



Congestion Isolation Backwards Implicit Notification

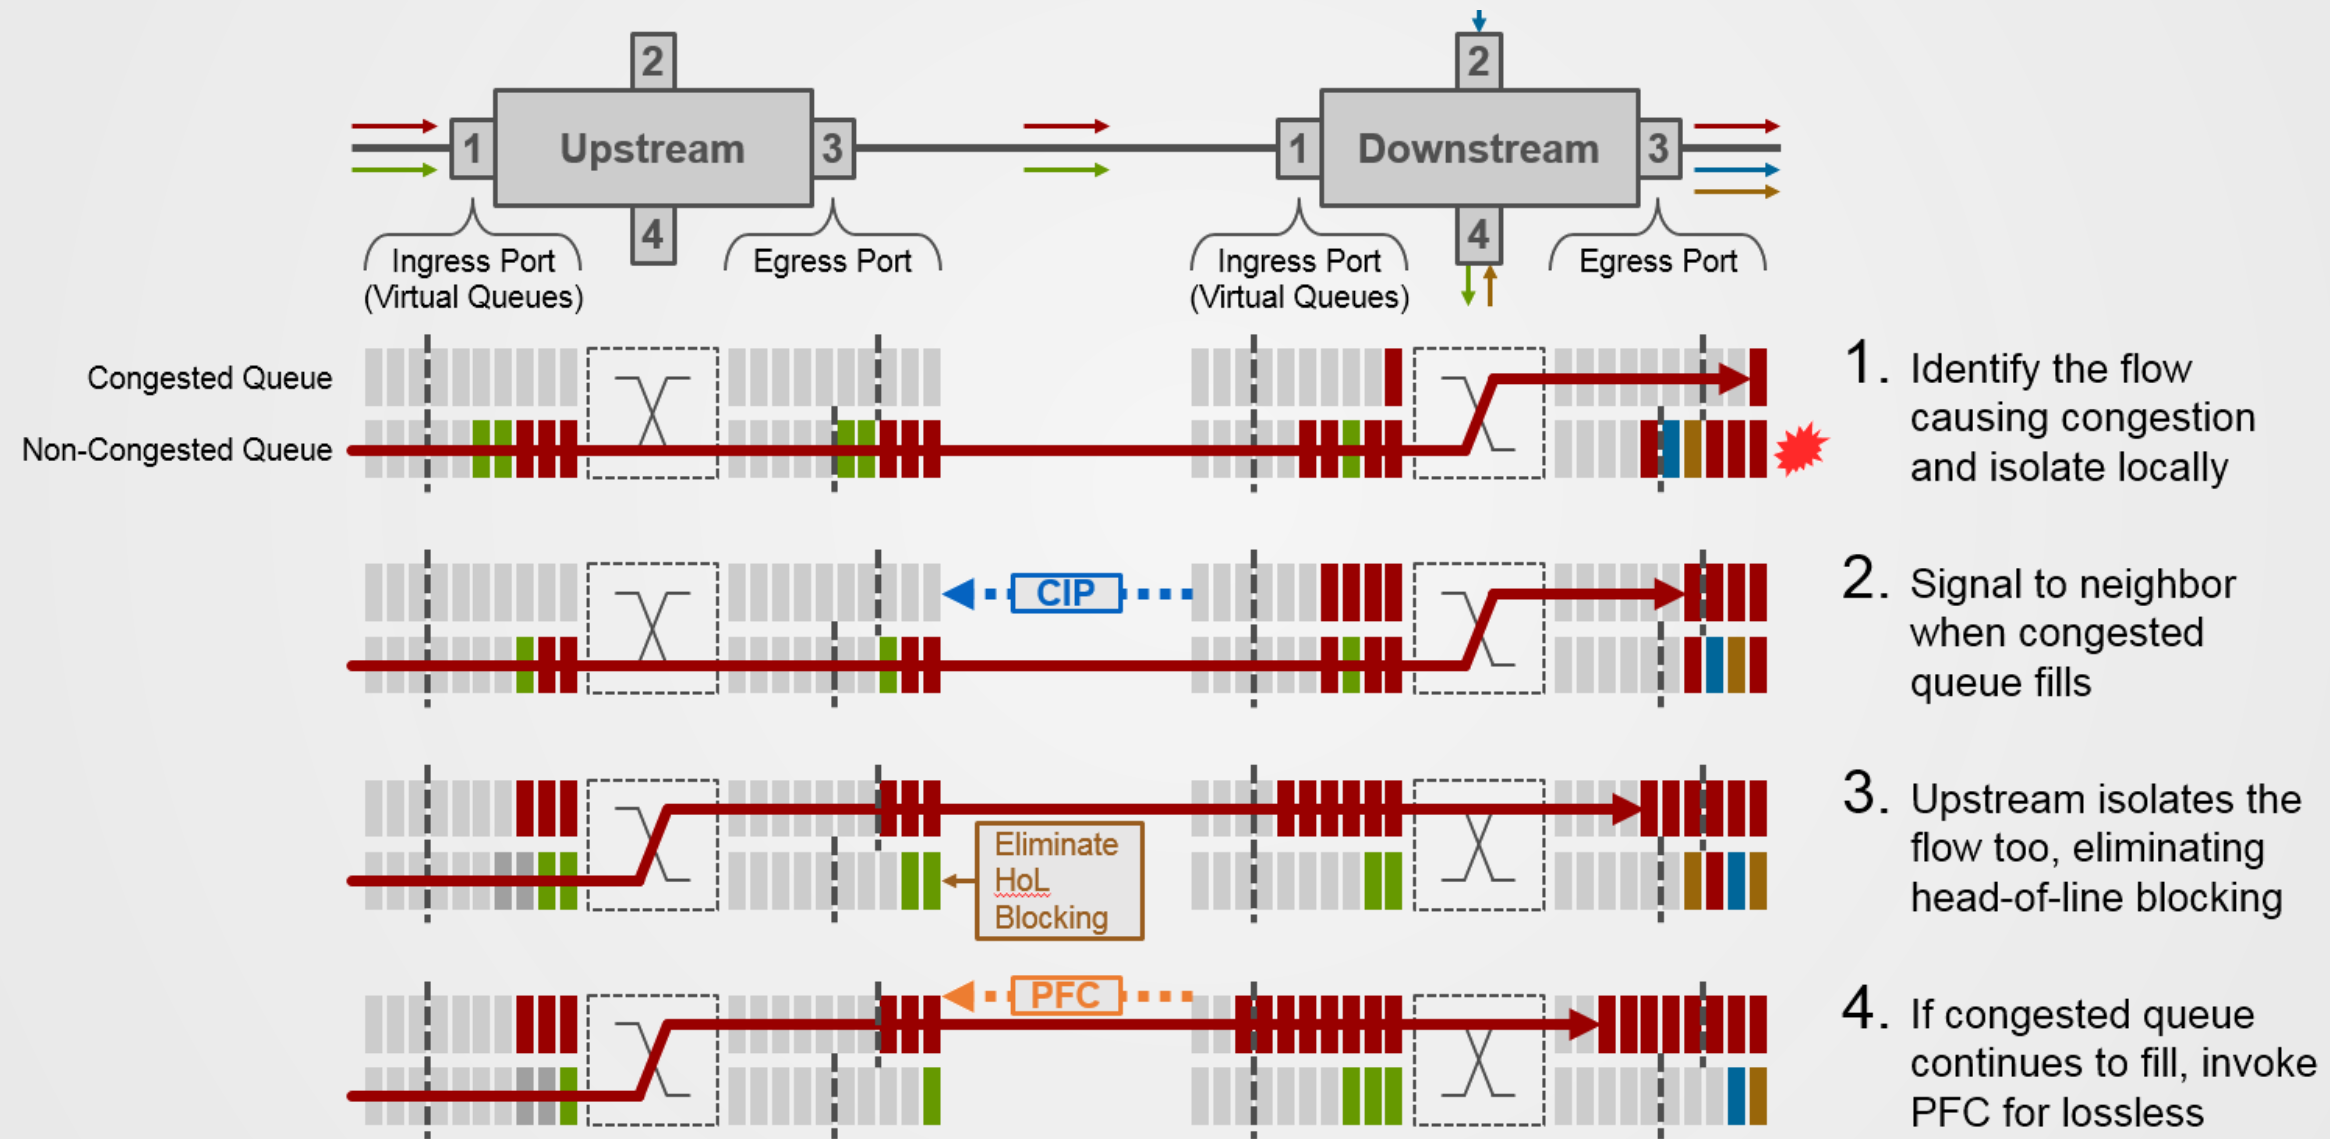
Barak Gafni

March 2018

The Need for Congestion Isolation

- As stated, some of the modern network require lossless operation in order to get the best performance for the applications
- The industry recognize situations were congestion may spread across multiple networking devices
 - As part of the architecture done in IEEE to enable these deployments
- As part of the mitigation, congestion isolation was proposed
 - This is in addition to some other technologies, such as congestion control algorithms etc.
- We believe this technology is needed in order to lower the effect of potential congestion spreading within such networks
- We get supportive feedback from operators to the statement above

Current Suggestion of the Solution



Problem Statement

- We would like to suggest that CIP may add too much complexity to the solution
- This includes, but not limited, to
 - Proprietary definitions of the flows
 - Proprietary implementations of hash functions
 - Potential race events, which may impact the implementation or the architecture

Suggested Solution – Implicit Backwards Notifications

- We would like to suggest Implicit Backwards Notifications, I.e. PFC
 - Assume two ingress buffers have been reserved in the downstream switch, as seen in previous slide
 - Once downstream switch implemented congestion isolation, it's default ingress buffer may still get over its threshold
 - In this case, “Congestion” occurs in the upstream switch, which in turn implement congestion isolation
 - Once the upstream switch recognize the big flows, it isolate them and apply new priority to them
 - This, in turn, will result in freeing the ingress buffer of the downstream switch, while using the second ingress buffer for the isolated flows
- The end result of the above technology seems to converge to the same network state as with CIP
 - One may argue about convergence time. This is indeed depends on implementation...



Thank You

