

Troubleshoot Nexus 7000: F3 Input Discards and LACP PDU drops

Contents

[Introduction](#)

[Prerequisites](#)

[Background information](#)

[Common Causes](#)

[Input discards](#)

[Loss of LACP PDU:](#)

[Troubleshooting](#)

[Solution](#)

[Known bugs](#)

[Reference](#)

Introduction

This document describes how to troubleshoot input discards on the port-channel on nexus 7000.

Prerequisites

Cisco recommends to have knowledge about following topics:

[Nexus 7000 Series switches](#)

[F series line cards](#)

[Link aggregation control protocol](#)

Background information

The F3 line card queues packets on ingress instead of egress and implements virtual output queues (VOQs) on all ingress interfaces, so that a congested egress port does not affect traffic directed to other egress ports. The extensive use of VOQs in the system helps ensure maximum throughput on a per-egress basis. Congestion on one egress port does not affect traffic destined for other egress interfaces, which avoids head-of-line blocking(HOLB) that otherwise causes congestion to spread.

In burst-optimized mode, we should see drops in PL if IB gets exhausted. In mesh-optimized mode, drops moves to VQ due to exceeded threshold. Mesh-optimized avoids HOLB drops.

VOQs also use the concept of credited and uncredited traffic. Unicast traffic is classified as credited traffic; broadcast, multicast, and unknown unicast traffic are classified as uncredited traffic. Uncredited traffic does not utilize VOQs, and traffic is queued on egress rather than

ingress. If an ingress port has no credit to send traffic to an egress port, the ingress port buffers until it gets credit. Since the ingress port buffers are not deep, input drops might occur.

Common Causes

Input discards

- The most common cause of input discards occurs when you have a Switched Port Analyzer (SPAN) with the destination port on an F2 linecard and with SPAN traffic that exceeds the line rate. Eventually the ingress port buffers the packets, which leads to input discards.

Note: Next-Gen I/O modules such as F2E, F3, and M3 are not susceptible to SPAN destination port oversubscription scenarios causing indiscards and HOLB on ingress ports. This is also noted in [Guidelines and Limitations for SPAN](#)

- Inappropriate design (such as 10G of input bandwidth and 1G of output bandwidth) triggers the F2 hardware limitation (HOL blocking).
- If traffic from multiple ports egresses out of same interface (1G to 1G or 10G to 10G interfaces), if you exceed the line rate, it might result in input discards on ingress ports.
- A VLAN mismatch may cause input discards. Use the **show interface trunk** command in order to verify that both switches forward the same VLAN.

Loss of LACP PDU:

A port-channel gets suspended when it does not receive any LACP PDUs from the neighbor. The line card queues packets on ingress instead of egress and an input discard indicates the number of packets dropped in the input queue because of congestion.

- Port Logic (PL) is a buffer before the decision engine and is after the front panel ports. Any congestion or flow control on Port Logic on the ingress would prevent or delay the LACP PDU from going any further causing the interface to be suspended. The VL is a high priority virtual lane. If there is a scenario where high priority VL 5 traffic is head-of-line blocking from a congested port, we will have a back pressuring in PL on VL 5 which can result in LACP PDU drop.

Troubleshooting

```
`show module`
```

Mod	Ports	Module-Type	Model	Status
5	0	Supervisor Module-2	N7K-SUP2E	active *
6	0	Supervisor Module-2	N7K-SUP2E	ha-standby
7	6	100 Gbps Ethernet Module	N7K-F306CK-25	ok
8	12	10/40 Gbps Ethernet Module	N7K-F312FQ-25	ok

In this example, input discards on port-channel 10 (7/1,7/2 and 7/5) and port-channel 20 (7/3,7/4 and 7,6) caused by congestion on the egress interface 8/6. These drops are caused by HOL blocking.

```
`show port-channel summary`
```

```
-----  
Group Port-      Type      Protocol  Member Ports  
Channel  
-----  
<snip>  
10    Po10(RU)    Eth       LACP      Eth7/1(P)  Eth7/2(P)  Eth7/5(P)  
20    Po20(RU)    Eth       LACP      Eth7/3(P)  Eth7/4(P)  Eth7/6(P)
```

```
switch# show interface counter errors
```

```
-----  
Port          InDiscards  
-----  
<snip>  
Eth7/1        253323164  
Eth7/2        253682395  
Eth7/3        66785160      >>>> input discards on interfaces 7/1-6 are incrementing  
continuously. These interfaces belong to Po10 and Po20 which eventually goes into suspended  
state with reason "no LACP PDUs received"  
Eth7/4        64770521  
Eth7/5        258650104  
Eth7/6        66533418  
<snip>  
Eth8/6        0  
<snip>  
Po10          765655663  
Po20          198089099
```

To determine the congested port:

On the VQI, non-zero counters were on the move constantly. On congested ports, the counters usually stay high most of the time

```
switch# attach mod 7  
Attaching to module 7 ...  
To exit type 'exit', to abort type '$.'
```

```
module-7# show hardware internal qengine voq-status | ex "0 0 0 0 0 0 0 0 0 0 0 0"
```

```
-----  
| VOQ Status for Queue Driver  
| ports 1-48  
-----  
VQI:CCOS INST0 INST1 INST2 INST3 INST4 INST5  
-----  
0:0      0    0    0    0    0    0  
0:1      0    0    0    0    0    0  
145:6    0    0    0    0    0    0  
145:7    0    0    0    0    0    0  
146:0    0    0    0    0    0    0  
146:1    14d  130  533  79b  258  447  
146:2     5   44    7   12   1a    2  
146:3   2325 2277 1ae8 1a39 27bc 1902  
146:4     0    0    0    0    0    0  
146:5     0    0    0    0    0    0  
146:6     0    0    0    0    0    0  
146:7     0    0    0    0    0    0  
147:0     0    0    0    0    0    0  
147:1     0    0    0    0    0    0
```

```
147:2    0    0    0    0    0    0
147:3    0    0    0    0    0    0
```

The VQI is 146

VQI === 146 has a non-zero counter and keeps incrementing

Convert to Hex:

```
switch# hex 146
0x92
```

```
switch# show system internal ethpm info module | egrep -i vqi
LTL(0x36), VQI(0x42), LDI(0), IOD(0x14c)
LTL(0x37), VQI(0x43), LDI(0x1), IOD(0x14d)
LTL(0x38), VQI(0x44), LDI(0x2), IOD(0x14e)
LTL(0x39), VQI(0x45), LDI(0x3), IOD(0x14f)
<snip>
LTL(0x72), VQI(0x8a), LDI(0xc), IOD(0x62)
LTL(0x76), VQI(0x8e), LDI(0x10), IOD(0x63)
LTL(0x7a), VQI(0x92), LDI(0x14), IOD(0xe6) >>>>>> VQI 0x92 maps to LTL 0x7a
LTL(0x7e), VQI(0x96), LDI(0x18), IOD(0xe7)
LTL(0x82), VQI(0x9a), LDI(0x1c), IOD(0xe8)
LTL(0x86), VQI(0x9e), LDI(0x20), IOD(0xe9)
<snip>
```

Convert the LTL to physical interface using pixm mapping

PIXM Manages LTL and FPOE mapping to build the hardware forwarding path through the switch

```
switch# show system internal pixm info ltl 0x7a
Member info
-----
Type          LTL
-----
PHY_PORT      Eth8/6          >>>> congested egress interface.
```

To determine if LACP PDU are dropped

LACP PDU is a high priority traffic and hence should not expect LACP PDU to be dropped and the port-channel to go down because of input discards unless there is high priority **VL 5** traffic is head-of-line-blocking from the congested port.

In order to confirm if high priority VL 5 traffic is getting dropped, run the command “**show hardware queuing drops ingress**” and this would show PL drops for VL 5 on the affected interface

```
switch# show hardware queuing drops ingress
slot 7
=====
Device: Flanker Queue
PL drops:
SOURCE INTERFACE          VL          COUNT
-----
Eth7/1                    5           24437734
Eth7/2                    5           24289997
Eth7/3                    5           24449567
Eth7/4                    5           26084373
```

