Pre-Conference Seminar F
Flash Storage Networking

Rob Davis, Ilker Cebeli, Brian Pan, Abdel Sadek, Rupin Mohan, Steve McQuerry, and Alan Weckel

Why Network Flash Based Storage?

- There are advantages to shared storage
  - Better utilization:
    - capacity, rack space, power
  - Scalability
  - Manageability
  - Fault isolation, and recovery
- Shared storage requires a Network

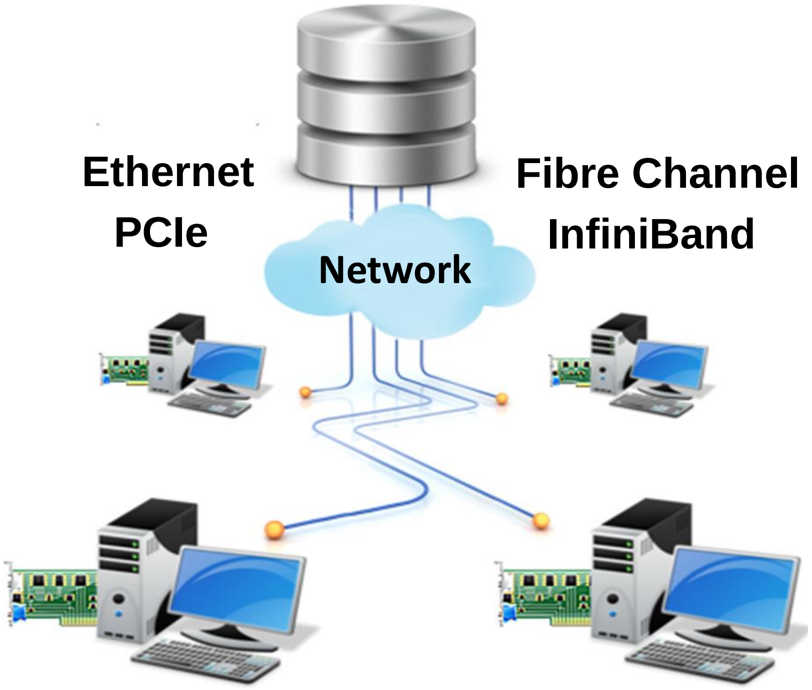

Agenda

- Networked Flash Storage Overview – 1:00 to 1:20
  - Rob Davis, Mellanox, VP Storage Technology
- ***Ethernet*** Networked Flash Storage – ~1:20 to 1:40
  - Steve McQuerry, **Pure Storage**, Senior Technical Marketing Engineer
- ***InfiniBand*** Networked Flash Storage – ~1:40 to 2:00
  - Abdel Sadek, **NetApp**, Technical Program Manager
- ***PCIe*** Networked Flash Storage – ~2:00 to 2:20
  - Brian Pan, **H3 Platform**, GM
- ***Fibre Channel*** Networked Flash Storage – ~2:20 to 2:40
  - Rupin Mohan, **HPE**, Chief Technologist

Agenda (cont.)

- Conference Break – 2:45 to 3:00
- How Networking Affects Flash Storage Systems – 3:00 to 3:20
  - Ilker Cebeli, **Samsung**, Sr. Dir. Product Planning
- Flash Storage Networking, How the market is evolving – ~3:20 to 3:40
  - Alan Weckel, **650 Group**, Technology Analyst/Co-Founder
- Q/A and Panel Discussion – ~3:40 to 4:00
  - All Presenters

Flash Makes Networking More Difficult

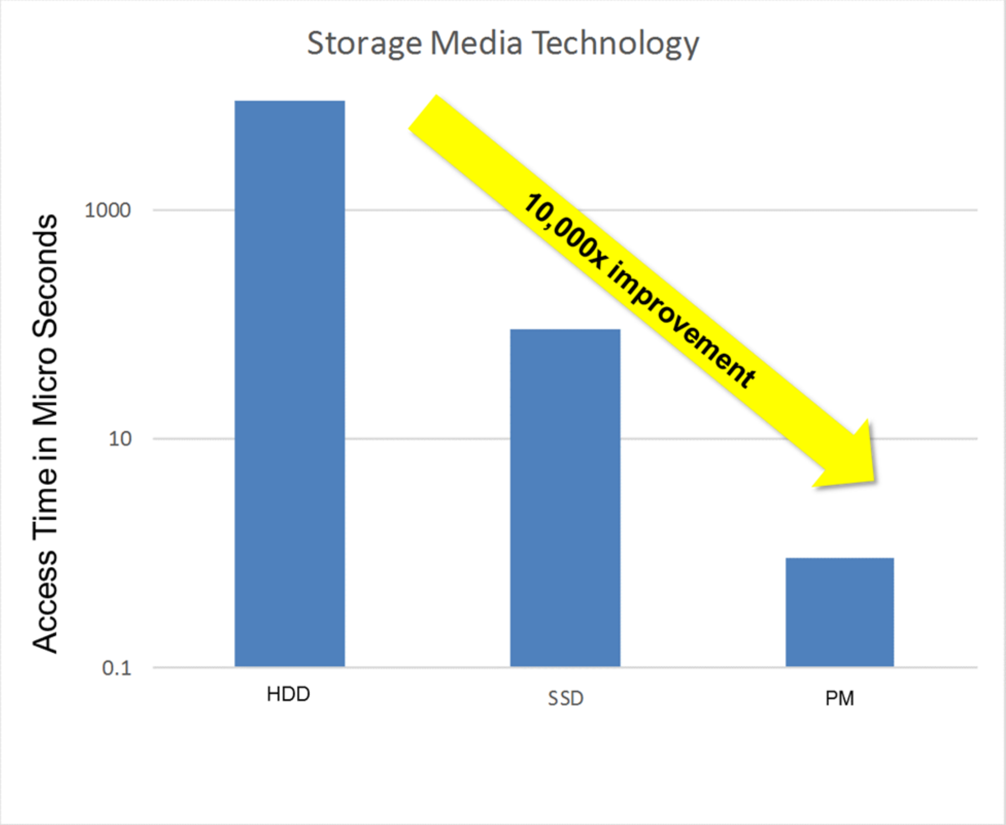

Faster Storage Needs a Faster Network

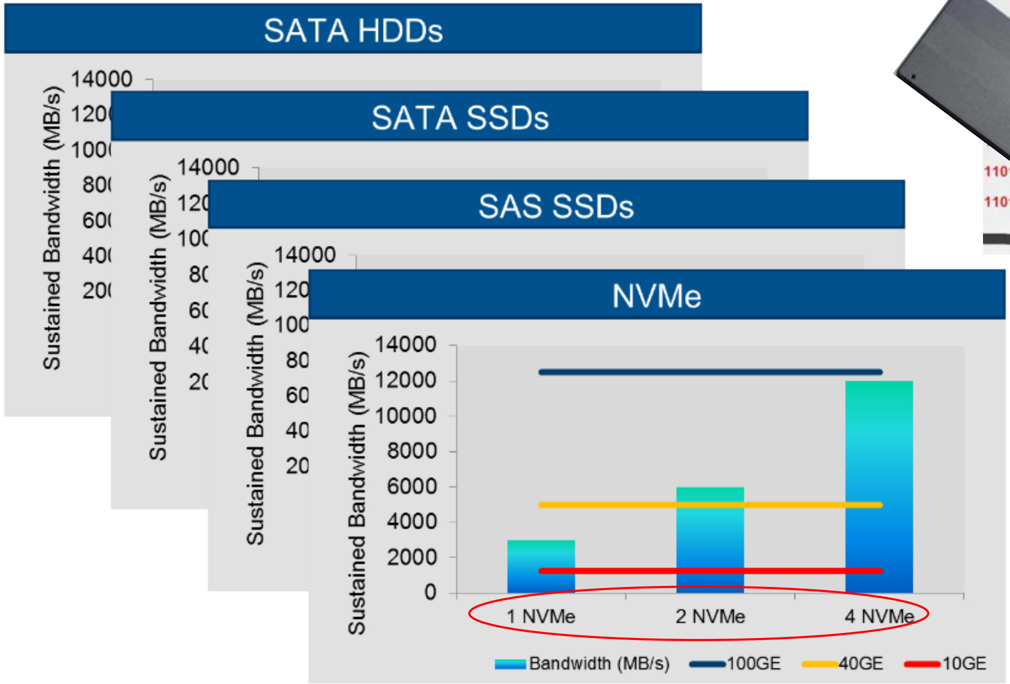

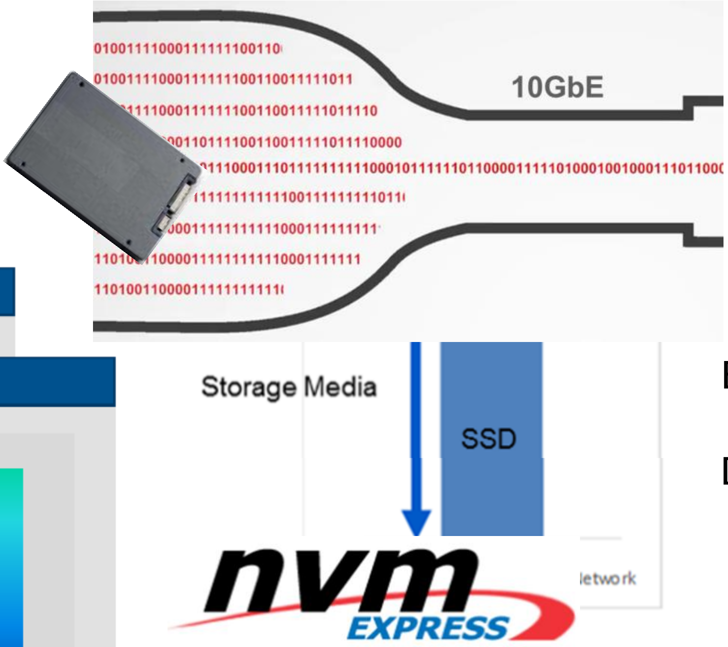

Flash SSDs move the Bottleneck from the Disk to the Network

What is the solution?

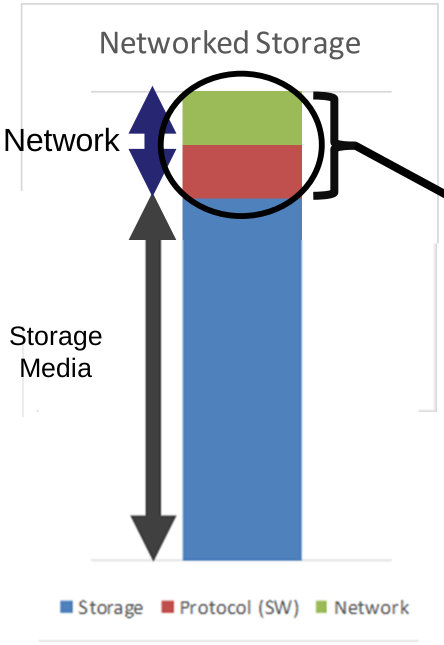

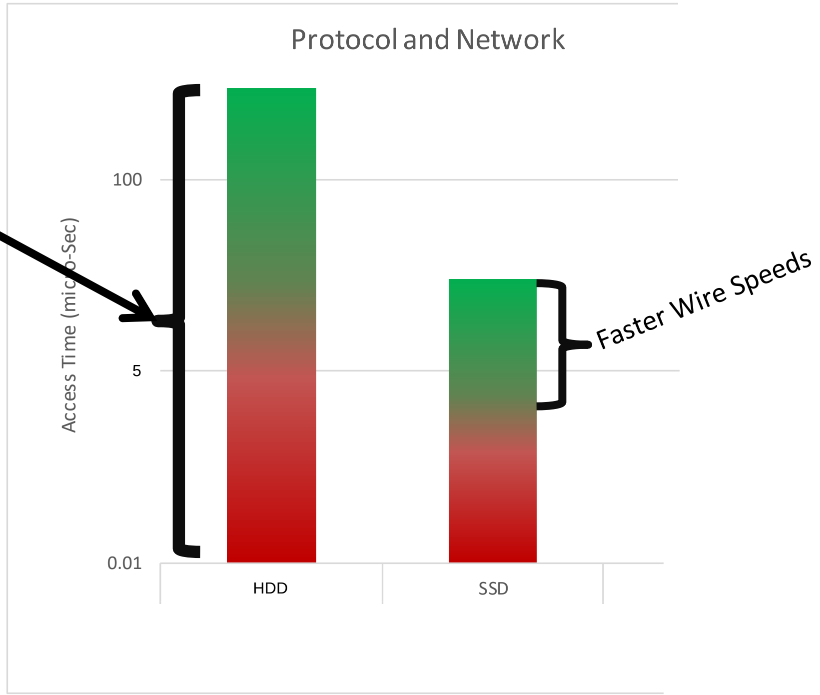

Faster Network Wires are Available

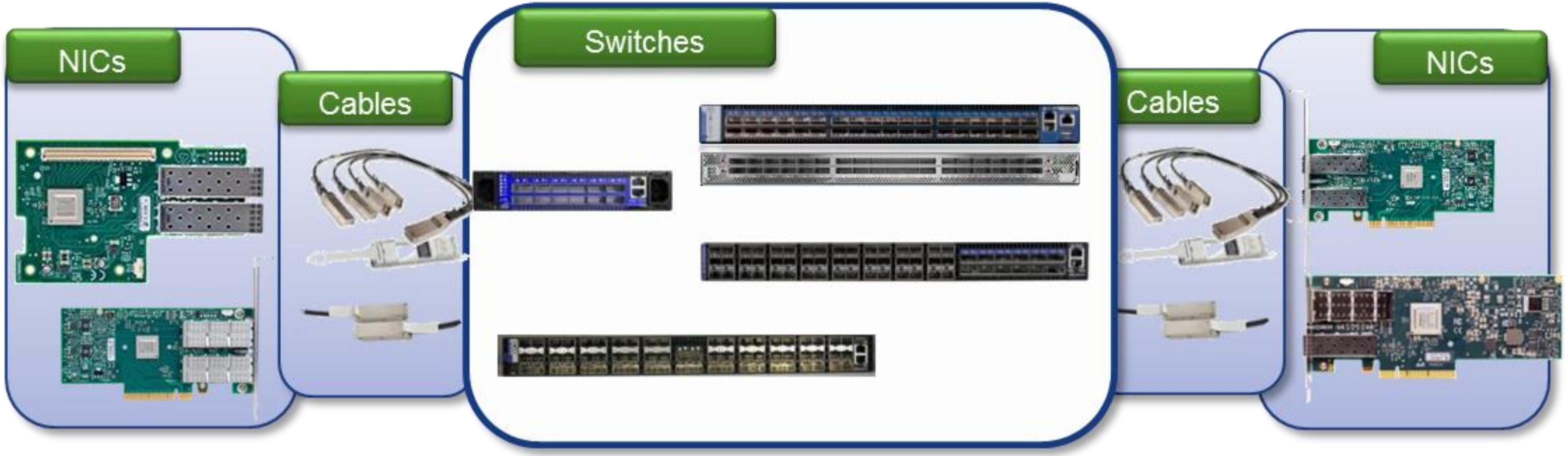

Ethernet & InfiniBand – 200Gb, going to 400Gb…
PCIe – Gen3(8Gb/lane), going to Gen4(16Gb/lane)…
FC – 32Gb, going to 128Gb…

Flash Needs Speed and Low Latency

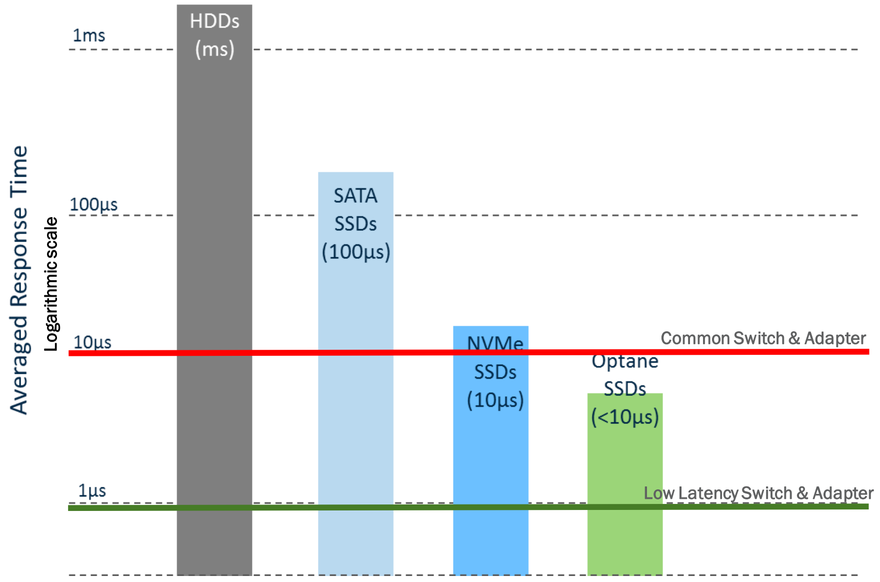

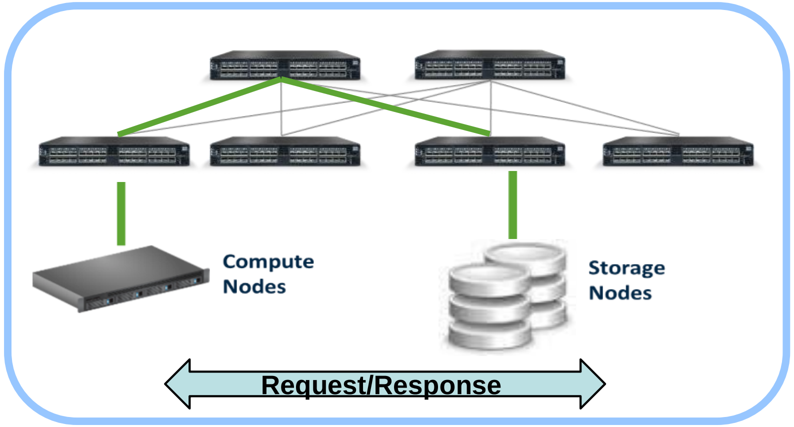

Network hops multiply latency

Faster Network Components Solves Some of the Problem…

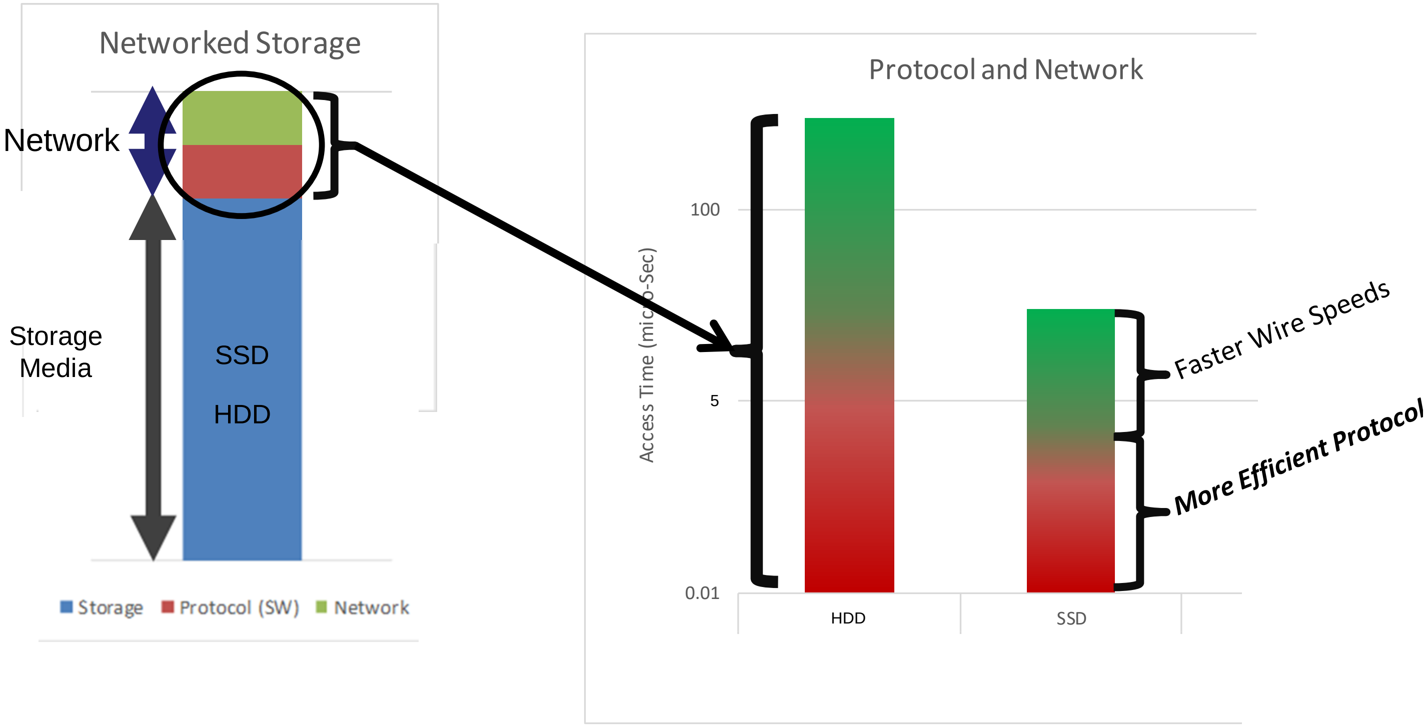

Flash Memory Summit

Faster Protocols

- NVMe-oF
  - RDMA(RoCE, IB)
  - Fibre Channel
  - PCIe
  - TCP
- RDMA
  - SMB Direct
  - VSAN over RDMA
  - iSER
  - NFSoRDMA
  - Ceph o RDMA

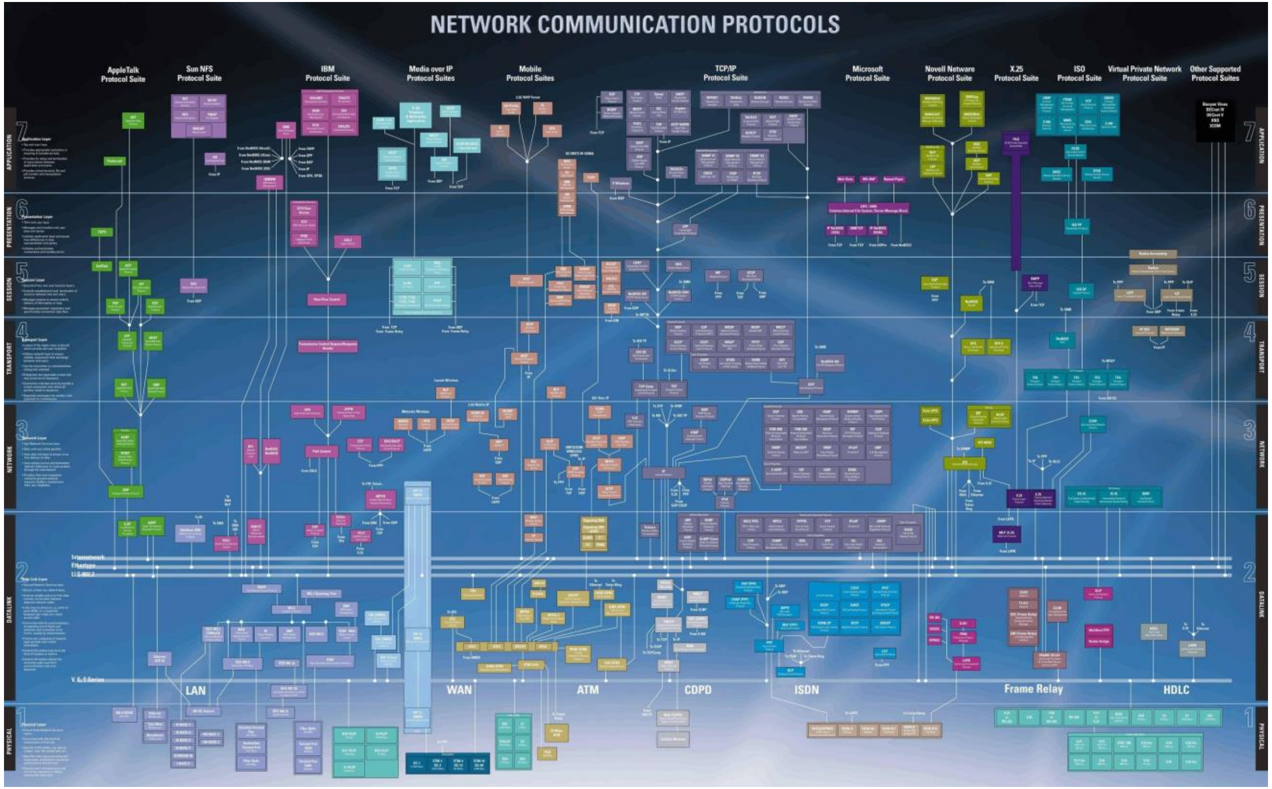

Where best to plug in?

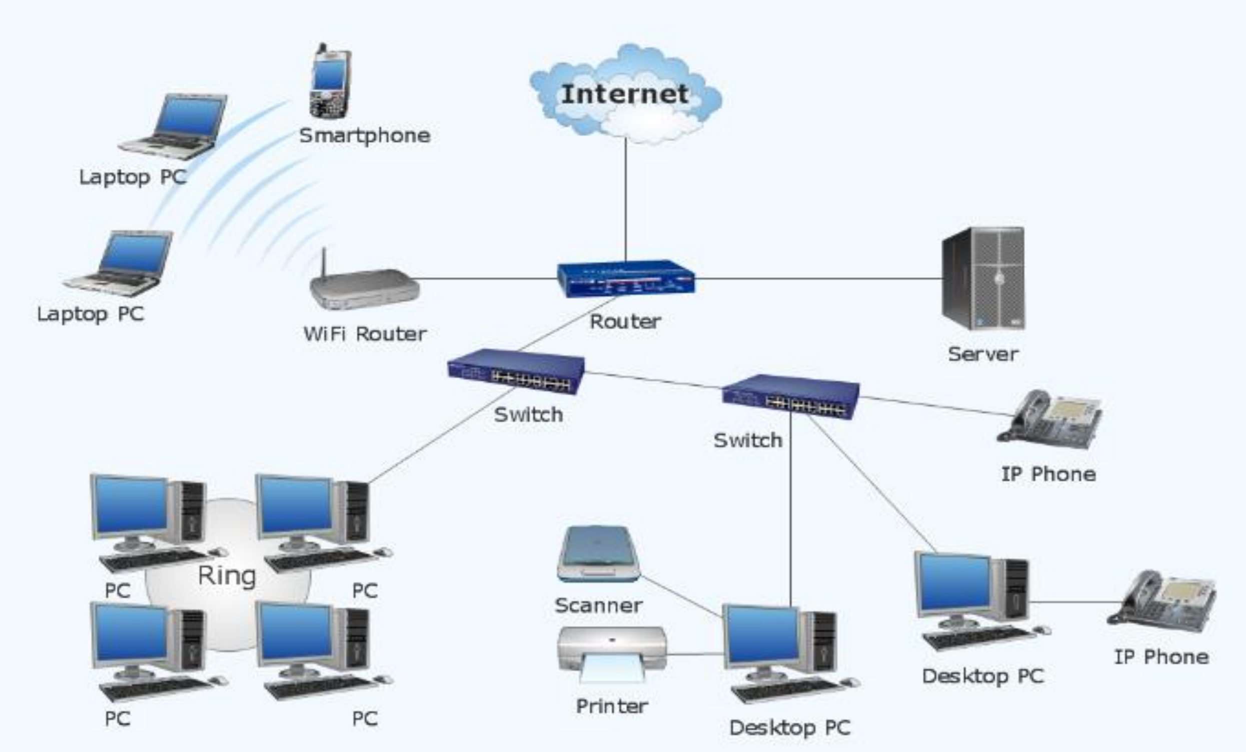

Flash Storage – Closer to Servers

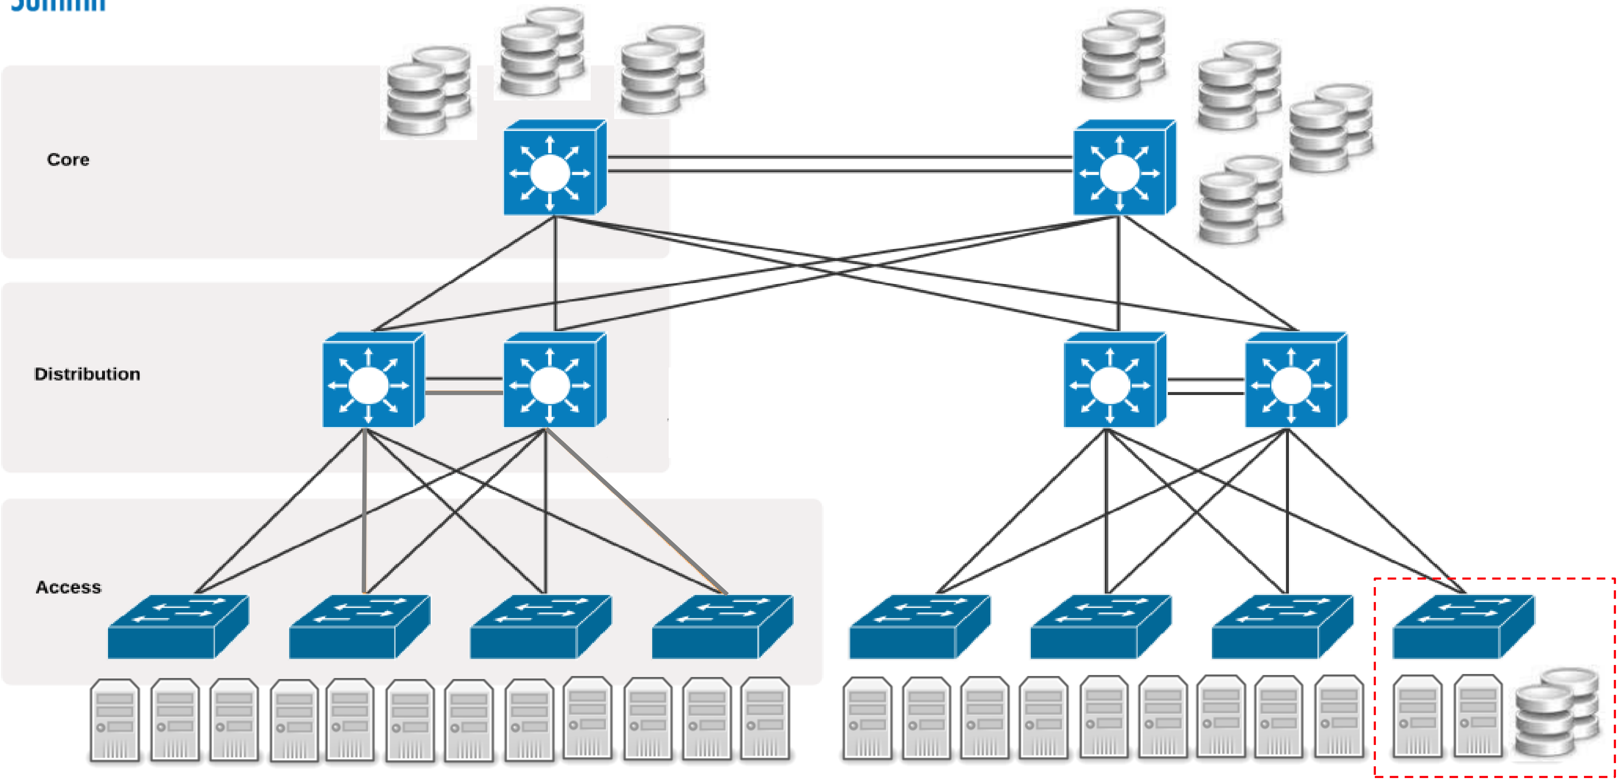

Match the Network to the Solution

- The Solution will often drive the protocol and the network technology
  - All technologies support Block
  - All technologies do not
-

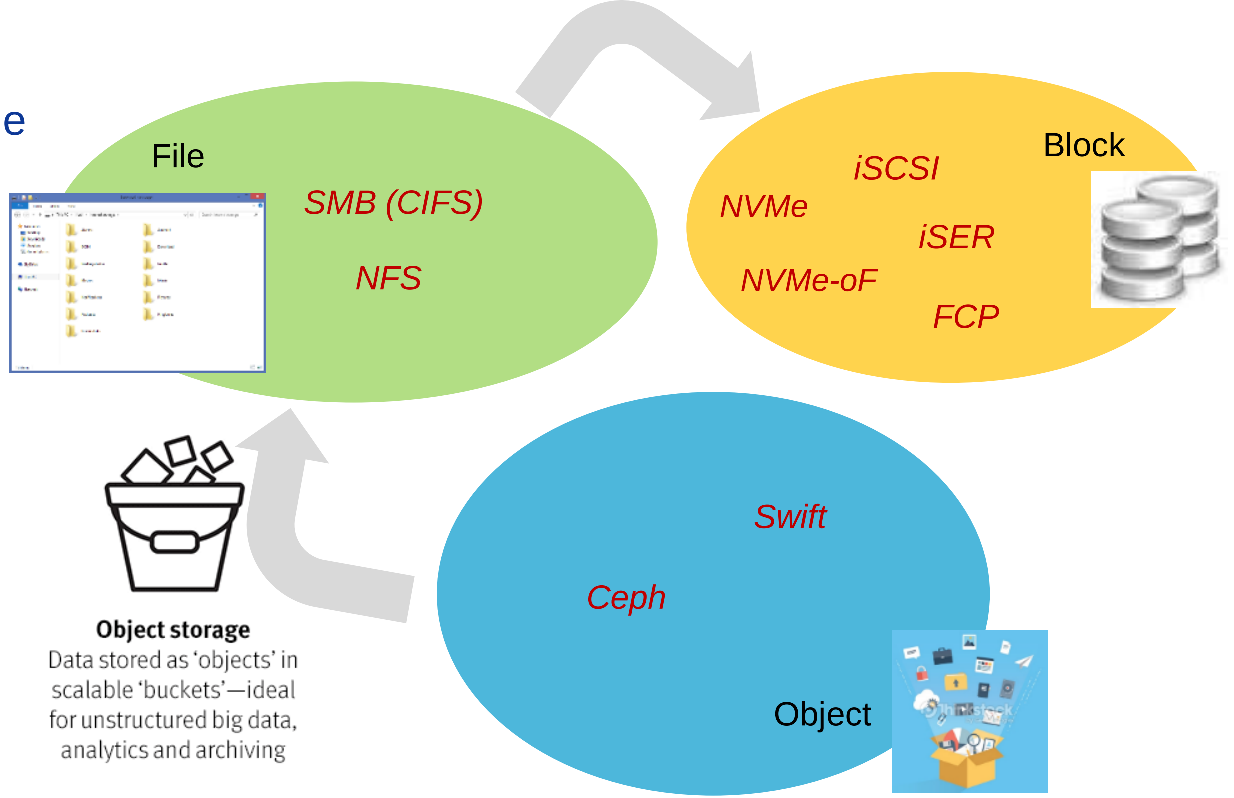

Conclusions

- There are tried and true reasons for networking your storage
- Networking flash requires special considerations
  - Faster Storage needs Faster Networks!
  - And protocols
- For the next few hours this team will present the different options and trade offs
- Then you get to question us

Thank You

Steve McQuerry

Steve McQuerry is a Senior Technical Marketing Engineer on the platform team at Pure Storage. One of his areas of focus is helping customers understand the use cases and best practices for deploying NVMe-oF with FlashArray.

Steve is a CCIE Emeritus (CCIE #6108) is a 20+ year data center veteran. For the last 16 years he has held both field and product positions for storage, networking, and compute OEMs. Steve has published multiple networking books and has been recognized as distinguished speaker at industry events. Steve holds a Bachelor of Science in Engineering Physics from Eastern Kentucky University.

Ethernet Networked Flash Storage

Flash Storage Networking

Steve McQuerry, Sr. TME PureStorage

- Background
- Common Transport Use Cases
- Customer profiles
- Observed Trends
- Pros & Cons

STORAGE ICEBERG

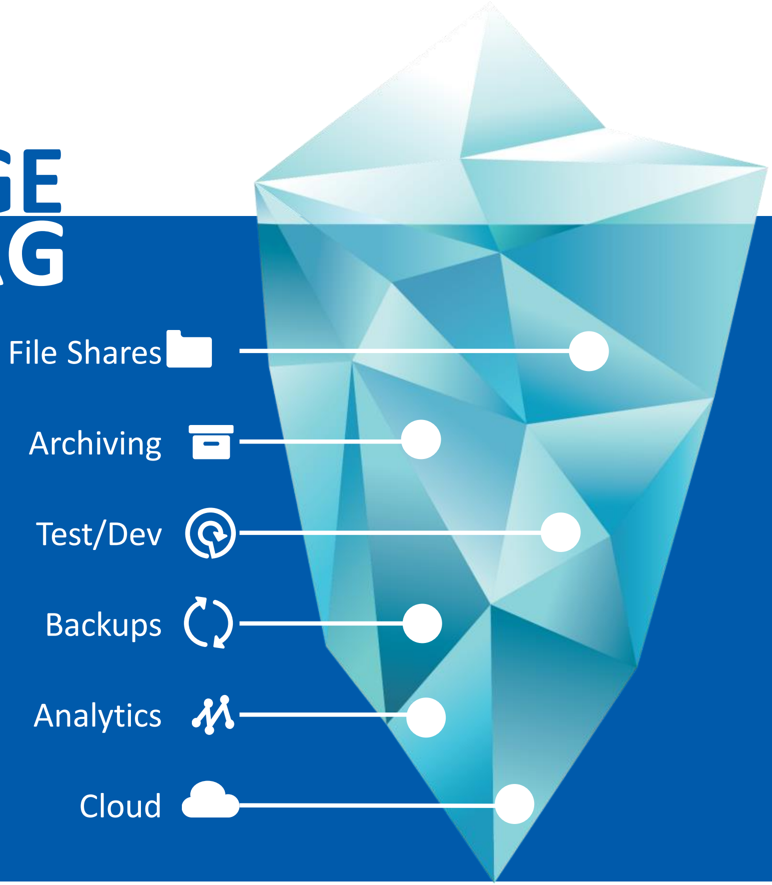

PRIMARY STORAGE

- Traditional SAN (FC, iSCSI)
- 20% of capacity

SECONDARY STORAGE

- Mostly Ethernet
- 80% of capacity
- Rapid growth
- Diverse data types
- Scale-out
- Tiered data

Ethernet Background & Roadmap

- Ethernet is the dominate transport in the Data Center from a port count perspective
- Ethernet speeds have increased exponentially over the last 10 years
- File and Software Defined Storage are predominantly Ethernet based transports
- Ethernet has had challenges when it comes to guaranteed throughput and deterministic latency

Common Transport Use Cases

- iSCSI – Front end block, not typically focused on performance
- iSER – Front end block focused on performance leverages RDMA
- iWARP – Front end; block, file, and SDS
- RoCE – Front/Back end; block, and SDS
- Other IP – Front end; block, file and SDS,

Customer Ethernet Transport Mix

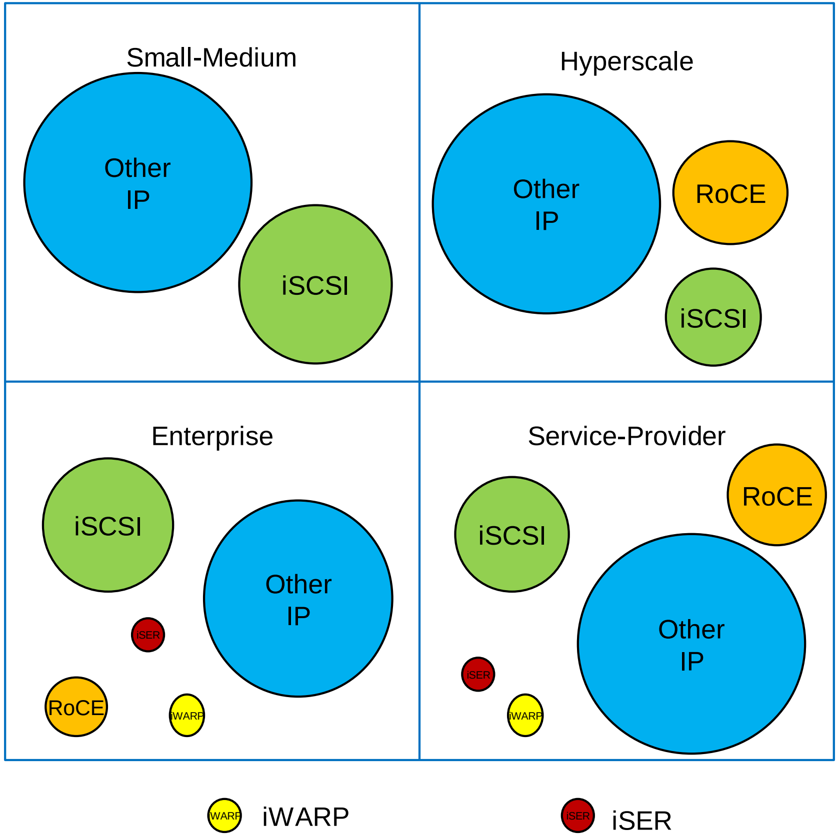

Ethernet Performance Roadmap

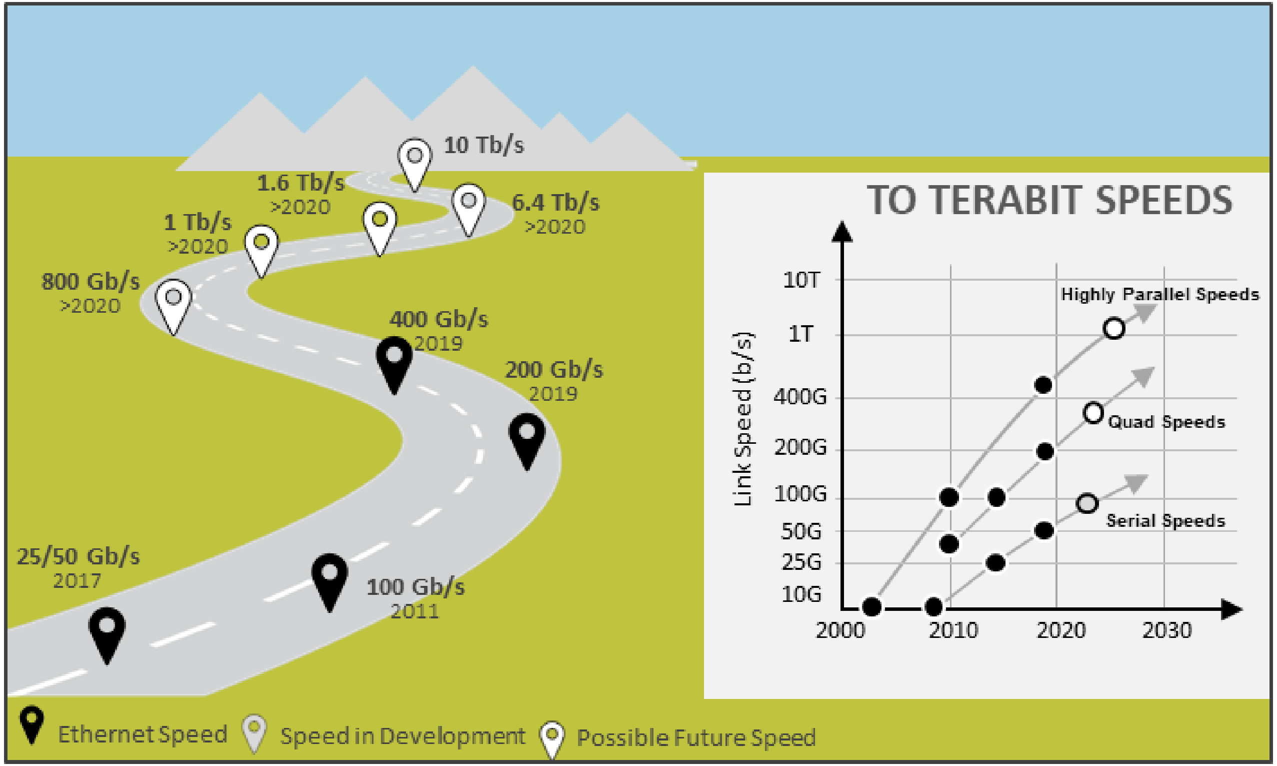

Observed Trends

- Majority of block storage continues to be on Fibre Channel
- Ethernet is becoming more interesting to customers from an operational standpoint
- File has been predominately IP/Ethernet and that trend continues
- SDS leverages IP/Ethernet and DAS
- Ethernet has become an interesting transport for disaggregating storage from compute

Pros & Cons

- Well understood
- Common Infrastructure
- Multiple Vendors
- Rapidly evolving
- Lower Cost

- Lossy/Unpredictable performance
- QoS can be complex
- Sometimes still a separate infrastructure

Thank You

Abdel Sadek Bio

Abdel has been in the storage industry for over 15 years. Working for LSI then NetApp. He has a wide experience in different SAN protocol including NVMe-oF, IB, FC, SAS and iSCSI.

He's part of the NetApp E-Series Technical Marketing team focusing on High Performance Computing, data analytics and Media and entertainment.

Agenda

- Background
- Storage types
- Latency vs. Scalability
- InfiniBand solution
- Use Cases
- Pros and Cons

Background

- Storage is moving to be “Memory-like”
  - Economics based on NAND, not on “spinning rust” and mechanicals
  - Supply line is now linked to Fabs and Semiconductor technology
  - Memory-like semantics supported by RDMA
- Software and supporting H/W is evolving to address the moving bottleneck
  - Demands lighter weight (I/O) software stack
  - Network now a critical component
  - Persistent memory becoming a market reality
  - Persistent memory Filesystems are emerging
  - Rapid adoption of In-Memory computing

Different Types of Storage

	Layer 0	Layer 1	Layer 2	Layer 3	Layer 4
Latency	<20ns	2µs–10µs	10µs–200µs	300µs–10ms	100ms–100sec+
IOPS	00s millions	0s millions	0s millions	00,000s	000s
Bandwidth	00s GBps	0s GBps	100s GBps	100s GBps	1s MBps
Capacity	1s TB	10s TB	100s TB	10s PB	100s PB
Cost	Very high	High	Mid	Mid/low	Low
					
Type	DRAM	SCM	SSD with SCM cache	HDD with SSD cache	HDD
	DDR	DDR/PCIe	PCIe/SAS	PCIe/SAS	SAS
Access	Memory bus	Memory bus	Fabric	Ethernet/FC	Ethernet

Challenge: Latency vs. Scalability

Scalability with low latency

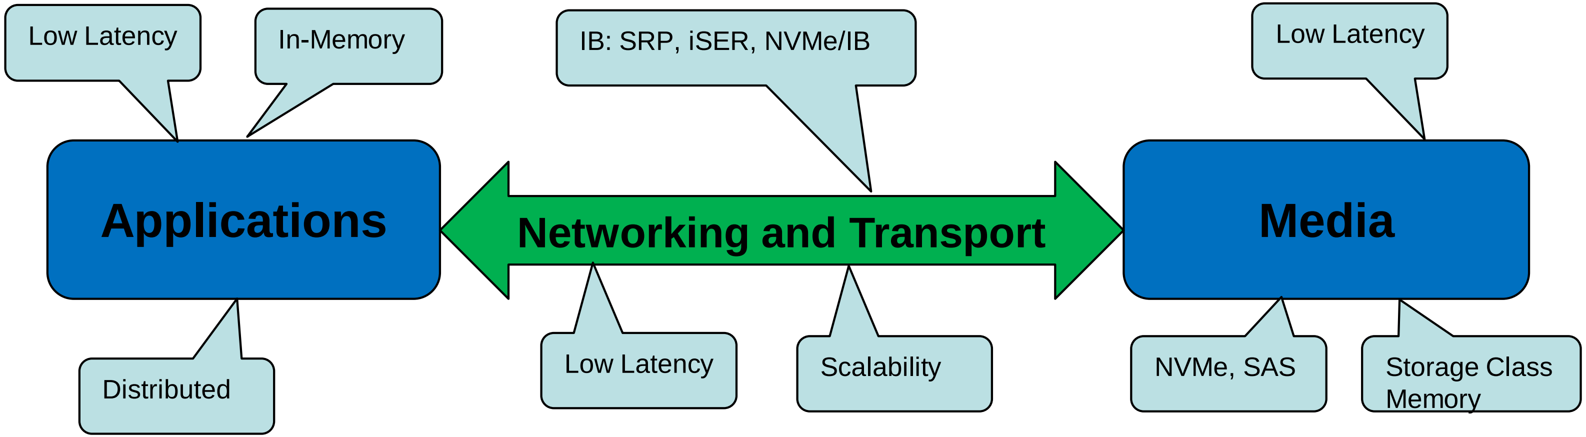

InfiniBand Architecture Highlights

- Reliable, lossless, self-managed fabric
- Hardware based transport protocol- Remote Direct Memory Access (RDMA)
- Centralized fabric management – Subnet Manger (SM)

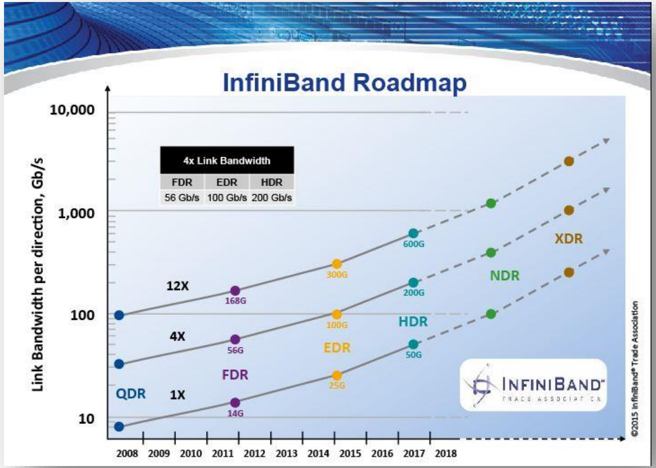

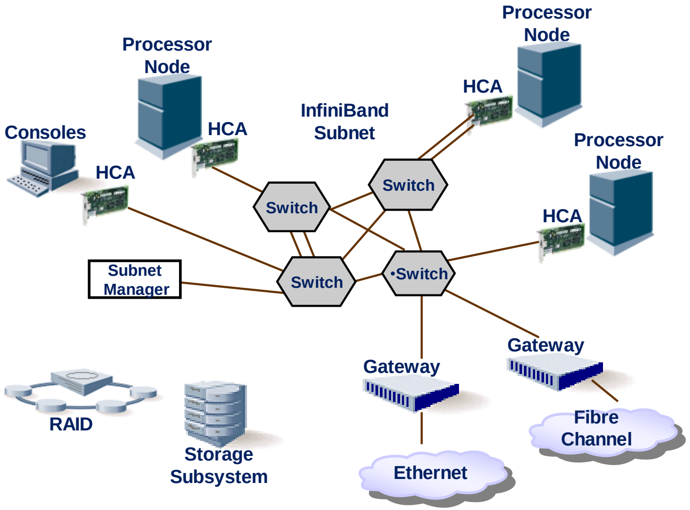

Reliable, Lossless, Self-Managed Fabric

- Credit-based link-level flow control
  - Link Flow control assures **NO packet loss** within fabric even in the presence of congestion
  - Link Receivers grant packet receive buffer space credits per Virtual Lane
  - Flow control credits are issued in 64 byte units
- Separate flow control per Virtual Lanes provides:
  - Alleviation of head-of-line blocking
  - Virtual Fabrics – Congestion and latency on one VL does not impact traffic with guaranteed QOS on another VL even though they share the same physical link

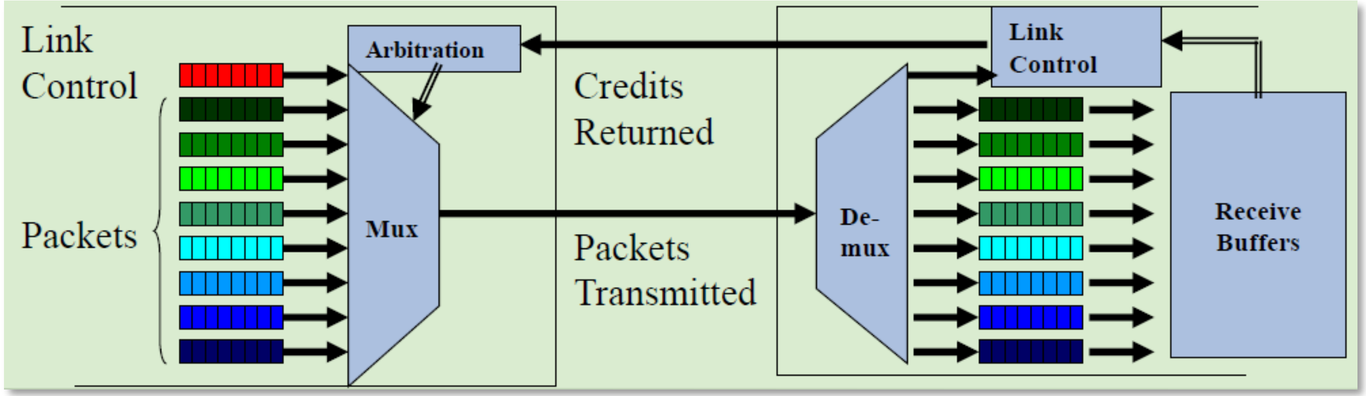

Use Cases

- High performance databases
- Parallel Filesystems
- AI/ML/DL workloads.
- Financial trading.
- Realtime modeling (metrology, logistics)

Case 1: RDMA Hand Shake with NVMe/IB

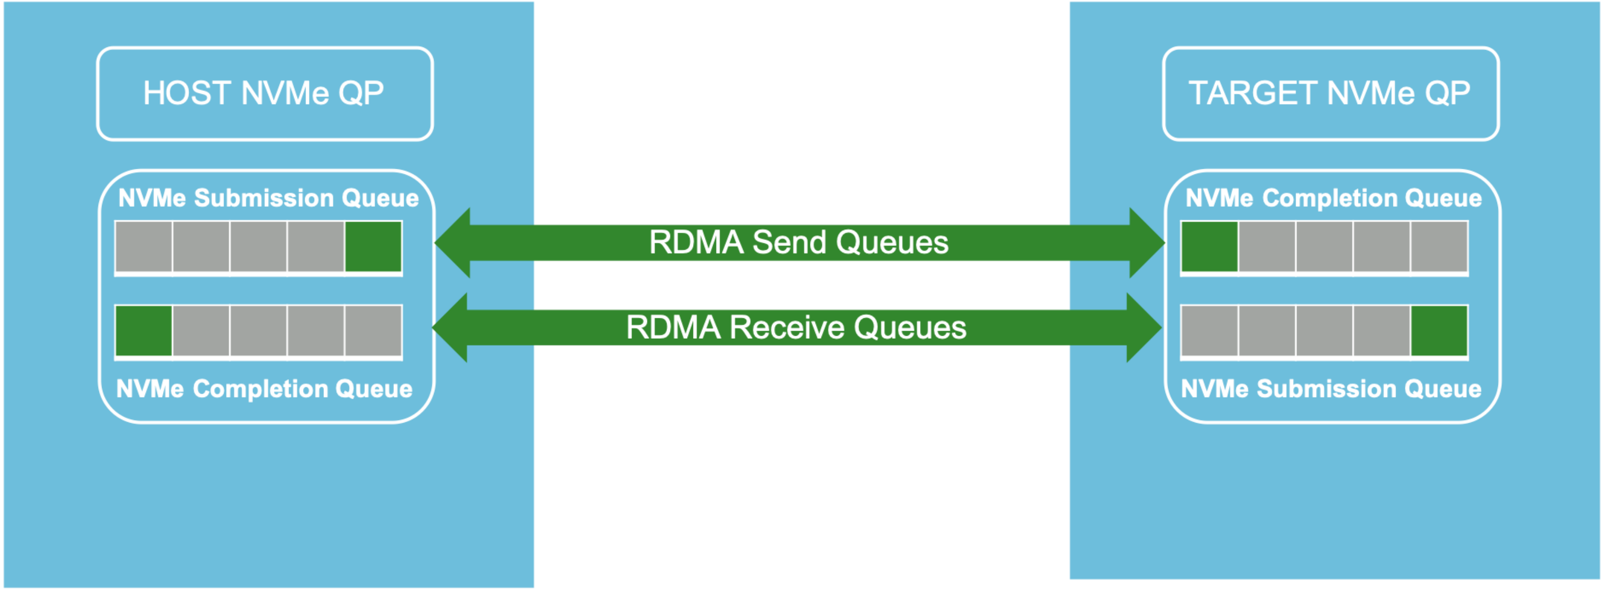

Case 2: Parallel Filesytems (HPC)

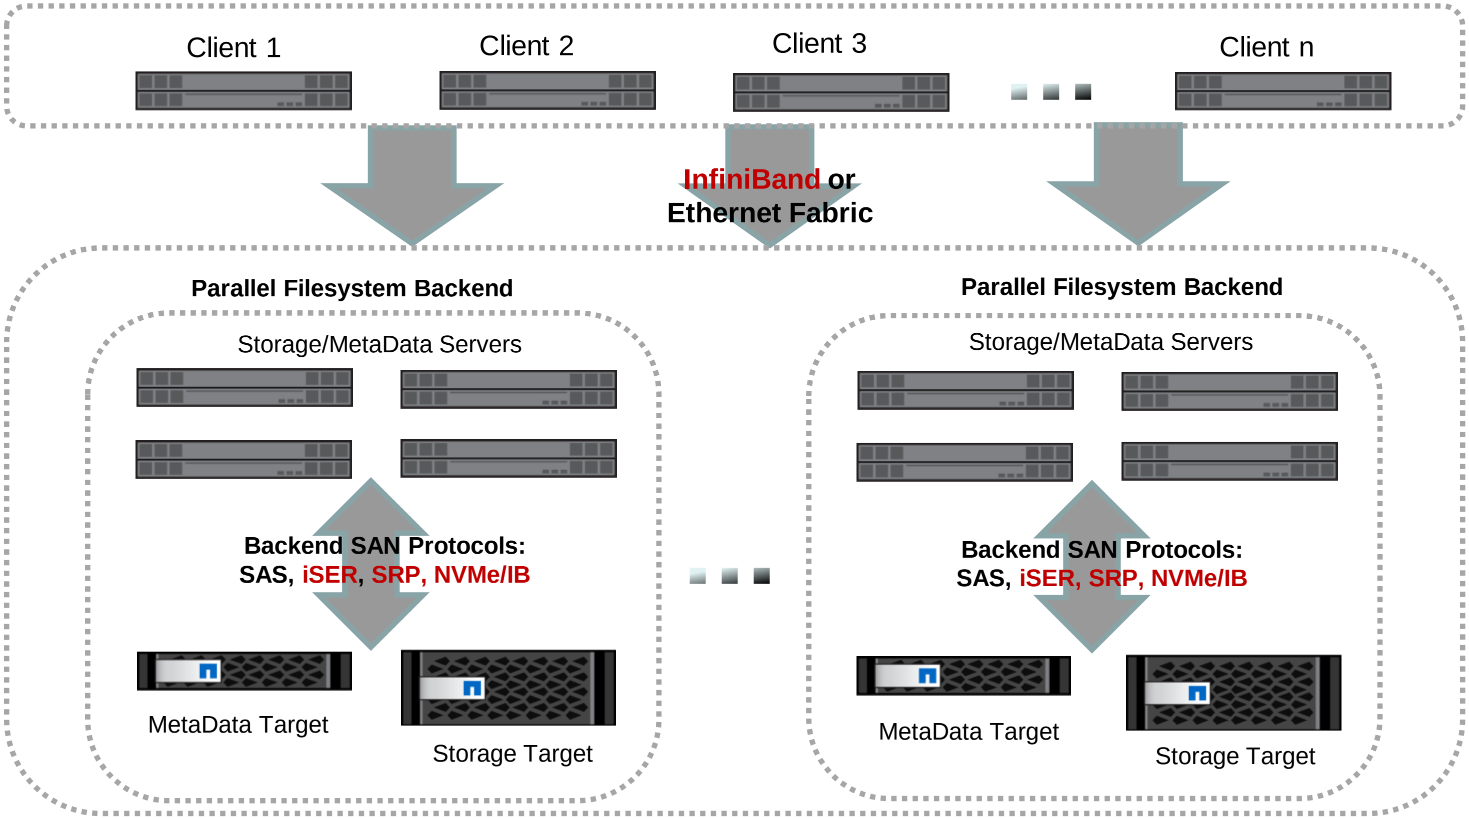

Flash Memory Summit

Public References

- Simula Research Lab: https://www.simula.no/news/simula-signs-contract-nextron-delivery-first-procurement-ex3-infrastructure

- Autralian National University: https://www.cio.com.au/article/664690/australia-fastest-ever-supercomputer-go-live-november/

Pros and Cons of InfiniBand

Pros	Cons
Low Latency	Distance with low latency
RDMA support	H/W offload adoption limited
High Bandwidth	Optimization required for IOPs
Open source drivers	Limited OS support
Reliable	Different cables/connectors for different speeds
Lossless	

Thank You

Brian Pan | H3

GM

huaiyangpan

www.h3platform.com

brian.pan@h3platform.com

+886 2 2698 3800#110

PCIe® Networked Flash Storage

Brian Pan
H3 Platform

PCI Express® (PCIe®)

- Specification defined by PCI-SIG®
  - www.pcisig.com
- Packet-based protocol over serial links
  - Software compatible with PCI and PCI-X
  - Reliable, in-order packet transfer
- High performance and scalable from consumer to Enterprise
  - Scalable link speed (2.5 GT/s, 5.0 GT/s, 8.0 GT/s, 16 GT/s, and 32 GT/s)
    - Gen5 (32 GT/s) is still being standardized
  - Scalable link width (x1, x2, x4, …. x32)
- Primary application is as an I/O interconnect

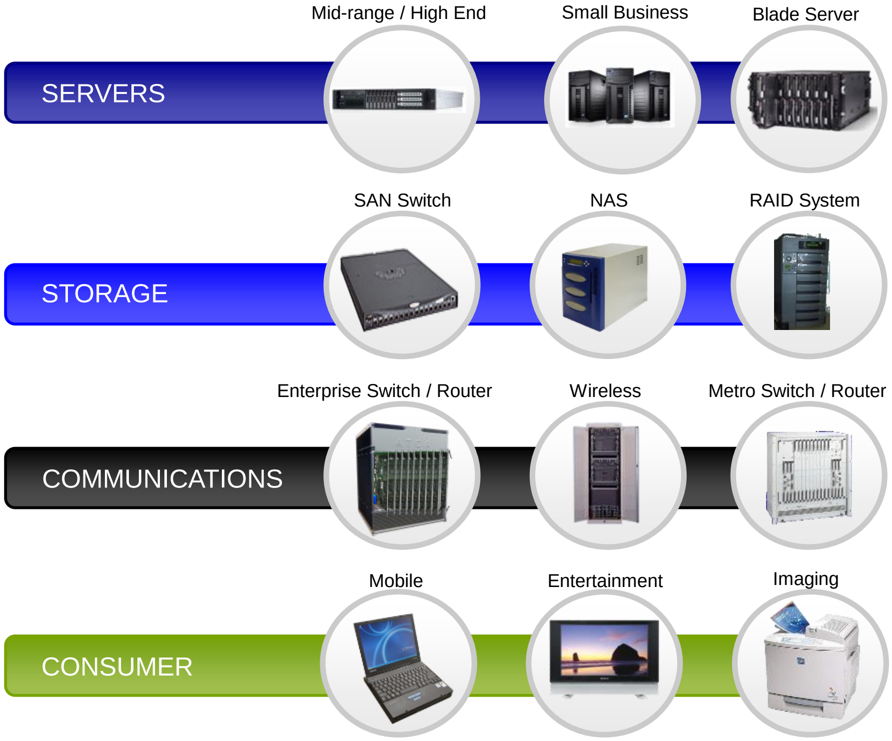

PCIe Characteristics

- Scalable speed
  - Encoding
    - 8b10b: 2.5 GT/s (Gen 1) and 5 GT/s (Gen 2)
    - 128b/130b: 8 GT/s (Gen 3), 16 GT/s (Gen4) and 32 GT/s (Gen5)
- Scalable width: x1, x2, x4, x8, x12, x16, x32

Generation	Raw Bit Rate	Bandwidth Per Lane Each Direction	Total x16 Link Bandwidth
Gen 1*	2.5 GT/s	~ 250 MB/s	~ 8 GB/s
Gen 2*	5.0 GT/s	~500 MB/s	~16 GB/s
Gen 3*	8 GT/s	~ 1 GB/s	~ 32 GB/s
Gen 4	16 GT/s	~ 2 GB/s	~ 64 GB/s
Gen 5	32 GT/s	~4 GB/s	~128 GB/s

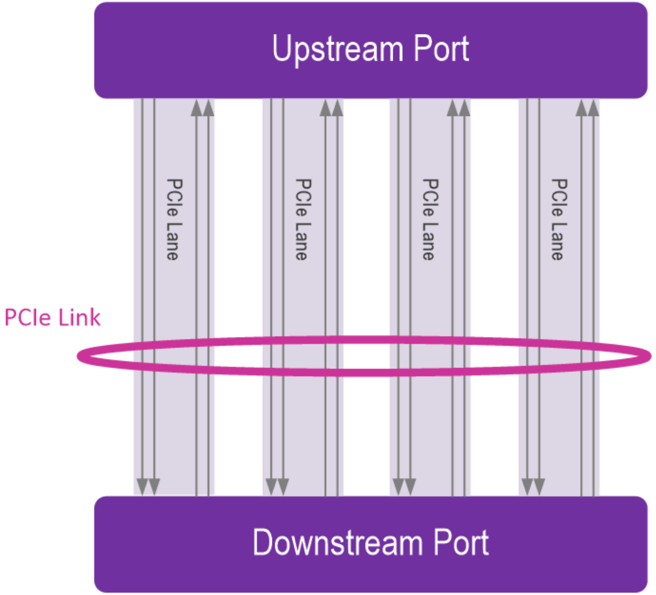

Note
** Source – PCI-SIG PCI Express 3.0 FAQ*

NVMe Through PCIe Fabric

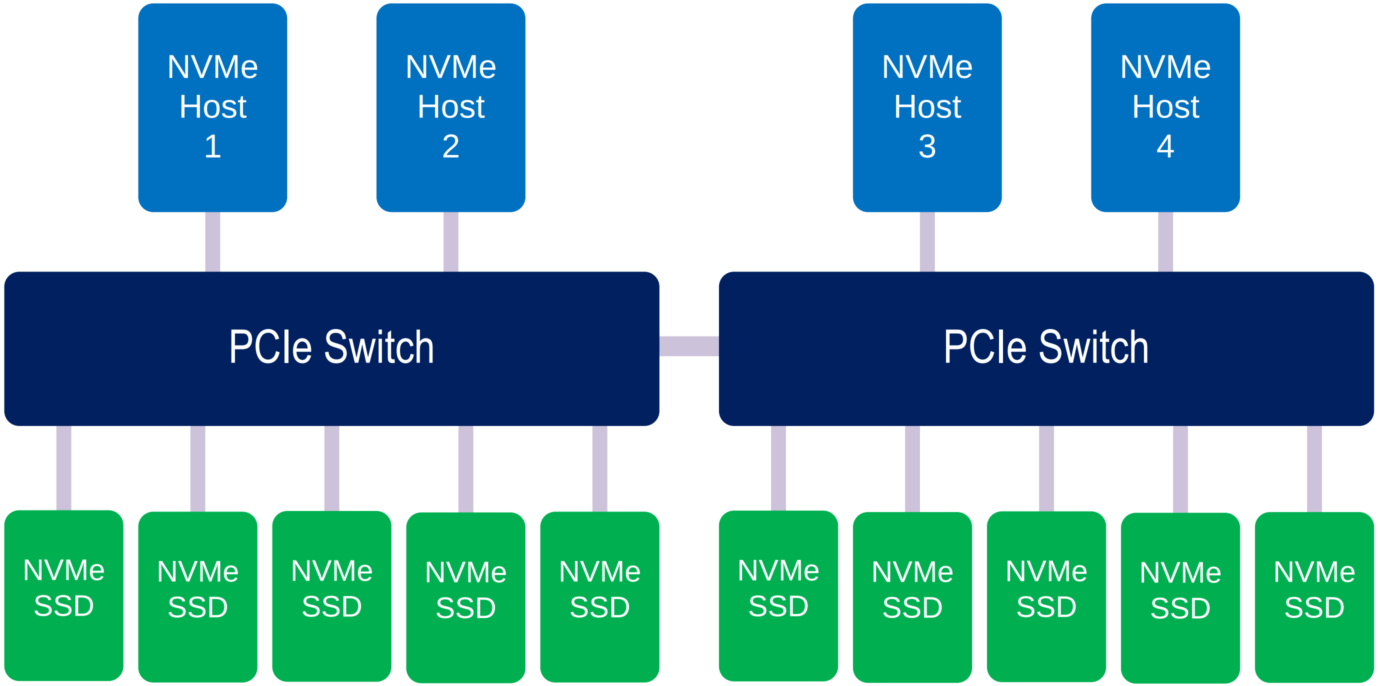

Features of PCIe Fabric

- NVMe SSD sharing
  - NVMe can be assigned to any connected host in PCIe fabric
  - NVMe SSD can be shared by using NVMe virtual function (The VF can be assigned to VM on the host)
- Namespace and VF mapping
  - Users can create/delete namespace and manage the mapping of VF and namespace

Architecture of NVMe SSD SR-IOV

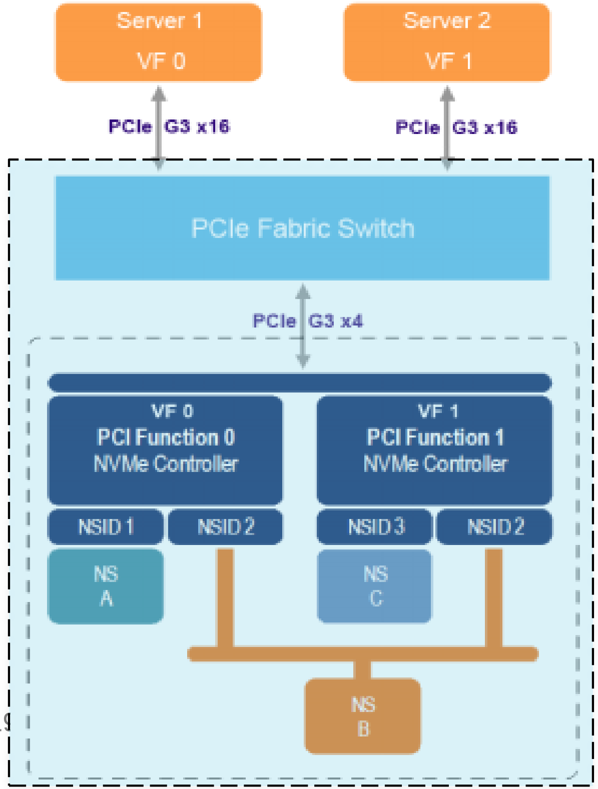

x86 server

- CentOS 7.5
- VMware

JBOf

- Broadcom ARM mCPU
- Broadcom PCIe switch
- PCIe switch driver

NVMe SSD

- Samsung 1725a
- Samsung PF/ VF driver

JBOf Specification

- PCIe switch
  - Broadcom 9797 PCIe switch
  - Broadcom ARM 58522 mCPU
- Host connection
  - 2x PCIe Gen3 x16 for host connection
- NVMe SSD
  - 16x U.2 Samsung 1725b NVMe SSD

Flash Memory Summit

Software Features

- NVMe SSD management
  - Create namespace of NVMe SSD
  - Create maximum VF of NVMe SSD (15+1 VF)
  - Assign VF to namespace or vise versa
- Host and NVMe SSD hotplug
  - Manage surprise removal and plug-in
- GUI and API
  - GUI or API for JBOf management

Multi-host SR-IOV vs Pass Throught vs Hyper+ SR-IOV

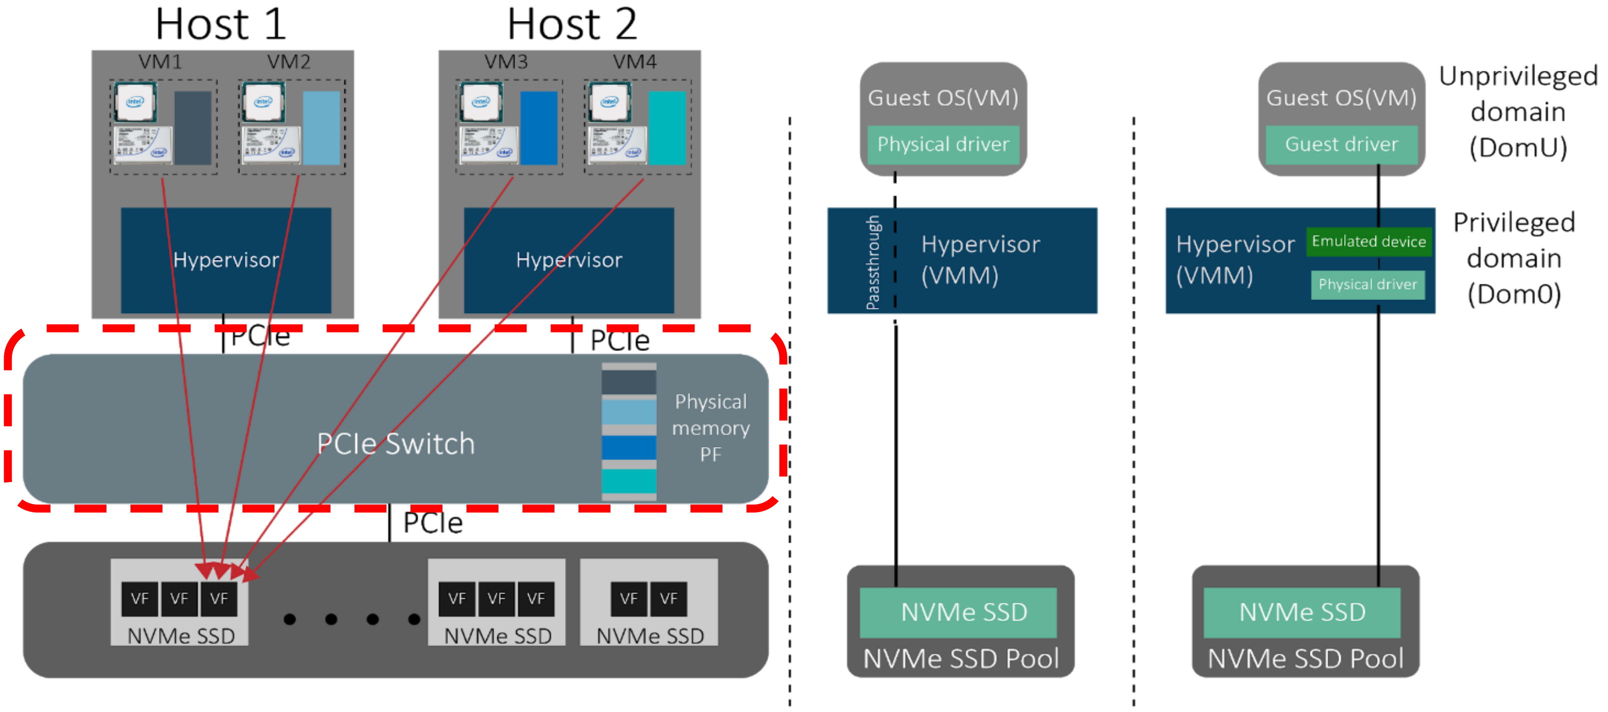

NVMe with SR-IOV. VM talk to VF directly. PF is sit on PCIe switch.

NVMe without SR-IOV. Passthrough model

NVMe without SR-IOV. Hypervisor manage VMs to NVMe SSD

SR-IOV Preparation Process

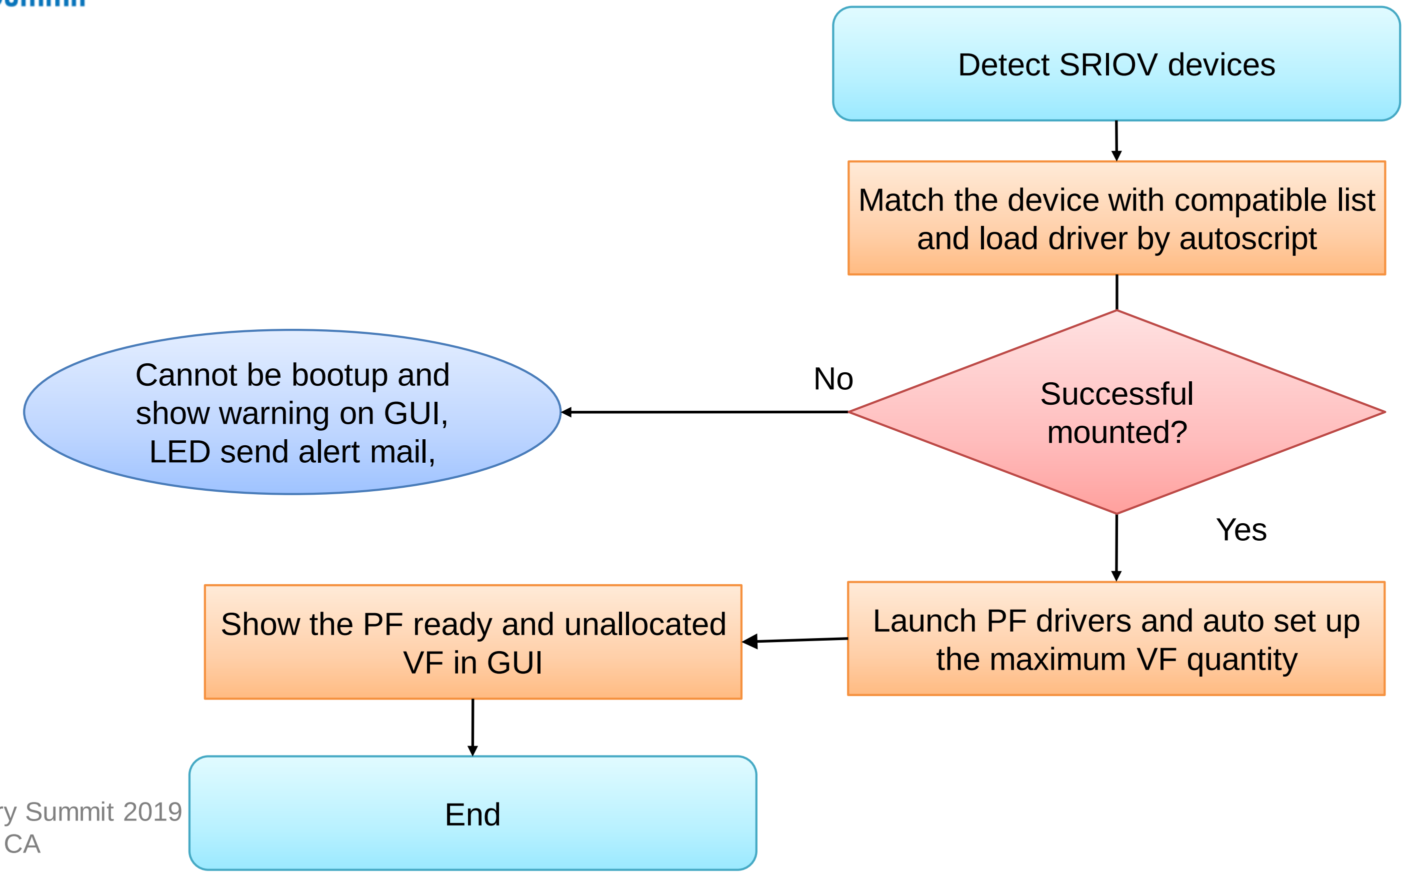

VF and Namespace Assignment

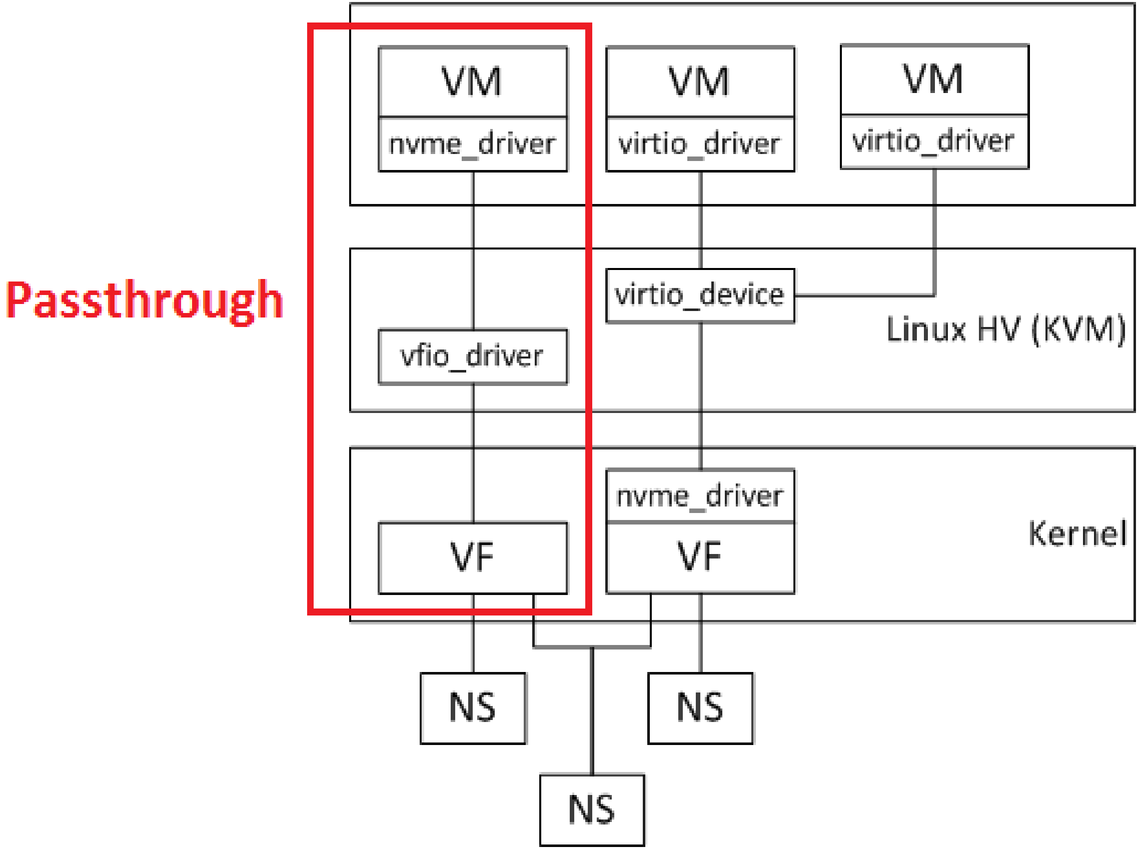

Namespace on VM

```
Direct-Attach VF to VM Passthruogh
Dump NS from VM
===============
[root@localhost ~]# nvme list
Node                 SN                    Model
Namespace Usage                           Format            FW Rev
---------------- -------------------- ----------------------------
-------- ------------------------- ---------------- --------
/dev/nvme0n1     S3H9NX0J700013       SAMSUNG MZWLL6T4HMLS-00003  2
53.69  GB /  53.69  GB    512   B +  0 B   GPNA6B3T
/dev/nvme0n2     S3H9NX0J700013       SAMSUNG MZWLL6T4HMLS-00003  5
134.22  GB / 134.22  GB    512   B +  0 B   GPNA6B3T
```

FIO Test Result– Direct vs SR-IOV

Performance	Throughput	Latency (4K R)
Directly attached	**3180 MB/s**	**91 usec**
Aggregate SR-IOV	**3058 MB/s**	**90 usec**

Multi-host SR-IOV vs Pass Throught vs Hyper+ SR-IOV

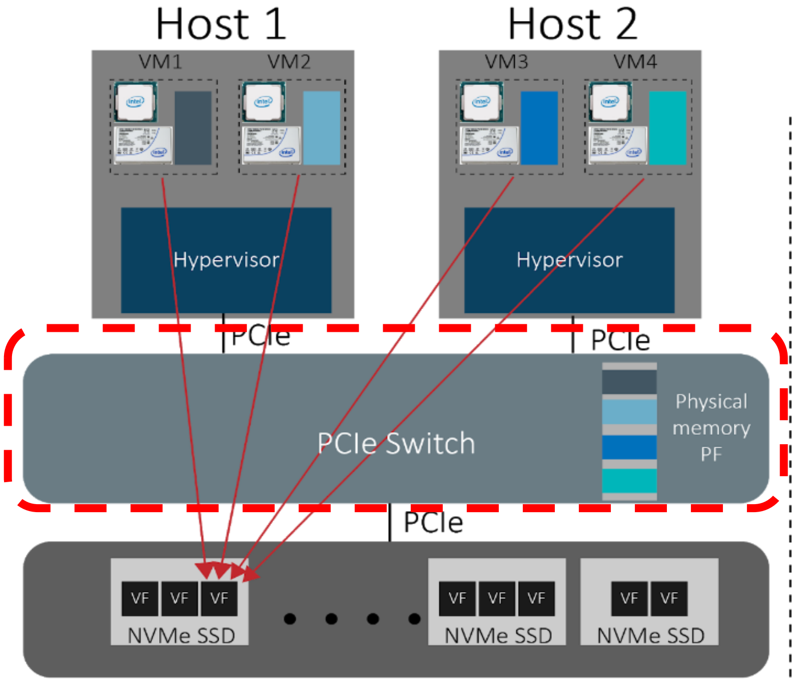

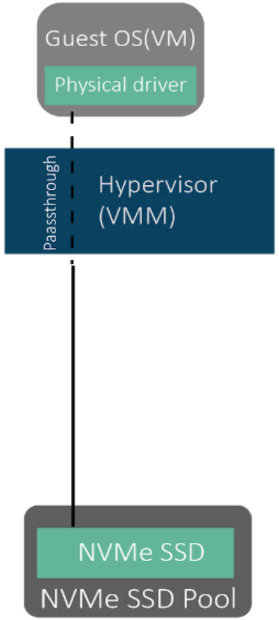

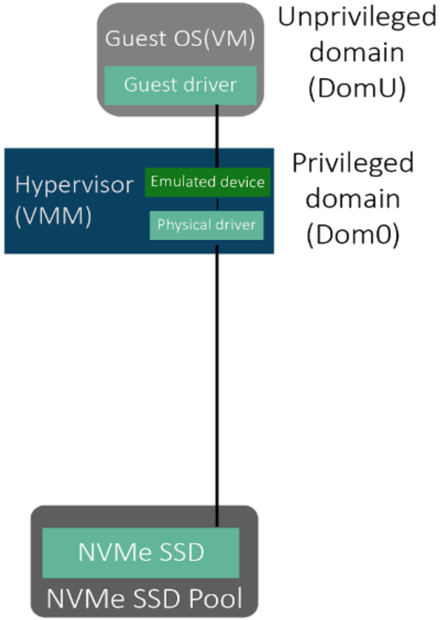

NVMe with SR-IOV. VM talk to VF directly. PF is sit on PCIe switch.

NVMe without SR-IOV. Passthrough model

NVMe without SR-IOV. Hypervisor manage VMs to NVMe SSD

Benefits of Sharing NVMe Through SR-IOV

- Performance and latency
  - VF latency is only 1/3 of PF latency in multi-VMs
- Cost saving
  - Tens of VFs associated with a single PF, extending the capacity of a device and lowering the hardware cost
  - With better latency and performance, the utilization rate will be higher to further reduce the hardware cost

Benefits of Sharing NVMe Through SR-IOV

- Multi-path IO via PCIe
  - The namespace on NVMe can be accessed by different hosts through PCIe connection
- Flexibility configuration
  - Dynamical control by the mCPU to assign VF or create namespace, users have the flexibility to manage NVMe SSD

Comparison of Different Interconnect

Technology	PCIe	Ethernet	Fiber	Infiniband
Latency (us)	0.25	10	5	3
Perf. Gbps, single port	256	400	64	200
Applications	I/O Area Network (IAN)	Local Area Network (LAN)	Storage Area Network (SAN)	I/O Area Network (IAN)

Comparison of Different Interconnect

Technology	PCIe	Ethernet	Fiber	Infiniband
Scalibility	Low	Very high	Medium	High
Ease of management	Low	High	Medium	Medium

Performance Results with SR-IOV

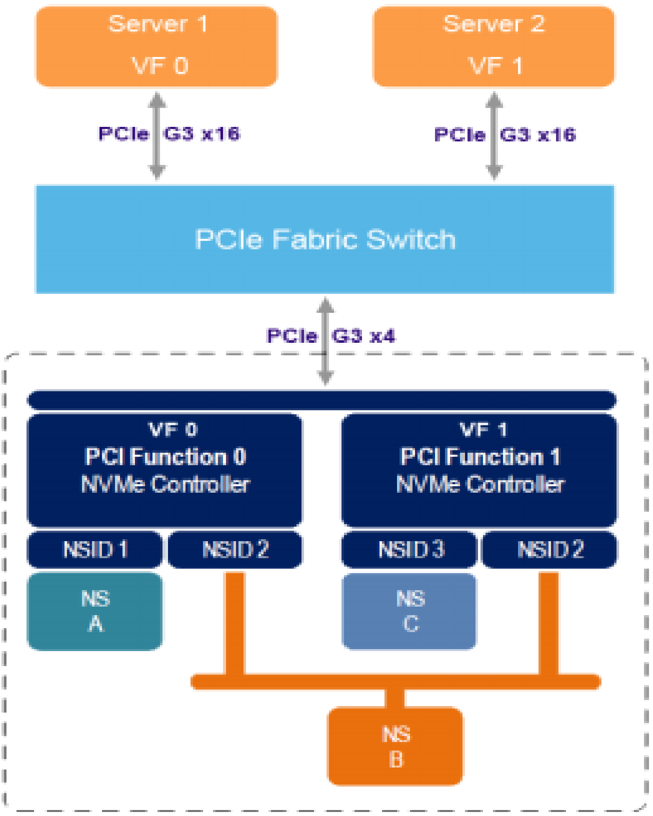

Latency

Server_1 access to VF_0 (NS_A)

4K Read (Random)	
Latency (usec) avg. and distribution	
90	50-100=97.70%, 100-250=2.30%

Server_1 access to VF_0 (NS_A)

4K Write (Random)	
Latency (usec) avg. and distribution	
25	10-20=85.11%, 20-50=14.84%

Note

The Latency measured using Fio in CentOS 7.5, with queue depth 1 by 1 worker and CPU core allowed 1.

The performance measured using Fio in CentOS 7.5, with queue depth 32 by 16 workers and CPU core allowed 8.

Performance Results with SR-IOV

Read Performance

Tested Functions	4K Read (Random)		4K Read (Sequential)	
	MB/s	IOPS	MB/s	IOPS
Server1 access to VF0 (NS_A)	1,496	382,990	1,414	362,011
Server2 access to VF1 (NS_C)	1,562	400,016	1,832	469,080
Server1 access to VF0 (NS_B)	1,494	382,601	1,402	359,052
Server2 access to VF1 (NS_B)	1,560	399,519	1,815	464,725

Write Performance

Tested Functions	4K Write (Random)		4K Write (Sequential)	
	MB/s	IOPS	MB/s	IOPS
Server1 access to VF0 (NS_A)	924	230,984	849	212,427
Server2 access to VF1 (NS_C)	979	244,881	1,319	337,891
Server1 access to VF0 (NS_B)	953	238,382	851	212,874
Server2 access to VF1 (NS_B)	988	253,093	1,366	349,872

Performance Results without SR-IOV

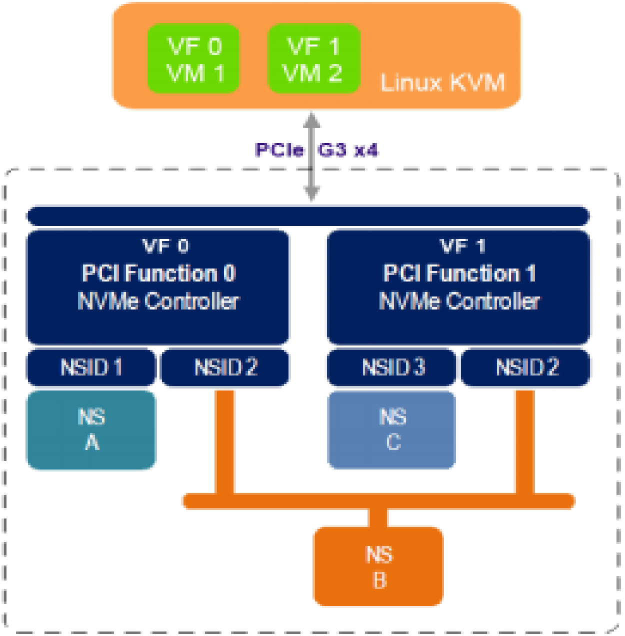

Latency

Linux access to VF_0 (NS_A)

4K Read (Random)	
Latency (usec) avg. and distribution	
91	50-100=92.42%, 100-250=7.57%

Linux access to VF_0 (NS_A)

4K Write (Random)	
Latency (usec) avg. and distribution	
18	10-20=96.71%, 20-50=3.24%

Note

The Latency measured using Fio in CentOS 7.5, with queue depth 1 by 1 worker and CPU core allowed 1.

The performance measured using Fio in CentOS 7.5, with queue depth 32 by 16 workers and CPU core allowed 8.

Performance Results without SR-IOV

Read Performance

Tested Functions	4K Read (Random)		4K Read (Sequential)	
	MB/s	IOPS	MB/s	IOPS
Linux access to PF	3,180	814k	3,420	874k
Linux access to VF0	3,130	801k	3,400	872k
VM 1 access to VF0 (NS_A)	2,822	722k	3,126	800k
VM 1 access to VF0 (NS_B)	1,611	415k	--	--
VM 2 access to VF1 (NS_B)	1,570	403k	--	--

Write Performance

Tested Functions	4K Write (Random)		4K Write (Sequential)	
	MB/s	IOPS	MB/s	IOPS
Linux access to PF	1,716	439k	--	--
Linux access to VF0	1,783	457k	--	--
VM 1 access to VF0 (NS_A)	1,782	456k	--	--

Thank You

Rupin Mohan

Rupin Mohan is a Director of R&D and CTO of Storage Networking (SAN) at HPE Storage. Rupin leads a global engineering team responsible for development of Storage Networking products. Rupin has filed 30+ patents at HPE. He is a Board Member and Marketing Chairman for FCIA. Rupin completed his MBA from MIT Sloan School of Management as a Sloan Fellow. He also holds a MS in Engineering from Tufts University and BE in Computer Engineering from Delhi Institute of Technology.

NVMe over Fibre Channel

Rupin Mohan
Director R&D, Chief Technologist (SAN)
HPE Storage

Fibre Channel: Timeline

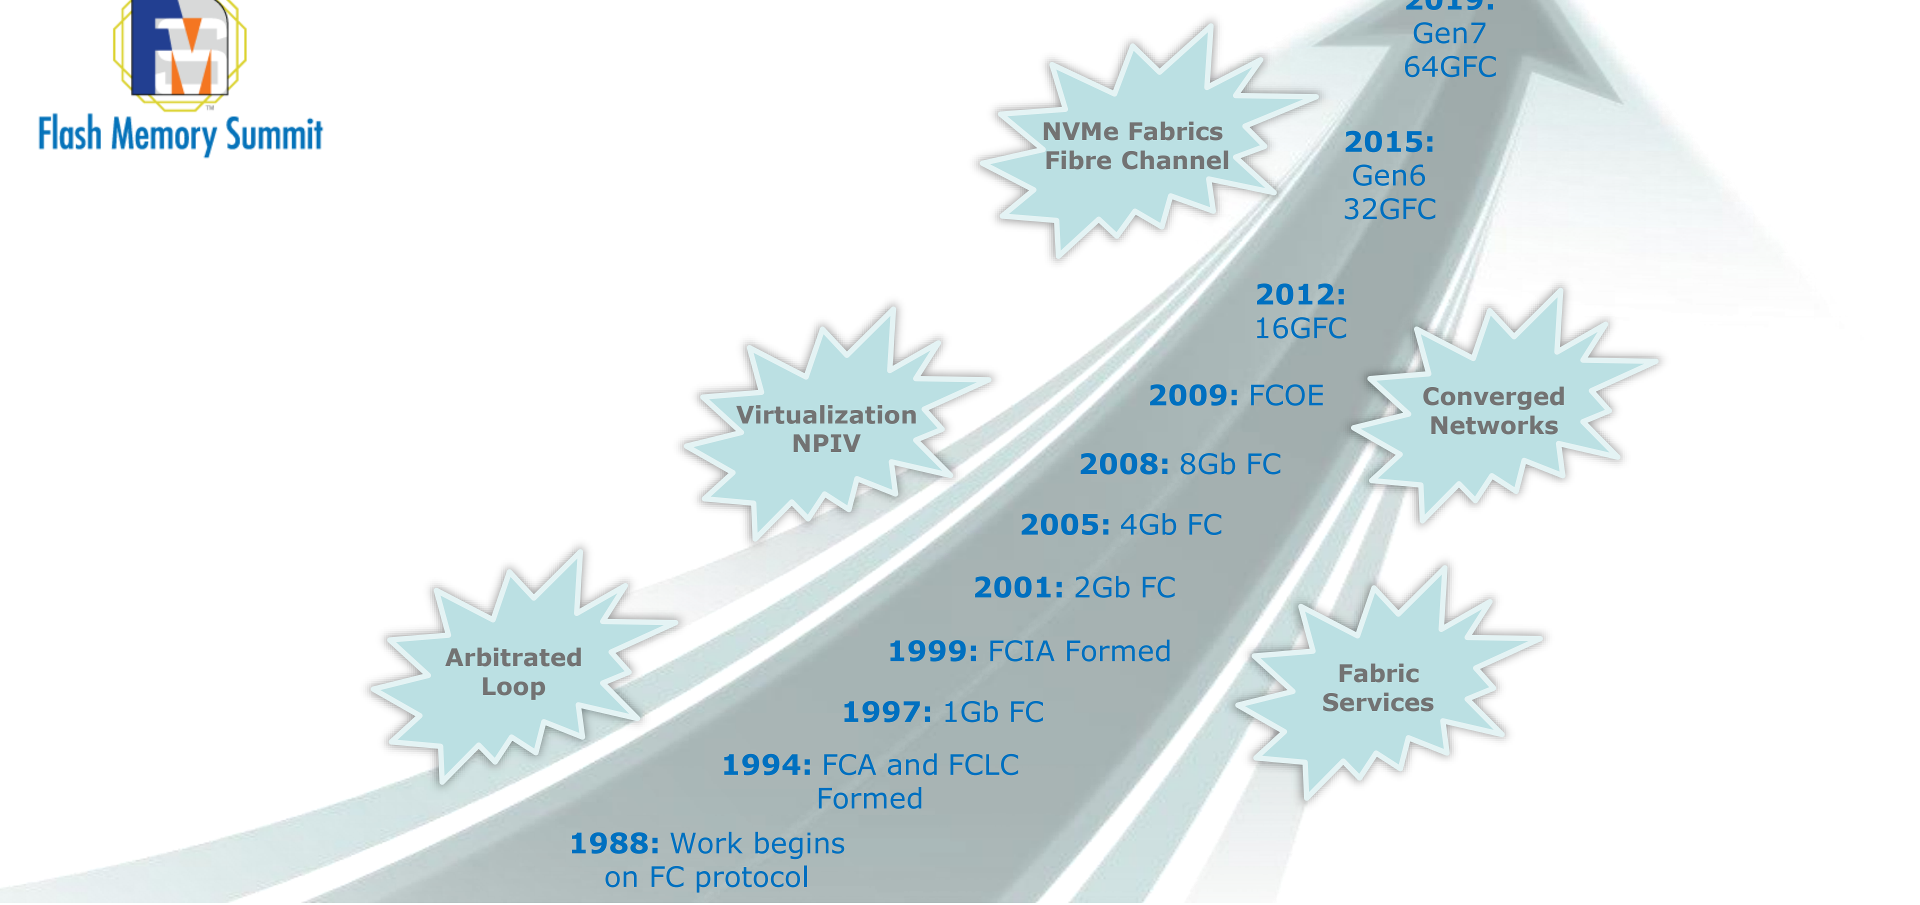

2019 Data Center Storage Trends

Flash Memory Summit

Workloads are being brought back from the public cloud

- Lack of understanding performance or sensitivity requirements
- Last year 41% of business brought at least one workload back*

Rapid growth of all-flash arrays, NVMe & NVMe-oF

- NVMe over FC expected to outpace NVMe over Ethernet starting this year*
- NVMe over FC is easy to deploy, high-performance, extremely reliable

Analytics/Artificial Intelligence Requirements

- Requires great volumes of data & fast access to it. Storage has become disaggregated due to public cloud & edge. Finding data takes too long. Lost time = diminished value of analytics and AI initiatives.

*ESG, 1/7/19, 2019 Data Storage Predictions: More Cloud Missteps, FC Is Back, and Finding Data Holds Back AI

Fibre Channel Return to Growth

- According to Dell'Oro Group's latest Storage Area Networks (SAN) 5-Year Forecast Report (final data will be available in Dell'Oro Group's 4Q18 report):
  - Total Fibre Channel SAN port shipments (Fibre Channel switch and Fibre Channel adapter) for 2018 is expected to approach 7.7 million, up more than 11 percent over 2017 (6.8 million)
  - Total Fibre Channel SAN revenue (Fibre Channel switch and Fibre Channel adapter) is expected to approach $2.5 billion in 2018, up nearly 22 percent over 2017

127M Total Ports Shipped (since 2001)

40M Still in Service (last 5 years)

~9M Average Per Year During "Peak" years (2007-2014)

Flash Memory Summit

FC-NVMe-2

NVMe over Fabrics

WHATS NEW?

- Refinements to existing FC-NVMe standard
- Enhances error recovery to more granular level
- Improves performance predictability
- Prevents connection disruption

HOW

- Sequence level error recovery – SLER - (exchanges are given more time before getting terminated, re-transmissions, new commands)
- A method to respond/correct minor error conditions to avoid NVMe subsystem disconnect/reconnect
- Will auto-negotiate for compatibility with FC-NVMe connections

WHEN

- Now a published standard.

NVMe-oF deployment (FC)

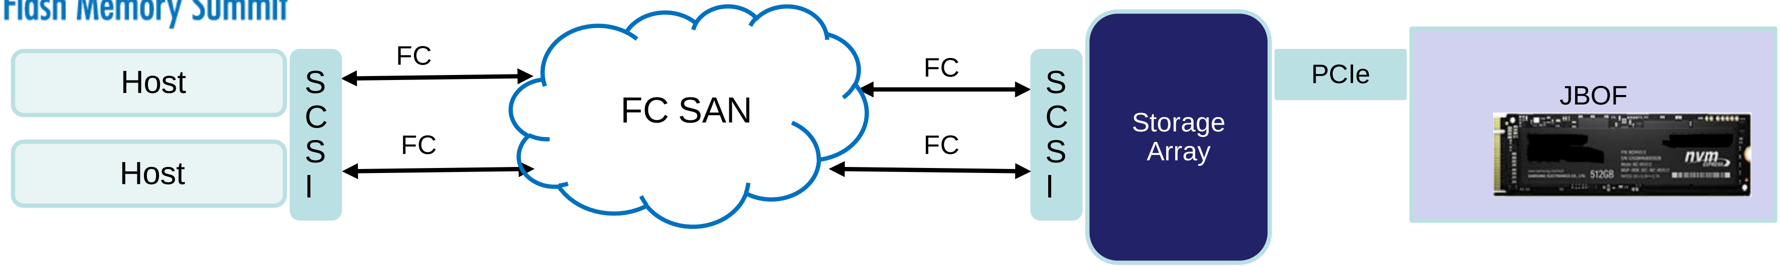

NVMe storage attached in the backend

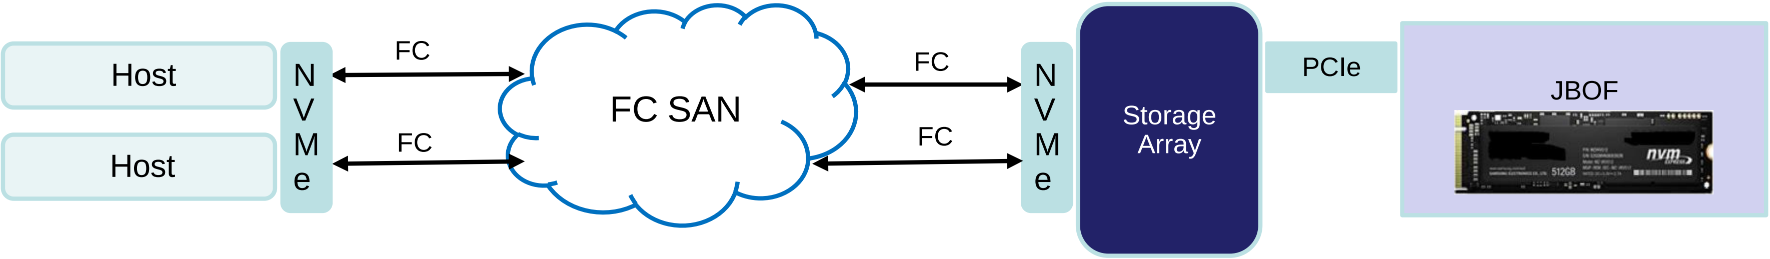

NVMe end to end using FC

The landscape today….

Protocol	Latency	Scalable	Performance	Enterprise Footprint
Fibre Channel	Lower	Yes	High	Reliable Storage Fabric
RoCEv2	Lowest	Yes	High	Negligible
iWARP	Medium	Yes	Medium	Negligible
TCP	High	Yes	Medium	Medium with iSCSI
InfiniBand	Lowest	Limited	High	None

Flash Memory Summit

TABLE OF CONTENTS

Download: https://fibrechannel.org/

Hardcopy: FCIA booth

Thank You

Ilker Cebeli

Ilker is a Senior Director of Product Planning at Samsung. He is responsible for leading the Emerging memory, SSD, and All-Flash-Array related storage solutions and technologies. He has spent 25 years in enterprise computing, storage, and networking working in various roles. Prior to joining to Samsung, Ilker worked at Micron, and was leading and directing emerging memory projects in memory division. Ilker also spent 15 years at Intel and he was responsible for Intel's Xeon™ product planning and server platform architecture definition.

Flash Memory Summit

How Networking Affects Flash Storage Systems

Ilker Cebeli

Senior Director of Product Planning, Samsung

Disclaimer

Flash Memory Summit

This presentation and/or accompanying oral statements by Samsung representatives collectively, the “Presentation”) is intended to provide information concerning the SSD and memory industry and Samsung Electronics Co., Ltd. and certain affiliates (collectively, “Samsung”). While Samsung strives to provide information that is accurate and up-to-date, this Presentation may nonetheless contain inaccuracies or omissions. As a consequence, Samsung does not in any way guarantee the accuracy or completeness of the information provided in this Presentation.

This Presentation may include forward-looking statements, including, but not limited to, statements about any matter that is not a historical fact; statements regarding Samsung’s intentions, beliefs or current expectations concerning, among other things, market prospects, technological developments, growth, strategies, and the industry in which Samsung operates; and statements regarding products or features that are still in development. By their nature, forward-looking statements involve risks and uncertainties, because they relate to events and depend on circumstances that may or may not occur in the future. Samsung cautions you that forward looking statements are not guarantees of future performance and that the actual developments of Samsung, the market, or industry in which Samsung operates may differ materially from those made or suggested by the forward-looking statements in this Presentation. In addition, even if such forward-looking statements are shown to be accurate, those developments may not be indicative of developments in future periods.

SAMSUNG

Flash Memory Summit

NVMe Technology – Background

- Optimized for flash
  - Traditional SCSI designed for disk
  - NVMe bypasses unneeded layers
  - Dramatically reducing latency

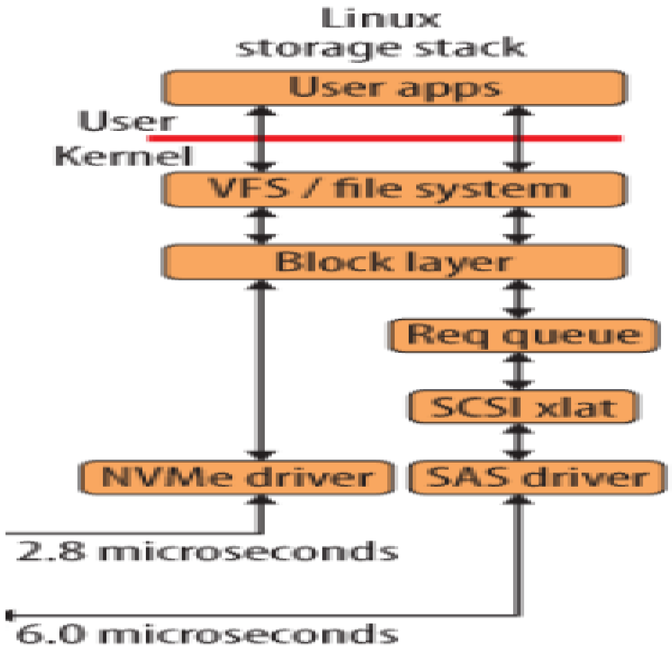

NVMe Design Advantages

- Lower latency
  - Direct connection to CPU's PCIe lanes
- Higher bandwidth
  - Scales with number of PCIe lanes
- Best in class latency consistency
  - Lower cycles/IO, fewer cmds, better queueing
- Lower system power
- No HBA required

NVMe Technology – Background

- NVMe outperforms SATA SSDs
  6X-7X more bandwidth,
  40-50% lower latency
  Up to 8x more IOPS

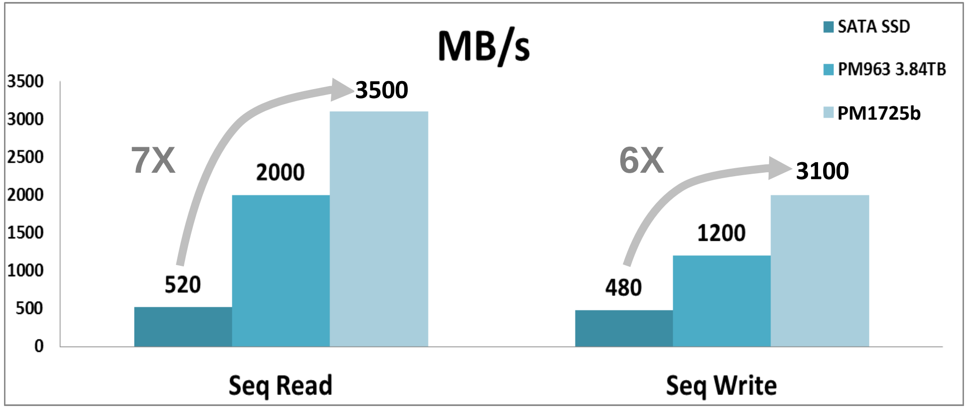

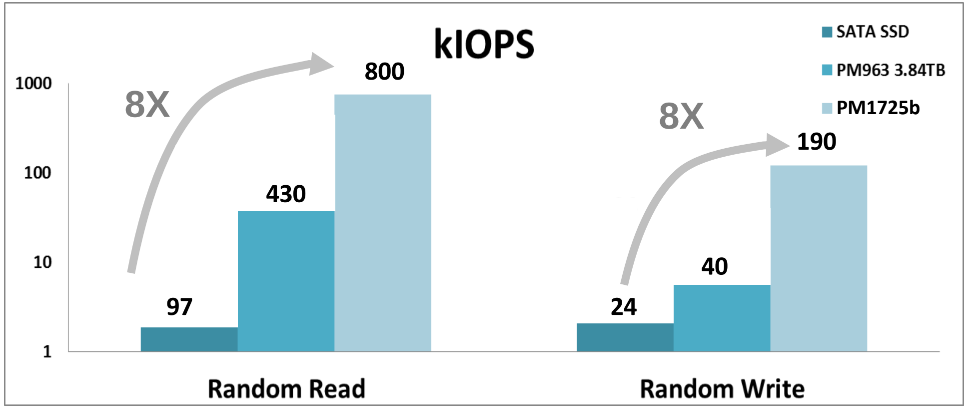

What is NVM Express Over Fabrics?

- A protocol interface to NVMe that enable operation over other interconnects (e.g., Ethernet, InfiniBand™, Fibre Channel).
- Shares the same base architecture and NVMe Host Software as PCIe
- Enables NVMe Scale-Out and low latency (<10µS latency) operations on Data Center Fabrics
- Avoids protocol translation (avoid SCSI)

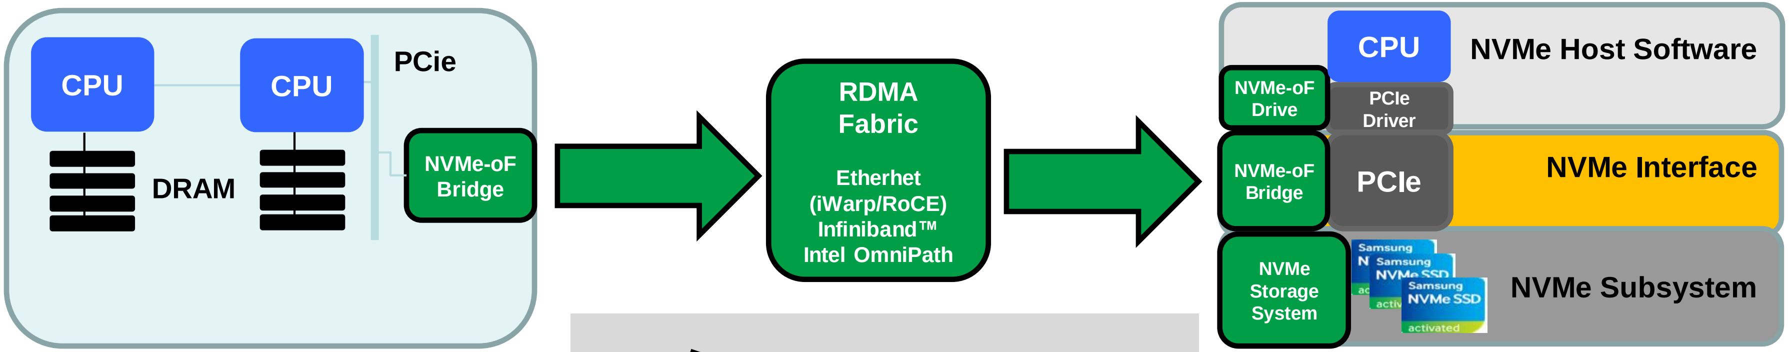

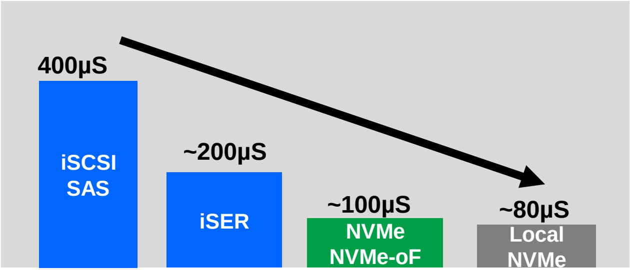

Some of the use cases for NVMe Over Fabrics

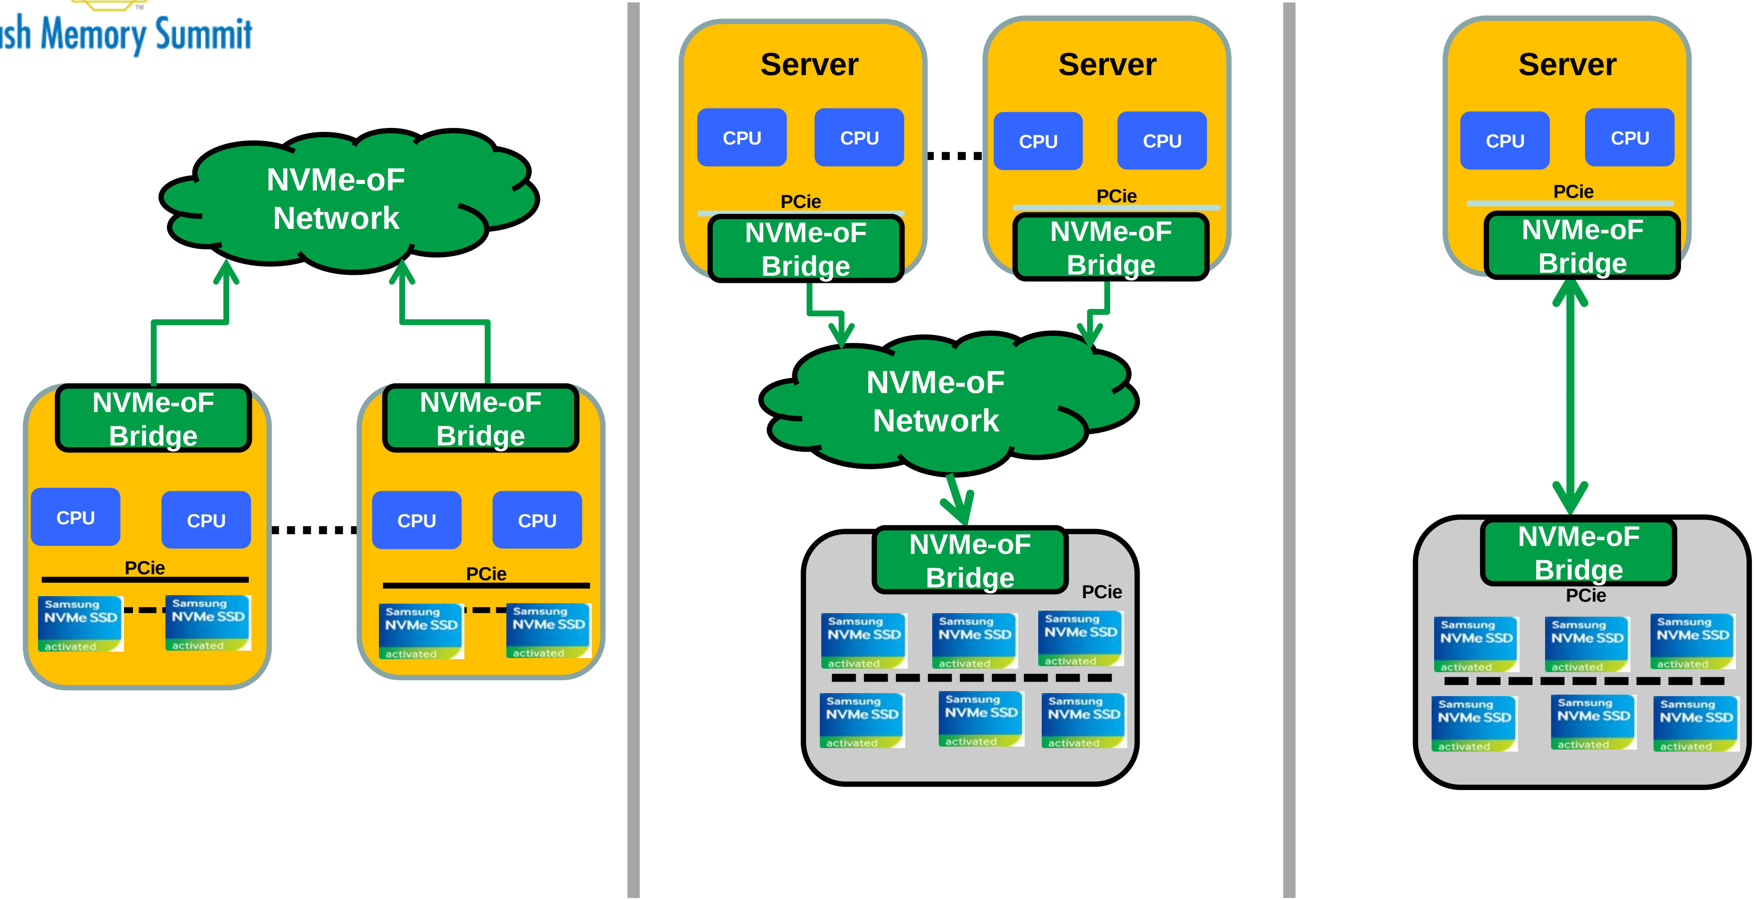

Performance Test Configuration – 2016

Flash Memory Summit

- **1x NVMe-oF target**
  - 24x NVMe 2.5” SSDs
  - 2x 100GbE NICs
  - Dual x86 CPUs
- **4x initiator hosts**
  - 2x25GbE NICs each
- **Open Source NVMe-oF kernel drivers**

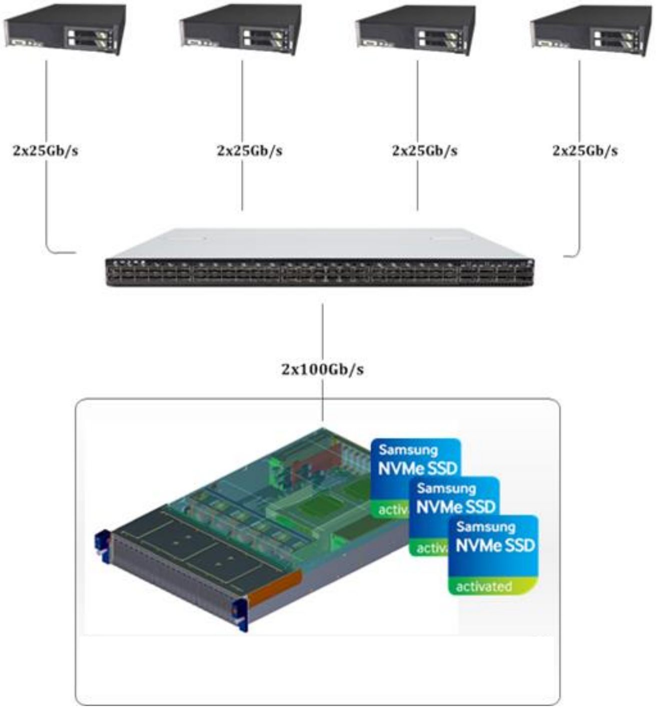

Local vs. Remote Latency Comparison – 2016

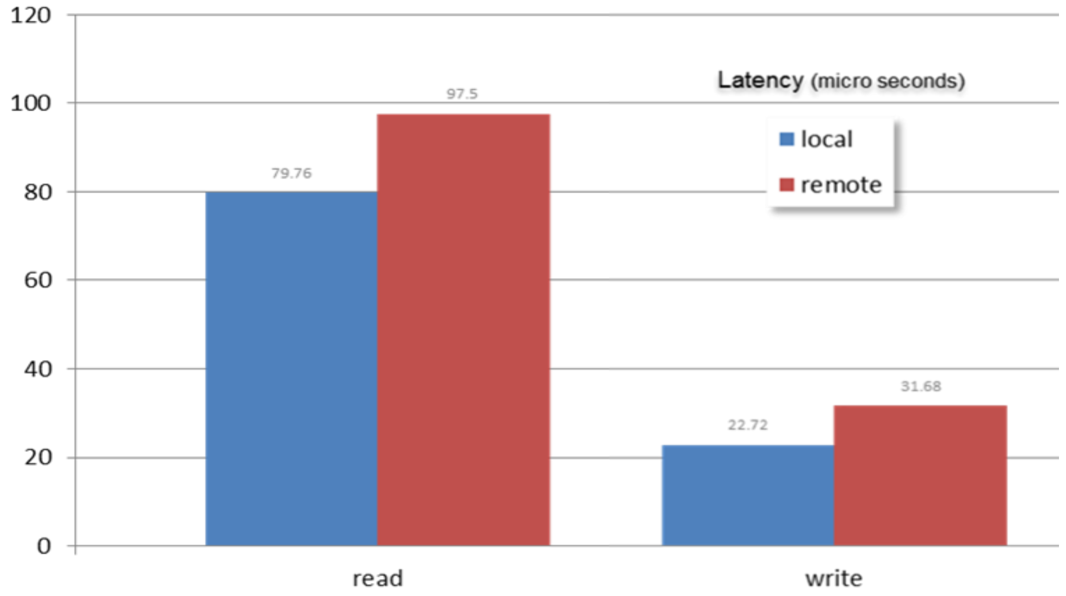

Read Gap	Write Gap
~17 us	~9 us

Performance Test Configuration – 2017

Flash Memory Summit

- **1x NVMeoF target**
  - 36x NF1 SSDs
  - 2x 100GbE NICs, 2x 50GbE NICs
  - Dual x86 CPUs
- **6x initiator clients**
  - 2x25Gb/s each
- **Open Source NVMe-oF kernel drivers**

- Ubuntu Linux 16.04/4.9 on Target

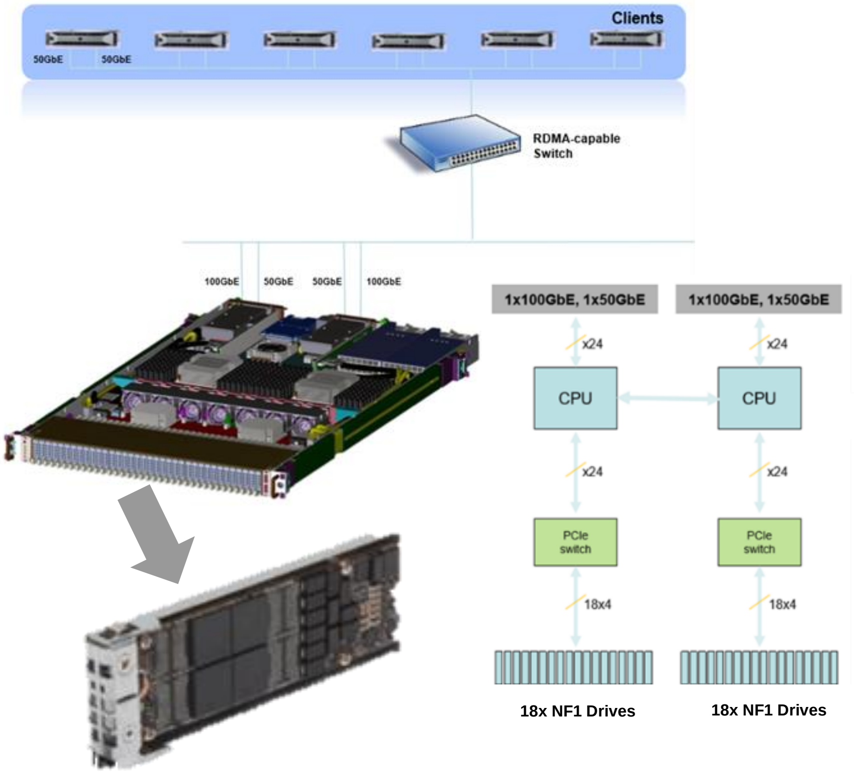

Flash Memory Summit

Local vs. Remote Latency Comparison - 2017

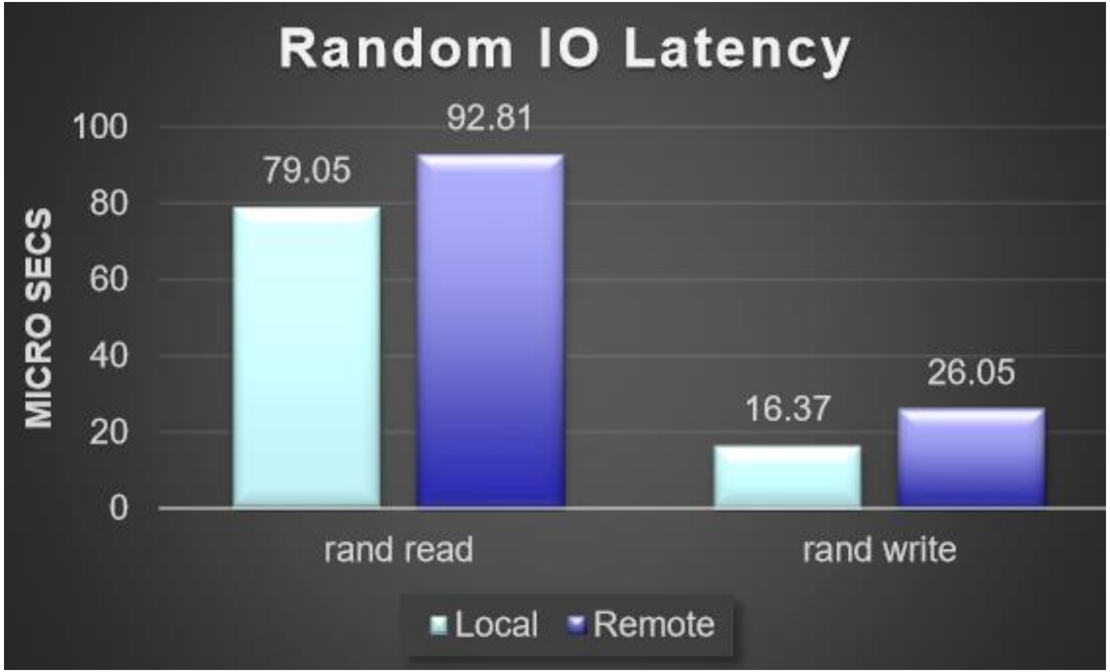

Read Gap	Write Gap
~14 us	~10 us

Read Gap	Write Gap
~17 us	~9 us

Flash Memory Summit

SSDs Will Continue to get Faster

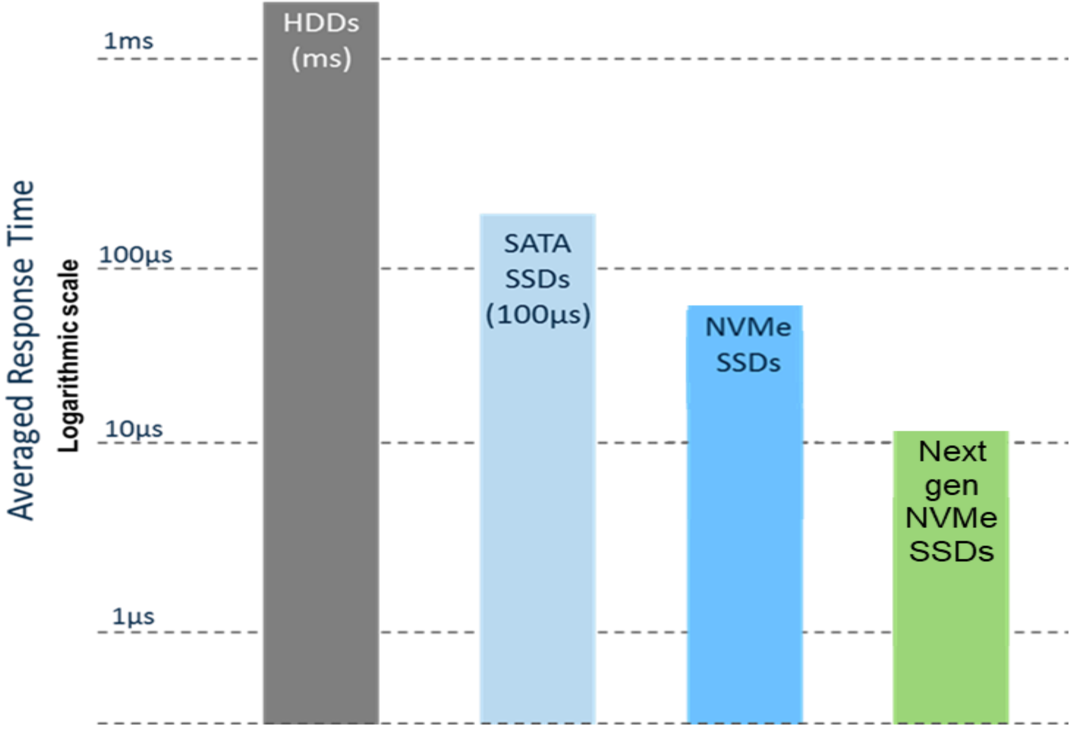

Read Gap	Write Gap
~14 us	~10 us

Read Gap	Write Gap
~17 us	~9 us

Thank You

Alan Weckel

Alan Weckel is Technology Analyst/Co-Founder at 650 Group, where he is in charge of Ethernet switch, Cloud and data center research. He has written many articles for the trade and technical press, and is frequently quoted in such leading publications as Bloomberg, Businessweek, Forbes, Network World, and the Wall Street Journal. Before co-founding 650 Group, he was VP/analyst at Dell’Oro Group and had engineering and software development experience at Raytheon, General Electric Power Systems, and Cisco. He holds a BSEE and an MS in Management from Rensselaer Polytechnic Institute.

Flash Storage Networking, How the market is evolving

Alan Weckel (alan@650group.com)

Trends changing how compute and storage are consumed

Storage: How and Where We Store Data is Changing

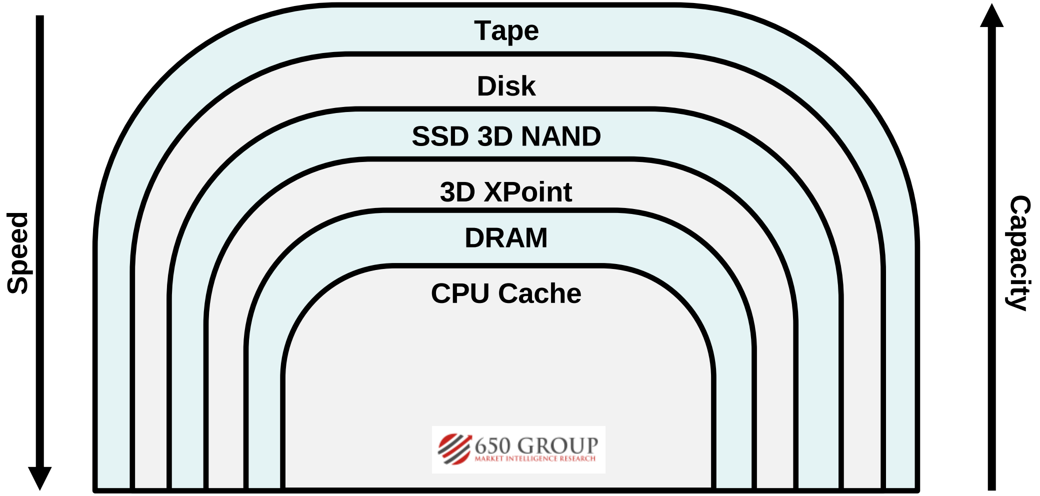

- Enterprise Storage Systems Market is Shrinking
  - Enterprises continue to buy systems
  - Enterprise market for converged and hyperconverged is growing
- Cloud Market is Growing
  - Hyperscalers buy components
  - Hyperscalers build their own software
- Areas of growth in Storage Systems Market
  - Cloud
  - All Flash Arrays
  - Hyperconverged

Server Shipments: Shipments into the Cloud

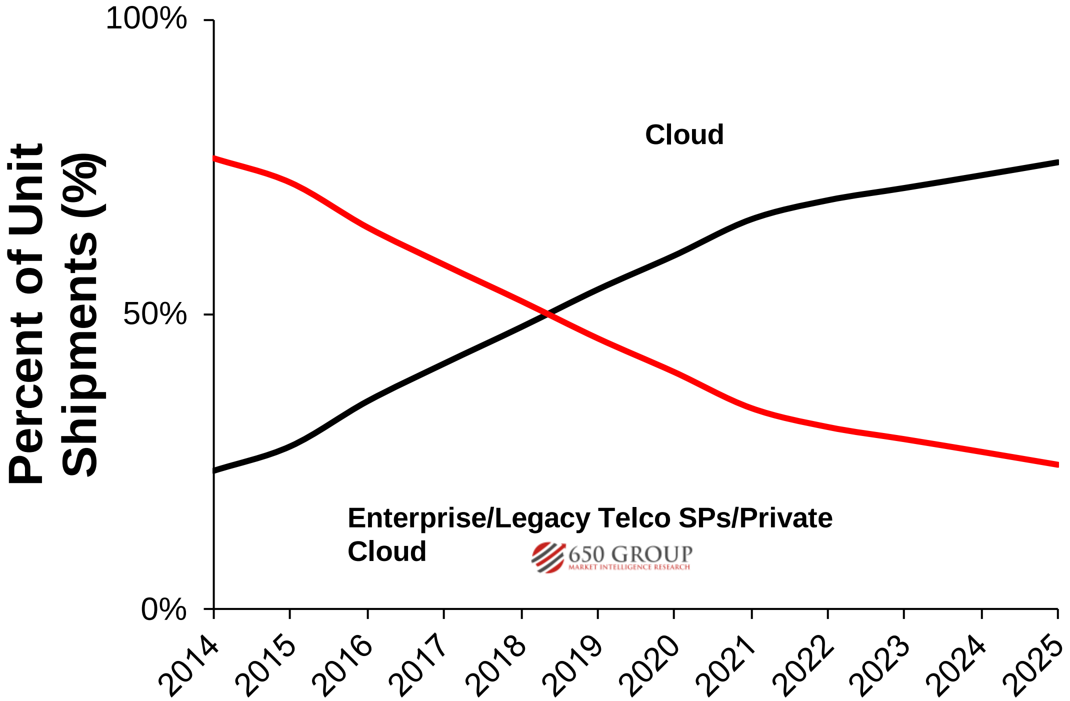

- Cloud servers will dominate compute
  - Higher-end processor
  - Smart NIC
  - Better software
  - Different type of storage
- Enterprise servers are being deployed in colocation facilities
- East/West traffic is no longer limited to one data center
  - Ethernet Based Architectures
  - Large amounts of data being moved across the world

Workloads: Installed Base by Deployment

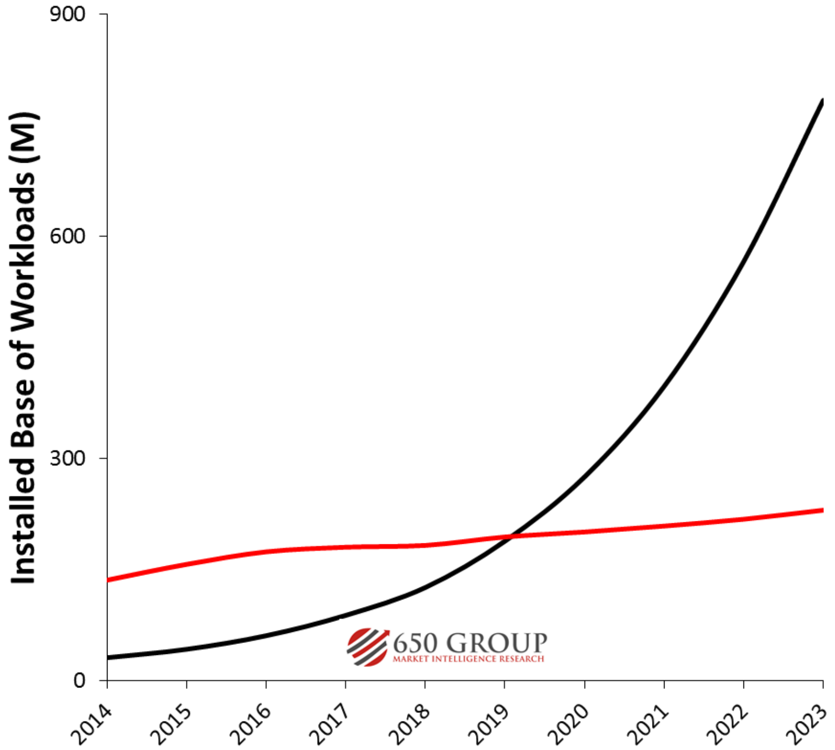

- Enterprise workloads continue to grow
  - More workloads per server
  - Type of application is changing
  - Colocation becoming common
- Cloud workload growth exploding
  - All types of applications are growing
  - IoT will be a major driving of workload growth
  - Edge Computing and AI changing DC design

Server and Smart NICs: Server Installed Base

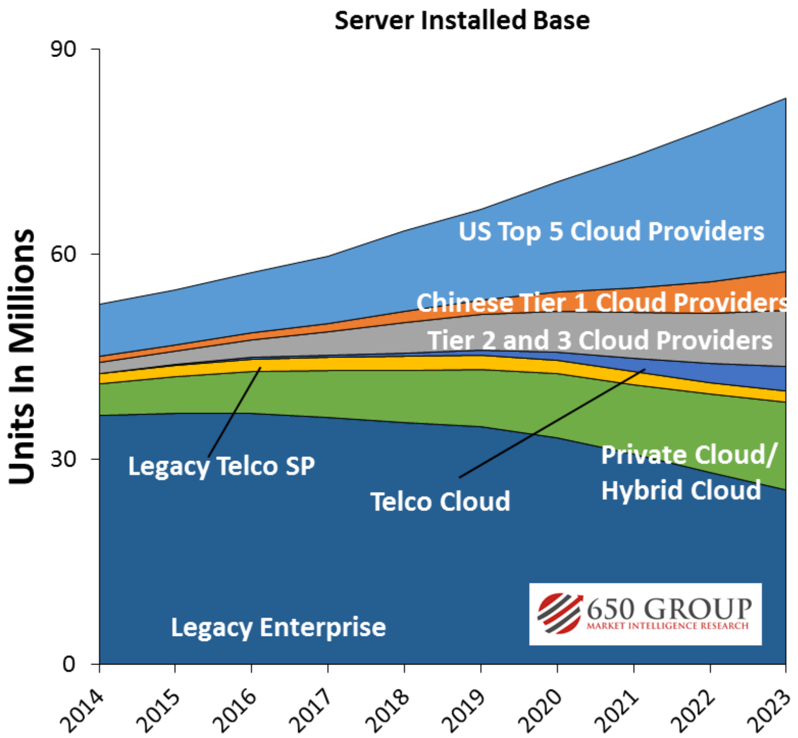

- Cloud is the new leader in technology transitions
  - Entire Telco market is smaller than Amazon
  - Cloud is moving from 2-3 to 3-4 technology generations ahead of the enterprise
- Tier 2 and 3 Clouds are increasingly riding on top of Tier 1 Cloud Infrastructure
- Clouds uses different architecture and buys different equipment then the enterprise

Ethernet Switch – Data Center

Ethernet Switch – Data Center: Total Market Revenue

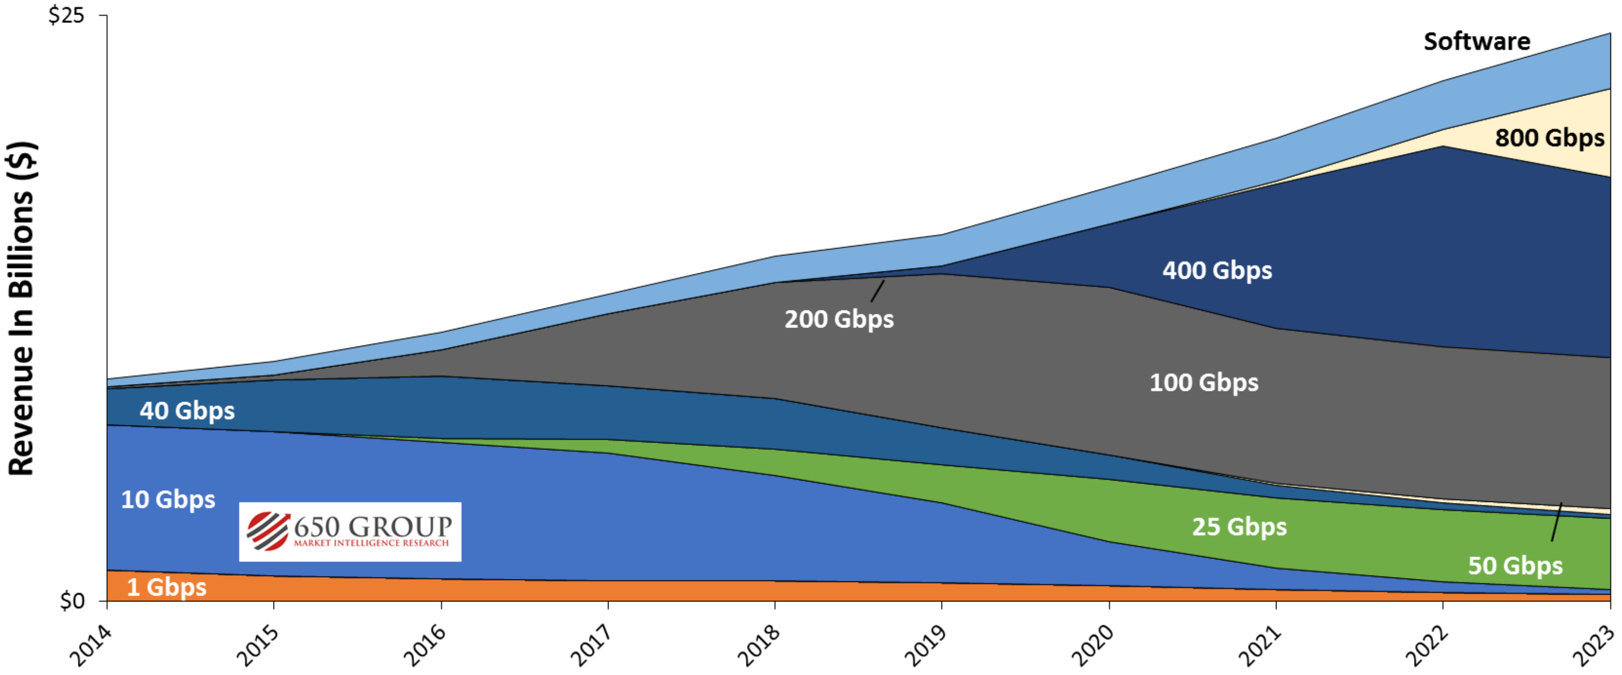

Ethernet Switch – Data Center: Total Market Revenue

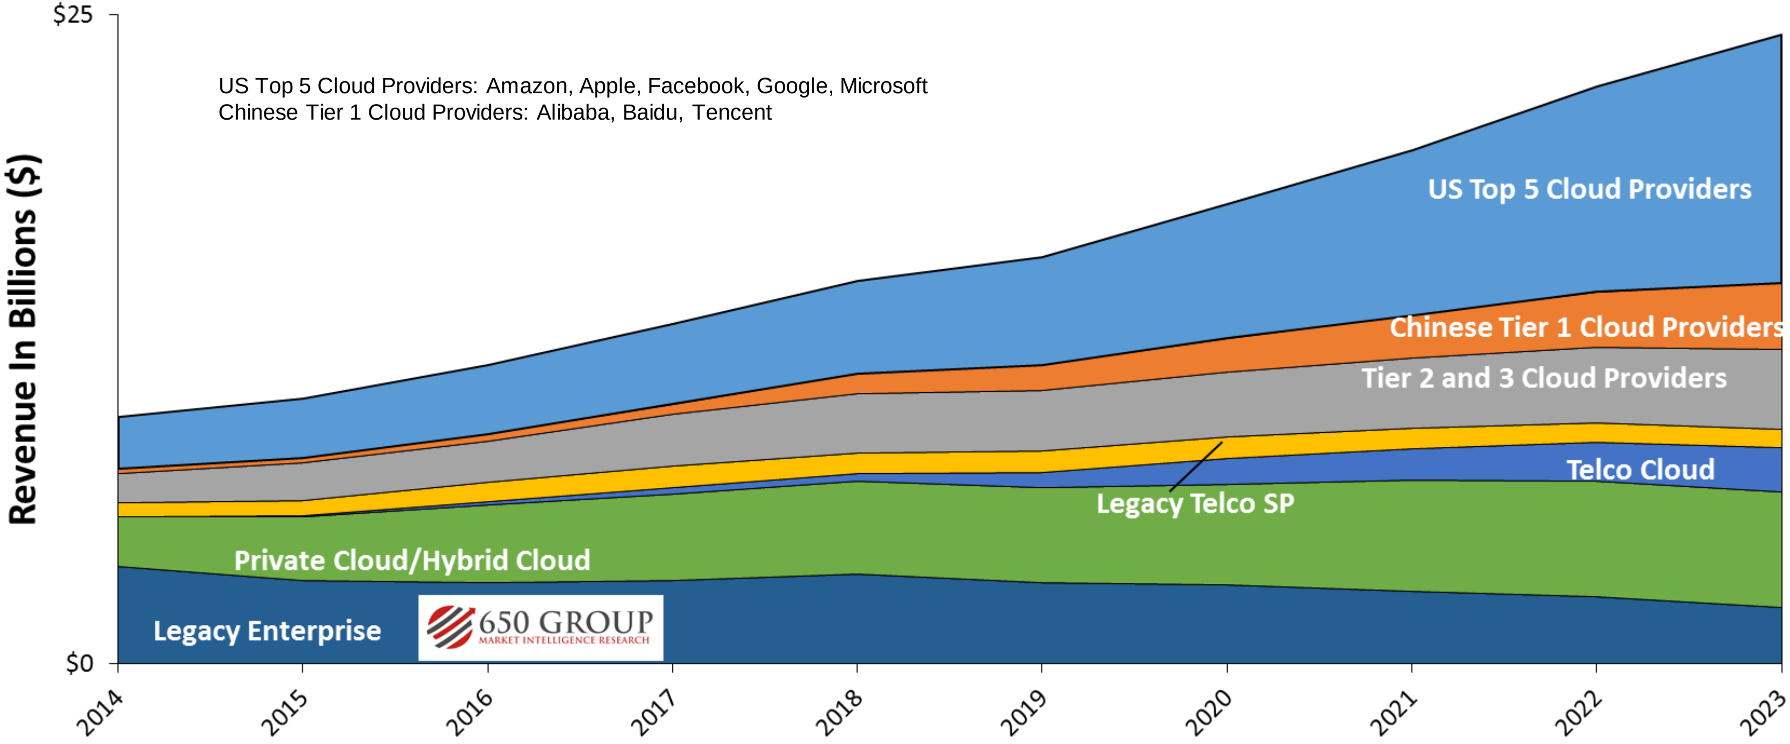

Fibre Channel

- FC Switches:
  - Shipments increased 15% in 2018 following six consecutive years of decline
- FC HBAs:
  - Shipments increased 10% in 2018
- Reasons for Fibre Channel switch and HBA growth:
  - Very strong economic growth for most of 2018 with Enterprise IT spending following suit.
  - Server and external storage array market recovery
  - Focus and execution following acquisitions
  - Some advance buying
- 2019 forecast assumes low single-digit shipment increase.

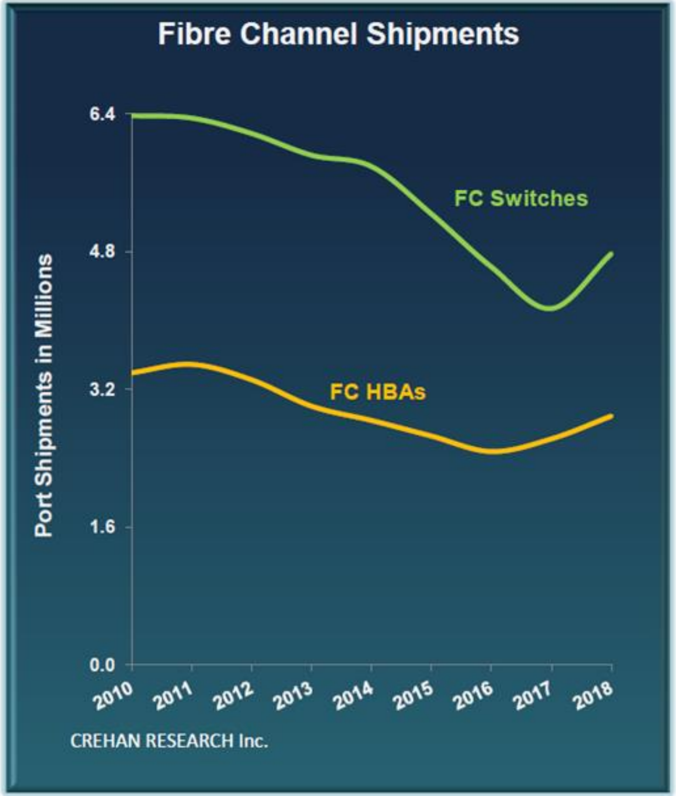

InfiniBand

- 2018 Shipments
  - Up 8% Y/Y to ~ 700K ports
    - Following steep 2017 decline
- 2018 Revenue
  - Up 6% Y/Y to $217M
- HDR/200Gbs HCA shipments very small in 2018
  - But starting to ramp very strongly for Switches
    - 3Q18: ~1K ports
    - 4Q18: ~20K ports

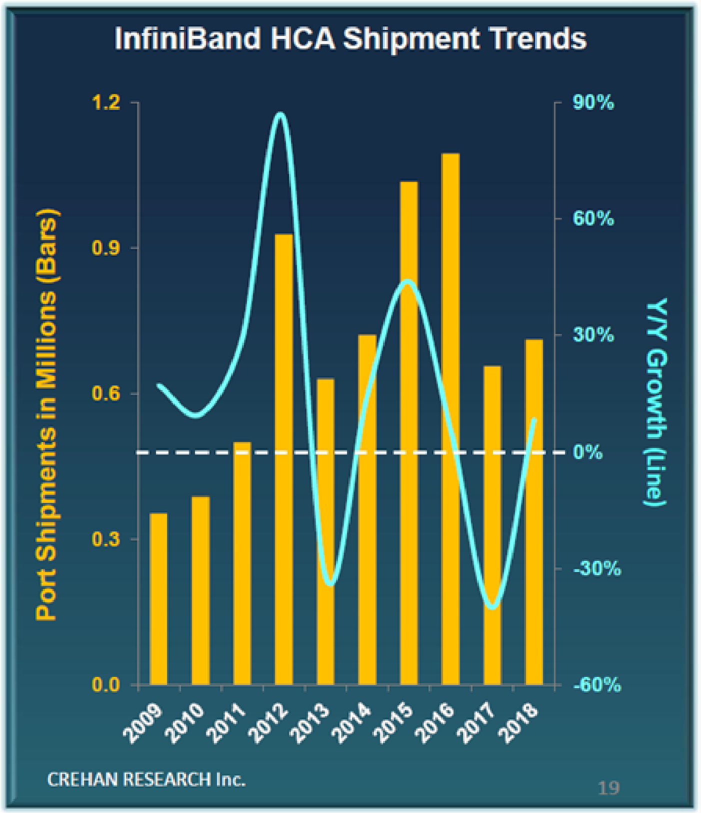

Merchant Silicon – Data Center Switching: Total SERDES Shipments

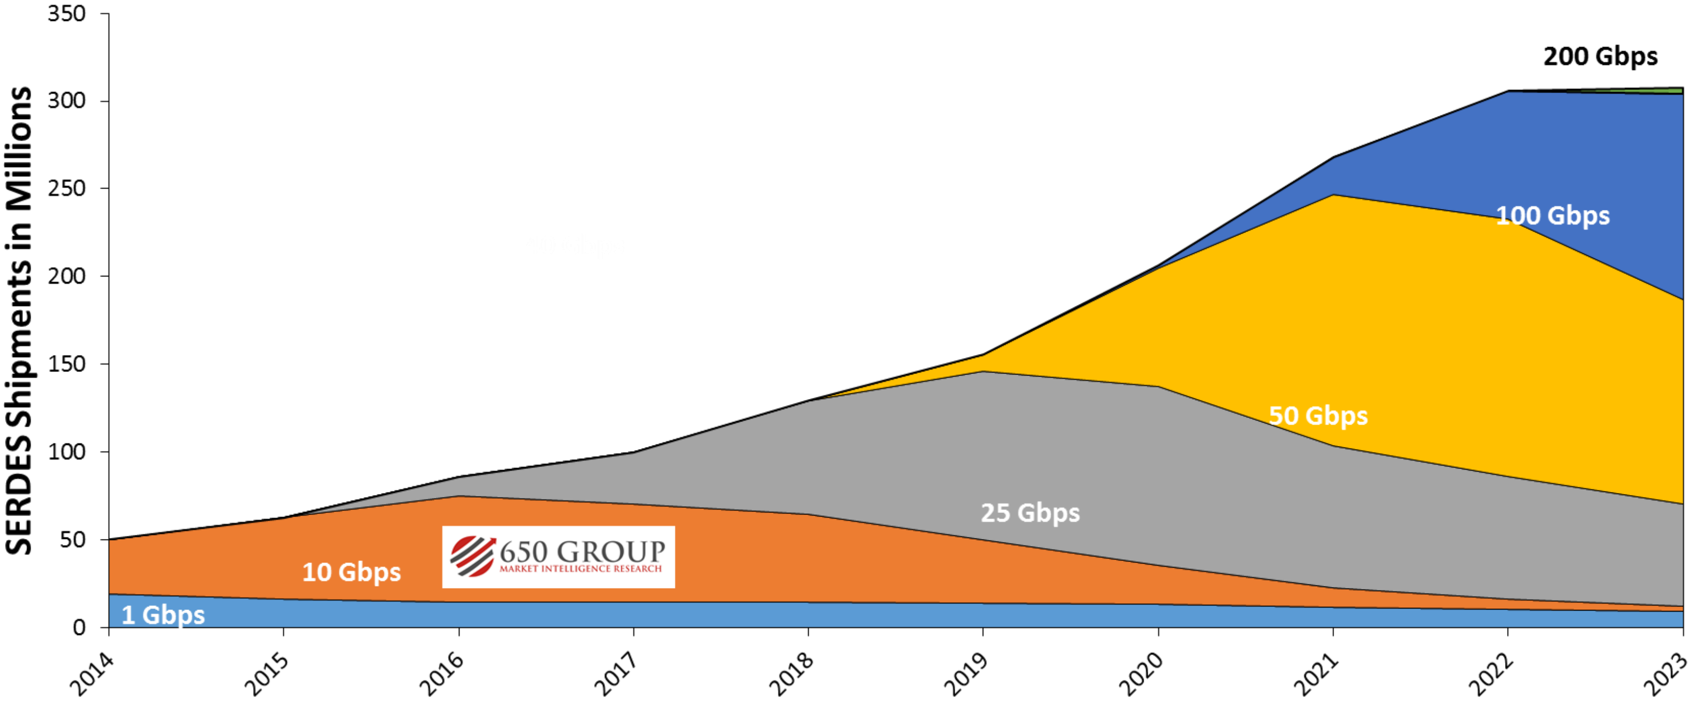

Merchant Silicon – Data Center Switching: ASIC Usage in the Tier 1 Cloud

650 GROUP MARKET INTELLIGENCE RESEARCH

Merchant Silicon's product cycles accelerating in the Cloud

ASIC Size	SERDES Technology	Leaf Port Speed	Spine/Core Port Speed
1.3 Tbps	10 Gbps	10 Gbps	10/40 Gbps
1.8 Tbps	25 Gbps	25 Gbps	100 Gbps
3.2 Tbps	25 Gbps	25/50 Gbps	100 Gbps
6.4 Tbps	25 Gbps	25/50 Gbps	100/200 Gbps
7.2 Tbps	100 Gbps	100 Gbps	400 Gbps
12.8 Tbps	50 Gbps	50/100 Gbps	200/400 Gbps
12.8 Tbps	100 Gbps	100 Gbps	400 Gbps
25.6 Tbps	100 Gbps	100 Gbps	800 Gbps

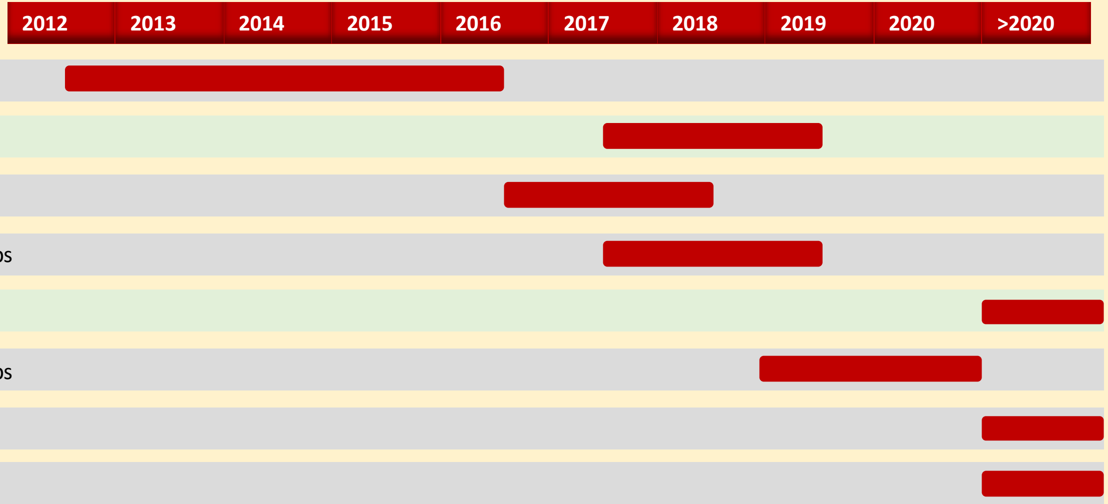

- Two waves of 400 Gbps
  - 8 X 50 Gbps
  - 4 X 100 Gbps
- Pace of Innovation Increasing
  - Four major silicon cycles in five years
  - Some technologies will get orphaned

Conclusion

- Speed of technology advancement is more rapid
- Ethernet is expanding into the Storage connectivity and Data Center transport markets at a rapid pace
- Cloud customers have different architectures and use different equipment then the enterprise
- 2019 will usher in Smart NICs and 200/400 Gbps which will expand the market for Ethernet

Thank You

Panel Q/A

Rob Davis, Ilker Cebeli, Brian Pan, Abdel Sadek, Rupin Mohan, Steve McQuerry, and Alan Weckel

Thank You