

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Problem](#)

[Troubleshoot & Verify](#)

[Solution](#)

[Related Cisco Support Community Discussions](#)

Introduction

EtherChannel bundle is used to provide high bandwidth interconnects. This article discusses a limitation which applies to Cisco EtherChannels on Catalyst 6500 switches running Supervisor 720 with PFC3A, PFC3B or PFC3BXL that can cause overrun to increment on Etherchannel member interfaces. This limitation is related to Layer 2 Forwarding Engine and hence applies only to layer 2 EtherChannels.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on Cisco Catalyst 6500 Series switches that run Supervisor Engine 720. WS-X6704-10GE has been used in this lab setup. WS-X6704-10GE is a Catalyst 6500 module with no oversubscription & has 2 fabric channel connections of 20 Gbps each.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Problem

A Catalyst 6500 might experience interface overrun when a Layer 2 Distributed Etherchannel (DEC) is configured. DEC is an etherchannel An Etherchannel with 2 member interfaces connected on the same linecard but on different fabric channels is not considered a DEC.

Overrun counter accounts number of times the receiver hardware was unable to hand received data to a hardware buffer. In other words, the input rate of traffic exceeded the ability of the receiver to handle the data.

In certain cases, DFC provides the capability to recirculate the packets. Recirculation can be used to perform additional lookups in the ACL or QoS Ternary Content Addressable memory (TCAM), the NetFlow table, or the Forwarding Information Base (FIB) TCAM table. Packet recirculation occurs only on a particular packet flow; other packet flows are not affected. The rewrite of the packet occurs on the modules; the packets are then forwarded back to the Policy Feature Card (PFC) for additional processing.

When using a layer 2 DEC, packet recirculation at ingress module is required during Catalyst 6500 is running in flow-through mode along with 3B/3BXL PFC mode.

More information about flow-through mode is available [here](#).

Overrun counter can start incrementing when the fabric utilization reaches about 50%.

Troubleshoot & Verify

- 1) Find out the member interfaces in the Etherchannel experiencing incrementing overruns.
- 2) Verify input rate and overrun counters on member interfaces.

```
6500#show interfaces tenGigabitEthernet 2/1
TenGigabitEthernet2/1 is up, line protocol is up (connected)
Hardware is C6k 10000Mb 802.3, address is 0002.fcc1.21ac (bia 0002.fcc1.21ac)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 251/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 10Gb/s, media type is 10Gbase-SR
input flow-control is on, output flow-control is off
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output 00:00:51, output hang never
Last clearing of "show interface" counters 00:00:13
Input queue: 0/2000/5597178/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
30 second input rate 9868906000 bits/sec, 822409 packets/sec
30 second output rate 3000 bits/sec, 5 packets/sec
10968368 packets input, 16452552000 bytes, 0 no buffer
Received 0 broadcasts (0 multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 5597195 overrun, 0 ignored
0 watchdog, 0 multicast, 0 pause input
0 input packets with dribble condition detected
79 packets output, 5596 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 PAUSE output
0 output buffer failures, 0 output buffers swapped out
```

```
6500#show interfaces tenGigabitEthernet 2/2
TenGigabitEthernet2/2 is up, line protocol is up (connected)
Hardware is C6k 10000Mb 802.3, address is 0002.fcc1.21ad (bia 0002.fcc1.21ad)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 251/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 10Gb/s, media type is 10Gbase-SR
```

```
input flow-control is on, output flow-control is off
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output 00:00:26, output hang never
Last clearing of "show interface" counters 00:00:03
Input queue: 0/2000/45043/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
30 second input rate 9868149000 bits/sec, 822345 packets/sec
30 second output rate 2000 bits/sec, 4 packets/sec
8823464 packets input, 13233496000 bytes, 0 no buffer
Received 0 broadcasts (0 multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 4575029 overrun, 0 ignored
0 watchdog, 0 multicast, 0 pause input
0 input packets with dribble condition detected
0 packets output, 0 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 PAUSE output
0 output buffer failures, 0 output buffers swapped out
```

3) Find out the modules on which these interfaces are present.

```
6500#show interfaces tenGigabitEthernet 2/1
TenGigabitEthernet2/1 is up, line protocol is up (connected)
Hardware is C6k 10000Mb 802.3, address is 0002.fcc1.21ac (bia 0002.fcc1.21ac)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 251/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 10Gb/s, media type is 10Gbase-SR
input flow-control is on, output flow-control is off
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output 00:00:51, output hang never
Last clearing of "show interface" counters 00:00:13
Input queue: 0/2000/5597178/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
30 second input rate 9868906000 bits/sec, 822409 packets/sec
30 second output rate 3000 bits/sec, 5 packets/sec
10968368 packets input, 16452552000 bytes, 0 no buffer
Received 0 broadcasts (0 multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 5597195 overrun, 0 ignored
0 watchdog, 0 multicast, 0 pause input
0 input packets with dribble condition detected
79 packets output, 5596 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 PAUSE output
0 output buffer failures, 0 output buffers swapped out
```

```
6500#show interfaces tenGigabitEthernet 2/2
TenGigabitEthernet2/2 is up, line protocol is up (connected)
Hardware is C6k 10000Mb 802.3, address is 0002.fcc1.21ad (bia 0002.fcc1.21ad)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 251/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 10Gb/s, media type is 10Gbase-SR
input flow-control is on, output flow-control is off
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output 00:00:26, output hang never
```

```

Last clearing of "show interface" counters 00:00:03
Input queue: 0/2000/45043/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
30 second input rate 9868149000 bits/sec, 822345 packets/sec
30 second output rate 2000 bits/sec, 4 packets/sec
8823464 packets input, 13233496000 bytes, 0 no buffer
Received 0 broadcasts (0 multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 4575029 overrun, 0 ignored
0 watchdog, 0 multicast, 0 pause input
0 input packets with dribble condition detected
0 packets output, 0 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 PAUSE output
0 output buffer failures, 0 output buffers swapped out

```

4) Find out the fabric interface utilization corresponding to these modules.

```

6500#show fabric utilization
slot    channel    speed    Ingress %    Egress %
  2      0          20G      0            0
  2     1         20G     49         0
  3      0          20G      0            0
  3      1          20G      0            50
  5      0          20G      0            0

```

5) As seen above, interface Tengigabitethernet 2/1 and Tengigabitethernet 2/2 are running at line rate, but overrun (drops) are happening on ingress due to ingress flow control. In order to check if switch is experiencing flow control and is running into the limitation described above use commands below.

```

6500#show platform hardware capacity rewrite-engine performance slot 2
slot channel perf_id description          packets          total overruns
-----+-----+-----+-----+-----+-----+
2  0  0x235  FAB RX 0          41083           0
2  0  0x237  FAB RX 1           0               0
2  0  0x27B  FAB TX 0         1904            0
2  0  0x27F  FAB TX 1           0               0
2  0  0x350  REPLICATION ML3   0               0
2  0  0x351  REPLICATION ML2   0               0
2  0  0x352  RECIRC L2         0               0
2  0  0x353  RECIRC L3         0               0
2  0  0x34C  SPAN TX 0         0               0
2  0  0x34D  SPAN TX 1         0               0
2  0  0x34E  SPAN RX 0         0               0
2  0  0x34F  SPAN RX 1         0               0
2  0  0x354  SPAN TERMINATION  0               0
2  1  0x235  FAB RX 0       759500888       0
2  1  0x237  FAB RX 1         0               0
2  1  0x27B  FAB TX 0     4545890286       0
2  1  0x27F  FAB TX 1         0               0
2  1  0x350  REPLICATION ML3   0               0
2  1  0x351  REPLICATION ML2   0               0
2  1  0x352  RECIRC L2     68615145        1047 <<< L2 Recirculation
2  1  0x353  RECIRC L3         0               0
2  1  0x34C  SPAN TX 0         0               0
2  1  0x34D  SPAN TX 1         0               0
2  1  0x34E  SPAN RX 0         0               0
2  1  0x34F  SPAN RX 1         0               0
2  1  0x354  SPAN TERMINATION  0               0

```

The above output indicates that the ingress replication engine is performing L2 recirculation due to the presence of a L2 DEC. Because of the recirculation, packets arrive at the replication engine twice the time which consumes twice the bandwidth. In other words, performance of single fabric channel is cut in half because each each packet is seen twice internally.

Solution

1) Configure a non distributed Etherchannel which is not affected by this limitation.

To validate this theory, an Etherchannel was configured on interfaces on same module (non DEC) and it was observed that at same packet rate as above, interfaces did not see any overruns incrementing. This can be a workaround to bypass this problem.

```
6500#show fabric utilization
```

slot	channel	speed	Ingress %	Egress %
2	0	20G	0	0
2	1	20G	99	0
3	0	20G	0	0
3	1	20G	0	99
5	0	20G	0	0

```
TenGigabitEthernet2/1 is up, line protocol is up (connected)
Hardware is C6k 10000Mb 802.3, address is 0002.fcc1.21ac (bia 0002.fcc1.21ac)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 251/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 10Gb/s, media type is 10Gbase-SR
input flow-control is on, output flow-control is off
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output 00:00:06, output hang never
Last clearing of "show interface" counters 00:36:12
Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
30 second input rate 9868487000 bits/sec, 822374 packets/sec
30 second output rate 3000 bits/sec, 6 packets/sec
 1783710310 packets input, 2675565466500 bytes, 0 no buffer
Received 0 broadcasts (0 multicasts)
 0 runts, 0 giants, 0 throttles
 0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
 0 watchdog, 0 multicast, 0 pause input
 0 input packets with dribble condition detected
13115 packets output, 946206 bytes, 0 underruns
 0 output errors, 0 collisions, 0 interface resets
 0 babbles, 0 late collision, 0 deferred
 0 lost carrier, 0 no carrier, 0 PAUSE output
 0 output buffer failures, 0 output buffers swapped out
```

```
TenGigabitEthernet2/2 is up, line protocol is up (connected)
Hardware is C6k 10000Mb 802.3, address is 0002.fcc1.21ad (bia 0002.fcc1.21ad)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 251/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 10Gb/s, media type is 10Gbase-SR
input flow-control is on, output flow-control is off
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output 00:00:11, output hang never
Last clearing of "show interface" counters 00:37:31
Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
30 second input rate 9868462000 bits/sec, 822371 packets/sec
30 second output rate 3000 bits/sec, 6 packets/sec
 1849499775 packets input, 2774249662500 bytes, 0 no buffer
```

Received 0 broadcasts (0 multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 watchdog, 0 multicast, 0 pause input
0 input packets with dribble condition detected
13599 packets output, 980928 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 PAUSE output
0 output buffer failures, 0 output buffers swapped out

2) Use Catalyst 6500 switch in PFC 3C/3CXL mode in case L2 DEC is required.

Note: DFC hardware upgrade would be required in case existing modules are running DFC3A/ DFC3B/ DFC3BXL.

3) Upgrade IOS version if your design and configuration applies to the conditions in [CSCti23324](#).

This bug fix relaxes the recirculation requirement for L2 DEC or multi-module EtherChannel for Catalyst 6500 switches with 67xx modules only. This bug is resolved in Cisco IOS release 12.2(33)SXJ1 and later. Be aware of following points that apply to this bug.

a) The bug fix relaxes the recirculation requirement for L2 DEC or multi-module EC for Catalyst 6500 switches 67xx modules only. In case the Catalyst 6500 switch has at least one L2 DEC across any older DFC module (e.g. 6516/6816) or combination of 67xx and 6516/6818 module, recirculation will be imposed for all L2 DEC's configured in the system. In case Catalyst 6500 switch has any older module and is configured with L2 DEC on 67xx modules only, recirculation will not be imposed.

b) Presence of all 67xx line cards is not enough to remove the recirculation requirement for DEC's. For example, if you have a DEC across 2 6704 DFC's and another port-channel configured on a 6748 CFC, the system will check the forwarding engine of supervisor (for the CFC module) and start using recirculation.

c) For VS-SUP720-10G, this bug fix does not work in scenarios where at least one port of L2 DEC is on CFC linecard / supervisor. In this scenario recirculation still happens. In addition, adjacency is not upgraded and recirculation is still in place even if you remove supervisor/CFC enabled port from the port-channel. In such scenario, reload is required to reprogram the hardware and removing & reconfiguring port-channel / redundancy switchover / removing L2 VLAN, etc. do not help.