P802.1Qcz / D1.0
WG Ballot Report

Paul Congdon

5/20/2020

First WG ballot stats

CATEGORY	All Respondents	
	TOTAL	%
Yes	9	81.82%
No	2	18.18%
Voting Yes or No	11	100.00%
Abs. Time	7	17.95%
Abs. Expertise	12	30.77%
Abs. Other	0	0.00%
Respondents	39	62.90%
Voting members	62	
Non-voting	3	
No. of commenters	6	15.38%
No. of comments	45	

D1.0 Changes

- Inclusion of SNMP MIB extensions and YANG for LLDP TLVs
- Framework to document YANG for all 802.1 Extension TLVs, with basicSet included
- Resolutions of comments from 3rd Task Group Ballot – mostly editorial

Disposition Plan

http://www.ieee802.org/1/files/private/cz-drafts/d1/802-1Qcz-d1-0-pdis-v01.pdf.

Priority items to discuss:

- 6, 11, 18, 20-21, 30, 35

Propose to Accept with discuss and/or modify the remedy:

- 12, 33, 36-37, 39, 42, 44, 45

Propose to Accept without:

- 1-5, 7-10, 13-17, 19, 22-29, 31-32, 34, 38, 40-41, 43

Material for discussion of comments 30 and 35

- Essence of the comments: Why are we changing the PCP on frames that are having their priority remapped?
- Material taken from previous contributions:
  - cz-congdon-ci-design-topics-1118-v01
  - cz-congdon-d0-4-status-1119-v02

Problem Statement

Once a flow has been isolated and a CIM is sent to the upstream switch to also isolate the same flow. The flow will be isolated to the same traffic class.

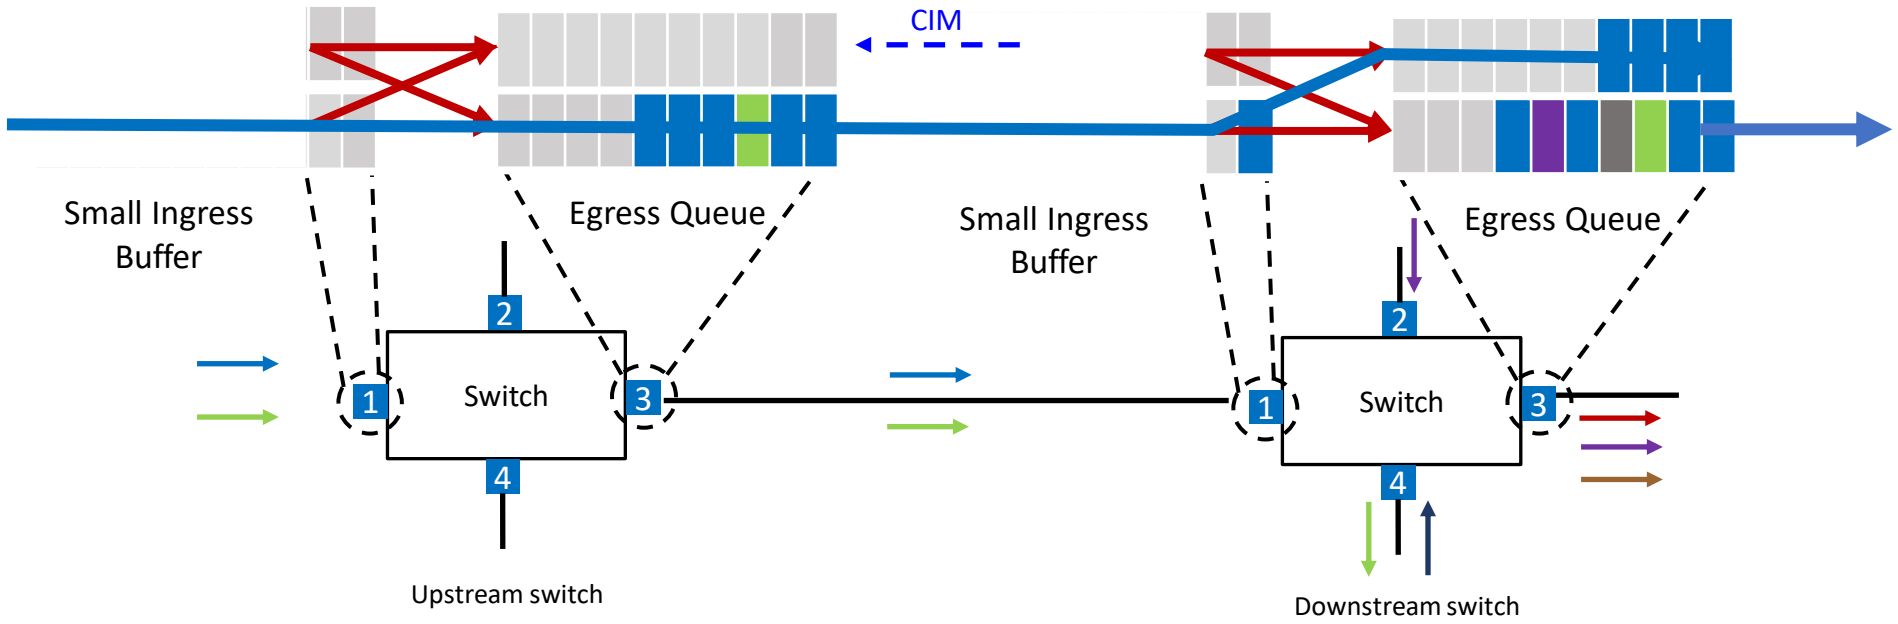

Problem Statement

After processing the CIM, subsequent packets of the isolated flow will be place in the congested queue upstream

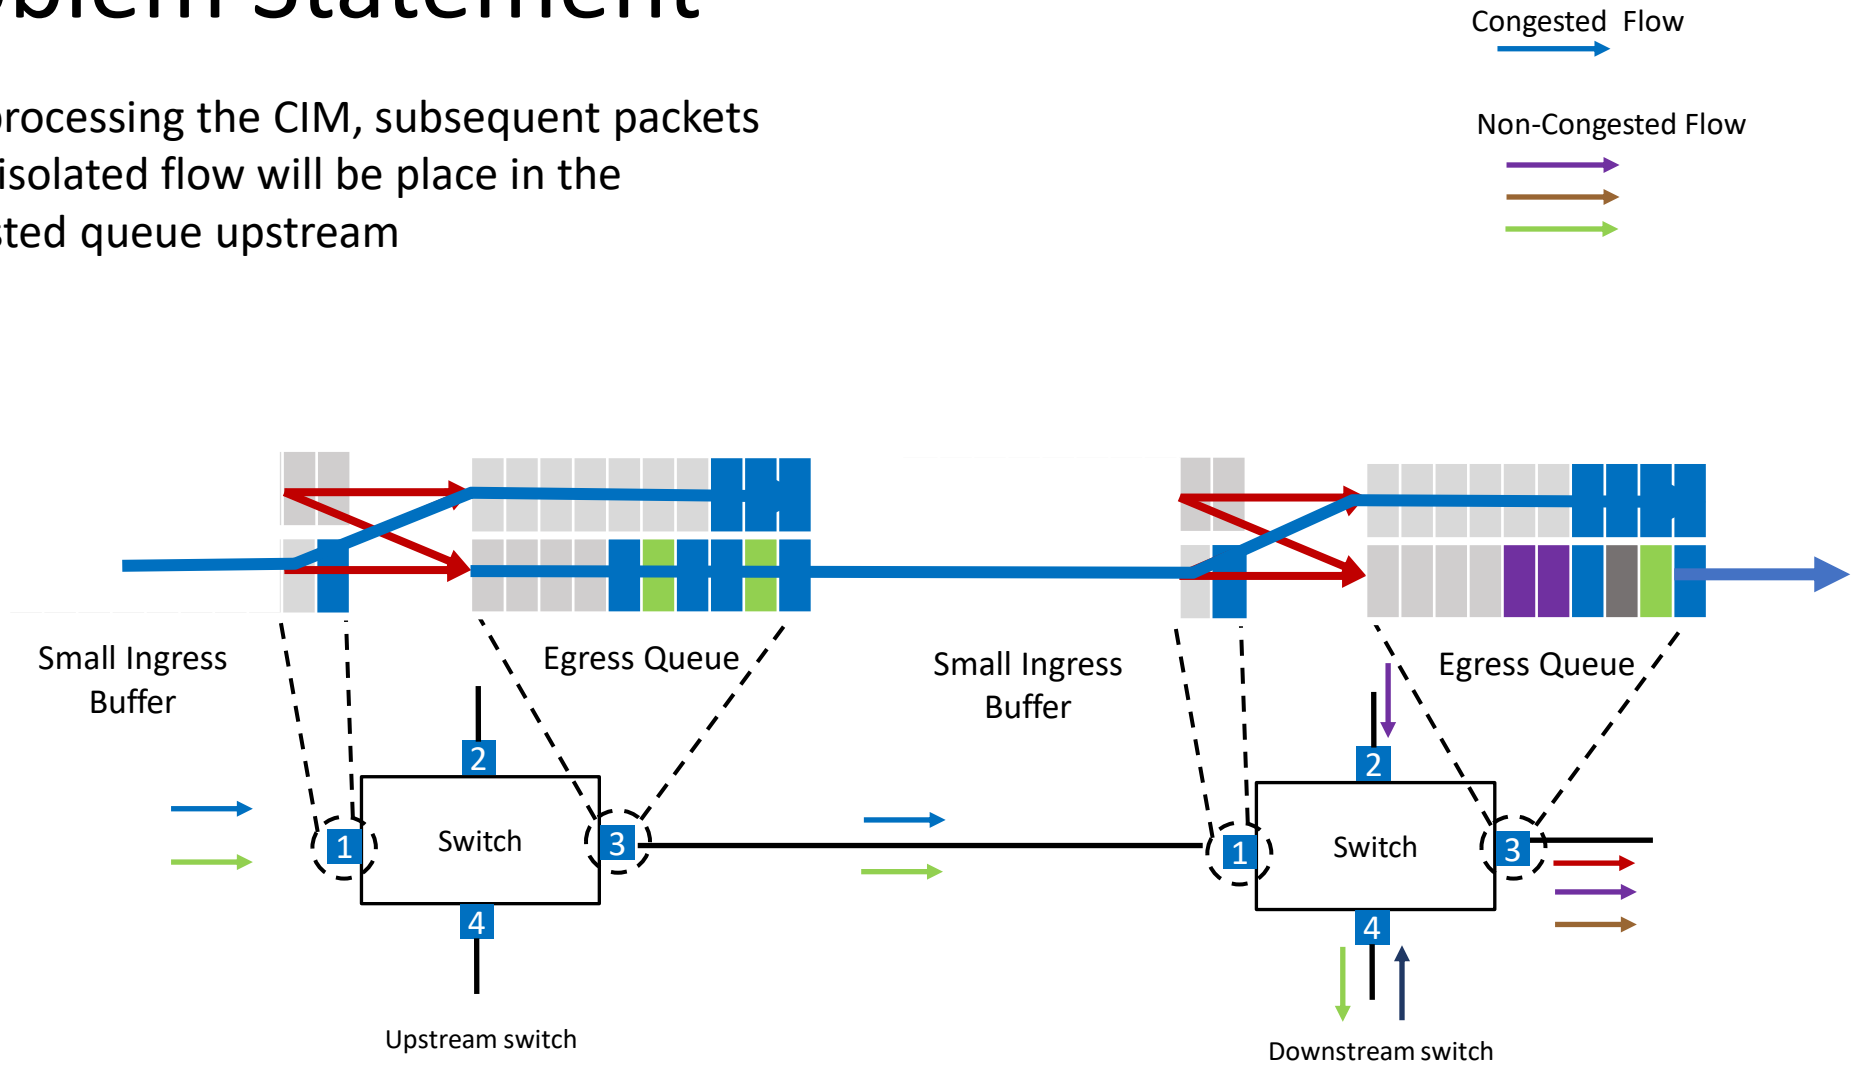

Problem Statement

It may occur that when the PFC Pause arrives at the upstream there are still packets of congested flow in the non-congested queue.

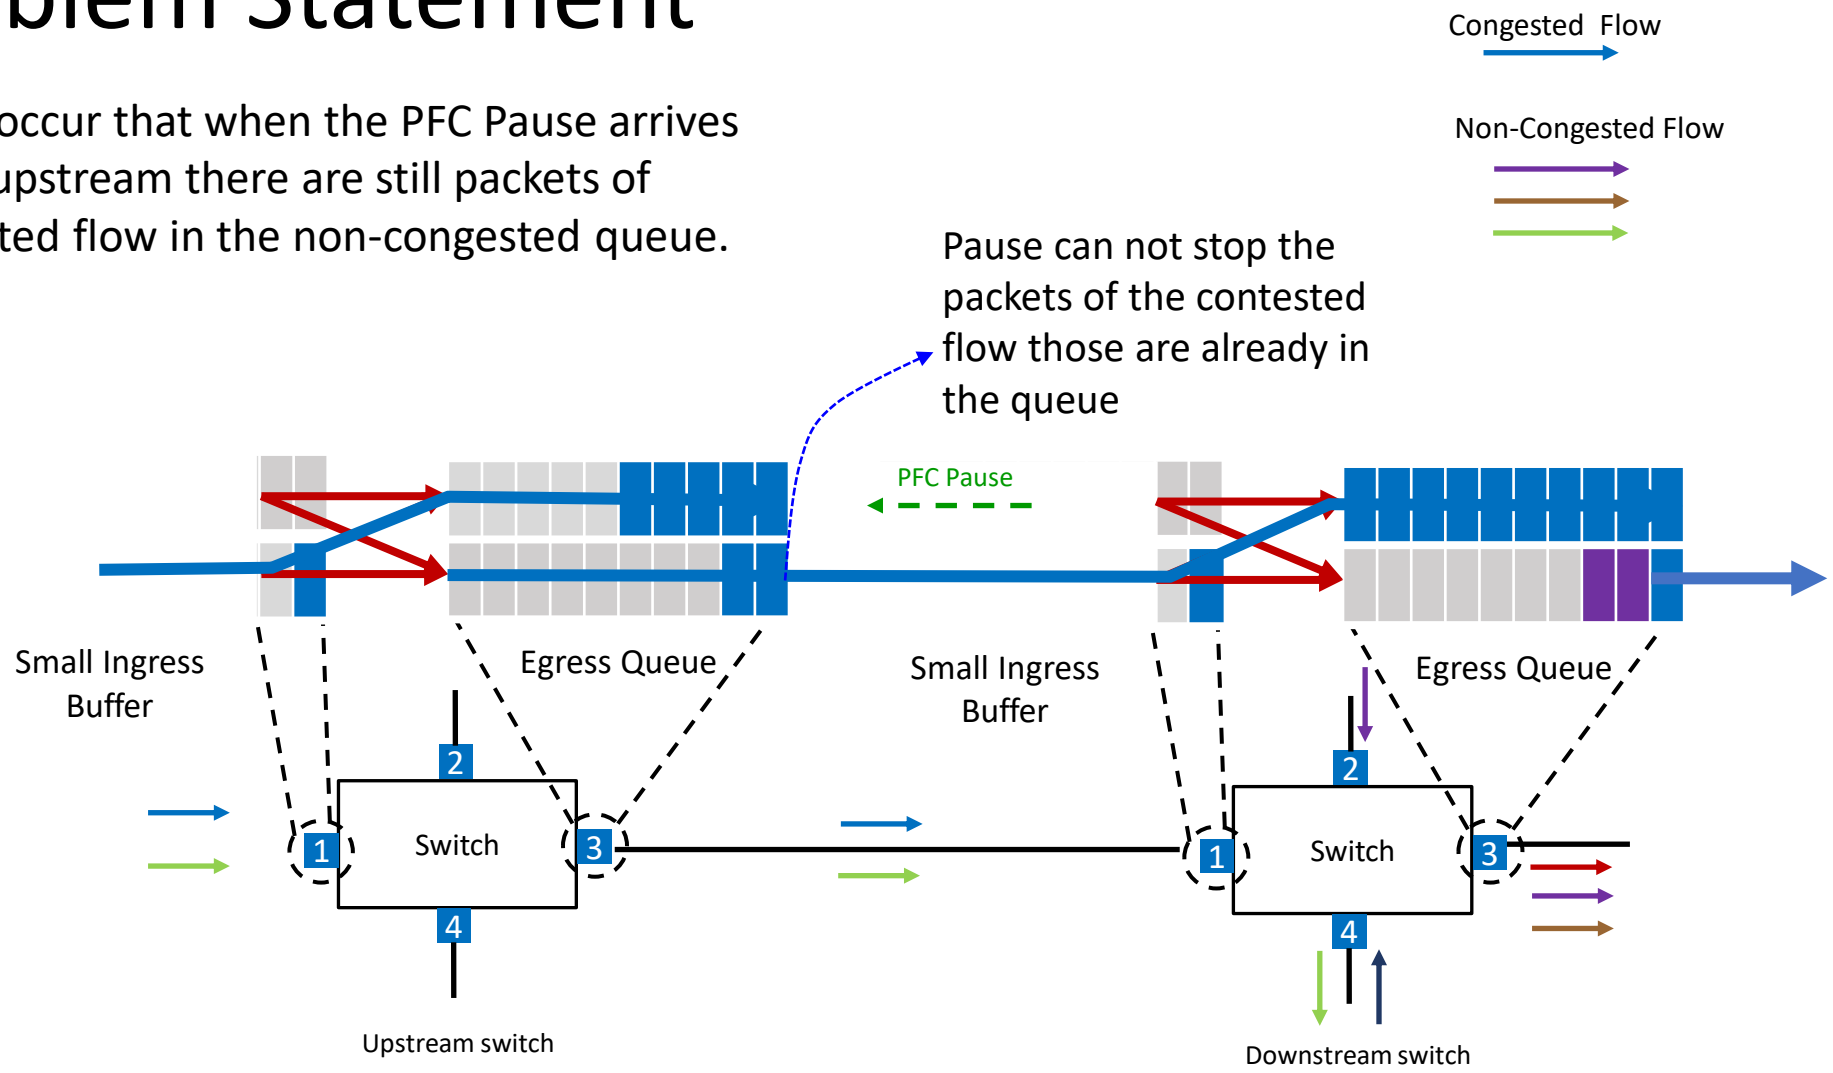

Observation

- Previous discussion of the complexity of interacting with PFC had assumed the downstream switch can not identify the traffic class used upstream to egress the packet.
- As long as the downstream switch ‘knows’ what traffic class the upstream switch egressed the packet, the downstream switch can Pause the correct traffic class

Proposed Solution

If using PFC, ***require*** that packets are priority or VLAN tagged on egress and the PCP contains an explicit indication of the traffic class used on egress

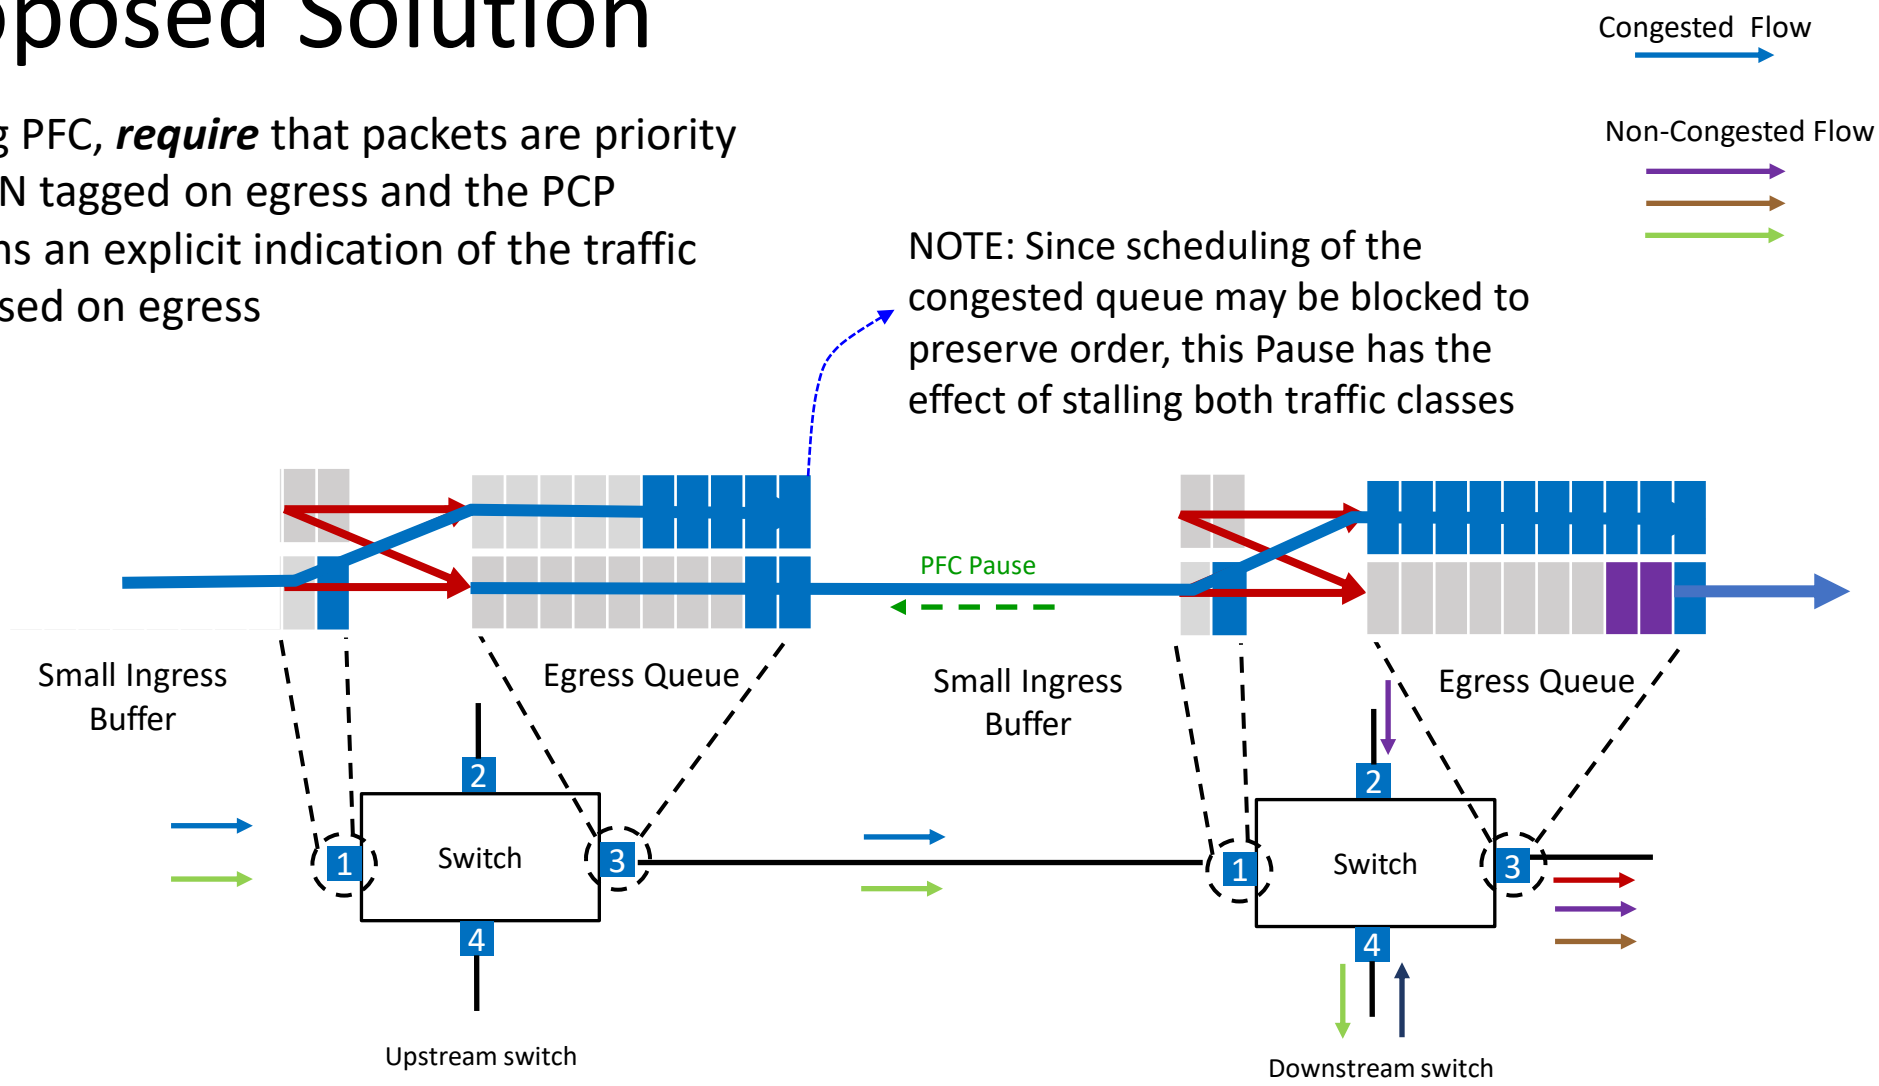

Congesting Flow Stickiness

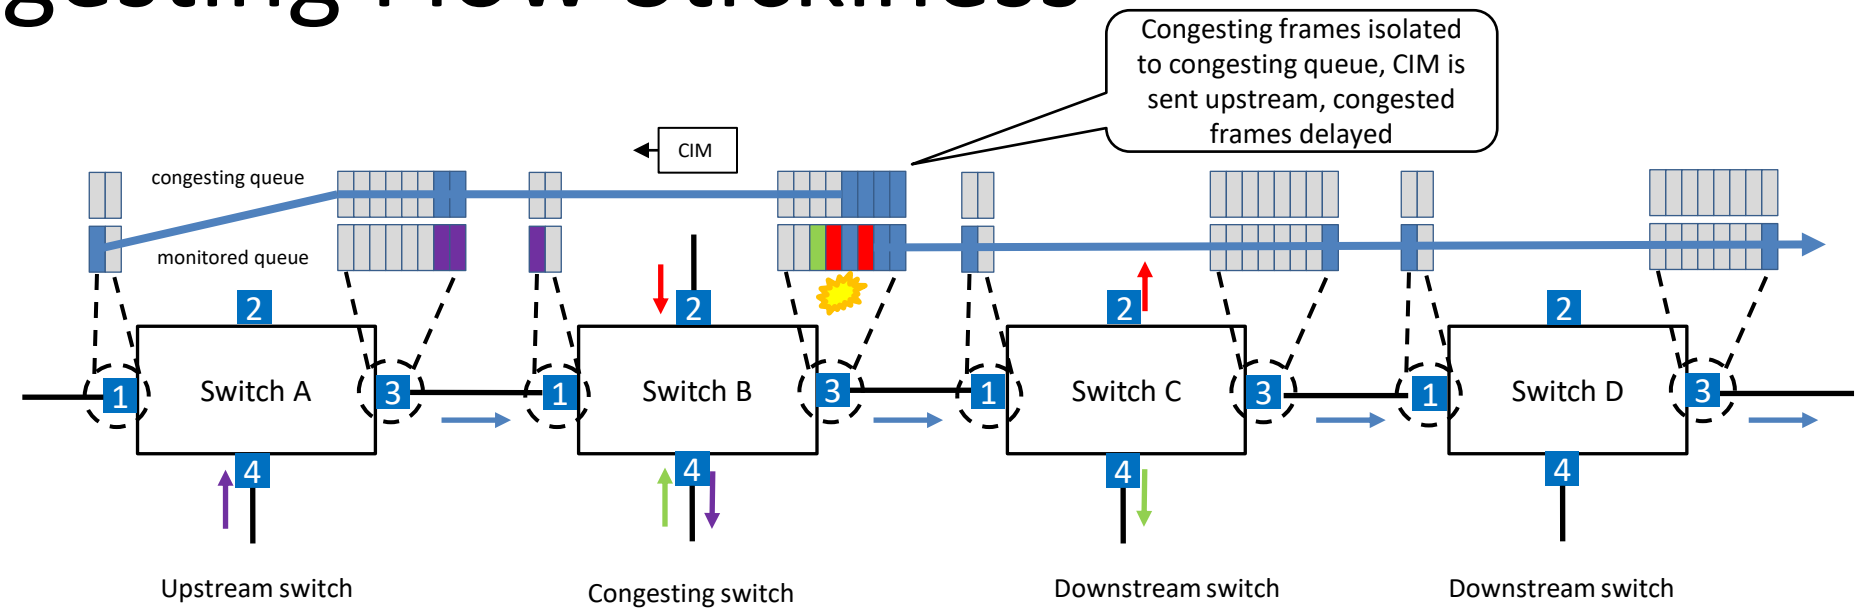

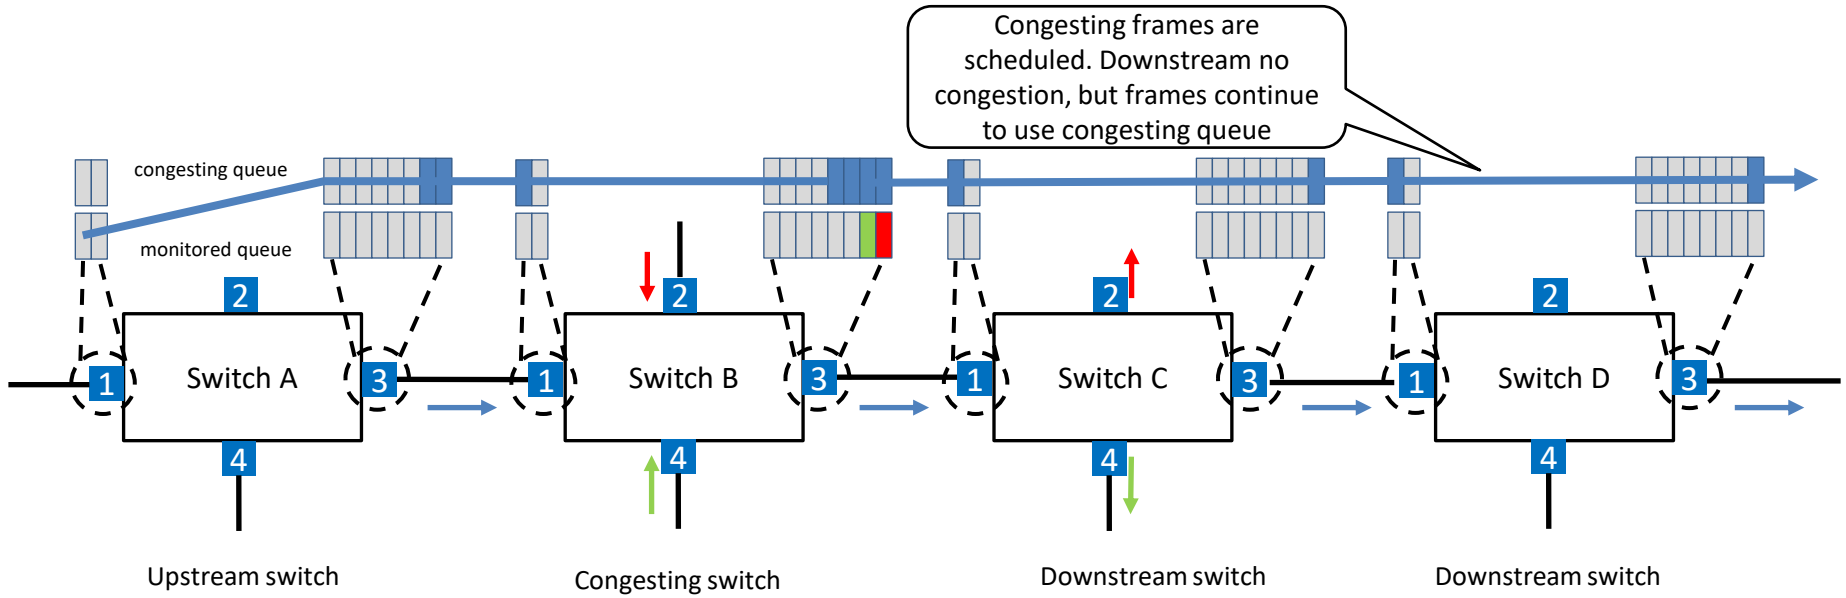

Solution to congesting stickiness

- Downstream switches should change the priority of these frames back to use the monitored queue.
  - We need a stream filter to match these frames
- Downstream switches do NOT have an entry in the stream identity table for these streams
  - Entries are created by detecting local congestion – not present downstream
  - Entries are created by receipt of CIM – not sent downstream
- We need a new stream_handle specification to match the absence of a stream_handle

Solution to congesting stickiness

- Downstream switches should change the priority of these frames back to use the monitored queue.
  - We need a stream filter to match these frames

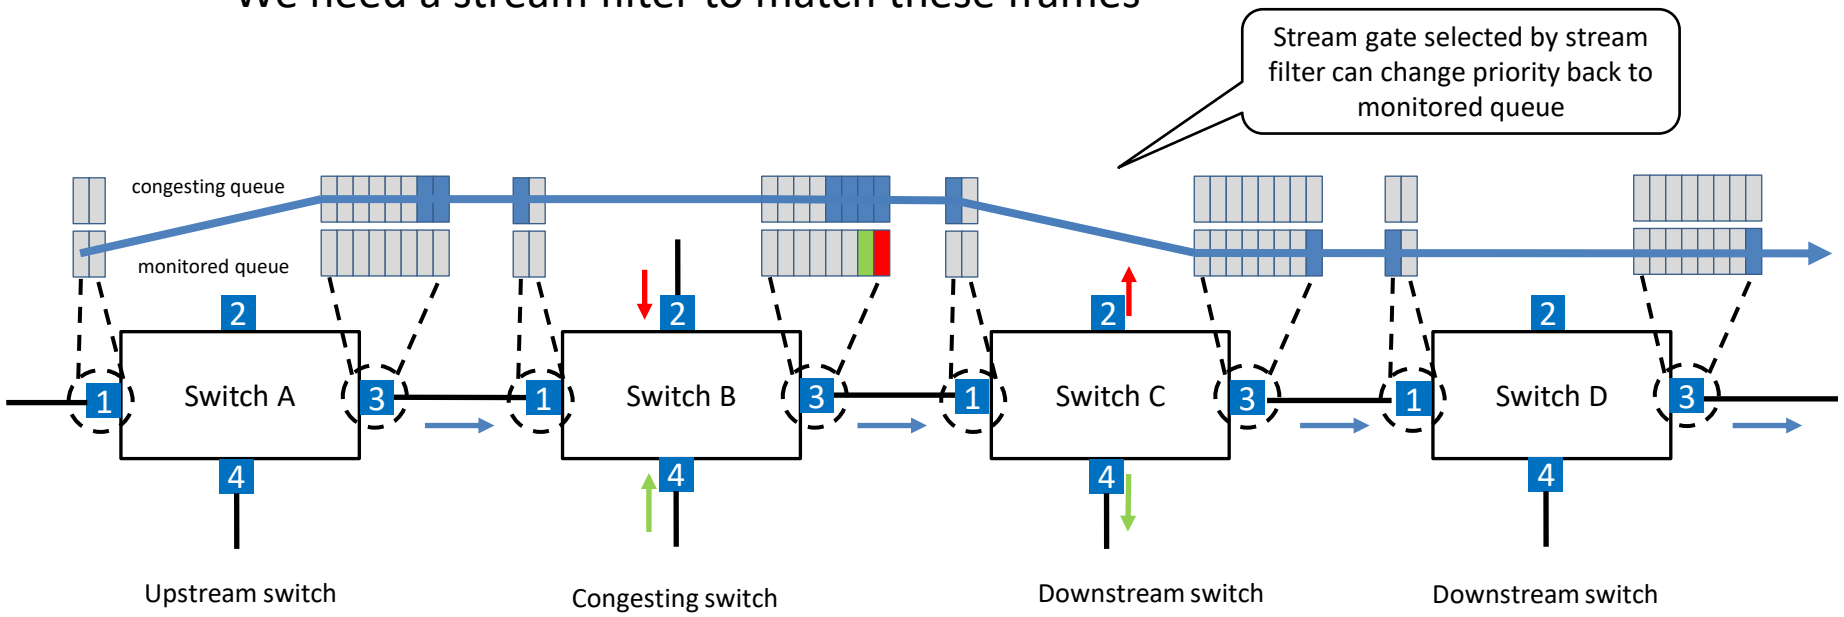

- However, downstream switches do NOT have an entry in the stream identity table for these streams
  - Entries are created by detecting local congestion – not present downstream
  - Entries are created by receipt of CIM – not sent downstream
- We need a new stream_handle specification to match the absence of a stream_handle!