

Requirements of the 'Data-Centric' Computing Network

Paul Congdon

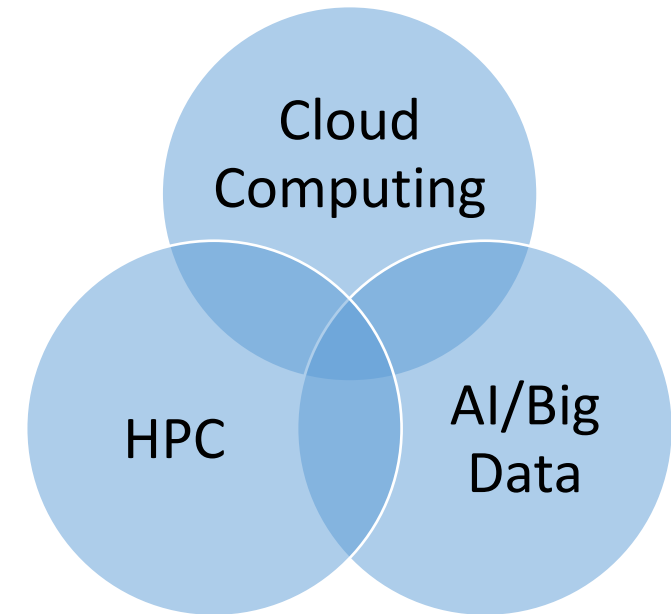
The 'Data-Centric' era is pushing change

Our daily lives are changing from the combination of mobility, cloud computing, high performance computing and AI/Big Data.

Their cause and effect has brought an explosive growth of data.


The exponential growth of data is forcing the evolution of computing systems to 'data-centric' computing systems.

- High performance is essential for storing, moving and processing data.
- Computing systems require breakthroughs in processing power, storage performance and network connectivity.



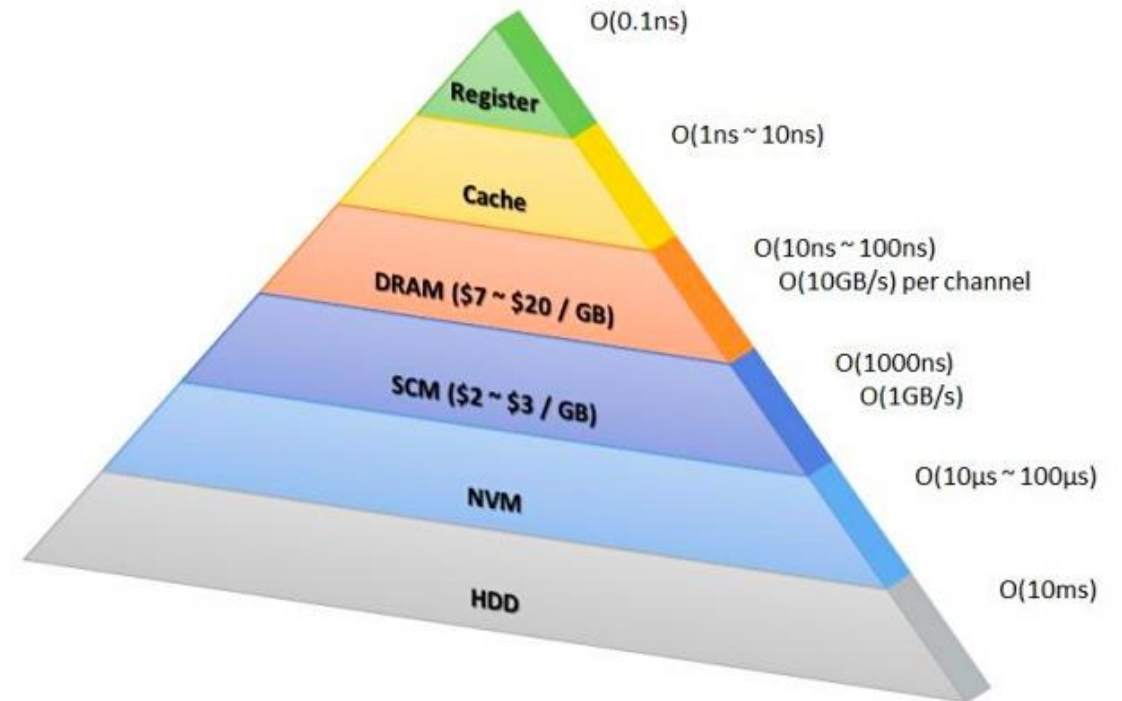
'Data-Centric' is enabled by Data

Big Data



IDC predicts that the Global Data will grow from 45 Zettabytes in 2019 to 175 Zettabytes by 2025

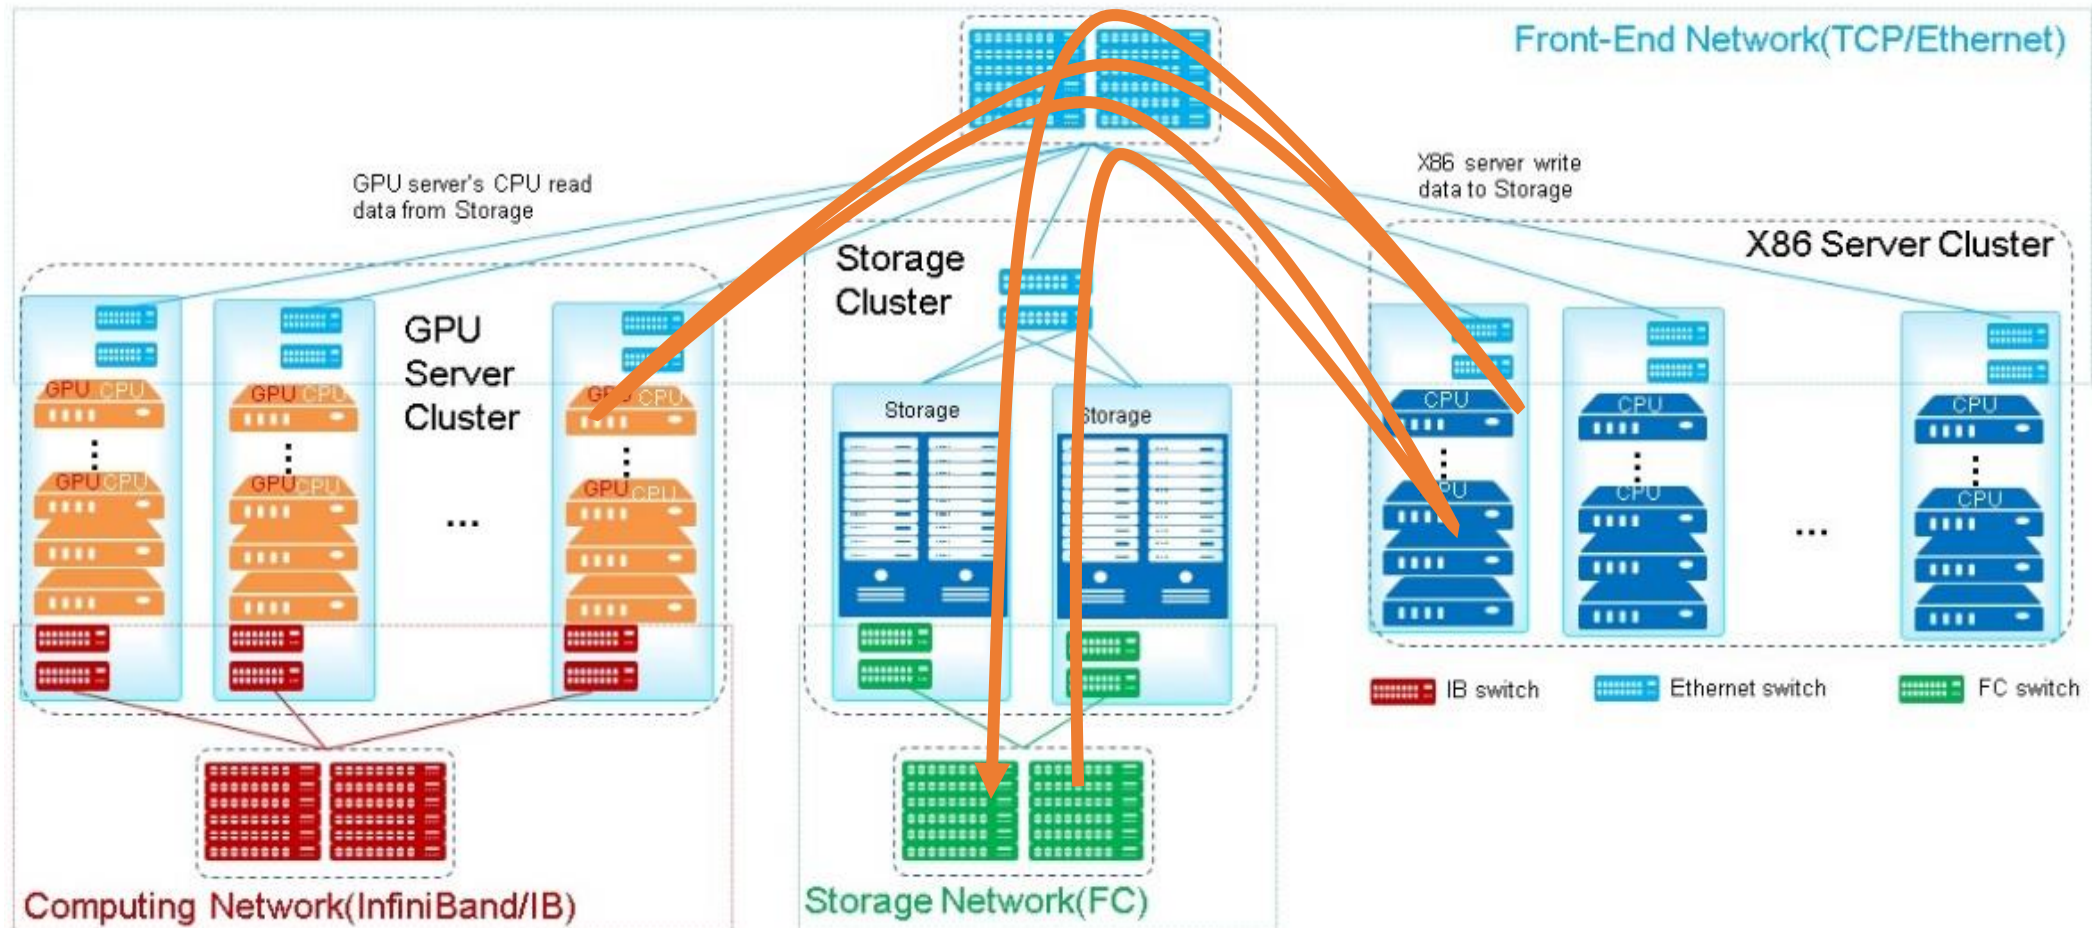
Fast Data



Persistent storage latencies are approaching memory latencies with the latest Storage class memory (SCM) technology. Accessed over the network using NVMeoF

The traditional approach can't keep up


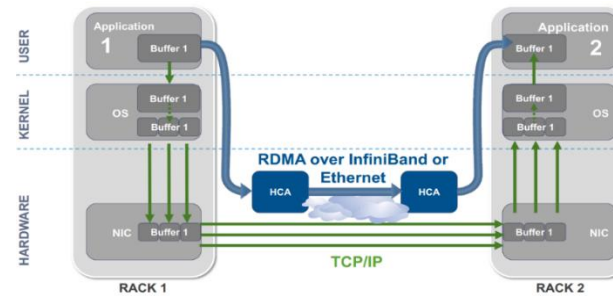
Moving the data to the compute doesn't scale or perform



New Data-Centric Technologies stress the Network

RDMA

Reduce memory bandwidth, CPU cycles and read/write time with efficiency and zero copy at end-points. **E2E latency** is mainly contributed by network.




(c) Mean Latency

(Sigcomm 2015: Congestion Control for Large-Scale RDMA Deployments)

NVMe


Faster storage media has higher IOPS and lower latency. **Network latency** become the biggest portion in the total latency.



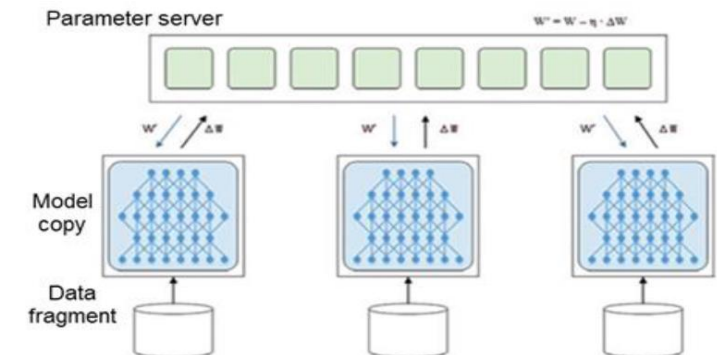
Distributed HW with Parallel SW

Distributed infrastructure using Data Processing Units (DPUs/SmartNICs) and parallel software architectures increase communication and require data transmission with **lower tail latency**.

Distributed Infrastructure with DPUs



Parallel AI Software



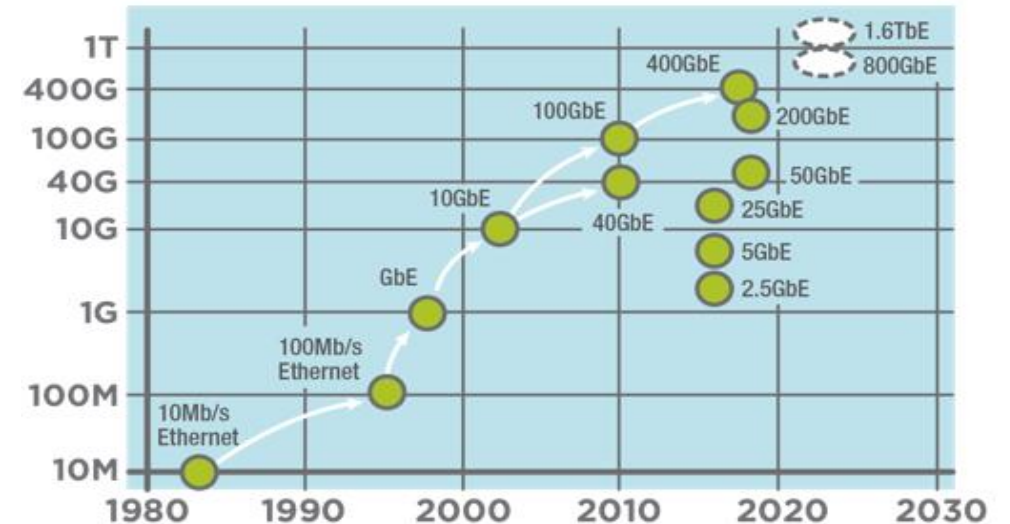
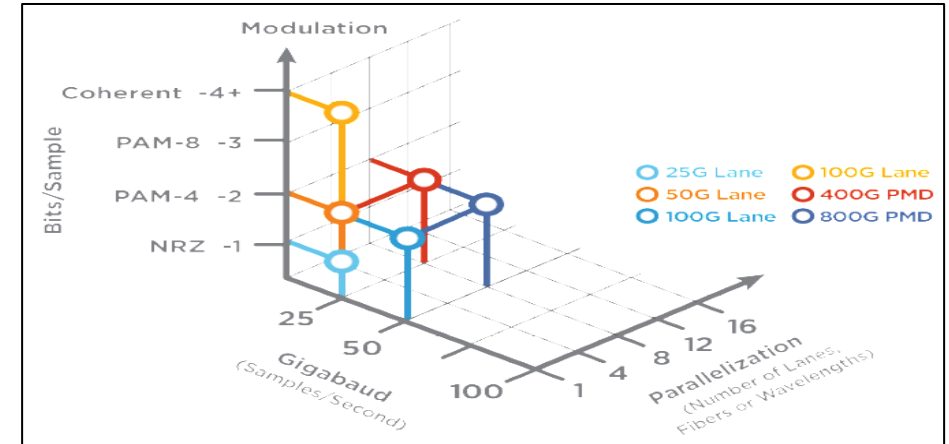
Ethernet for the 'Data-Centric' Network Fabric

Ethernet is **ubiquitous technology**.

- Cost-effective solution
- Relatively easy to deploy and manage
- Leading technology development

Ethernet provides **large bandwidth** connectivity

- up to 400 Gbps, 100G for single lane
- towards 800 Gbps, 200G for single lane




(Ethernet Alliance 2020 Roadmap)

Ethernet is the popular choice

Although dedicated protocols are designed for computing and storage networks, like IB for high performance compute clusters, FC storage networks, **people see the appeal of Ethernet due to Ethernet network advantages.**

- In **TOP500**, which is a list of the world's 500 most powerful computer systems, 51% supercomputers use high performance Ethernet fabrics, more percentage than IB and other proprietary technologies.

Interconnect Family System Share

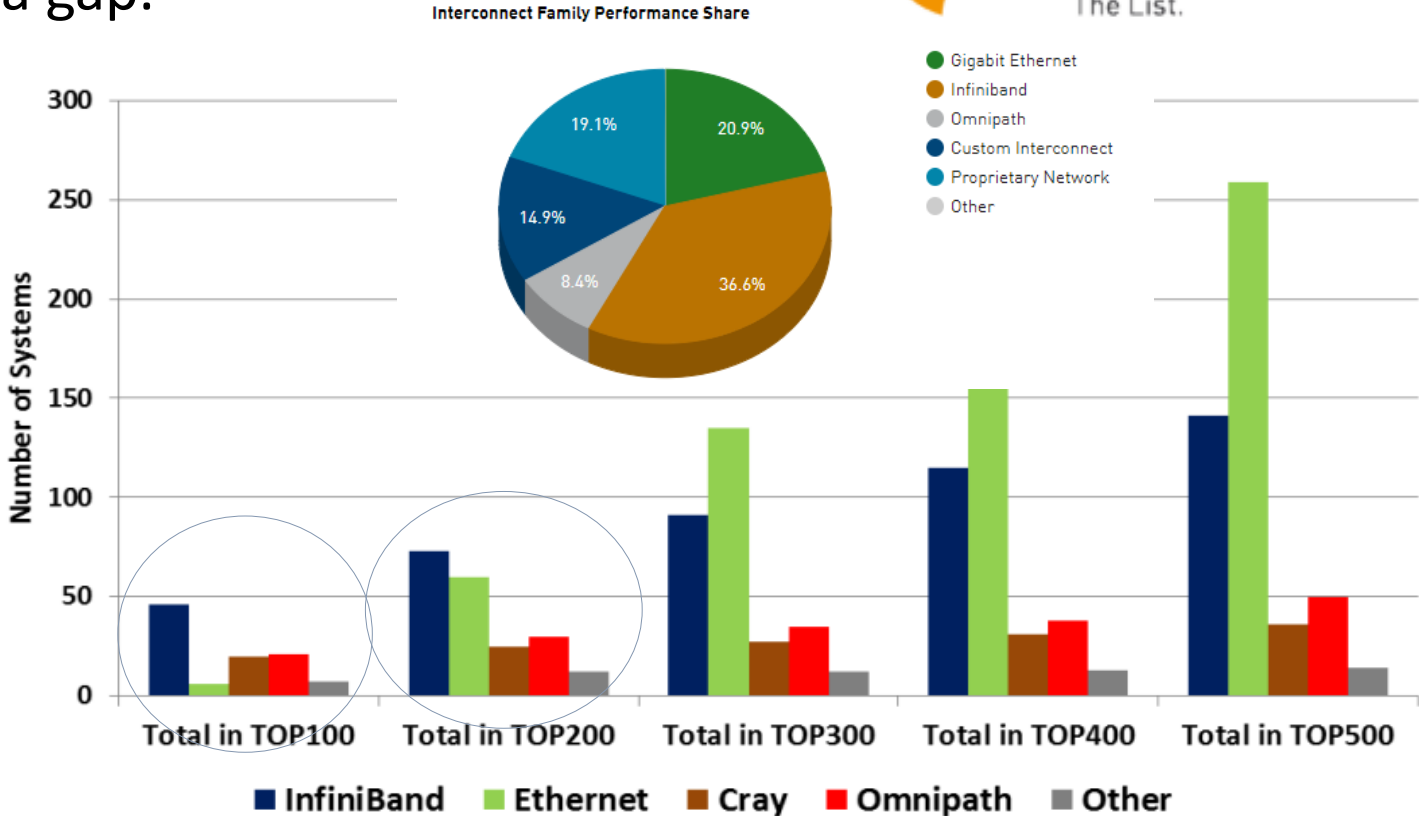


However, Improvements are still needed

In the Top500, the Ethernet advantage is most likely due to price-performance.

In terms of performance, there is still a gap.

- In Top500 performance statistic, Ethernet is only 21%.
- In Top100, Ethernet is far behind IB.



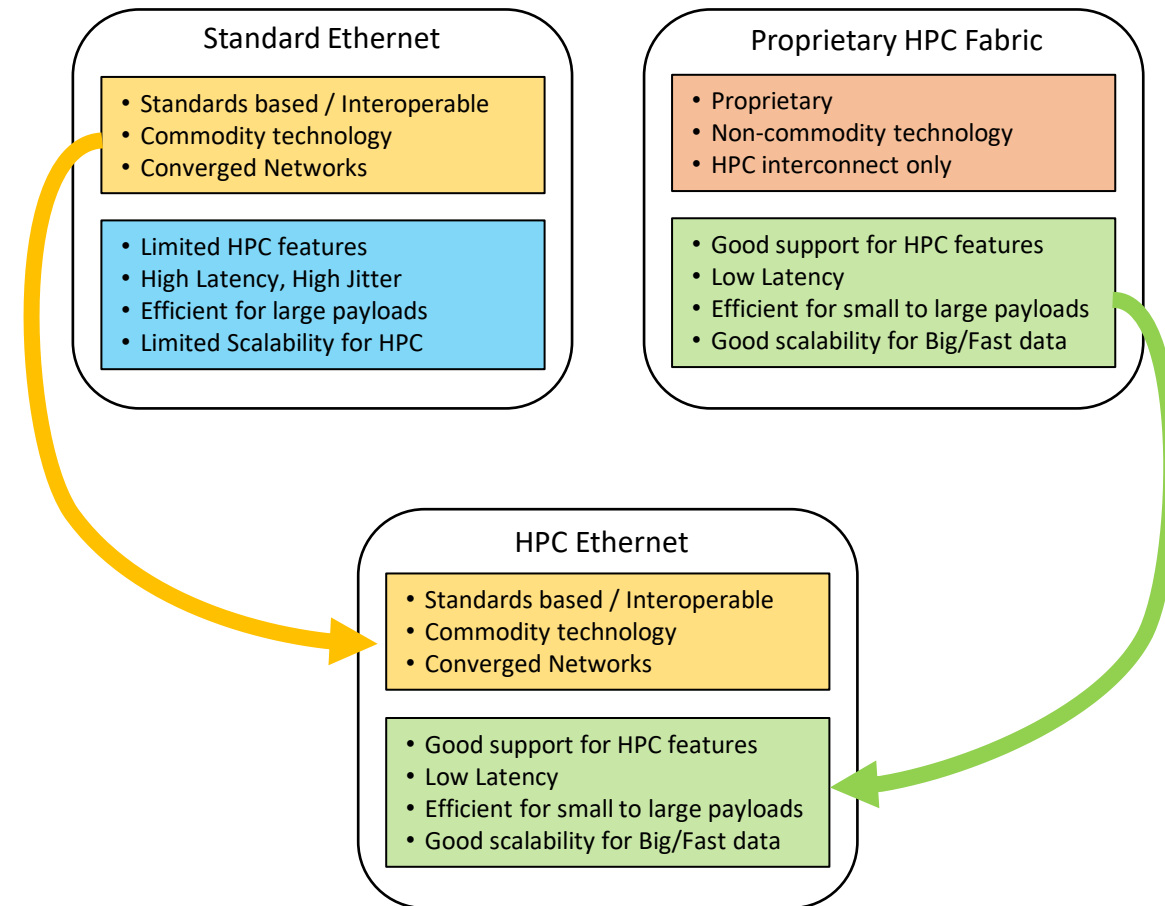
Perceived Infiniband/Custom fabric advantages

Some claims for why Infiniband is superior

- Guaranteed delivery at the HW level. Credit based flow control
- Hardware based re-transmission. Higher throughput
- Better lossless congestion management
- Cut through design with late packet invalidation
- 16 Virtual Lanes vs 8 Traffic Classes in Ethernet
- Preselected failover paths and switches for instant recovery
- Lower cost IB switch chips due to technical differences
- References
 - <http://www.informatix-sol.com/docs/EthernetvInfiniBand.pdf>
 - <https://www.slideshare.net/PatrikKristel/design-cloud-system-with-infiniband-or-ethernet>
 - <http://people.engr.ncsu.edu/efg/506/sum06/lectures/notes/lec12.ppt>

Some known proprietary tweaks to Ethernet for HPC

- Per-flow, credit basis congestion control
- Reduced minimum frame size (< 64B) with local addressing
- Auto-negotiation between Standard Ethernet and HPC Ethernet features
- Low-latency FEC, link-level retry to tolerate transient errors
- References:
 - <https://www.nextplatform.com/2019/08/16/how-cray-makes-ethernet-suited-for-hpc-and-ai-with-slingshot/>
 - <https://infoshri.com/viral-news/inside-rosetta-the-engine-behind-crays-slingshot-exascale-era-interconnect/>



Why not TSN for Data-Centric Computing Network?

- What is in common?
 - Desire for zero congestion loss
 - Lower latency with no tail
 - Deterministic performance
- What is different with the data-centric computing network?
 - Larger network scale
 - Higher network bandwidth
 - Unpredictable traffic patterns, not able to plan or configure in advance
 - Applications can tolerate small variances in latency and throughput
 - Workloads change frequently


DCB progressed data center ethernet

Standard	Description	Contribution
802.1Qau-2010	Congestion Notification	Layer-2 end-to-end congestion control
802.1Qaz-2011	Enhanced Transmission Selection	Bandwidth sharing between traffic classes
802.1Qbb-2011	Priority-based Flow Control	Lossless traffic classes
802.1Qcz-2021	Congestion Isolation	Avoid head-of-line blocking

However, more is needed...

- Avoid incast congestion
- Fine grain load balancing
- Further reduce network latency

Network Latency Breakdown



Dynamic Network Latency (>99%)

■ Queuing latency (50 μ s)

■ Packet loss latency (5000 μ s)

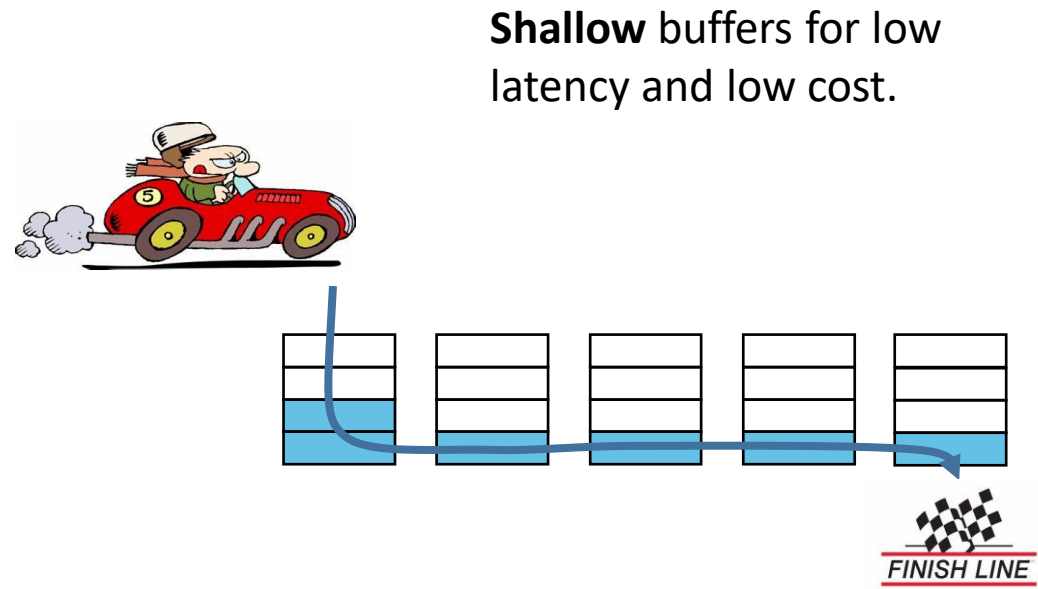
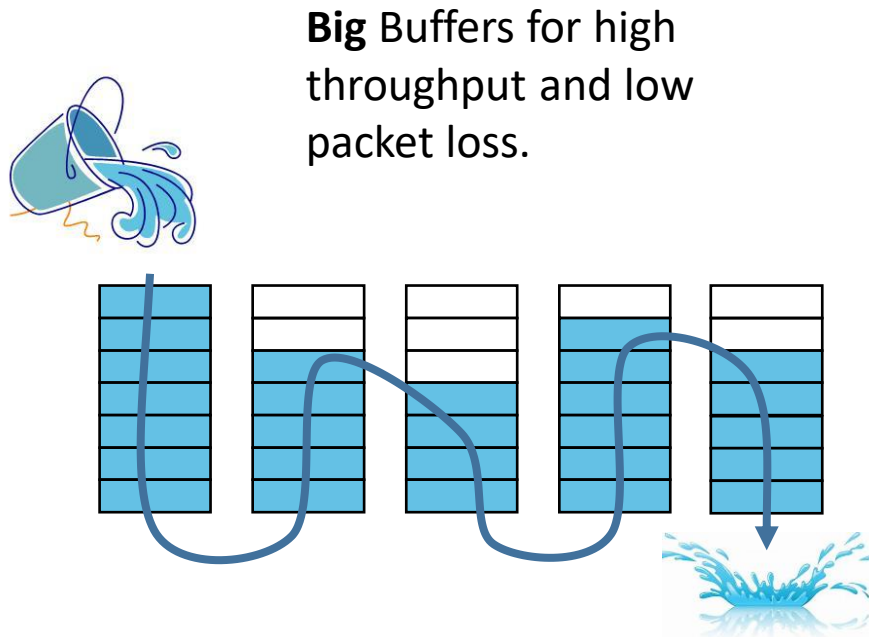
Static Network Latency (<1%)

■ Switching latency (3 μ s)

■ Transmission latency (0.3 μ s)

Addressing Dynamic Latency

- Avoid packet loss by deploying **Big** buffers everywhere
- However, to minimize queuing delay and lower latency, deploy **Shallow** buffers in switches




Shallow Buffers are important to ASICs

Switch capacity keeps increasing. However, the buffer size of commodity switching chips does not keep pace. **The bandwidth capacity significantly outpaces the buffer size.**

Reasons for shallow switch buffers.

(APNet'17, Congestion Control for High-speed Extremely Shallow-buffered Datacenter Networks)

- Switch buffers use high-speed SRAM. Compared to DRAM, SRAM is more expensive as it requires more transistors.
- The area increases with the memory size. When the area becomes large, the read/write latency will increase, making the memory access speed hard to match the link speed.
- The area increases with the memory size. Therefore, the power increases and the yield decreases.



Hardware trends for top-of-the-line data center switches manufactured by Broadcom. Buffer size is not keeping up with increases in switch capacity.

(BS '19: Backpressure Flow Control)

It's hard to throw memory at the problem!

Research in Congestion Management to Reduce Latency

Fan-in Congestion


In oversubscribed network topology, downlink bandwidth of TOR is bigger than uplink bandwidth. The traffic exceeding uplink bandwidth will cause congestion on the uplink egress port.

In-network Congestion

Improper load balancing of flows can oversubscribe an egress port in the Spine of the network fabric.

In-cast Congestion

While multiple senders transmit data to the same receiver, several flows converge from different path to the same TOR egress port, which exceeds capacity of egress port.



- Three types of congestion occur in the computing network:
 1. Fan-in
 2. In-network
 3. In-cast congestion
- Research in congestion management aims to reduce mean latency and/or tail latency.
- **Keeping buffer utilization low involves better coordination with transports.**

Sender-Driven Transport

Principle --- Push



Sender transmissions probe the channel for maximum available bandwidth and adjusts transmission rate based on feedback from receiver and the network..

Typical Protocol:

DCTCP (Sigcomm 10 Data Center TCP (DCTCP))

DCQCN (Sigcomm 15 Congestion Control for Large-Scale RDMA Deployments)

Conceptual Queue Depth Illustration



Pros

Well-known mature solution

Cons

Depends on creating and detecting congestion. Untimely adjustments cause queue depth, throughput and latency fluctuations. Congestion signaling delays

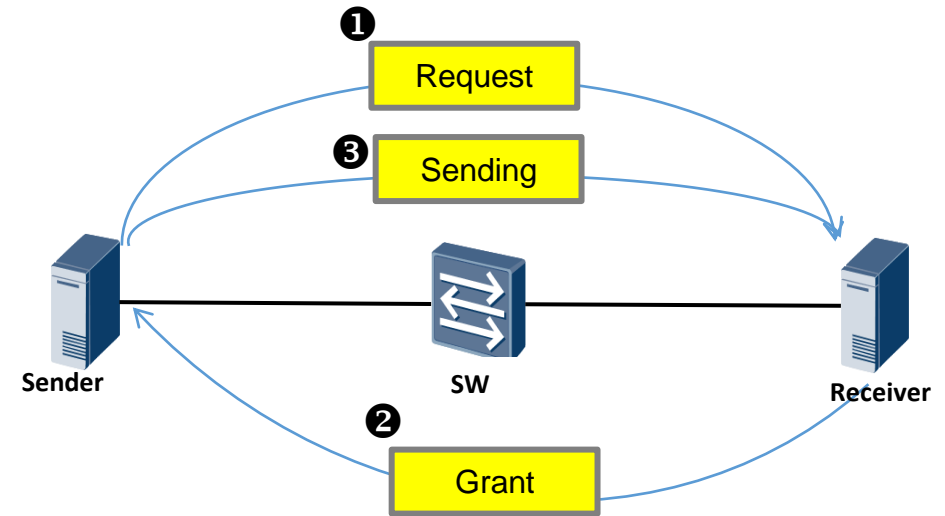
Receiver-Driven Transport

Principle --- Pull

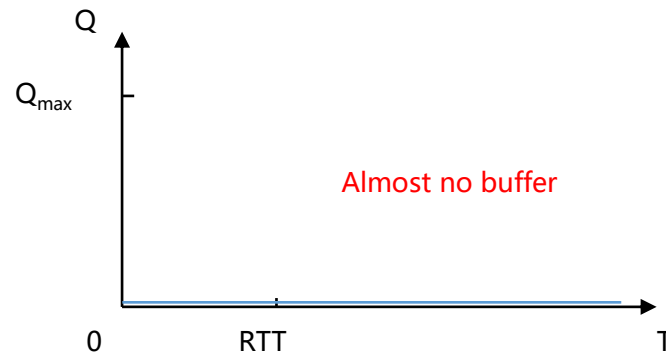
Sender transmissions are paced by the receivers scheduling. Request-Grant, credit-based protocol.

Typical Protocol:

ExpressPass (Sigcomm 17 Credit-Scheduled Delay-Bounded Congestion Control for Datacenters)



Conceptual Queue Depth Illustration



Pros

Fully utilize network bandwidth. Avoids incast congestion. Supports minimum buffering in the network

Cons

Requires 1st RTT which wastes throughput, increases latency and is unfriendly to the dominate number of small flows.

Hybrid-Driven Transport

Principle --- Push and Pull



Push initial data into network (normally in 1st RTT), then receiver schedules the remainder of the data. The initial pushed data is 'unscheduled', and the pulled data is 'scheduled'.

Typical Protocol:

NDP (Sigcomm 17 Re-architecting datacenter networks and stacks for low latency and high performance)

HOMA (Sigcomm 18 Homa: A Receiver-Driven Low-Latency Transport Protocol Using Network Priorities)

Conceptual Queue Depth Illustration



Pros

Increase throughput and lower latency with low buffer utilization

Cons

Unscheduled traffic creates jitter and uncertainty with latency and packet loss.

Next Steps

1. Call for interest of 802.1 support for lowest possible latency in the 'Data-Centric' computing network
 - How do bridges support congestion management research to reduce latency?
 - What are the actual deficiencies of Ethernet in the 'Data-Centric' computing environment?
 - What projects are within scope ensure Ethernet application in the 'Data-Centric' computing network?
2. Continue socializing ideas and expanding eco-system of interested parties in ensuring Ethernet success in the 'Data-Centric' computing network.
3. Articulate technical details of specific proposals in future contributions.