

# Troubleshoot "QM\_SANITY\_WARNING" Message on 12000 series router

## Contents

---

### [Introduction](#)

### [Prerequisites](#)

[Requirements](#)

### [Components Used](#)

### [Conventions](#)

### [Problem](#)

### [Solution](#)

[Background](#)

[Scenario 1:](#)

[Example:](#)

[Scenario 2:](#)

[Scenario 3:](#)

### [Troubleshoot Commands](#)

[Example:](#)

---

## Introduction

This Document describes the procedure to debug packet buffer depletion messages that can happen on different Line cards in a 12000 series cisco router running IOS. It is far too common to see valuable time and resources wasted replacing the hardware that actually functions properly due to lack of knowledge on GSR buffer management.

## Prerequisites

### Requirements

The reader should have an overview of the [Cisco 12000 Series Router architecture](#).

## Components Used

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

- Cisco 12000 Series Internet Router
- Cisco IOS® Software Release that supports the Gigabit Switch Router

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

# Conventions

Refer to [Cisco Technical Tips Conventions](#) for more information on document conventions.

## Problem

GSR or 12000 series cisco routers have a truly distributed architecture. This means that each LC runs its own copy of Cisco IOS software image and have the intelligence to complete the packet forwarding decision on its own. Each line card does its own

1. Forwarding lookup
2. Packet buffer management
3. QOS
4. Flow control

One of the most important operation during packet switching in GSR is the buffer management which is done by various Buffer Management ASICs ( BMA) located in the Line cards. Below are some messages related to GSR buffer management which could show up in the router logs while in production. In the following section we will discuss the different triggers that could cause these messages to appear on the router logs and what are the corrective action to be done to mitigate the problem. In some situation this could also lead to packet loss which could manifest as protocol flaps and cause network impact.

```
%EE48-3-QM_SANITY_WARNING: ToFab FreeQ buffers depleted
```

```
SLOT 1:Sep 16 19:06:40.003 UTC: %EE48-3-QM_SANITY_WARNING: Few free buffers(1) are available in ToFab FreeQ pool# 2
```

```
SLOT 8:Sep 16 19:06:45.943 UTC: %EE48-3-QM_SANITY_WARNING: Few free buffers(0) are available in ToFab FreeQ pool# 1
```

```
SLOT 0:Sep 16 19:06:46.267 UTC: %EE48-3-QM_SANITY_WARNING: Few free buffers(2) are available in ToFab FreeQ pool# 2
```


```
SLOT 8:Sep 16 19:06:47.455 UTC: %EE48-3-QM_SANITY_WARNING: ToFab FreeQ buffers depleted. Recarving the ToFab buffers
```

```
SLOT 8:Sep 16 19:06:47.471 UTC: %EE192-3-BM QUIESCE:
```

## Solution

### Background

To troubleshoot the QM-SANITY warning errors we need to understand the packet flow on a GSR line card. The Figure below explains the main blocks of a C12k line card and the packet flow path.



GSR Line Card

The Line Card (LC) on a Cisco 12000 Series Internet Router has two types of memory:

- Route or processor memory (Dynamic RAM - DRAM): This memory enables mainly the onboard processor to run Cisco IOS software and store network routing tables (Forwarding Information Base - FIB, adjacency)
- Packet memory (Synchronous Dynamic RAM - SDRAM): Line card packet memory temporarily stores data packets awaiting switching decisions by the line card processor.

As seen from the above image, GSR line card has specialised packet buffer ASIC( Application Specific Integrated Circuit), one in each direction of traffic flow which provides access to the packet memory. These ASICs also known as Buffer management ASIC ( BMA) does the packet buffering and buffer queue management function on the line card. To support high throughput /forwarding rates, the packet memory on either direction is carved into different size memory pools designed to forward packets of varying MTU sizes.

The frames received by the Physical Layer Interface Module (PLIM) cards are Layer 2 processed and DMAed to a local memory in the PLIM card. Once the received data unit is complete, an ASIC in the PLIM contacts the ingress BMA and requests a buffer of appropriate size. If the buffer is granted, the packet moves to the line card ingress packet memory. If there are no available buffers the packet is dropped and ignored interface counter will go up. The ingress packet processor does the features processing on the packet, makes the forwarding decision and moves the packet to the toFab queue corresponding to the egress line card. The Fabric Interface ASIC( FIA) segments the packet to cisco cells and the cella are transmitted to the switch fabric. The packets are then received from the switch fabric by the FIA on the egress line card and goes on to the Frfab queues where they are reassembled, then to the egress PLIM, and finally sent on the wire.

The Decision of the FrFab BMA to select the buffer from a particular buffer pool is based on the decision made by the ingress line card switching engine. Since all queues on the entire box are of the same size and in the same order, the switching engine tells the transmitting LC to put the packet in the same number queue from which it entered the router.

While the packet is being switched, queue size of a particular buffer pool at the ingress line card which was

used to move the packet will be decremented by one till the BMA in the egress line card returns the buffer. Here we should also note that the complete buffer management is done in hardware by the buffer management ASIC's and for flaw less operation it is necessary that the BMA's returns the buffers to the original pool from where it was sourced.

There are three scenarios where the GSR packet buffer management can experience stress or failure leading to packet loss. Below are the three Scenarios.

### **Scenario 1:**

The hardware queue management fails. This happens when the egress BMA fails to return the packet buffer or returns the packet buffer to incorrect buffer pool. If the buffers are returned to the incorrect pool, we will see some buffer pools growing and some buffer pools depleting over a period of time and eventually effecting packets with the depleting buffer pool size. We will start seeing the QM-Sanity warnings as the packet buffer depletes and crosses the warning threshold.

Use the QM sanity debugs and show controllers tofab queues command to check if you are impacted by this condition. Refer to the troubleshooting section to find how to enable QM sanity thresholds.

This condition is generally caused by faulty hardware. Check the below outputs on the router and looks for parity errors or line card crashes. The fix would be to replace the Line card.

```
show controllers fia
```

```
show context all
```

```
show log
```

### **Example:**

From the QM sanity debugs and show controller tofab queue we can see the Pool 2 is growing in size while Pool 4 is running low. This indicates Pool 4 is losing buffers and it is being returned to Pool 2.

#### QM sanity debugs:

```
<#root>
```

```
SLOT 5:Oct 25 04:41:03.286 UTC: Pool 1: Carve Size 102001: Current Size 73078
```

```
SLOT 5:Oct 25 04:41:03.286 UTC: Pool 2: Carve Size
```

```
78462
```

```
: Current Size
```

```
181569
```

```
SLOT 5:Oct 25 04:41:03.286 UTC: Pool 3: Carve Size 57539: Current Size 6160
```

```
SLOT 5:Oct 25 04:41:03.286 UTC: Pool 4: Carve Size 22870: Current Size 67
```

```
SLOT 5:Oct 25 04:41:03.286 UTC: IPC FreeQ: Carve Size 600: Current Size 600
```

show controllers tofab queues:

<#root>

<snip>

Qnum Head Tail #Qelem LenThresh

-----

4 non-IPC free queues:

102001/102001 (buffers specified/carved), 39.1%, 80 byte data size

1 13542 13448 73078 262143

78462/

**78462**

(buffers specified/carved), 30.0%, 608 byte data size

2 131784 131833

**181569**

262143

57539/57539 (buffers specified/carved), 22.0%, 1616 byte data size

3 184620 182591 6160 262143

23538/22870 (buffers specified/carved), 8.74%, 4592 byte data size

4 239113 238805 67 262143

<snip>

## Scenario 2:

Traffic congestion on the next hop device or the forward path. In this scenario the device to which the GSR feeds traffic cannot process at GSR's speed and as a result the next hop device is sending pause frames towards GSR asking it to slow down. If flow control is enabled on GSR PLIM cards, the router will honour the pause frames and will start buffering the packets. Eventually the router will run out of buffers causing the QM Sanity error messages and packet drops. We will start seeing the QM-Sanity warnings as the packet buffer depletes and crosses the warning threshold. Refer to the troubleshooting section on how to find the QM sanity thresholds.

Use the **show interface output** on the egress interface to check if the router is impacted by this Scenario. The below capture gives an example of an interface receiving pause frames. The action plan will be to look

at the cause of congestion in the next hop device.

GigabitEthernet6/2 is up, line protocol is up

Small Factor Pluggable Optics okay

Hardware is GigMac 4 Port GigabitEthernet, address is 000b.455d.ee02 (bia 000b.455d.ee02)

Description: Cisco Sydney Lab

Internet address is 219.158.33.86/30

MTU 1500 bytes, BW 500000 Kbit, DLY 10 usec, rely 255/255, load 154/255

Encapsulation ARPA, loopback not set

Keepalive set (10 sec)

Full Duplex, 1000Mbps, link type is force-up, media type is LX

**output flow-control is on, input flow-control is on**

ARP type: ARPA, ARP Timeout 04:00:00

Last input 00:00:02, output 00:00:02, output hang never

Last clearing of "show interface" counters 7w1d

Queueing strategy: random early detection (WRED)

Output queue 0/40, 22713601 drops; input queue 0/75, 736369 drops

Available Bandwidth 224992 kilobits/sec

30 second input rate 309068000 bits/sec, 49414 packets/sec

30 second output rate 303400000 bits/sec, 73826 packets/sec

143009959974 packets input, 88976134206186 bytes, 0 no buffer

Received 7352 broadcasts, 0 runts, 0 giants, 0 throttles

0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored

0 watchdog, 7352 multicast, **45 pause input**

234821393504 packets output, 119276570730993 bytes, 0 underruns

Transmitted 73201 broadcasts

0 output errors, 0 collisions, 0 interface resets

0 babbles, 0 late collision, 0 deferred

0 lost carrier, 0 no carrier, 0 pause output

0 output buffer failures, 0 output buffers swapped out

### Scenario 3:

At times of oversubscription due to poor network design/traffic bursts/DOS attack. QM Sanity warning can occur if there is sustained high traffic condition where more traffic is directed at the router than what the Line Cards can handle.

To rootcause this check the traffic rates on all the interfaces in the router. That will reveal if any of the high speed links are congesting slow links.

Use the **show interface output** command.

## Troubleshoot Commands

To Check the current QM sanity level for a LC

- Attach to LC
- Go to enable mode
- Run **test fab** command
- Collect the output of “qm\_sanity\_info”
- Option **q** to exit **test fab** command line
- Exit from LC

To configure QM Sanity parameters

- change to configuration mode
- Run **hw-module slot <slot#> qm-sanity tofab warning freq <>**

To enable/disable QM sanity debugs

- Attach to LC
- Go to enable mode
- Run **test fab** command
- Run “qm\_sanity\_debug”. Run again and it will stop the debugs
- Option **q** to exit **test fab** command line
- Exit from LC

To check the GSR fabric interface asic statistics

- show controller fia

To check the Tofab queues

- show controllers tofab queues

To check Frfab queus

- show controller frfab queues

### Example:

The below output is pulled from a working Lab router to demonstrate the command outputs.

```
GSR-1-PE-5#show controller fia
```

```
Fabric configuration: 10Gbps bandwidth (2.4Gbps available), redundant fabric
```

```
Master Scheduler: Slot 17      Backup Scheduler: Slot 16
```

```
Fab epoch no 0          Halt count 0
```

From Fabric FIA Errors

-----

```
redund overflow 0          cell drops 0
```

```
cell parity      0
```

```
Switch cards present 0x001F  Slots 16 17 18 19 20
```

```
Switch cards monitored 0x001F  Slots 16 17 18 19 20
```

```
Slot:      16      17      18      19      20
```

```
Name:      csc0      csc1      sfc0      sfc1      sfc2
```

-----

```
los      0      0      0      0      0
```

```
state Off      Off      Off      Off      Off
```

```
crc16 0      0      0      0      0
```

To Fabric FIA Errors

-----

```
sca not pres 0          req error      0          uni fifo overflow 0
```

```
grant parity 0          multi req      0          uni fifo undrflow 0
```

```
cntrl parity 0          uni req        0
```

```
multi fifo  0          empty dst req 0          handshake error  0
```



cell parity 0

GSR-1-PE-5#attach 1

Entering Console for Modular SPA Interface Card in Slot: 1

Type "exit" to end this session

Press RETURN to get started!

LC-Slot1>en

LC-Slot1#test fab

BFLC diagnostic console program

BFLC (? for help) [?]: qm\_sanity\_debug

QM Sanity Debug enabled

BFLC (? for help) [qm\_sanity\_debug]:

SLOT 1:02:54:33: ToFAB BMA information

SLOT 1:02:54:33: Number of FreeQs carved 4

SLOT 1:02:54:33: Pool 1: Carve Size 102001: Current Size 102001

SLOT 1:02:54:33: Pool 2: Carve Size 78462: Current Size 78462

SLOT 1:02:54:33: Pool 3: Carve Size 57539: Current Size 57539

SLOT 1:02:54:33: Pool 4: Carve Size 22870: Current Size 22870

SLOT 1:02:54:33: IPC FreeQ: Carve Size 600: Current Size 600

SLOT 1:02:54:33: Number of LOQs enabled 768

SLOT 1:02:54:33: Number of LOQs disabled 1280

SLOT 1:02:54:33: ToFAB BMA information

SLOT 1:02:54:33: Number of FreeQs carved 4

SLOT 1:02:54:33: Pool 1: Carve Size 102001: Current Size 102001  
SLOT 1:02:54:33: Pool 2: Carve Size 78462: Current Size 78462  
SLOT 1:02:54:33: Pool 3: Carve Size 57539: Current Size 57539  
SLOT 1:02:54:33: Pool 4: Carve Size 22870: Current Size 22870  
SLOT 1:02:54:33: IPC FreeQ: Carve Size 600: Current Size 600  
SLOT 1:02:54:33: Number of LOQs enabled 768  
SLOT 1:02:54:33: Number of LOQs disabled 1280  
QM Sanity Debug disabled

BFLC (? for help) [qm\_sanity\_debug]: qm\_sanity\_info

ToFab QM Sanity level Warning

FrFab QM Sanity level None

Sanity Check is triggered every 20 seconds

Min. buffers threshold in percentage 5

BFLC (? for help) [qm\_sanity\_info]: q

LC-Slot1#exi

Disconnecting from slot 1.

Connection Duration: 00:01:09

GSR-1-PE-5#config t

Enter configuration commands, one per line. End with CNTL/Z.

GSR-1-PE-5(config)#hw-module slot 1 qm-sanity tofab warning freq 10

GSR-1-PE-5(config)#end

GSR-1-PE-5#attach 1

02:57:25: %SYS-5-CONFIG\_I: Configured from console by console

GSR-1-PE-5#attach 1

Entering Console for Modular SPA Interface Card in Slot: 1

Type "exit" to end this session

Press RETURN to get started!

LC-Slot1>en

LC-Slot1#test fab

BFLC diagnostic console program

BFLC (? for help) [?]: qm\_sanity\_info

ToFab QM Sanity level Warning

FrFab QM Sanity level None

Sanity Check is triggered every 10 seconds

Min. buffers threshold in percentage 5

BFLC (? for help) [qm\_sanity\_info]: q

LC-Slot1#exit

Disconnecting from slot 1.

Connection Duration: 00:00:27

GSR-1-PE-5#execute-on all show controllers tofab queues

===== Line Card (Slot 0) =====

Carve information for ToFab buffers

SDRAM size: 268435456 bytes, address: E0000000, carve base: E0018000

268337152 bytes carve size, 4 SDRAM bank(s), 16384 bytes SDRAM pagesize, 2 carve(s)

max buffer data size 4592 bytes, min buffer data size 80 bytes

262141/262141 buffers specified/carved

265028848/265028848 bytes sum buffer sizes specified/carved

Qnum Head Tail #Qe Lem LenThresh

-----

4 non-IPC free queues:

107232/107232 (buffers specified/carved), 40.90%, 80 byte data size

601 107832 107232 262143

73232/73232 (buffers specified/carved), 27.93%, 608 byte data size

107833 181064 73232 262143

57539/57539 (buffers specified/carved), 21.94%, 1616 byte data size

181065 238603 57539 262143

23538/23538 (buffers specified/carved), 8.97%, 4592 byte data size

238604 262141 23538 262143

IPC Queue:

600/600 (buffers specified/carved), 0.22%, 4112 byte data size

155 154 600 262143

Raw Queue (high priority):

0 0 0 65535

Raw Queue (medium priority):

0 0 0 32767

Raw Queue (low priority):

0 0 0 16383

ToFab Queues:

Dest Slot Queue# Head Tail Length Threshold

pkts pkts

=====

0 0 0 0 0 262143

15 2191(hpr) 0 0 0 0

Multicast 2048 0 0 0 262143

2049 0 0 0 262143

===== Line Card (Slot 1) =====

Carve information for ToFab buffers

SDRAM size: 268435456 bytes, address: 26000000, carve base: 26010000

268369920 bytes carve size, 4 SDRAM bank(s), 32768 bytes SDRAM pagesize, 2 carve(s)

max buffer data size 4592 bytes, min buffer data size 80 bytes

262140/261472 buffers specified/carved

267790176/264701344 bytes sum buffer sizes specified/carved

Qnum Head Tail #Qelem LenThresh

-----

4 non-IPC free queues:

102001/102001 (buffers specified/carved), 39.1%, 80 byte data size

1 601 102601 102001 262143

78462/78462 (buffers specified/carved), 30.0%, 608 byte data size

2 102602 181063 78462 262143

57539/57539 (buffers specified/carved), 22.0%, 1616 byte data size

3 181064 238602 57539 262143

23538/22870 (buffers specified/carved), 8.74%, 4592 byte data size

4 238603 261472 22870 262143

IPC Queue:

600/600 (buffers specified/carved), 0.22%, 4112 byte data size

30 85 84 600 262143

Raw Queue (high priority):

27 0 0 0 65368

Raw Queue (medium priority):

28 0 0 0 32684

Raw Queue (low priority):

31 0 0 0 16342

ToFab Queues:

Dest Slot Queue# Head Tail Length Threshold

pkts pkts

=====

:::::::::::::::::::::

Hi Priority

0 2176(hpr) 0 0 0

1 2177(hpr) 0 0 0

2 2178(hpr) 0 0 0

3 2179(hpr) 0 0 0

4 2180(hpr) 553 552 0

5 2181(hpr) 0 0 0

6 2182(hpr) 0 0 0

7 2183(hpr) 0 0 0

8 2184(hpr) 0 0 0

9 2185(hpr) 0 0 0

10 2186(hpr) 0 0 0

11 2187(hpr) 0 0 0

12 2188(hpr) 0 0 0

13 2189(hpr) 0 0 0

14 2190(hpr) 0 0 0

15 2191(hpr) 0 0 0

Multicast

2048 0 0 0

2049 0 0 0

2050 0 0 0

2051 0 0 0

2052 0 0 0

2053 0 0 0

2054 0 0 0

2055 0 0 0

==== Line Card (Slot 3) =====

Carve information for ToFab buffers

SDRAM size: 268435456 bytes, address: E0000000, carve base: E0018000

268337152 bytes carve size, 4 SDRAM bank(s), 16384 bytes SDRAM pagesize, 2 carve(s)

max buffer data size 4112 bytes, min buffer data size 80 bytes

262142/262142 buffers specified/carved

230886224/230886224 bytes sum buffer sizes specified/carved

Qnum Head Tail #Qelem LenThresh

-----

3 non-IPC free queues:

94155/94155 (buffers specified/carved), 35.91%, 80 byte data size

601 94755 94155 262143

57539/57539 (buffers specified/carved), 21.94%, 608 byte data size

94756 152294 57539 262143

109848/109848 (buffers specified/carved), 41.90%, 1616 byte data size

152295 262142 109848 262143

IPC Queue:

600/600 (buffers specified/carved), 0.22%, 4112 byte data size

207 206 600 262143

Raw Queue (high priority):

0 0 0 65535

Raw Queue (medium priority):

0 0 0 32767

Raw Queue (low priority):

0 0 0 16383

ToFab Queues:

Dest Slot Queue# Head Tail Length Threshold

pkts pkts

=====

0 0 0 0 0 262143

1 0 0 0 262143

2 0 0 0 262143

3 0 0 0 262143

:::::::::::::::::::::::::::::::

2049 0 0 0 262143

2050 0 0 0 262143

2051 0 0 0 262143

2052 0 0 0 262143

2053 0 0 0 262143

2054 0 0 0 262143

2055 0 0 0 262143



===== Line Card (Slot 2) =====

Carve information for FrFab buffers

SDRAM size: 268435456 bytes, address: D0000000, carve base: D241D100

230567680 bytes carve size, 4 SDRAM bank(s), 16384 bytes SDRAM pagesize, 2 carve(s)

max buffer data size 4592 bytes, min buffer data size 80 bytes

235926/235926 buffers specified/carved

226853664/226853664 bytes sum buffer sizes specified/carved

Qnum Head Tail #Qelem LenThresh

-----

4 non-IPC free queues:

96484/96484 (buffers specified/carved), 40.89%, 80 byte data size

11598 11597 96484 262143

77658/77658 (buffers specified/carved), 32.91%, 608 byte data size

103116 103115 77658 262143

40005/40005 (buffers specified/carved), 16.95%, 1616 byte data size

178588 178587 40005 262143

21179/21179 (buffers specified/carved), 8.97%, 4592 byte data size

214748 235926 21179 262143

IPC Queue:

600/600 (buffers specified/carved), 0.25%, 4112 byte data size

66 65 600 262143

Multicast Raw Queue:

0 0 0 58981

Multicast Replication Free Queue:

235930 262143 26214 262143

Raw Queue (high priority):

78 77 0 235927

Raw Queue (medium priority):

11596 11595 0 58981

Raw Queue (low priority):

0 0 0 23592

Interface Queues:

Interface Queue# Head Tail Length Threshold

pkts pkts

=====

0 0 103107 103106 0 32768

3 178588 178587 0 32768

1 4 103110 103109 0 32768

7 11586 11585 0 32768

2 8 0 0 0 32768

11 0 0 0 32768

3 12 0 0 0 32768

15 0 0 0 32768

GSR-1-PE-5#