How Are The Networks Coping Up With Flash Storage

Saurabh Sureka

Sr. Product Manager,

Emulex, an Avago Technologies Company

www.emulex.com

Goals

- Data deluge – quick peek
- The flash landscape
- Fibre Channel: Roadmap, challenges and solutions
- Ethernet: Roadmap, challenges and solutions
- Q&A

The Data Deluge

- 2004 to 2014: Annual disk storage shipment
  - 1.5 Exabyte (EB) to ~100 EB[1]
  - 1 EB = 1 Giga GB
  - Flash is a toddler (accelerated growth) in memory age

- Google data network:
  - Serves 3.5 billion search queries per day [2]
  - YouTube has 300 hours of video content uploaded per hour [2]
  - Modest 480p frame, 5MB/min – requires 50 petabytes per year

[1] IDC WW Quarterly Enterprise Storage Systems Tracker, 1Q2015, June 2015

The Flash Landscape

Where there be a memory slot, there be thy Flash: thy savior

- Hard drive SSD
- PCIe SSD
- DDR DIMM
- Flash DRAM
- Flash arrays (hybrid, all flash) - Network Attached Storage (NAS) / Storage Attached Network (SAN)

Networking Storage: Fibre Channel

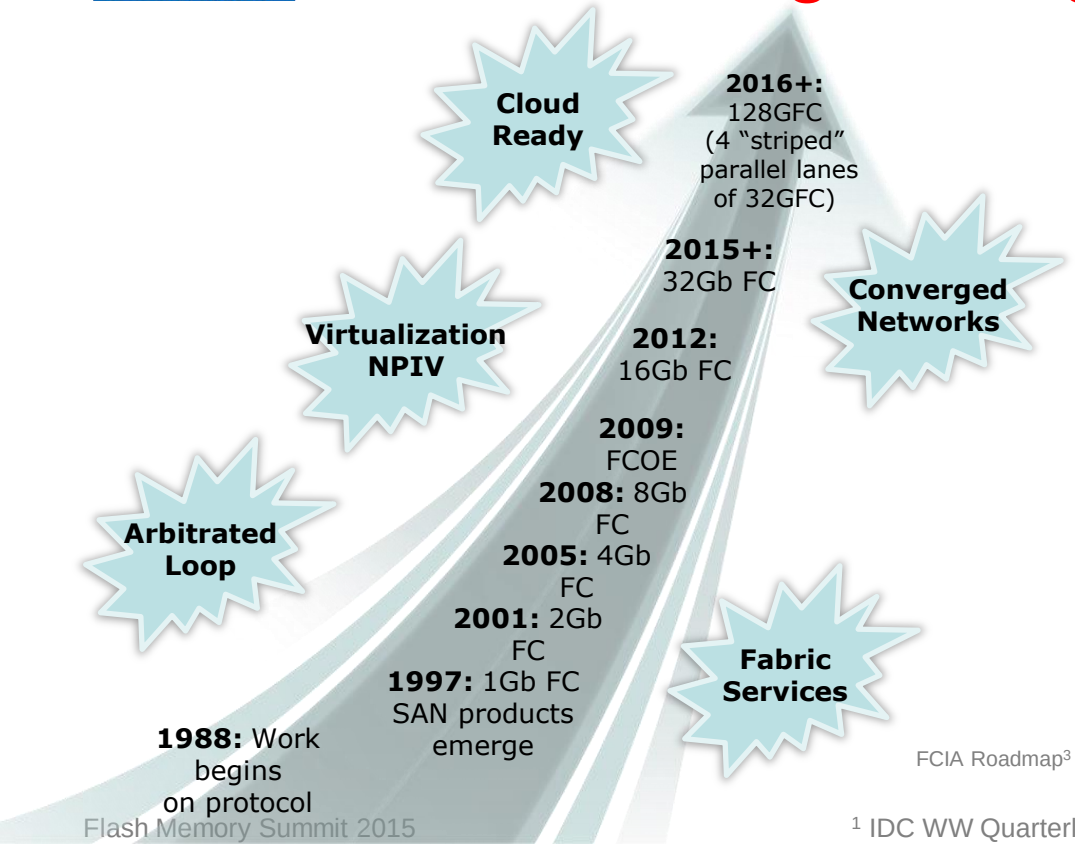

FCIA Roadmap[3]

- 15+ years deployment
- Lossless (guaranteed delivery – buffer to buffer credit)
- Reliable
- Secure & pervasive (SAN) – 30% market revenue share in 2014[1]

[1] IDC WW Quarterly Enterprise Storage Systems Tracker, 1Q2015, June 2015

Gen 5 (16Gb) Fibre Channel Advantage

- Flash & Gen 5 Fibre Channel (FC) intersection – low latency (ms to µs)[4]

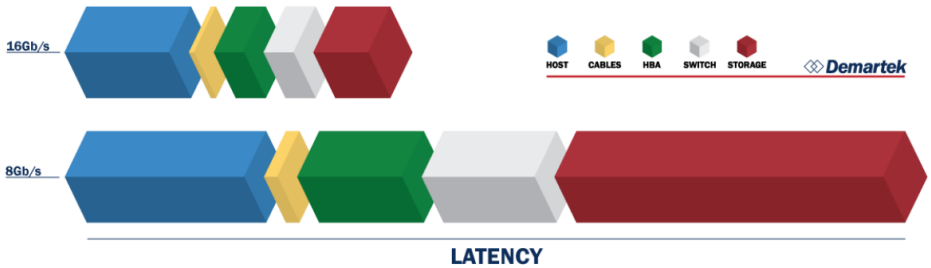

- Flash arrays (all flash or hybrid) – Gen 5 FC is a popular choice
  - Dell, EMC, HP 3PAR, IBM, Nimbus, NetApp, Pure Storage, Solidfire, EMC, Violin

Gen 5 Fibre Channel Advantage

- Hybrid implementation – “jitter” is a problem.

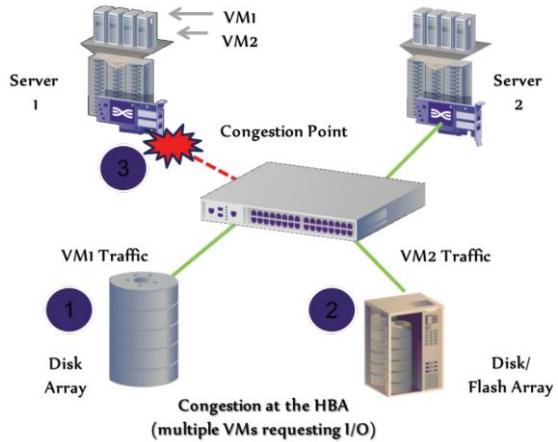

- How do we prioritize low latency flash traffic in hybrid network[5]
- ISL, Multi-Switch: Use CS_CTL bit in FC frame to mark traffic class (low priority)

Networking Storage: Ethernet

- Lossy (rely on upper protocols)
- Pervasive networking protocol
- Lower entry cost compared to FC

- Convergence (Network + Storage)
  - iSCSI, iSER
  - FCoE
  - RDMA over Ethernet

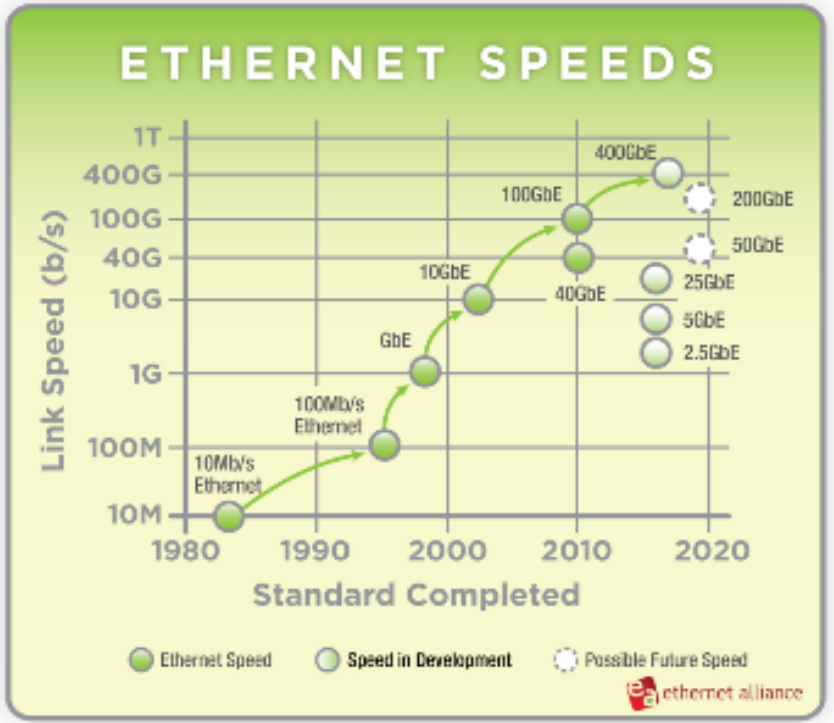

Ethernet Roadmap [6]

Ethernet Network Issues for Storage

- Ethernet is converged network: carries network + storage traffic
- Flow Control (IEEE 802.3x standard): “port wide” – affects “good guy”
  - Cause of horror stories in data center (enabled on any hop?)
- Priority Flow Control (IEEE 802.1Qbb): Use “priority bits” in Ethernet frame
  - Classify low-latency storage traffic with CoS (class of service)
  - Separate queue’s (host/switch - network) for specific CoS
- Can we guarantee bandwidth (QoS) for storage?
  - IEEE 802.1Qaz: Enhanced Transmission Selection (ETS)
  - Example: 50% of link bandwidth reserved for storage class

Ethernet: Example Deployment [7]

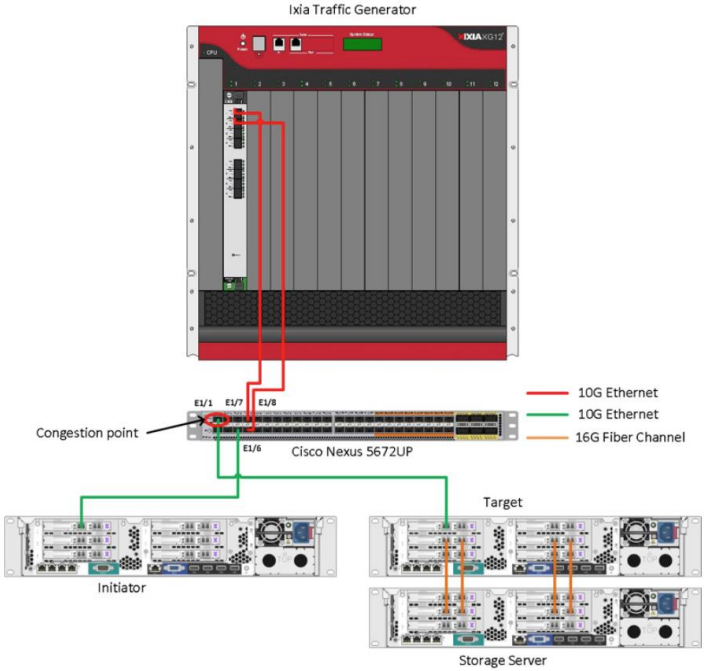

- Gen 5 (16Gb) FC back-end (target – storage server)
- RoCE (10GbE) – SMB Direct from initiator to target
- PFC (CoS 5) configured for RoCE traffic – with ETS of 5 Gbps (QoS)

Observations (chart):

- TCP traffic (CoS 0) throughput suffers under congestion, compared to ~5 Gbps for CoS5 (RoCE)
- High CoS0 latency under congestion

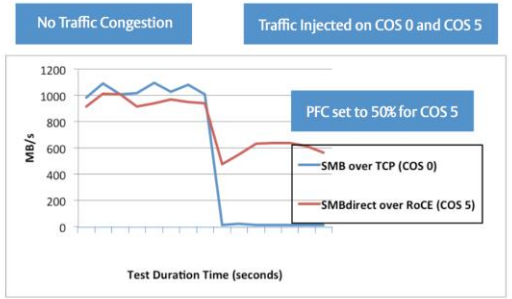

Figure 8. Bandwidth with Congested SMB Traffic

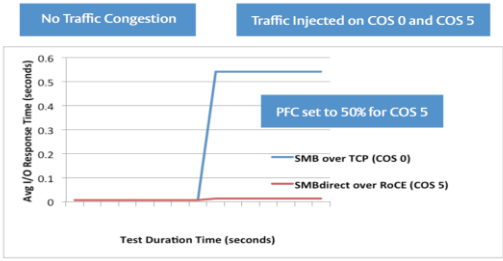

Figure 9. Latency with Congested SMB Traffic

Deployment & Management

- PFC/ETS limitations:
  - Each hop needs to support DCB and be explicitly configured
  - PFC benefit is topology dependent

- IEEE 802.1Qau Quantified Congestion Notification (QCN): Early congestion indication?

- Network management in multi-vendor deployment?
- Software Defined Networking (SDN)

Conclusion

- Flash brings in latency, prioritization and delivery challenges for networks
  - Fibre Channel (FC) is enterprise hardened
  - 16G Fibre Channel (Gen 5) offers compelling solutions
  - FC for NVMe over Fabrics
  - Ethernet - converged infrastructure solution: iSCSI, FCoE, RoCE, iSER
  - Scale UP and Scale OUT – Both Required

References

1. IDC Worldwide Storage Report, June 2015
2. Amin Vahdat, Google, Open Networking Summit 2015
3. Fibre Channel Industry Alliance Roadmap: http://fibrechannel.org/fibre-channel-roadmaps.html
4. Demartek, Evaluation Report: All-Flash HP 3PAR StoreServ 7450 Storage System and Generation 5 (Gen 5) 16 Gb/s Fibre Channel, March 2015
5. Emulex, Flash Storage Gets Priority with Emulex ExpressLane, August 2014
6. Ethernet Alliance Roadmap: http://www.ethernetalliance.org/roadmap/
7. Emulex/Cisco, Best Practices for Deployments using DCB and RoCE, July 2015
8. IEEE 802.1 Data Center Bridging Task Group (material referenced within)
9. IEEE 802.1Qau Congestion Notification (material referenced within)

Questions?

saurabh.sureka@avagotech.com

EMULEX
An Avago Technologies Company