Network Infrastructure Challenges

The Switching Silicon Perspective

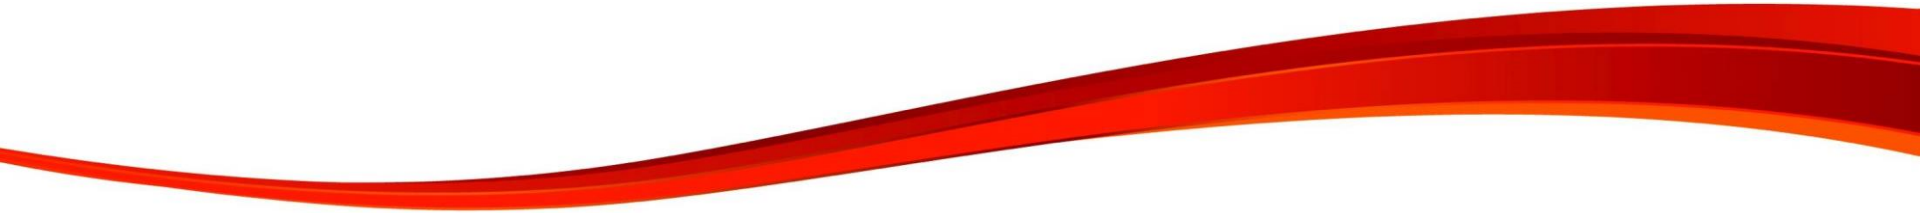

Carmi Arad
carmi@marvell.com

We will talk about

- The change in the network infrastructure
- The challenges
  - As seen by a switching silicon vendor
- Study some of the missing technologies

Network Infrastructure

Client

802.11ac Wave 2
802.11ax

HD Cameras
OTT Streaming
AI based services

More end-nodes/services

- More IP’s
- More management overhead
- More bandwidth to/from clients
- Lower latency

Cloud

- More throughput
- Lossless Ethernet
- Tight management / Analytics
- Low power
- Programmability

Edge Infrastructure

- More throughput
- Strict latency/jitter guarantees

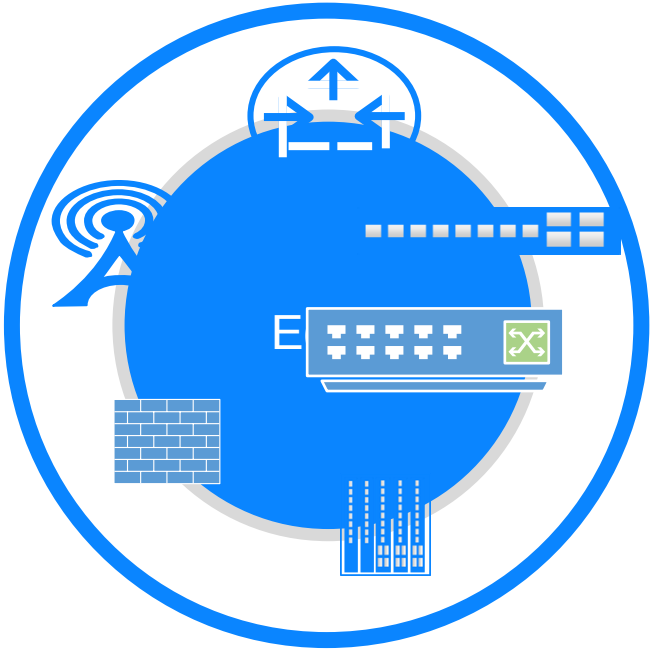

Edge Data Centers

- Tunneling
- NAT
- Throughput
- Low latency

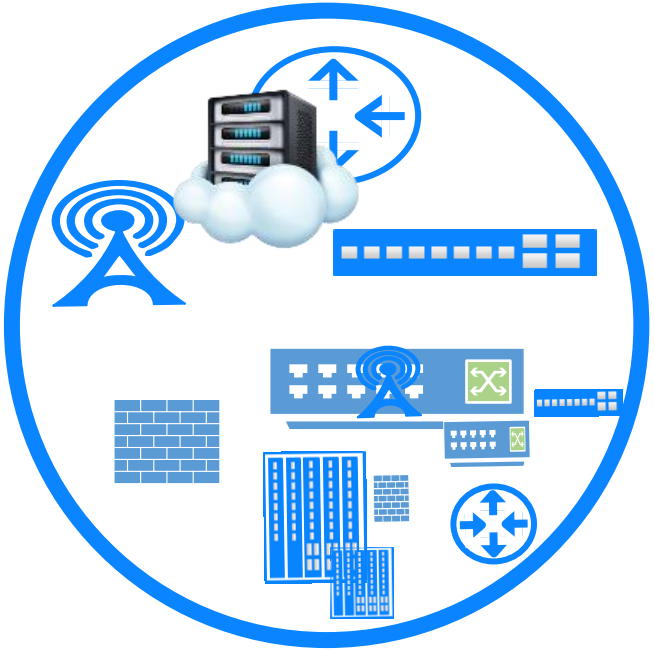

Challenges

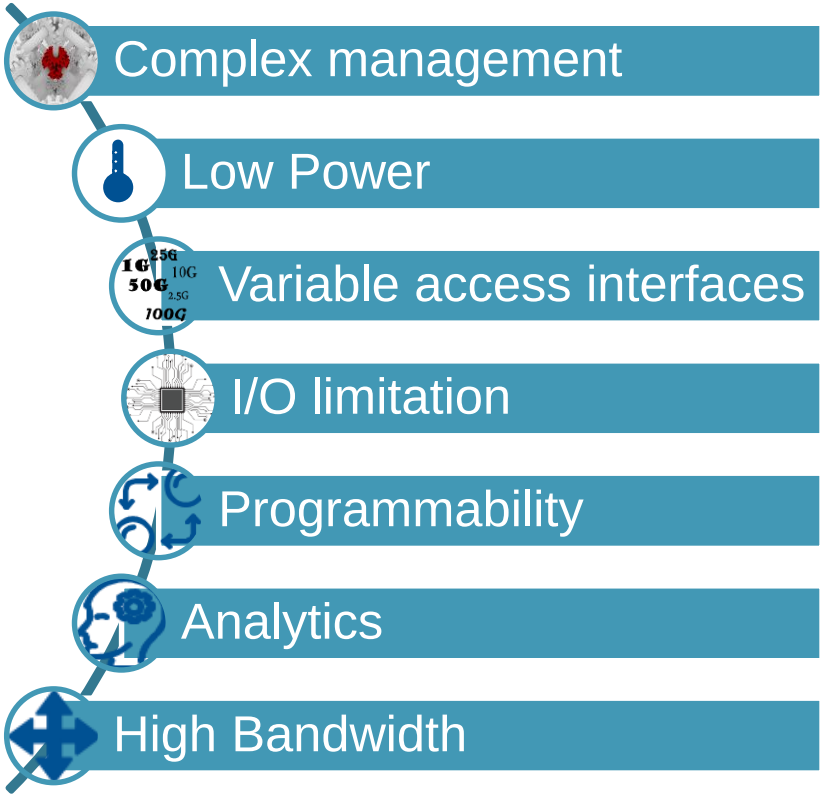
Complex management
Low Power
1G
25G
50G
10G
2.5G
100G
Variable access interfaces
I/O limitation
Programmability
Analytics
High Bandwidth

System Level Modularity

- Simpler management
- Port extender $ < Switch $
- Port extender W < Switch W
- Plug & play among vendors

- Addressed by 802.1BR and 802.1Qbg

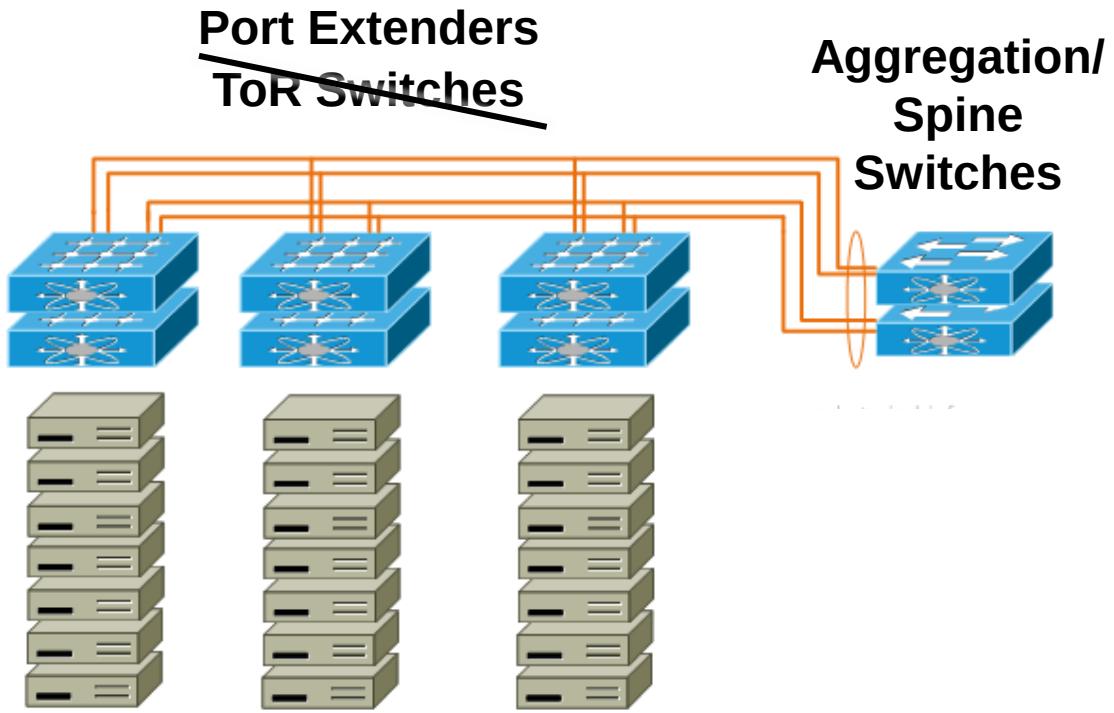

System Level Modularity

- Has the work been completed by 802.1BR/802.1Qbg?

- Centralized congestion avoidance and transmission selection
  - Reduce PE cost by reducing PE buffer size
  - Small buffer → PE can't be a congestion point → Congestion should be handled by the CB
  - How would the PE signal congestion to the CB?
    - 802.3x and PFC are blocking

- Time for standard inband channelized flow control over Ethernet?

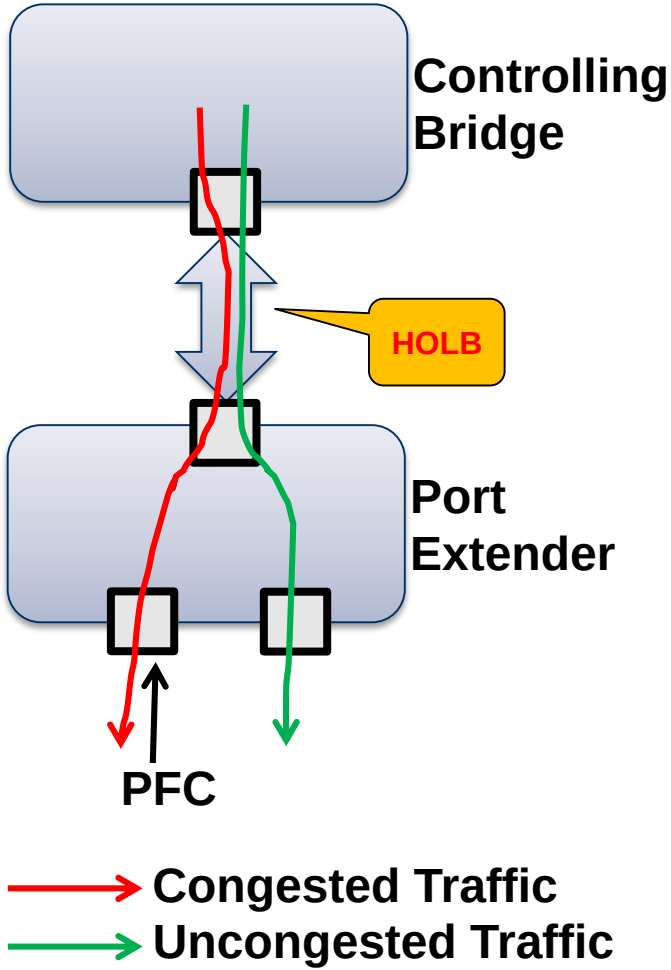

Semiconductor Reality Take #1: Silicon Is More Complex Than Ever

- Process cost
- High investment
- Analog IP
- Many and sometimes conflicting requirements

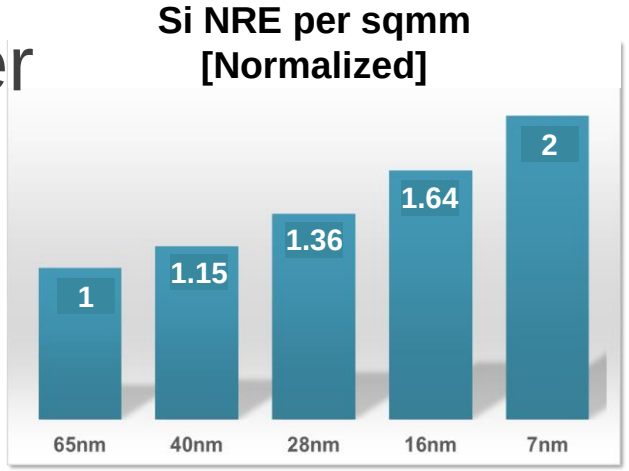

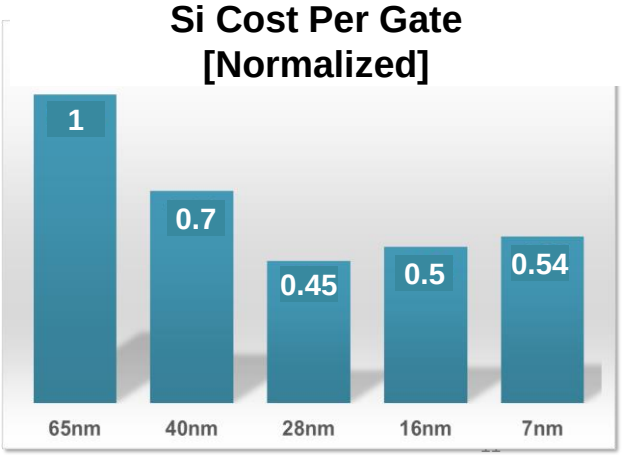

Modularity at the Silicon Level

- Separate the I/O from the Main Die
- Separate Serdes Process from Packet Processor Process
- Ability to mix and match processes
- Ability to mix and match FAB Technology

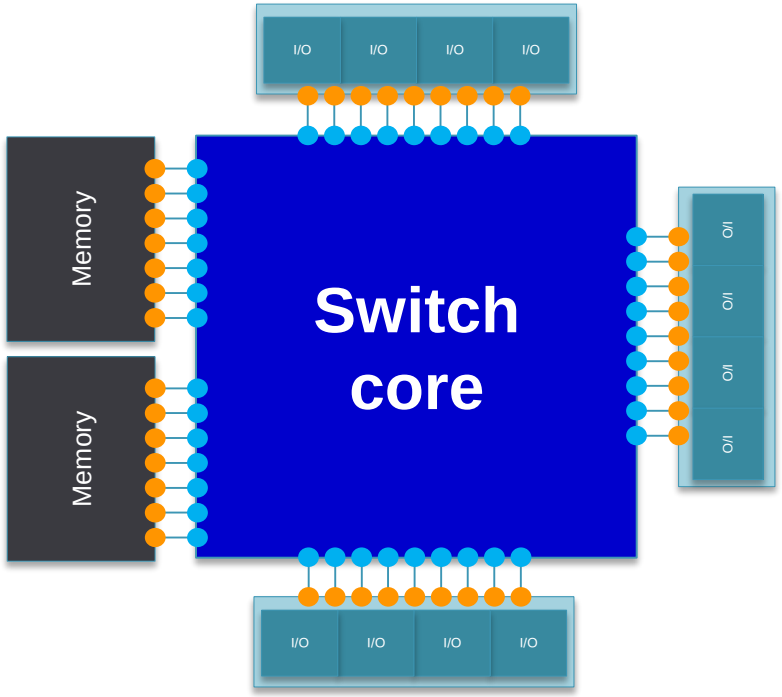

USR ALLIANCE

- USR – Ultra Short Reach Link
- Up to 500Gbps of throughput with low power and small footprint
- Designed for organic substrate and standard MCM package
- Working toward standardization of the USR interface
- Join the Alliance: www.usr-alliance.org

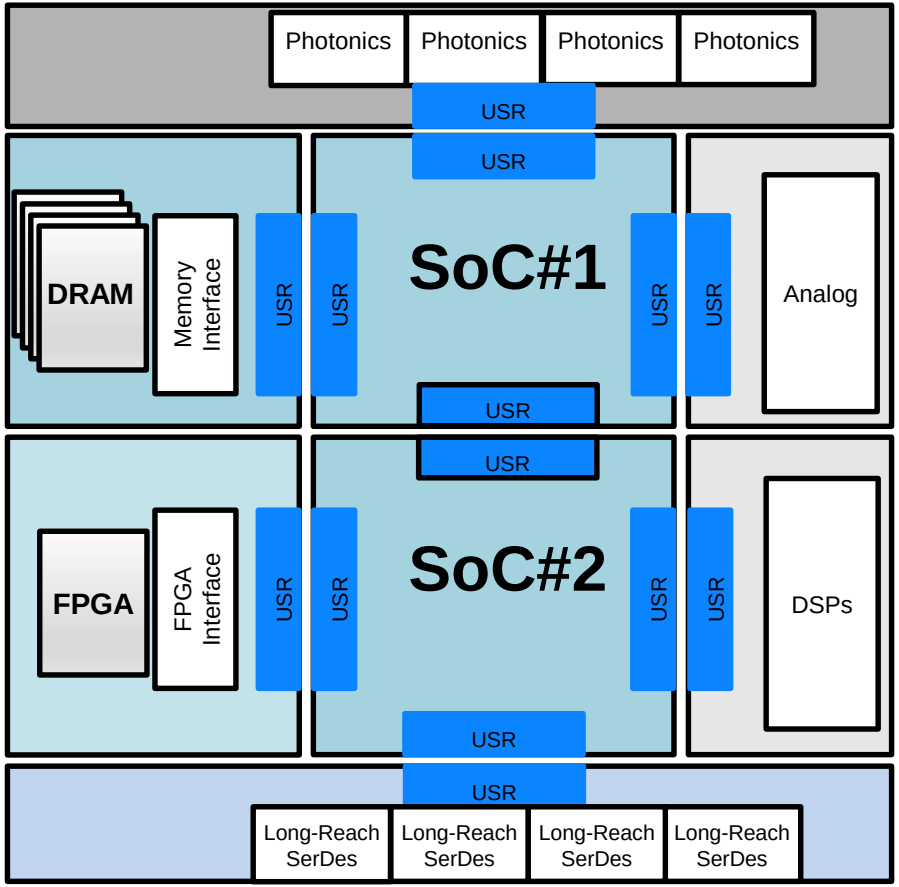

Semiconductor Reality Take #2: Switch Packet Buffer Trends (DCN Switches)

- What happens?
  - Switch throughput scale >> Si technology shrink
- Challenges
  - High throughput
  - Short average FCT
  - Lossless operation

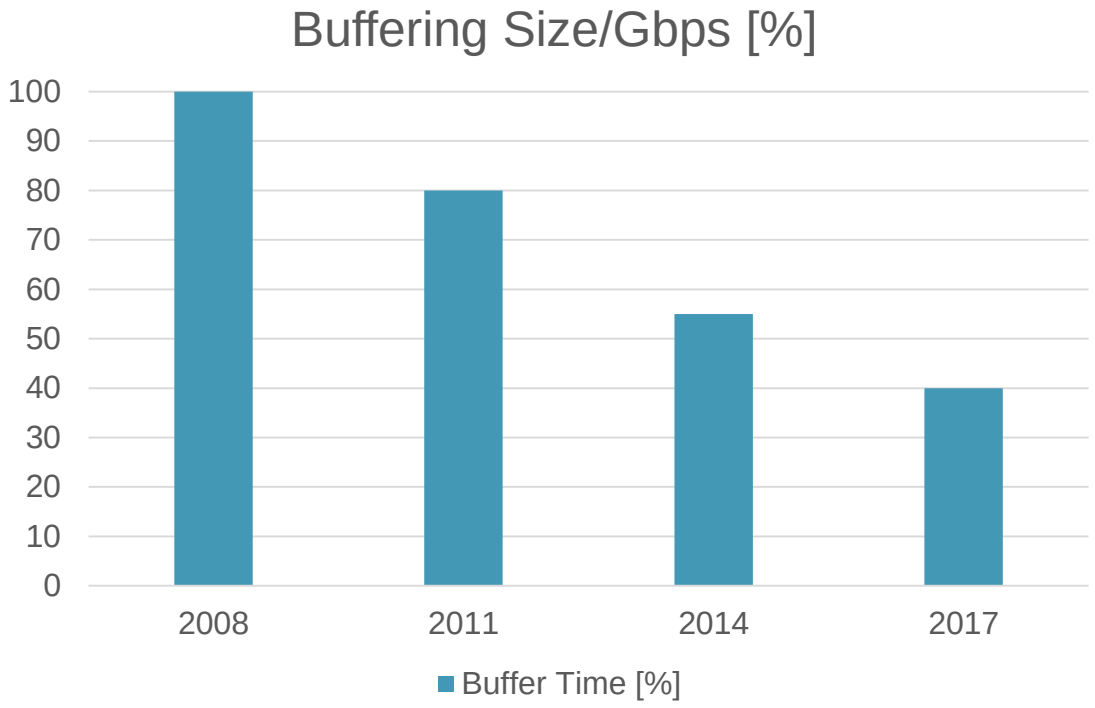

Throughput: Dynamic Load Balancing

- Many ideas
  - Flow based
  - Flowlet based
  - Packet based
- Some are implemented…
- But all are proprietary/ not interoperable
  - Congestion sensing
  - Congestion signaling
  - Load balancing policy
  - Capability exchange

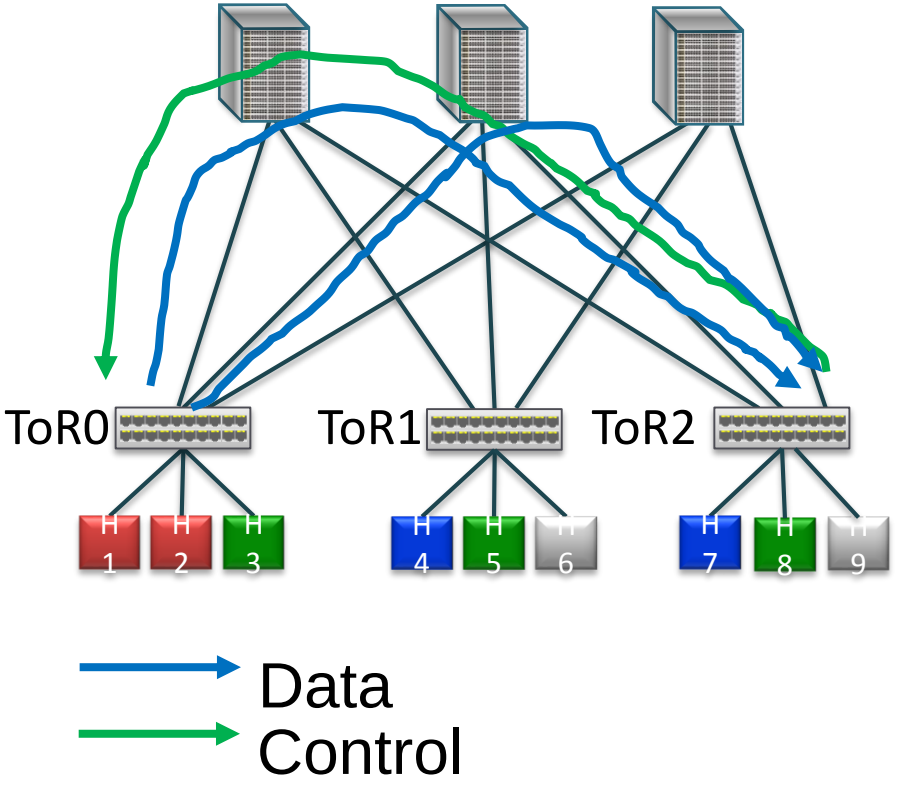

Short FCT in Lossless Networks: Congestion Isolation

Congested Flow
Victim Flow
Congested Queue
Normal Queue

Today – Without Congestion Isolation

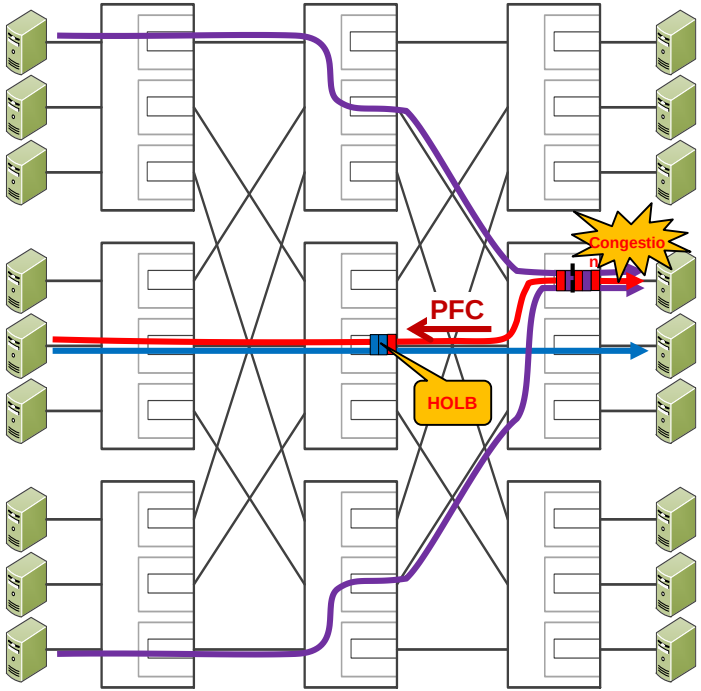

Congestion Isolation

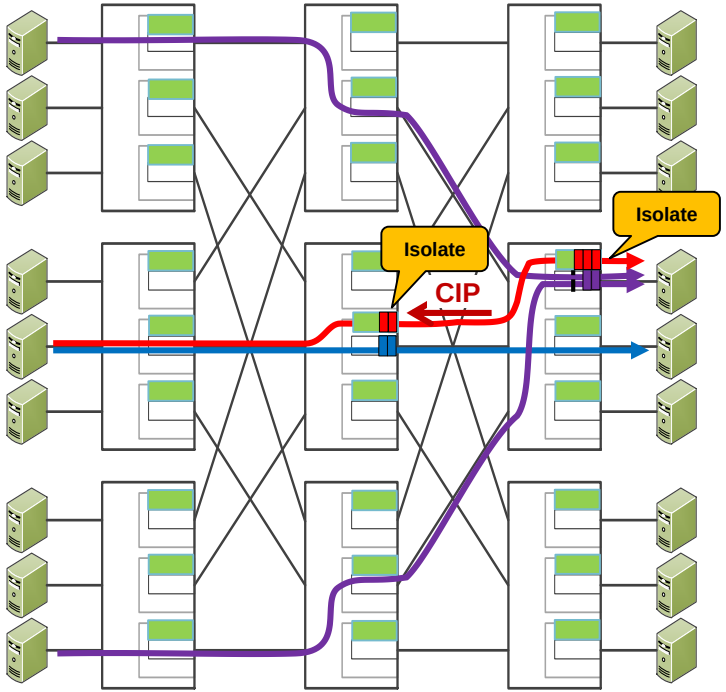

MARVELL

Summary

- Network infrastructure is evolving and introducing new challenges
- System/Si modularity is driven by requirements for
  - High throughput +
  - Low cost/power +
  - Simpler management
- Sophisticated congestion avoidance/control is driven by
  - Reduced buffering time in switches
- State-of-art is proprietary

Thank you for listening

This is the time for questions ☺


MARVELL

The information contained in this presentation is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided “AS IS”, without warranty of any kind, express or implied. This information is based on Marvell’s current product roadmap, which are subject to change by Marvell without notice. Marvell assumes no obligation to update or otherwise correct or revise this information. Marvell shall not be responsible for any direct, indirect, special, consequential or other damages arising out of the use of, or otherwise related to, this presentation or any other documentation even if Marvell is expressly advised of the possibility of such damages. Marvell makes no representations or warranties with respect to the contents of the presentation and assumes no responsibility for any inaccuracies, errors or omissions that may appear in this presentation.