

Debugging Packet LOSS and QMAN Enqueue Rejection Issues on DPAA

Platforms

This document will take an in-depth look at QorIQ Data Path Acceleration Architecture (DPAA) and how each component interacts with each other and the core. Discussion will focus on debugging common issues such as packet loss and queue manager enqueue rejections. This document guide users to find where missing packets are being dropped or lost in receive path and determine reason for Queue Manager enqueue rejection.

1. Debugging Packet Loss Issue

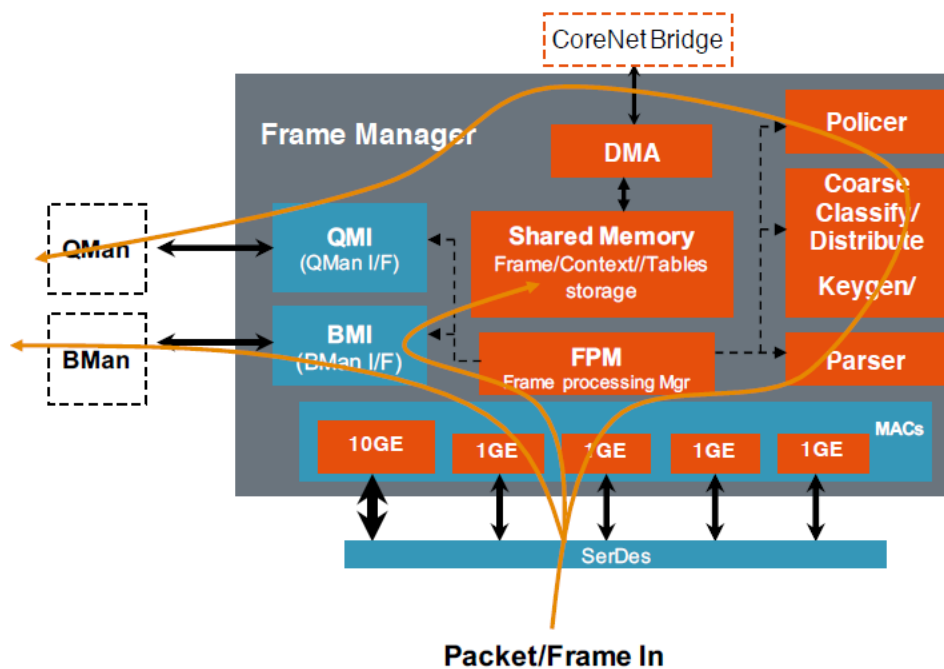
Packets sent to the SoC but we have lost track of where they ended up. Missing some or all packets sent to the SoC.

1.1 Frame Manager(FMan) Introduction

FMan supports a flexible pipeline of packet processing elements.

Frame data and per frame context is stored in internal memory while frame is processed Frame Processor Manager (FPM) “schedules” frames for processing by different elements to create appropriate pipeline.

Default pipeline configured for each port



1.2 Frame Manager Buffer Manager Interface (BMI) Rx Port Statistics Counters

FMBM_RFRC – Receive Frame Counter

– Total number of frames received on the Rx port

- FMBM_RBFC – Bad Frames Counter
 - Number of frames received on the Rx port with an error indication.
 - Error cause could be FCS error, MAC FIFO overflow, code error, etc.
 - These frames are discarded and not shown to receive queues, unless FMBM_RCFG[FDOVR] is set.
 - Rx FD Status Field “FPE” (frame physical error) bit will be set.
- FMBM_RFFC – Filter Frames Counter
 - Number of frames received on the Rx port that were filtered out by the parse and classify modules of the Fman.
 - See Rx FD Status Field.
 - These frames are discarded and not shown to receive queues, unless FMBM_RCFG[FDOVR] is set.
- FMBM_RFDC – Frames Discard Counter
 - Number of frames received on the Rx port that were not able to enter the receive queue system due to WRED algorithm. Other reasons for enqueue reject may be tail drop, out of service FQ, etc.
- FMBM_RODC – Out of Buffers Discard Counter
 - Number of received frames that were discarded due to lack of external buffers.
- FMBM_RDBC – Buffer Deallocate Counter
 - Number of buffer deallocate operations. The counter increments whenever a buffer is returned to BMan pools.
- FMBM_RLFC – Large Frames Counter
 - Over size indication is marked when frame size exceeds the maximum configured in the corresponding MAC configuration register.

1.3 Linux Sysfs Support for Fman Rx Port Statistics

Example: FMan1 dTSEC5 on LS1046ARDB

```
$ ls /sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/*
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_dealloc_buf
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_discard_frame
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_enq_total
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_frame
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_rx_bad_frame
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_rx_filter_frame
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_rx_large_frame
/sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/port_rx_out_of_buffers_discard

$ cat /sys/devices/platform/soc/1a00000.fman/1a8c000.port/statistics/*
fm0-port-rx4 counter: 0
fm0-port-rx4 counter: 0
fm0-port-rx4 counter: 0
fm0-port-rx4 counter: 0
fm0-port-rx4 counter: 0
fm0-port-rx4 counter: 0
```

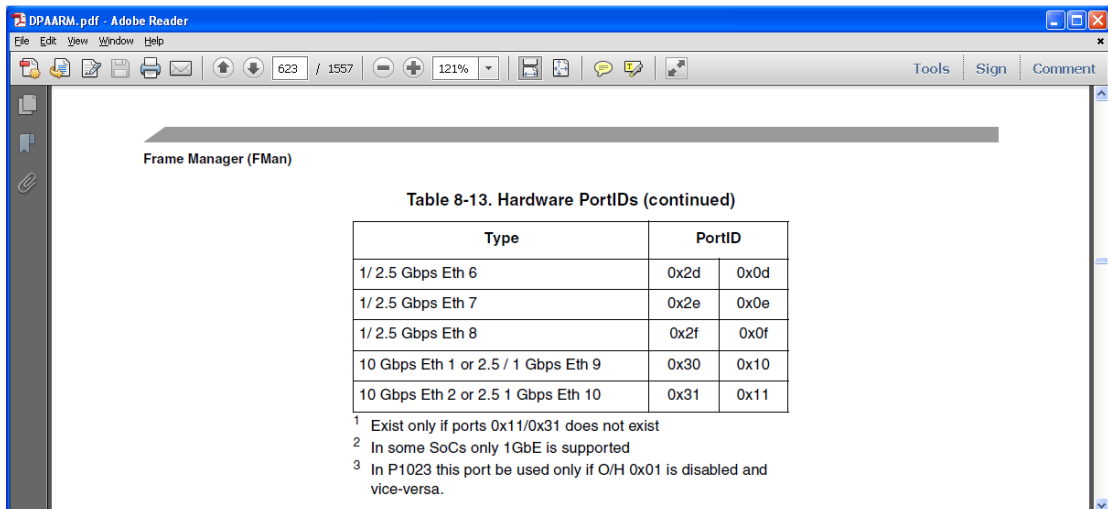
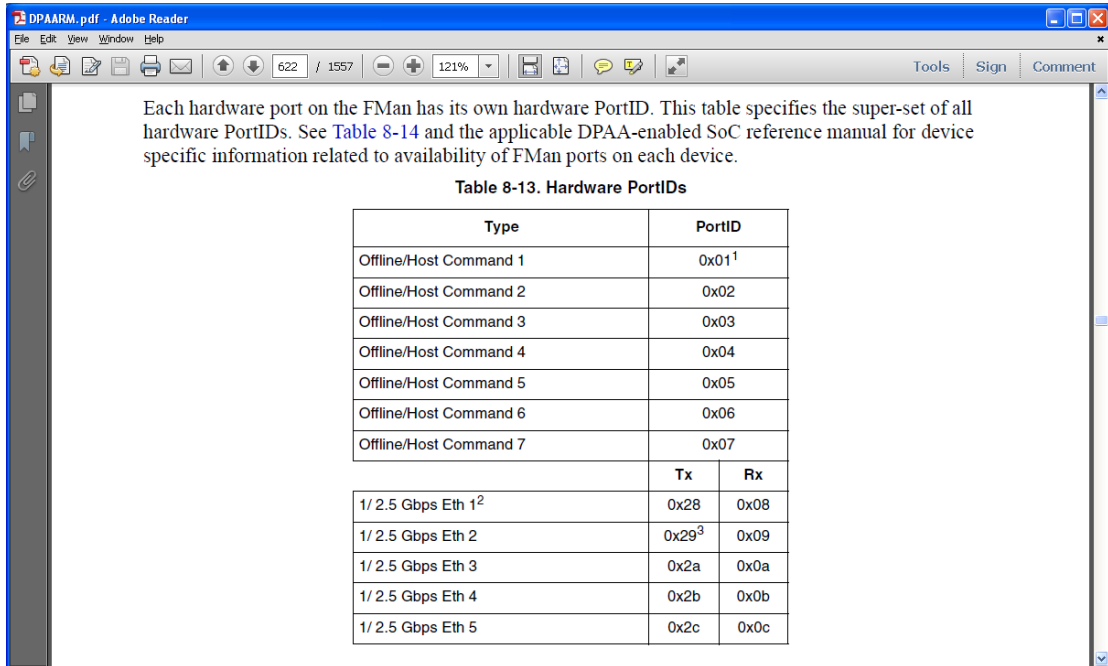
fm0-port-rx4 counter: 0

fm0-port-rx4 counter: 0

FMan1 Base Address + HW Ports Base Address + HW Port ID

1a00000 + 0x8_0000 + 0x0_C000 = 0x1a8c000

Frame Manager Hardware PortIDs



Example: FMAN Rx Port Statistics – Discarded Frames

```
# cat /sys/devices/platform/soc/1a00000.fman/1a88000.port/statistics/*
```

```
FM 0 Port 0 counter: 71 port_dealloc_buf
```

```
FM 0 Port 0 counter: 71 port_discard_frame
```

```
FM 0 Port 0 counter: 495255 port_enq_total
```

```
FM 0 Port 0 counter: 495255 port_frame
```

```
FM 0 Port 0 counter: 0 port_rx_bad_frame
```

FM 0 Port 0 counter: 0 port_rx_filter_frame
FM 0 Port 0 counter: 0 port_rx_large_frame
FM 0 Port 0 counter: 0 port_rx_out_of_buffers_discard

FMBM_RFDC – Frames Discard Counter

– Number of frames received on the Rx port that were not able to enter the receive queue system due to WRED algorithm. Other reasons for enqueue reject may be tail drop, out of service FQ, etc.

Example: FMAN Rx Port Statistics – Out of Buffers

```
cat /sys/devices/platform/soc/1a00000.fman/1a90000.port/statistics/*  
fm1-port-rx6 counter: 0  
fm1-port-rx6 counter: 0  
fm1-port-rx6 counter: 71836626 port_enq_total  
fm1-port-rx6 counter: 71862438 port_frame  
fm1-port-rx6 counter: 0  
fm1-port-rx6 counter: 0  
fm1-port-rx6 counter: 0  
fm1-port-rx6 counter: 25628 port_rx_out_of_buffers_discard  
Running out of buffers points to congestion in the system
```

Example: FMAN Rx Port Statistics – No Dropped Packets

```
# cat /sys/devices/platform/soc/1a00000.fman/1a89000.port/statistics/*  
FM 0 Port 1 counter: 0  
FM 0 Port 1 counter: 0  
FM 0 Port 1 counter: 714541 port_enq_total  
FM 0 Port 1 counter: 714541 port_frame  
FM 0 Port 1 counter: 0  
FM 0 Port 1 counter: 0  
FM 0 Port 1 counter: 0  
FM 0 Port 1 counter: 0
```

All packets received on this Rx Port were enqueued to Queue Manager.

After checking FMan Port statistics, if packets are still unaccounted for, then check MAC counter registers.

DPAA Ethernet Driver exports a series of information in Sysfs such as buffer pool IDs and frame queue IDs used by the interface and MAC registers.

```
root@localhost:~# ls /sys/devices/platform/soc/soc:fsl,dpaa/soc:fsl,dpaa:ethernet@8/net/fm1-  
mac9  
addr_assign_type  carrier_down_count  device_type          ifindex  
name_assign_type  proto_down          type  
addr_len          carrier_up_count    dormant              iflink              netdev_group  
queues           uevent
```

```

address          consumers          duplex            link_mode        operstate
speed
bpids          dev_id            flags             mac_regs         phys_port_id
statistics
broadcast        dev_port          fqids            mac_rx_stats
phys_port_name  subsystem
carrier          device            gro_flush_timeout mac_tx_stats     phys_switch_id
suppliers
carrier_changes  device_addr       ifalias           mtu               power
tx_queue_len

```

```

/sys/devices/platform/soc/soc:fsl,dpaa/soc:fsl,dpaa:ethernet@8/net/fm1-mac9/statistics# ls
collisions      rx_crc_errors    rx_frame_errors  rx_over_errors   tx_carrier_errors
tx_fifo_errors
multicast       rx_dropped       rx_length_errors rx_packets        tx_compressed
tx_heartbeat_error rs
rx_bytes        rx_errors        rx_missed_errors tx_aborted_errors tx_dropped
tx_packets
rx_compressed  rx_fifo_errors  rx_nohandler     tx_bytes          tx_errors
tx_window_errors

```

Please refer to the following to use Sysfs to check FQIDs and BPIDs.

```

/sys/devices/platform/soc/soc:fsl,dpaa/soc:fsl,dpaa:ethernet@8/net/fm1-mac9# cat fqids
Rx error: 403
Rx default: 404
Rx PCD: 15360 - 15487
Rx PCD High Priority: 80896 - 81023
Tx confirmation (mq): 405 - 420
Tx error: 421
Tx default confirmation: 422
Tx: 423 – 438

```

```

/sys/devices/platform/soc/soc:fsl,dpaa/soc:fsl,dpaa:ethernet@8/net/fm1-mac9# cat bpids
32

```

2. Queue Manager(Qman) Enqueue Rejections

2.1 Reasons for an Enqueue Rejection

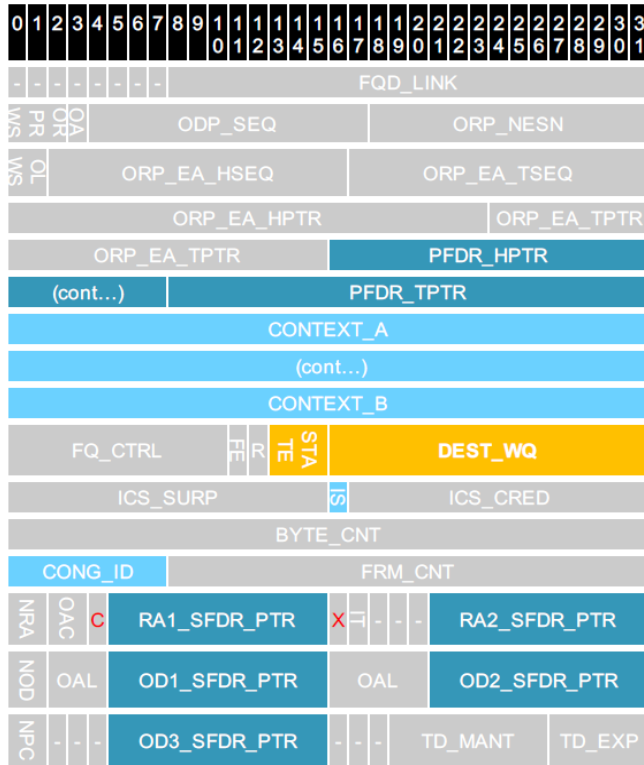
- Congestion Group tail drop threshold exceeded
- WRED congestion avoidance

- FQ tail drop threshold exceeded
- Error condition, indicated in QMAN_ERR_ISR
- Order Restoration enabled and frame arrived in late or early rejection window. Plus a couple others related to order restoration

2.2 Frame Queue Descriptor

FQD Selected Field Description:

- FQD_LINK: Link to the next FQD in a queue of FQDs, used for Work Queues
- ORPRWS: ORP Restoration Window Size
- OA: ORP Auto Advance NESN Window Size
- ODP_SEQ: ODP Sequence Number
- ORP_NESN: ORP Next Expected Sequence Number.
- ORP_EA_HPTR, ORP_EA_TPTR: ORP Early Arrival Head and Tail Pointer
- PFDR_HPTR, PFDR_TPTR : PFDR Head and Tail Pointer
- CONTEXT_A, CONTEXT_B: Frame Queue Context A and B
- STATE: FQ State
- DEST_WQ: Destination Work Queue
- ICS_SURP: Intra-Class Scheduling Surplus or Deficit.
- IS: Intra-Class Scheduling Surplus or Deficit identifier
- ICS_CRED: Intra-Class Scheduling Credit
- CONG_ID: Congestion Group ID
- RA[1-2]_SFDR_PTR: SFDR Pointer for Recently Arrived frame # 1 and 2
- TD_MANT, TD_EXP : Tail Drop threshold Exponent and Mantissa
- C: FQD in external memory or in cache (Qman 1.1)
 - X: XON or XOFF for flow control command (Qman1.1)



2.3 Qman Debugfs

```

root@localhost:/sys/kernel/debug/qman#
root@localhost:/sys/kernel/debug/qman# echo 1 > query_fq_fields
root@localhost:/sys/kernel/debug/qman# cat query_fq_fields
Query FQ Programmable Fields Result fqid 0x1
orprws: 0
oa: 0
olws: 0
cgid: 0
fq_ctrl: None
dest_channel: 0
dest_wq: 0
ics_cred: 0
td_mant: 0
td_exp: 0
ctx_b: 0
ctx_a: 0x0
ctx_a_stash_exclusive: None
ctx_a_stash_annotation_cl: 0
ctx_a_stash_data_cl: 0
ctx_a_stash_context_cl: 0

```

In the following example, check DPAARM Work Queue Channel Assignments
 -Dedicated Channel serviced by software portal 49

-Work Queue number 3

```
root@localhost:/sys/kernel/debug/qman # echo 0x3c00 > query_fq_fields
```

```
root@localhost:/sys/kernel/debug/qman # cat query_fq_fields
```

Query FQ Programmable Fields Result fqid 0x3c00

orprws: 0

oa: 0

olws: 0

cgid: 3

fq_ctrl:

Prefer in cache

Avoid Blocking

Context-A stashing

Congestion Group Enable

dest_channel: 49

dest_wq: 3

ics_cred: 0

td_mant: 0

td_exp: 0

ctx_b: 0xd5

ctx_a: 0x175d70018

ctx_a_stash_exclusive:

FQ Ctx Stash

Frame Annotation Stash

ctx_a_stash_annotation_cl: 1

ctx_a_stash_data_cl: 2

ctx_a_stash_context_cl: 2

FQ in tentatively scheduled state, Snap shot shows that FQ is empty

```
root@localhost:/sys/kernel/debug/qman # echo 0x3c00 > query_fq_np_fields
```

```
root@localhost:/sys/kernel/debug/qman # cat query_fq_np_fields
```

Query FQ Non Programmable Fields Result fqid 0x3c00

force eligible pending: no

retirement pending: no

state: Tentatively Scheduled

fq_link: 0x0

odp_seq: 0

orp_nesn: 0

orp_ea_hseq: 0

orp_ea_tseq: 0

orp_ea_hptr: 0x0

orp_ea_tptr: 0x0

pfdr_hptr: 0x0

pfdr_tptr: 0x0

is: ics_surp contains a surplus

ics_surp: 0
byte_cnt: 0
frm_cnt: 0
ra1_sfd: 0x800
ra2_sfd: 0x0
od1_sfd: 0x0
od2_sfd: 0x0
od3_sfd: 0x0

2.4 Buffer Manager (BMan) Debugfs

root@localhost:/sys/kernel/debug/bman# cat query_bp_state

bp_id	free_buffers_avail	bp_depleted
0	no	no
1	no	no
2	no	no
3	no	no
4	no	no
5	no	no
6	no	no
7	no	no
8	no	no
9	no	no
10	no	no
11	no	no
12	no	no
13	no	no
14	no	no
15	no	no
16	no	no
17	no	no
18	no	no
19	no	no
20	no	no
21	no	no
22	no	no
23	no	no
24	no	no
25	no	no
26	no	no
27	no	no
28	no	no
29	no	no
30	no	no
31	no	no

32	<i>yes</i>	<i>no</i>
33	<i>no</i>	<i>no</i>
34	<i>no</i>	<i>no</i>
35	<i>no</i>	<i>no</i>
36	<i>no</i>	<i>no</i>
37	<i>no</i>	<i>no</i>
38	<i>no</i>	<i>no</i>
39	<i>no</i>	<i>no</i>
40	<i>no</i>	<i>no</i>
41	<i>no</i>	<i>no</i>
42	<i>no</i>	<i>no</i>
43	<i>no</i>	<i>no</i>
44	<i>no</i>	<i>no</i>
45	<i>no</i>	<i>no</i>
46	<i>no</i>	<i>no</i>
47	<i>no</i>	<i>no</i>
48	<i>no</i>	<i>no</i>
49	<i>no</i>	<i>no</i>
50	<i>no</i>	<i>no</i>
51	<i>no</i>	<i>no</i>
52	<i>no</i>	<i>no</i>
53	<i>no</i>	<i>no</i>
54	<i>no</i>	<i>no</i>
55	<i>no</i>	<i>no</i>
56	<i>no</i>	<i>no</i>
57	<i>no</i>	<i>no</i>
58	<i>no</i>	<i>no</i>
59	<i>no</i>	<i>no</i>
60	<i>no</i>	<i>no</i>
61	<i>no</i>	<i>no</i>
62	<i>no</i>	<i>no</i>
63	<i>no</i>	<i>no</i>

Note that on earlier slide, sysfs showed that BPID 32 was being used