

Contents

[Introduction](#)

[Background Information](#)

[Navigate the CLI](#)

[Attach to the FEX](#)

[Enter Debug Exec Mode](#)

[Exit Debug Exec Mode](#)

[Exit the FEX](#)

[Terminology](#)

[Host Interface \(HI\)](#)

[Network Interface \(NI\)](#)

[FEX Fabric Port](#)

[FEX ASIC Names](#)

[Front Port Mapping](#)

[N2K-C2148T-1GE](#)

[N2K-C2224TP-1GE / N2K-C2248TP-1GE](#)

[N2K-C2232PP-10GE / N2K-C2232TM-10GE](#)

[N2K-C2248TP-E-1G](#)

[N2K-C2248PQ-10GE & N2K-C2348UPQ-10GE](#)

[Verify SFP](#)

[Find Loss](#)

[View HI Port Counters](#)

[View NI Port Counters](#)

[View Historical Drops](#)

[View Recent Drops and Interrupts](#)

[View the Port Traffic Rate in Real Time](#)

[Mitigate Loss](#)

[Reposition Servers](#)

[Add Additional Uplinks](#)

[Share HI Buffers](#)

[Nexus 6000 FEX Load Balance Enhancement](#)

Introduction

This document describes how to troubleshoot performance on the Fabric Extenders (FEX) that can attach to Nexus 5000 or 6000 Series Switches.

Note: None of the commands introduced in this document are disruptive. You must have a Nexus 2000 switch connected to a 5000 or 6000 series switch.

Background Information

Navigate the CLI

Attach to the FEX

Attach to the FEX to run show commands on the FEX command line:

```
Nexus# attach fex fex  
fex>
```

Enter Debug Exec Mode

Enter the debug mode on the FEX in order to run advanced commands and specify the FEX asic name. Refer to Table 1. for the FEX asic names.

```
fex# dbgexec [prt/woo/red/pri]
```

Exit Debug Exec Mode

In order to exit the Debug Exec Mode use CTRL+C keyboard sequence:

```
fex> [CTRL+C]
```

Exit the FEX

In order to exit the fex, use the command **exit**:

```
fex# exit
```

Terminology

Host Interface (HI)

Hi's are the ports that face the servers on the FEX. These are commonly known as front ports. Every front port on a FEX has a HI number. This number is usually different than the port number, but it is used to troubleshoot commands to refer to a port. Each asic tabulates front ports differently.

Network Interface (NI)

NI's are the FEX Control Ports on the FEX that connect back to the parent switch. These are also referred to as Network Uplinks. These also have a unique NI number dependent on the model.

FEX Fabric Port

These ports are the parent switch side of the unique link to the FEX. These ports are configured with **switchport mode fex-fabric** and a **fex association** commands.

FEX ASIC Names

Each FEX is designed with a different ASIC. The abbreviation of the ASIC name is used in the debugging mode to run commands.

Most models of the FEX have one ASIC, however the 2148 has 6, each with 8 front ports. These are referred to as **rmon** in troubleshoot commands.

The ASIC names and associated abbreviations are listed for reference:

Table 1.

FEX Model	ASIC Name	Abbreviation
N2K-C2148T-1GE	redwood	rw
N2K-C2224TP-1GE	portola	prt
N2K-C2248TP-1GE	woodside	woo
N2K-C2232PP-10GE	princeton	pri
N2K-C2232TM-10GE	woodside	woo
N2K-C2248TP-E-1GE	woodside	woo
B22	woodside/belmont	woo
N2K-C2232TM-E-10GE	tiburon	tib

Front Port Mapping

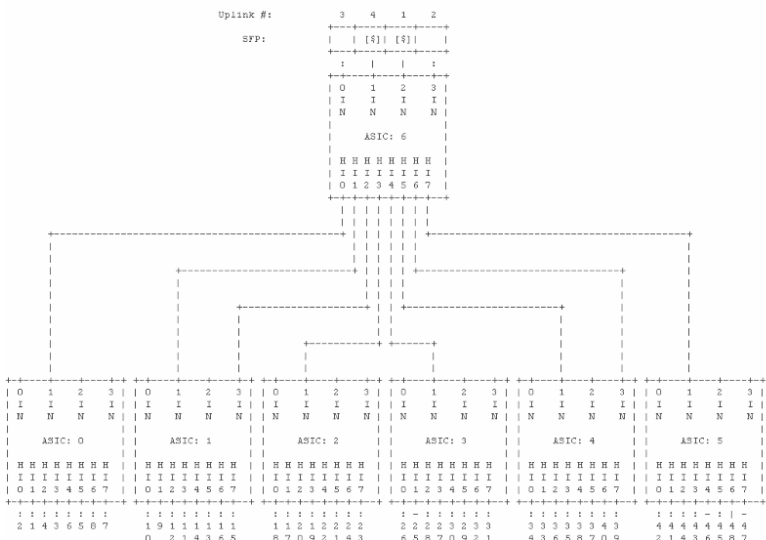
To interpret interface counter output it can be necessary to convert the front port number to a HI number. The conversion is dependent on the FEX chassis model.

N2K-C2148T-1GE

In this example, front port 26 (chassis-id/1/26) has been assigned rmon 3 HI 0:

```
switch# attach fex chassis_id
```

```
fex-[chassis_id]# show platform software redwood sts
```



N2K-C2224TP-1GE / N2K-C2248TP-1GE

In this example, front port 10 (135/1/10) has been assigned HI 9:

```
switch# attach fex chassis_id
```

```
fex-[chassis_id]# dbgexec portola
```

```
prt> fp
```

```
fex-135# dbgexec prt
prt> fp
Fabric port map:
Fabric port map:
  1      3
  |      :
+-----+-----+
| NI1 | NIO |
+-----+-----+
| NI2 | NI3 |
+-----+-----+
  :      |
  2      4
Front port map:
  1  3  5  7  9 11 13 15 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47
  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
HIF | 3 | 7 | 2 | 6 |11 |16 | |10 |15 |17 |20 |21 |23 | |26 |30 |27 |31 |35 |39 | |34 |38 |42 |46 |43 |47 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
HIF | 1 | 5 | 0 | 4 | 9 |13 | | 8 |12 |14 |18 |19 |22 | |24 |28 |25 |29 |32 |37 | |33 |36 |40 |44 |41 |45 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :
  2  4  6  8 10 12 14 16 18 20 22 24 26 28 30 32 34 36 38 40 42 44 46 48
prt> █
```

N2K-C2232PP-10GE / N2K-C2232TM-10GE

In this example, front port 20 (135/1/20) has been assigned HI 19:

```
switch# attach fex chassis_id
```

```
fex-[chassis_id]# show platform software woodside sts
```

```
(FINAL POSITION TBD)      Uplink #:      1  2  3  4  5  6  7  8
                          Link status:   :  :  :  :  :  :  :  |  |
                          SFP:            [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ]
                          +-----+-----+-----+-----+-----+-----+-----+
                          | N N N N N N N N N |
                          | I I I I I I I I I |
                          | 0 1 2 3 4 5 6 7 |
                          |
                          |           NI (0-7)           |
                          +-----+-----+-----+-----+
                          |
                          +-----+-----+-----+-----+
                          |
                          |           HI (0-7)           | |           HI (8-15)           | |           HI (16-23)           | |           HI (24-31)           |
                          | H H H H H H H H H | | H H H H H H H H H | | H H H H H H H H H | | H H H H H H H H H |
                          | I I I I I I I I I | | I I I I I I I I I | | I I I I I I I I I | | I I I I I I I I I |
                          | 0 1 2 3 4 5 6 7 | | 8 9 1 1 1 1 1 1 | | 1 1 1 1 2 2 2 2 | | 2 2 2 2 2 2 3 3 |
                          |           0 1 2 3 4 5 | | 6 7 8 9 0 1 2 3 | | 4 5 6 7 8 9 0 1 |
                          +-----+-----+-----+-----+
                          [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ]
                          +-----+-----+-----+-----+
                          - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
                          1  2  3  4  5  6  7  8      9  1  1  1  1  1  1  1      1  1  1  2  2  2  2  2      2  2  2  2  2  3  3  3
                          0  1  2  3  4  5  6      7  8  9  0  1  2  3  4      5  6  7  8  9  0  1  2
```

N2K-C2248TP-E-1G

```
fex-111# dbgexec pri
pri> fp
Fabric port map:
Fabric port map:
    1      3
    |      :
+-----+
|  NI1  | NI0 |
+-----+
|  NI2  | NI3 |
+-----+
    2      4
Front port map:
    1  3  5  7  9 11  13 15 17 19 21 23  25 27 29 31 33 35  37 39 41 43 45 47
    |  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :
HIF | 3 | 7 | 2 | 6 |11 |16 | |10 |15 |17 |20 |21 |23 | |26 |30 |27 |31 |35 |39 | |34 |38 |42 |46 |43 |47 |
+-----+
HIF | 1 | 5 | 0 | 4 | 9 |13 | | 8 |12 |14 |18 |19 |22 | |24 |28 |25 |29 |32 |37 | |33 |36 |40 |44 |41 |45 |
+-----+
    :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :  :
    2  4  6  8 10 12  14 16 18 20 22 24  26 28 30 32 34 36  38 40 42 44 46 48
```

N2K-C2248PQ-10GE & N2K-C2348UPQ-10GE

In this example, HI28 maps to Front Port 29:

```
tib> fp
-----
+-----+-----+-----+
| 1 3 5 7 9 1 1 1 | 1 1 2 2 2 2 2 3 | 3 3 3 3 4 4 4 4 |
|                   | 7 9 1 3 5 7 9 1 | 3 5 7 9 1 3 5 7 |
+-----+-----+-----+
| H                   | H                   | H                   |
| I                   | I                   | I                   |
| 0 2 4 6 8 1 1 1   | 1 1 2 2 2 2 2 3   | 3 3 3 3 4 4 4 4   |
|                   | 0 2 4   6 8 0 2 4 6 8 0 | 2 4 6 8 0 2 4 6 |
+-----+-----+-----+
| H                   | H                   | H                   |
| I                   | I                   | I                   |
| 1 3 5 7 9 1 1 1   | 1 1 2 2 2 2 2 3   | 3 3 3 3 4 4 4 4   |
|                   | 1 3 5   7 9 1 3 5 7 9 1 | 3 5 7 9 1 3 5 7 |
+-----+-----+-----+
| 2 4 6 8 1 1 1 1   | 1 2 2 2 2 2 3 3   | 3 3 3 4 4 4 4 4   |
|                   | 0 2 4 6   8 0 2 4 6 8 0 2 | 4 6 8 0 2 4 6 8 |
+-----+-----+-----+
-----
|-----+-----+
| 1-4   | | 9-12 |
|-----+-----+
|-----+-----+
| 5-8   | | 13-16 |
|-----+-----+
|-----+-----+
| NI2,3 | | NI6,7 |
|-----+-----+
```

Verify SFP

This command shows the Small Form-Factor Pluggable (SFP)

```
fex# show platform software woodside sfp rmon 0 HI5
```

In this example, you see that the SFP in HI5 is a 10G-Base-SR (LC) made by CISCO-AVAGO:



Note: If you run this command on a FEX that uses copper ports, then you will notice the

command errors. This is expected as there is no SFP to query. The prompt will return to **no SFP found** when that port is fiber, but does not currently contain an SFP.

Find Loss

Show commands can be run at the FEX prompt for the HI and NI ports in order to view interface counters on the FEX side of the FEX Fabric port links.

View HI Port Counters

This command shows port counter verification, similar to a **show int**:

Note: **rmon 0** is used only when the FEX has one host asic. The 2224, 2248 and 2232 models have only one asic. The 2148 model has six asics, so rmon 0 through 5 will be used. See the Front Port Mapping section for further details.

View NI Port Counters

This command will show you the port counters for the network uplinks similar to a **show int**. This command shows you the FEX side of the link. This command does not show you the parent switch side of the link.

View Historical Drops

Historical Drops can be viewed with the **drops** command. This shows you all drops on the FEX since it was turned on.

This command also shows you drops to the FEX CPU that will not represent data traffic drops with DROP8 counters. These can be safely ignored.

Note: **tail drop [8]** and **TAIL_DROP8** represents tail drops to the FEX CPU and is not relevant to troubleshoot performance as this happens under normal conditions.

View Recent Drops and Interrupts

Interrupts sent to the CPU include tail-drops, which are drops due to congestion and lack of buffer space. These can be viewed with the **show new_ints** command:

Note: 6.0 and later code uses **show new_ints all**

This example shows that frames tail drop in the SS1 buffer:

This example shows that NI 3 receives symbol errors:

This example shows that the FEX tail drops frames that ingress NI3:

View the Port Traffic Rate in Real Time

The rate command outputs real time traffic rate statistics for a port. Unlike **show int**, its not an average, its the raw current data rate that second. In this example, NI 3 currently receives 2.96kbps in the Network to Host direction. A **show int** on the corresponding parent Nexus switch shows 2.96Kbps in the TX direction on the FEX Fabric uplink connected to NI 3.

```
prt> rate
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Port    || Tx Packets | Tx Rate | Tx Bit  || Rx Packets | Rx Rate | Rx Bit  | Avg Pkt| Avg Pkt| |
|         ||            | (pkts/s) | Rate    ||            | (pkts/s) | Rate    | (Tx)   | (Rx)   | Err|
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| O-CI    ||          11 |         2 | 4.80Kbps ||          12 |         2 | 8.64Kbps | 252    | 430    | |
| O-NI3   ||           6 |         1 | 4.32Kbps ||           6 |         1 | 2.96Kbps | 430    | 289    | |
| O-NI1   ||           6 |         1 | 4.32Kbps ||           5 |         1 | 1.89Kbps | 430    | 217    | |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

Mitigate Loss

Tail drops are caused by buffer exhaustion. Typically the buffer becomes exhausted when multiple servers burst to the HIFs at once, or the host egress buffers cannot empty their outbound traffic fast enough to replenish the credits on the NIFs.

There are several options available to mitigate that loss.

Reposition Servers

Move any servers with bursty traffic flows such as storage arrays and video endpoints off of the FEX and connect them directly to the base ports of the parent switch. This will prevent the bursty servers from exhausting the buffer and starving out traffic from less chatty hosts.

Nexus 5000 and 6000 Series Switches have larger buffers than FEX models, to connect bursty servers to the base ports mitigates the loss because the base port buffers are able to handle a much greater burst.

Add Additional Uplinks

Some models of FEX can unlock additional buffer space when more uplinks from the FEX to the parent switch are added. This can potentially cease drops on the Network Uplinks.

Table 2.

Model	Buffer Increase When Adding Uplinks
2148	none
2224	buffer increase up to 2 uplinks
2248TP	buffer increase up to 4 uplinks
2232	buffer increase up to 4 uplinks
2248TP-E	none

2248PQ none

Share HI Buffers

Most models of FEX can benefit from sharing the HI buffer across all host ports. If drops are seen on the HI, sharing the buffer might mitigate those drops.

Modify the FEX queue limit globally:

5k(config)# no fex queue-limit (globally applies to all fexes on that 5k)

Modify the FEX queue limit on individual FEX:

Fex queue

5k(config)# fex 100

5k(config-fex)# no hardware [model] queue-limit

Nexus 6000 FEX Load Balance Enhancement

The Nexus 6000 has an additional option to change the load balancing algorithm from HIF to NIF. By default, even if packets arrive on different HIF ports, they might still be queued to the same NIF. With uplink-load-balance-mode enabled, they are distributed across multiple NIFs, and allow for more even usage of NIF egress buffers.

6k(config)# hardware N2248PQ uplink-load-balance-mode