# Contents



# Introduction

This document describes the causes of and solutions for input discards for the Cisco Nexus 7000 F248 Series (F2/F2e) line card. An input discard indicates the number of packets dropped in the input queue because of congestion. This number includes drops that are caused by tail drop and weighted random early detection (WRED).

# Background Information

The F2 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 (HOL) blocking that otherwise causes congestion to spread.

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

These are common causes of 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.

# Command Solutions

This section provides information you can use to troubleshoot your configuration.

> **Notes**: Use the Command Lookup Tool (registered customers only) in order to obtain more information on the commands used in this section.
> The Output Interpreter Tool (registered customers only) supports certain **show** commands. Use the Output Interpreter Tool in order to view an analysis of **show** command output.

In these examples, Ethernet 2/1 (Eth2/1) has a host connected that receives two 1Gbps streams. Eth2/1 runs at 1G. The two streams ingress on Eth2/5 and Eth2/9.

## show interface <ingress interface>

Use this command in order to check the speed of the interfaces. If the ingress interface runs at 10Gbps and the egress interface runs at 1Gbps, the drops are likely caused by HOL blocking.

```
N7K1# show int eth2/5
Ethernet2/5 is up
admin state is up, Dedicated Interface
------
full-duplex, 1000 Mb/s
------
30 seconds input rate 588237960 bits/sec, 73524 packets/sec
30 seconds output rate 216 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
input rate 588.56 Mbps, 73.52 Kpps; output rate 156.11 Mbps, 19.45 Kpps
RX
221333142 unicast packets 0 multicast packets 0 broadcast packets
221333128 input packets 221333169400 bytes
0 jumbo packets 0 storm suppression packets
0 runts 0 giants 0 CRC 0 no buffer
0 input error 0 short frame 0 overrun 0 underrun 0 ignored
0 watchdog 0 bad etype drop 0 bad proto drop 0 if down drop
0 input with dribble 11590977 input discard <-------
0 Rx pause
```

## show hardware internal statistics module <x> pktflow dropped

Run this command several times in order to determine whether the value of congestion_drop_bytes increments; x is the module number of the ingress port.

## attach module <x> and show hardware internal qengine

Run these commands several times in order to identify the virtual queue index (VQI) number:

**attach module <x>**

**module-x# show hardware internal qengine voq-status | ex "0 0 0 0 0 0 0 0 0 0 0"**
or
**module-x# show hardware internal qengine inst 2 voq-status non-empty**

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

```
N7K1# attach module 2
Attaching to module 2 ...
To exit type 'exit', to abort type '$.'

module-2# show hardware internal qengine inst 2 voq-status non-empty
VQI:CCOS BYTE_CNT PKT_CNT TAIL HEAD THR
-------- -------- ------- ---- ---- ---
0036:3 6154 3077 6804 14168 1 <------- VQI is 36 here

module-2# show hardware internal qengine voq-status | ex "0 0 0 0 0 0 0 0 0 0 0"
VQI:CCOS CLP0 CLP1 CLP2 CLP3 CLP4 CLP5 CLP6 CLP7 CLP8 CLP9 CLPA CLPB
-------- ---- ---- ---- ---- ---- ---- ---- ---- ---- ---- ---- ----
0036:3 0 0 0 0 0 0 0 1 0 0 0 0
VQI === 36
```

Once you have the VQI number, use the **show hardware internal qengine vqi-map** command in order to look for the VQI map table. Review the slot number and low-speed data interface (LDI) number in order to determine the egress interface. (The slot is also known as the module, and the LDI is also known as the port.) The module is zero-based and a mapping function can be used to determine the LDI.

```
module-2# show hardware internal qengine vqi-map
VQI   SUP  SLOT  LDI  EQI  FPOE  NUM  XBAR   IN   ASIC  ASIC  SV  FEA_
NUM   VQI  NUM   NUM  NUM  BASE  DLS  MASK   ORD  TYPE  IDX   ID  TURE
----  ---  ----  ---  ---  ----  ---  -----  ---  ----  ----  --  ----
--snip
36     no   1     0    0    8     1    0x155  0    CLP   0     0   0x81
--snip
```

LDI to Physical Port Mapping:

| LDI | Port |
|---|---|
| 0 | 2 |
| 1 | 1 |
| 2 | 3 |
| 3 | 4 |
| 4 | 6 |
| 5 | 5 |
| 6 | 7 |
| 7 | 8 |
| 8 | 10 |
| 9 | 9 |
| 10 | 11 |
| 11 | 12 |
| 12 | 14 |
| 13 | 13 |
| 14 | 15 |

15 16
16 18
17 17
18 19
19 20
20 22
21 21
22 23
23 24
24 26
25 25
26 27
27 28
28 30
29 29
30 31
31 32
32 34
33 33
34 35
35 36
36 38
37 37
38 39
39 40
40 42
41 41
42 43
43 44
44 46
45 45
46 47
47 48

**Physical Port = Eth 2/2**

Validate VQI and LDI via **show system internal ethpm info interface Eth2/2 | include VQI**

The congested port from the test description was 2/1 but the VQI listed is e2/2. The reason for the discrepenacy is that the egress buffers are shared by a port group which is a group of 4 ports for a F2/F2e module. Ports 1-4, 5-8 and so on are part of each port groups. If any single port in the port group gets congested in the egress direction then it can cause back pressure on the ingress port resulting in input discards.

## Additional Commands

If you continue to notice input discards, run these commands several times:

- **show interface | in Mbps|Ethernet**
- **show hardware internal statistics pktflow dropped**
- **show hardware internal statistics pktflow dropped congestion**
- **show hard internal statistics pktflow all**

- **show hardware internal error**
- **show hardware internal statistics device qengine**
- **show hard internal mac port 38 qos config**
- **show hard internal statis device mac all port 38**
- **attach module 1**
- **show hardware internal qengine voq-status**
- **show hardware internal qengine vqi-map**

# Information for F2e

On an F2e, there is a hardware internal error counter that points to the VQI of the first port in the port group/asic with the congested egress interface.

## show hardware internal errors mod <x>

Use this command in order to check the module for the number of times congestion is detected.

```
N7K2# show hardware internal errors mod 1

|-----------------------------------------------------------------------|
| Device:Clipper XBAR Role:QUE Mod: 1 |
| Last cleared @ Wed Jul 10 14:51:56 2013
| Device Statistics Category :: CONGESTION
|-----------------------------------------------------------------------|
Instance:1
ID Name Value Ports
-- ---- ----- -----
16227 Num of times congestion detected on VQI 48 0000000000001296 5-8 -


Instance:2
ID Name Value Ports
-- ---- ----- -----
16227 Num of times congestion detected on VQI 48 0000000000000590 9-12 -


Instance:3
ID Name Value Ports
-- ---- ----- -----
16227 Num of times congestion detected on VQI 48 0000000000001213 13-16 -
```

## show hardware internal qengine vqi-map

Use this command in order to map the VQI to the physical interface. This example uses VQI 48 from the previous example. Review the slot number and LDI number in order to determine the egress interface. The module is zero-based and a mapping function can be used to determine the LDI.

```
module-1# show hardware internal qengine vqi-map
VQI SUP SLOT LDI EQI FPOE NUM XBAR IN ASIC ASIC SV FEA_
NUM VQI NUM NUM NUM BASE DLS MASK ORD TYPE IDX ID TURE
---- --- ---- --- --- ---- --- ----- --- ---- ---- -- ----
--snip
48 no 0 12 0 3 1 0x155 0 CLP 3 0 0x1
--snip
```

```
Module Number = SLOT NUM + 1 (zero-based)
Module Number = 0 + 1 = 1
```

**Physical Port = Eth 1/14 (check the LDI to physical port mapping table)**

```
Validate VQI and LDI via "show system internal ethpm info interface Eth1/14 | include VQI"
```

Although VQI 48 maps to Eth1/13, note that congestion on the first port in the port group/asic is reported. Because there are four ports in one port group/asic, use the next command in order to show the actual interface within that port group/asic that sees the congestion.

## show hardware queuing drops egress module <x> (applicable for F2e only)

Use this command in order to show the actual egress interface that sees congestion on the port group/asic that is a part of VQI 48 from the previous example.

```
N7K2# show hardware queuing drops egress module 1

VQ Drops
--------------------------------------------------------------------------------
| Output | VQ Drops | VQ Congestion | Src | Src | Input |
| Interface | | | Mod | Inst | Interface |
--------------------------------------------------------------------------------
| Eth1/14 | 0000000000000000 | 0000000000001296 | 1 | 1 | Eth1/5-8 |
| Eth1/14 | 0000000000000000 | 0000000000000590 | 1 | 2 | Eth1/9-12 |
| Eth1/14 | 0000000000000000 | 0000000000001213 | 1 | 3 | Eth1/13-16 |
| Eth1/14 | 0000000000000000 | 0000000000000536 | 2 | 1 | Eth2/5-8 |
| Eth1/14 | 0000000000000000 | 0000000000000009 | 2 | 2 | Eth2/9-12 |
| Eth1/14 | 0000000000000000 | 0000000000000262 | 2 | 3 | Eth2/13-16 |
--------------------------------------------------------------------------------
```