Contents

Introduction

The purpose of this document is to show how to modify ingress buffers on the Cisco Nexus 7000 (N7k) Cisco Nexus 7000 48-Port 1 and 10 Gigabit Ethernet F2-Series Module (F2) and Cisco Nexus 7000 Enhanced F2-Series 48-Port Fiber 1 and 10 Gigabit Ethernet Module (F2e) linecards for Virtual Lane 3 (VL3).

Also, you will see the amount of ingress buffering capacity you gain for VL3 after modifying these values.

Problem

Using Fiber Channel over Ethernet (FCoE) multihop connections between Datacenters over distances greater than 2 kilometers can result in input drops. By default, the F2/F2e linecards have 0 pages in the latency buffer to queue packets after pause is sent and this will lead to input drops on long distance FCoE multihop interfaces.

The latency buffer is defined as follows:

PL_STOP - HWM (PL_Pause) = LB (Latency Buffer)

You'll notice the values mentioned above are displayed as pages. Each page is roughly 384 bytes.

Notice below, the ingress buffer capacity of VL3 with the default FCoE QoS policy:

EX

```
module-10# show hardware internal mac port 1 qos configuration | begin IB | end EB
  IB
    Port page limit : 3584 (1376256 Bytes)
    VL#  HWM pages(bytes)  LWM pages(bytes) Used PL_STOP(HWM & LWM) SPAN
                                      pages                      THR
     0    1107 (  425088)   1035 (  397440)     0     1107   1035    100
     1       2 (     768)      1 (     384)     0        2      1      1
     2       2 (     768)      1 (     384)     0        2      1      1
     3    1053 (  404352)   1029 (  395136)     0     1053   1029    100
     4    1107 (  425088)   1083 (  415872)     0     1107   1083    100
     5     231 (   88704)    159 (   61056)     0      231    159     57
     6       2 (     768)      1 (     384)     0        2      1      1
     7       2 (     768)      1 (     384)     0        2      1      1
    Credited DWRR WT: 216 (0xd8) Uncredited DWRR WT: 144 (0x90)
    DWRR honor UC = FALSE
     Leak Lo weight = 0xd8, enabled = FALSE
  EB
```

PL_STOP and High Water Mark (HWM) is of same value. Here you can see the latency buffer has 0 pages by default. To support long distance FCoE these values will need to be modified.

Solution

First you will need to duplicate the 'default-4q-7e-in-policy' Quality of Service (QoS) policy-map:

```
module-10# show hardware internal mac port 1 qos configuration | begin IB | end EB
  IB
    Port page limit : 3584 (1376256 Bytes)
    VL#  HWM pages(bytes)  LWM pages(bytes) Used PL_STOP(HWM & LWM) SPAN
                                            pages                   THR
     0    1107 (  425088)   1035 (  397440)     0     1107   1035    100
     1       2 (     768)      1 (     384)     0        2      1      1
     2       2 (     768)      1 (     384)     0        2      1      1
     3    1053 (  404352)   1029 (  395136)     0     1053   1029    100
     4    1107 (  425088)   1083 (  415872)     0     1107   1083    100
     5     231 (   88704)    159 (   61056)     0      231    159     57
     6       2 (     768)      1 (     384)     0        2      1      1
     7       2 (     768)      1 (     384)     0        2      1      1
    Credited DWRR WT: 216 (0xd8) Uncredited DWRR WT: 144 (0x90)
    DWRR honor UC = FALSE
     Leak Lo weight = 0xd8, enabled = FALSE
  EB
```

Below you will see the amount of bytes allocated to the latency buffer of VL3, after modifying the service policies.

Note: You will not see a latency buffer until you allocate ATLEAST 60% of the queue-limit to the "ndrop" policy.

Policies will be modified in increments of 10, up to 99%

```
60/40 ingess buffer allocation
==============================
policy-map type queuing 7I_4q-7e-in
  class type queuing c-4q-7e-drop-in
    service-policy type queuing 7I_4q-7e-drop-in
    queue-limit percent 40
  class type queuing c-4q-7e-ndrop-in
    service-policy type queuing 7I_4q-7e-ndrop-in
    queue-limit percent 60

interface Ethernet2/5
  service-policy type queuing input 7I_4q-7e-in

module-2# show hardware internal mac port 5 qos configuration | begin IB | end EB
  IB
    Port page limit : 3584 (1376256 Bytes)
    VL#  HWM pages(bytes)  LWM pages(bytes) Used PL_STOP(HWM & LWM) SPAN
                                            pages                   THR
     0     624 (  239616)    576 (  221184)     0      624    576    100
     1       2 (     768)      1 (     384)     0        2      1      1
     2     624 (  239616)    576 (  221184)     0      624    576    100
     3    1913 (  734592)   1889 (  725376)     0     2126   1889    100
     4       2 (     768)      1 (     384)     0        2      1      1
     5     124 (   47616)     52 (   19968)     0      124     52     31
     6       2 (     768)      1 (     384)     0        2      1      1
     7       2 (     768)      1 (     384)     0        2      1      1
    Credited DWRR WT: 216 (0xd8) Uncredited DWRR WT: 144 (0x90)
```

```
    DWRR honor UC = FALSE
     Leak Lo weight = 0xd8, enabled = FALSE
  EB
```

60/40 will allocate 81792 bytes to vl3 latency buffer.

PL_STOP - HWM * 384 bytes
2126 - 1913 = 213pages * 384 = 81792 bytes

```
70/30 ingress buffer allocation
==============================
policy-map type queuing 7I_4q-7e-in
  class type queuing c-4q-7e-drop-in
    service-policy type queuing 7I_4q-7e-drop-in
    queue-limit percent 30
  class type queuing c-4q-7e-ndrop-in
    service-policy type queuing 7I_4q-7e-ndrop-in
    queue-limit percent 70

interface Ethernet2/5
  service-policy type queuing input 7I_4q-7e-in

module-2# show hardware internal mac port 5 qos configuration | begin IB | end EB
  IB
    Port page limit : 3584 (1376256 Bytes)
    VL#  HWM pages(bytes)  LWM pages(bytes) Used PL_STOP(HWM & LWM) SPAN
                                            pages                   THR
     0     463 (  177792)    415 (  159360)     0      463    415    100
     1       2 (     768)      1 (     384)     0        2      1      1
     2     463 (  177792)    415 (  159360)     0      463    415    100
     3    1987 (  763008)   1963 (  753792)     0     2484   1963    100
     4       2 (     768)      1 (     384)     0        2      1      1
     5      88 (   33792)     16 (    6144)     0       88     16     22
     6       2 (     768)      1 (     384)     0        2      1      1
     7       2 (     768)      1 (     384)     0        2      1      1
    Credited DWRR WT: 216 (0xd8) Uncredited DWRR WT: 144 (0x90)
    DWRR honor UC = FALSE
     Leak Lo weight = 0xd8, enabled = FALSE
  EB
```

70/30 allocates 190848 bytes to VL3 latency buffer.

```
policy-map type queuing 7I_4q-7e-in
  class type queuing c-4q-7e-drop-in
    service-policy type queuing 7I_4q-7e-drop-in
    queue-limit percent 20
  class type queuing c-4q-7e-ndrop-in
    service-policy type queuing 7I_4q-7e-ndrop-in
    queue-limit percent 80

interface Ethernet2/5
  service-policy type queuing input 7I_4q-7e-in

module-2# show hardware internal mac port 5 qos configuration | begin IB | end EB
  IB
    Port page limit : 3584 (1376256 Bytes)
    VL#  HWM pages(bytes)  LWM pages(bytes) Used PL_STOP(HWM & LWM) SPAN
                                            pages                   THR
     0     302 (  115968)    254 (   97536)     0      302    254     75
     1       2 (     768)      1 (     384)     0        2      1      1
     2     302 (  115968)    254 (   97536)     0      302    254     75
     3    1875 (  720000)   1851 (  710784)     0     2841   1851    100
```

```
     4       2 (     768)      1 (     384)     0        2      1      1
     5      52 (   19968)     46 (   17664)     0       52     46     13
     6       2 (     768)      1 (     384)     0        2      1      1
     7       2 (     768)      1 (     384)     0        2      1      1
    Credited DWRR WT: 216 (0xd8) Uncredited DWRR WT: 144 (0x90)
    DWRR honor UC = FALSE
     Leak Lo weight = 0xd8, enabled = FALSE
  EB
```

80/20 allocates 370944 bytes to VL3 latency buffer.

```
policy-map type queuing 7I_4q-7e-in
  class type queuing c-4q-7e-drop-in
    service-policy type queuing 7I_4q-7e-drop-in
    queue-limit percent 10
  class type queuing c-4q-7e-ndrop-in
    service-policy type queuing 7I_4q-7e-ndrop-in
    queue-limit percent 90


interface Ethernet2/5
  service-policy type queuing input 7I_4q-7e-in

module-2# show hardware internal mac port 5 qos configuration | begin IB | end EB
  IB
    Port page limit : 3584 (1376256 Bytes)
    VL#  HWM pages(bytes)  LWM pages(bytes) Used PL_STOP(HWM & LWM) SPAN
                                           pages                    THR
     0     141 (   54144)     93 (   35712)     0      141     93     35
     1       2 (     768)      1 (     384)     0        2      1      1
     2     141 (   54144)     93 (   35712)     0      141     93     35
     3    1055 (  405120)   1031 (  395904)     0     3199   1031    100
     4       2 (     768)      1 (     384)     0        2      1      1
     5      16 (    6144)     10 (    3840)     0       16     10      4
     6       2 (     768)      1 (     384)     0        2      1      1
     7       2 (     768)      1 (     384)     0        2      1      1
    Credited DWRR WT: 216 (0xd8) Uncredited DWRR WT: 144 (0x90)
    DWRR honor UC = FALSE
     Leak Lo weight = 0xd8, enabled = FALSE
  EB
```

90/10 allocates 823296 bytes to VL3 latency buffer

```
policy-map type queuing 7I_4q-7e-in
  class type queuing c-4q-7e-drop-in
    service-policy type queuing 7I_4q-7e-drop-in
    queue-limit percent 10
  class type queuing c-4q-7e-ndrop-in
    service-policy type queuing 7I_4q-7e-ndrop-in
    queue-limit percent 90


interface Ethernet2/5
  service-policy type queuing input 7I_4q-7e-in

module-2# show hardware internal mac port 5 qos configuration | begin IB | end EB
  IB
    Port page limit : 3584 (1376256 Bytes)
    VL#  HWM pages(bytes)  LWM pages(bytes) Used PL_STOP(HWM & LWM) SPAN
                                           pages                    THR
     0     141 (   54144)     93 (   35712)     0      141     93     35
     1       2 (     768)      1 (     384)     0        2      1      1
     2     141 (   54144)     93 (   35712)     0      141     93     35
     3    1055 (  405120)   1031 (  395904)     0     3199   1031    100
     4       2 (     768)      1 (     384)     0        2      1      1
```

```
   5      16 (    6144)     10 (    3840)     0        16     10      4
   6       2 (     768)      1 (     384)     0         2      1      1
   7       2 (     768)      1 (     384)     0         2      1      1
  Credited DWRR WT: 216 (0xd8) Uncredited DWRR WT: 144 (0x90)
  DWRR honor UC = FALSE
   Leak Lo weight = 0xd8, enabled = FALSE
EB
```

99/1 allocates 906240 bytes to VL3 latency buffer

Note: Each clipper asic has 6MB of buffer capacity. There are 4 ports per clipper, therefore this equates to ~1.5MB buffer capacity per port. With 99/1 you will see ~.9MB is allocated to VL3 latency buffer and the remainder is used by HWM for each VL(majority to VL3). When adding each VLs HWM with the LB of VL3 you'll see it equates to ~1.35MB buffer capacity.