SNIA

Flash: Plan for the Disruption

Ethernet Storage Forum Members

The SNIA Ethernet Storage Forum (ESF) focuses on educating end-users about Ethernet-connected storage networking technologies.

Housekeeping

SNIA
Advancing storage & information technology

- Presentation with live Q&A at the end
- Questions submitted via web tool will be answered verbally
- Unanswered questions will be placed on www.sniaesfblog.org
- Request complete the 10 second feedback form

Today's Panel

- David Fair, SNIA–ESF Board of Directors
- Paul Feresten, SNIA Member

SNIA Webinar Agenda

- Introduction to flash storage
- How flash is being deployed in storage systems
- Considerations and tradeoffs
- Performance benefits of flash in storage
- Trends in non-volatile memory
- Network impact of flash in storage

Performance Gap Challenge

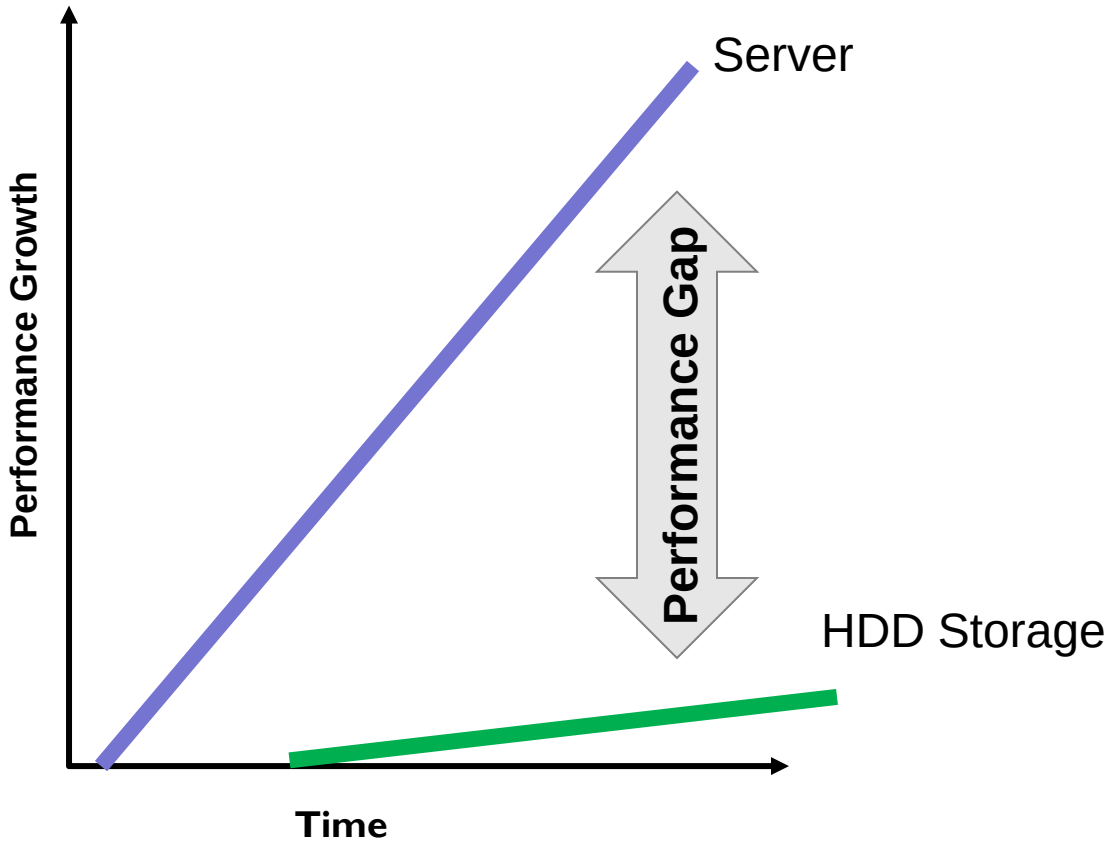

- Huge gap between CPU and Storage
- Relatively small differences between HDD types
- Latency impacts applications

Flash is a Game Changer

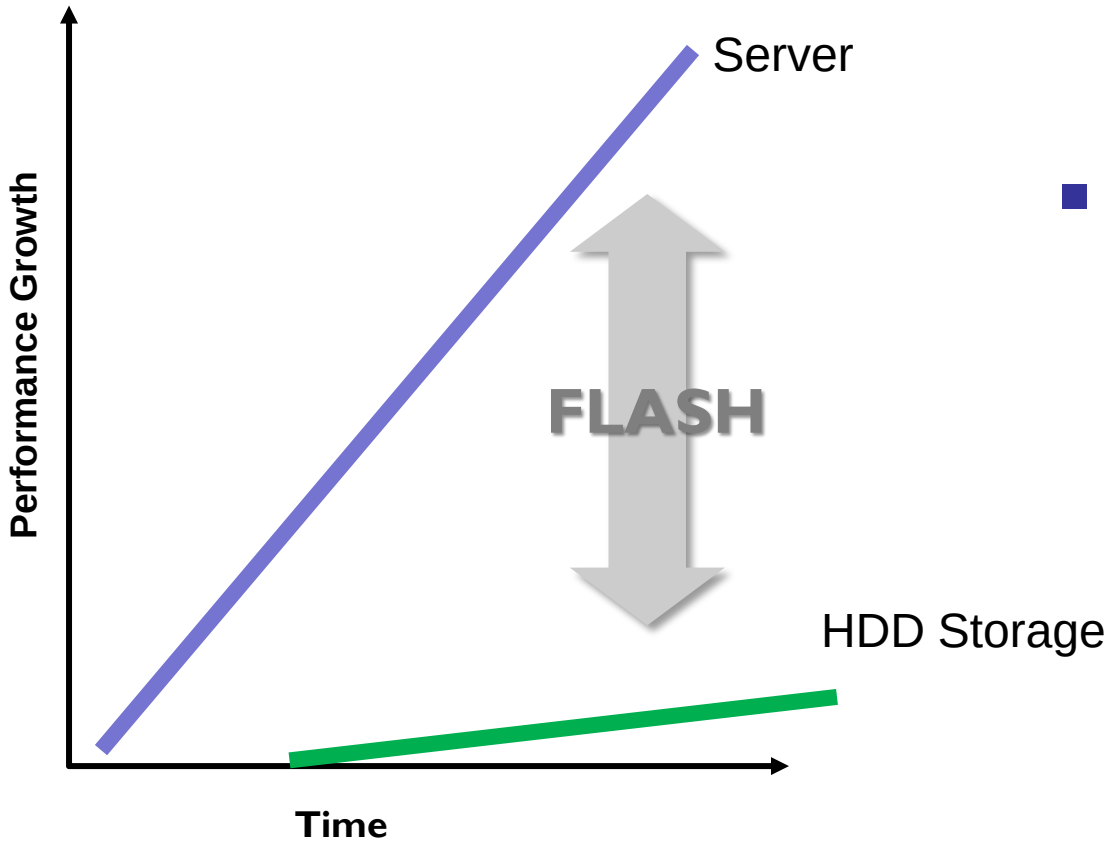

- Flash is a game changer
  - Lower Latency
  - Higher IOPS
  - Lower $ per IOPS

Why Flash in the Data Center Now?

- Why flash?
  - Capacity efficiency versus DRAM
    - ~5x better $ per GB
    - ~40x better power per GB
  - IOPS efficiency versus HDDs
    - ~40x better $ per IOPS
    - ~600x better power per IOPS
- Why now?
  - Period of rapid density advancements led to HDD-like bit density at lower $/GB than DRAM
  - Innovations in SSD and tiering technology

Flash Implementations

Server Centric Server Integrated	Controllers / Accelerators Boot Flash DIMM cards with Flash PCIe cards with Flash
Network Centric Server-attached appliances via PCIe, SAS, FC, iSCSI	SSD Arrays in non-HDD form factors SSD Arrays in HDD form factors
Storage Centric	SSD/PCIe/Hybrid Solutions

NAND-Flash Technologies

- Single-Level Cell (SLC) – One bit per cell
- Multi-Level Cell (MLC) – Two or more bits per cell
  - Triple Level Cell (TLC) – Three bits per cell
  - First announcements of MLC-3 & MLC-4 were made in 2009

	SLC	MLC-2	MLC-3	MLC-4
Bits per cell	1	2	3	4
Performance	Fastest			Slowest
Endurance	Longest			Shortest
Capacity	Smallest			Largest
Error Prob.	Lowest			Highest
Price per GB	Highest			Lowest
Applications	Enterprise	Mostly Consumer	Consumer	Consumer

Source: Demartek 2012

Performance vs. Cost

SNIA
Advancing storage & information technology

	$/GB	$/IOPS	IOPS/watt
SSD (SLC)	$5 – $40	$0.005 – $0.15	1000 – 15000
SSD (MLC)	$0.63 – $4	$0.004 – $0.05	1000 – 15000
HDD (enterprise)	$0.50 – $1	$1 – $3	10 – 30
HDD (desktop)	$0.05 – $0.37	$1 – $4	10 – 40

- SSDs are dollars per gigabyte and pennies per IOPS
- HDDs are pennies per gigabyte and dollars per IOPS
- Notes:
  - Prices sampled in early September 2012 and are subject to change
  - SSD pricing includes drive and PCIe card form factors
  - MLC pricing includes eMLC and cMLC
  - The HDD supply chain appears to have recovered from the Thailand flooding that occurred in the Fall of 2011. HDD availability and prices are beginning to return to pre-flood levels.

Source: Demartek 2012

Flash Deployment

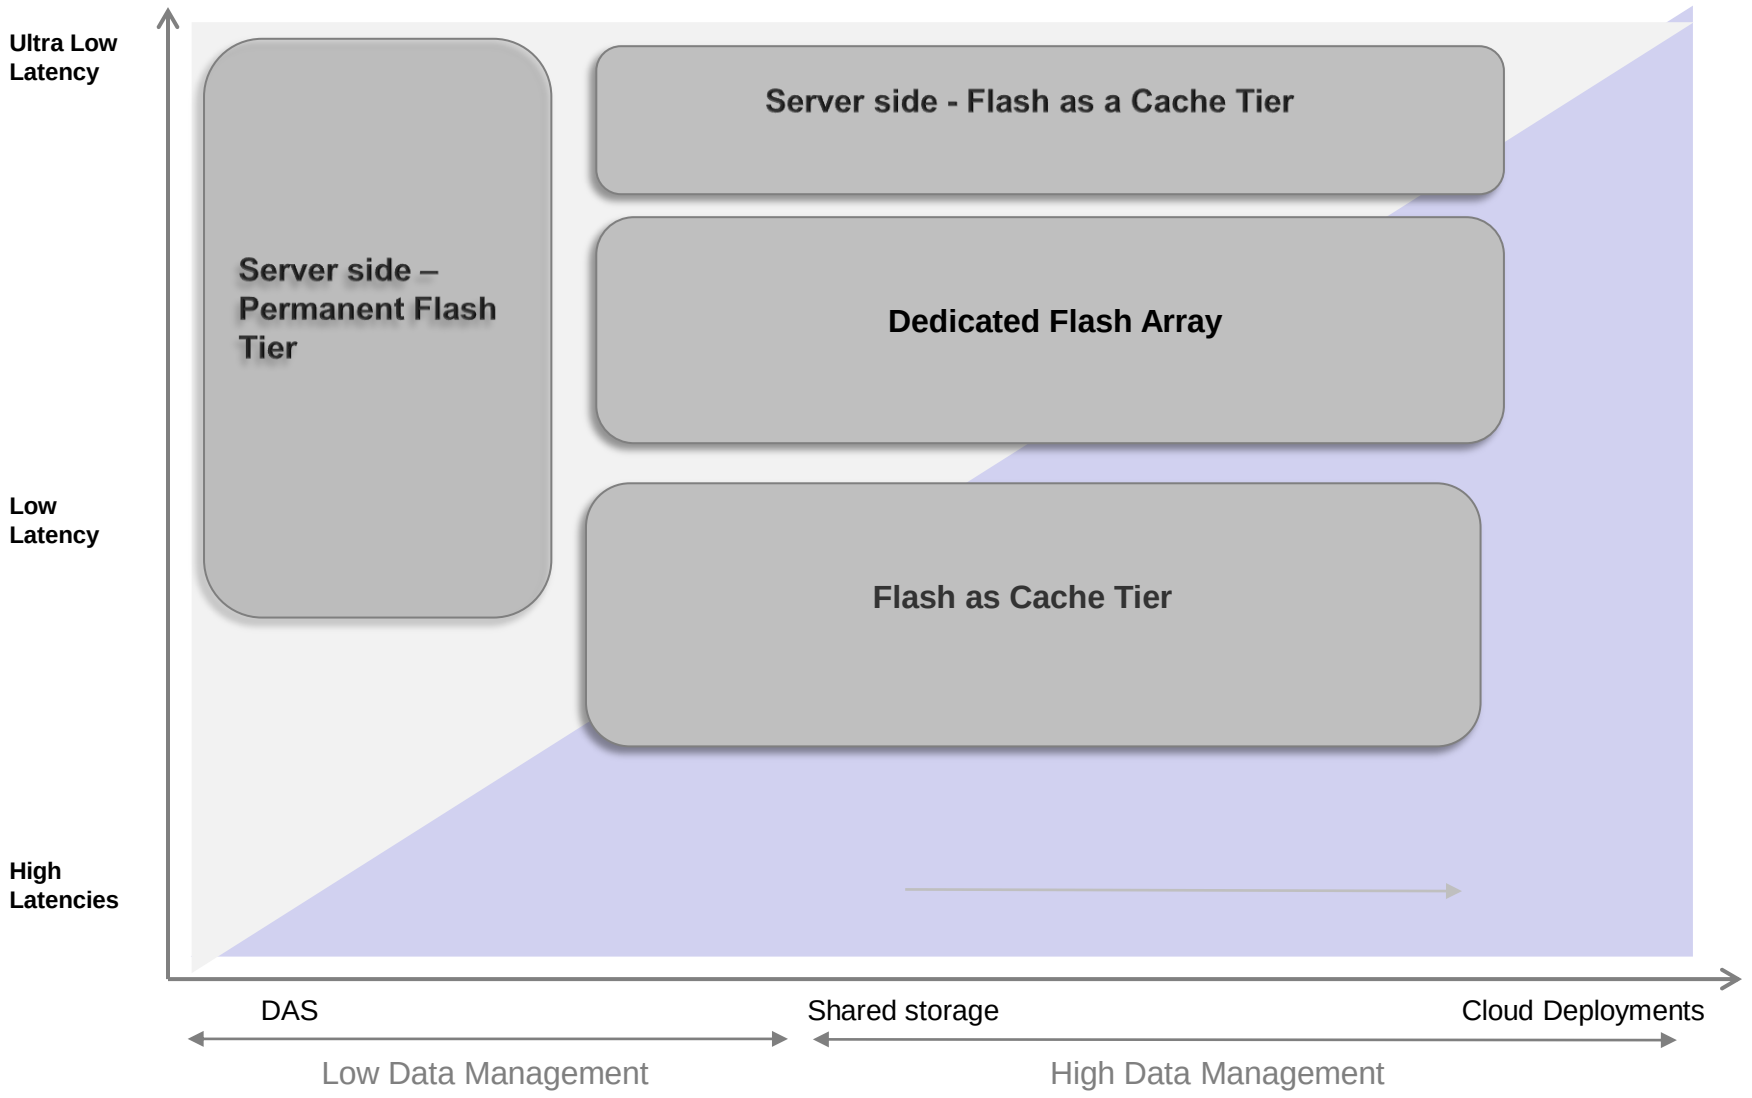

Flash and the Network

SNIA
Advancing storage & information technology

Flash Above the Network

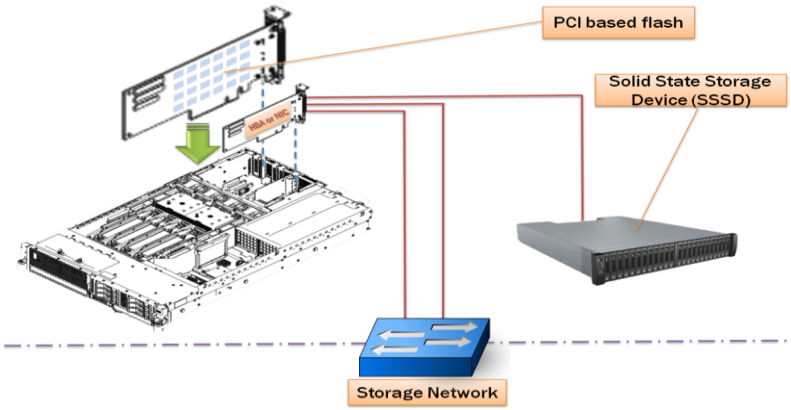

Flash Below the Network

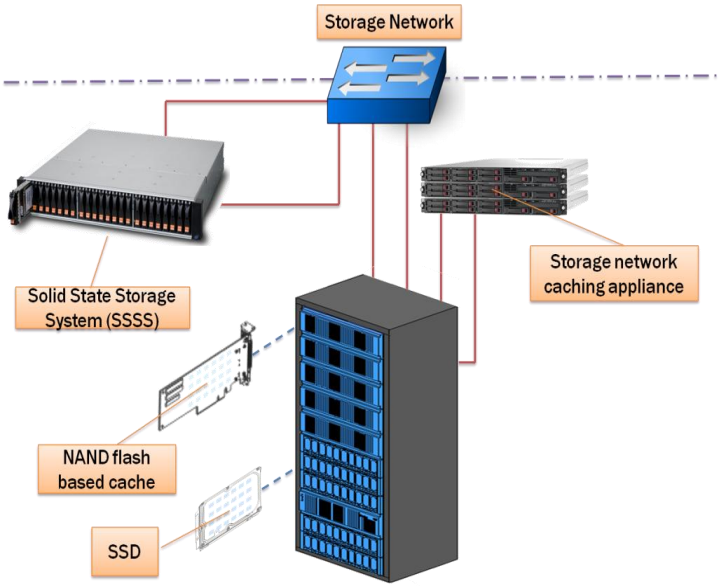

Performance and Efficiency Impact

SNIA
Advancing storage & information technology

Before:
FAS 6210 HA Pair with 144TB

240 SAS 600GB10K RPM Disks

After:
FAS 6210 HA Pair with 168TB

168 1TB 7.2K RPM SATA Disks

1TB Flash Cache

Cost/Efficiency Impact

- Entire workload moved from SAS to SATA (file services workload)
- 34.1% lower cost per TB
- 40.2 % lower $/IOPS
- 40.5% less power

Flexibility Impact

- 16.7% more storage capacity
- 28.5% more IOPS
- 18.5% improvement in average response time

NAND Brings Transformational Change To Compute Platforms

SNIA
Advancing storage & information technology

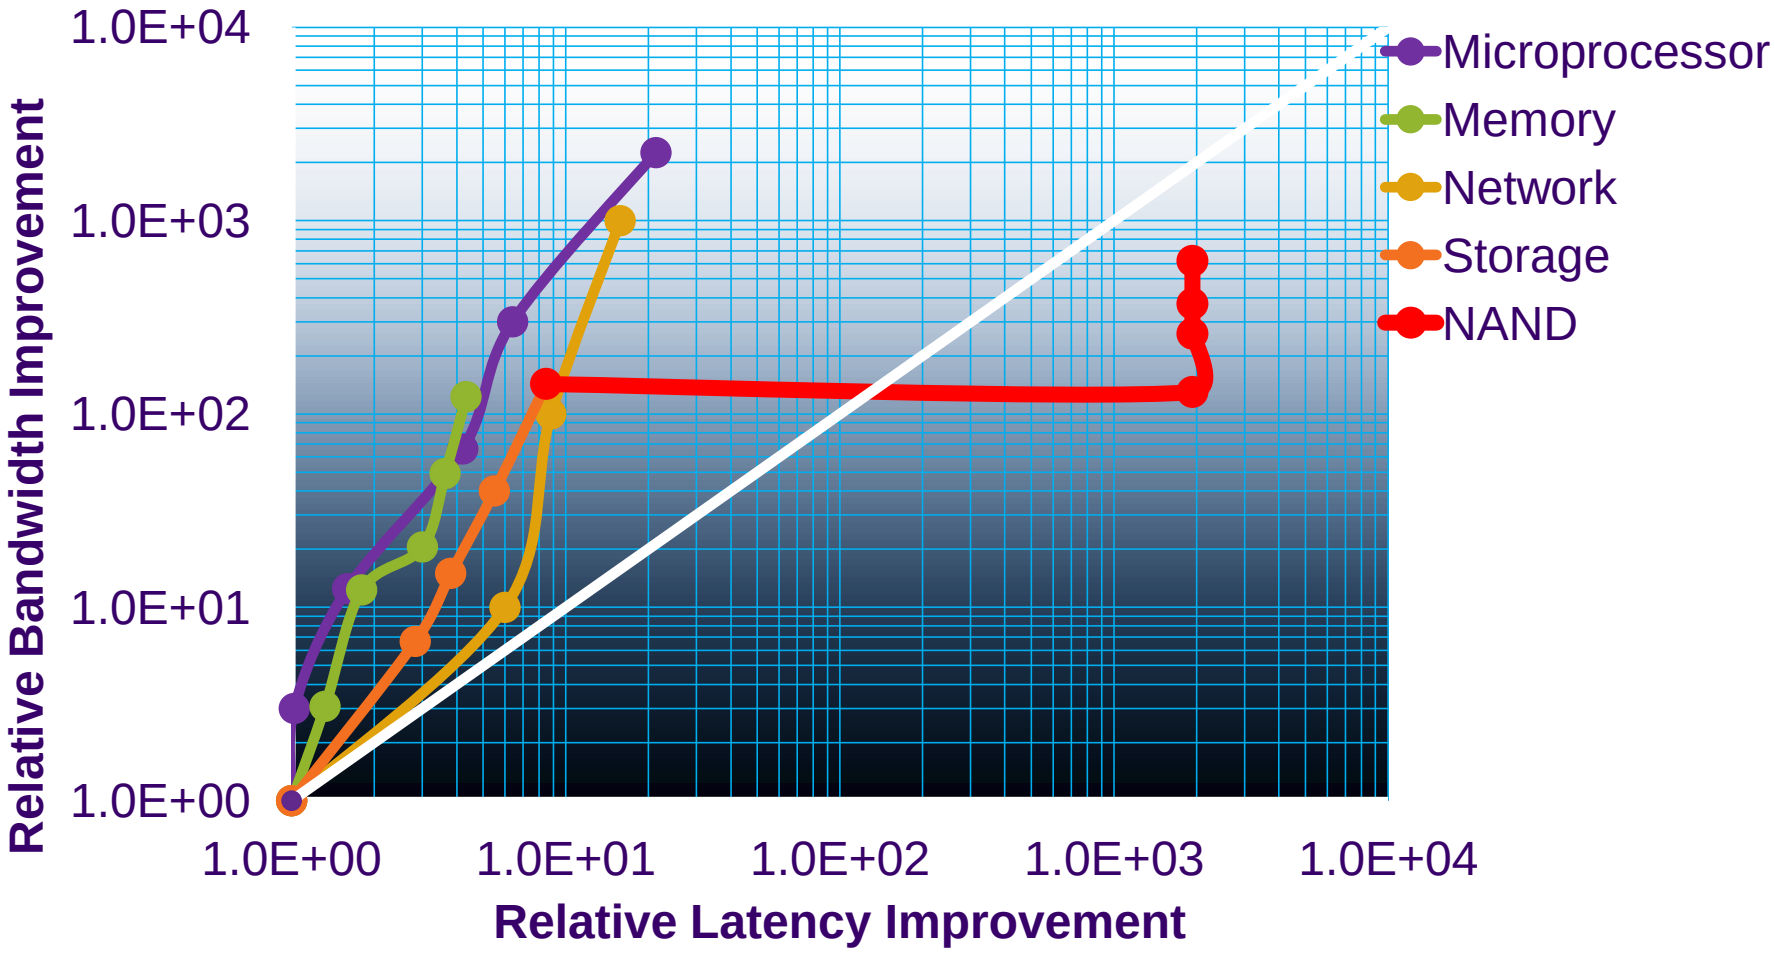

Source: Intel & David Patterson "Latency Lags Bandwidth" COMMUNICATIONS OF THE ACM October 2004/Vol. 47, No. 10 & Intel, 2011

NAND Is Forecast To Continue Sharply Declining In Price

SNIA
Advancing storage & information technology

Blended Average Price per GB Comparison (December 2011 vs. May 2012) for All Market Segments, 2010-2016

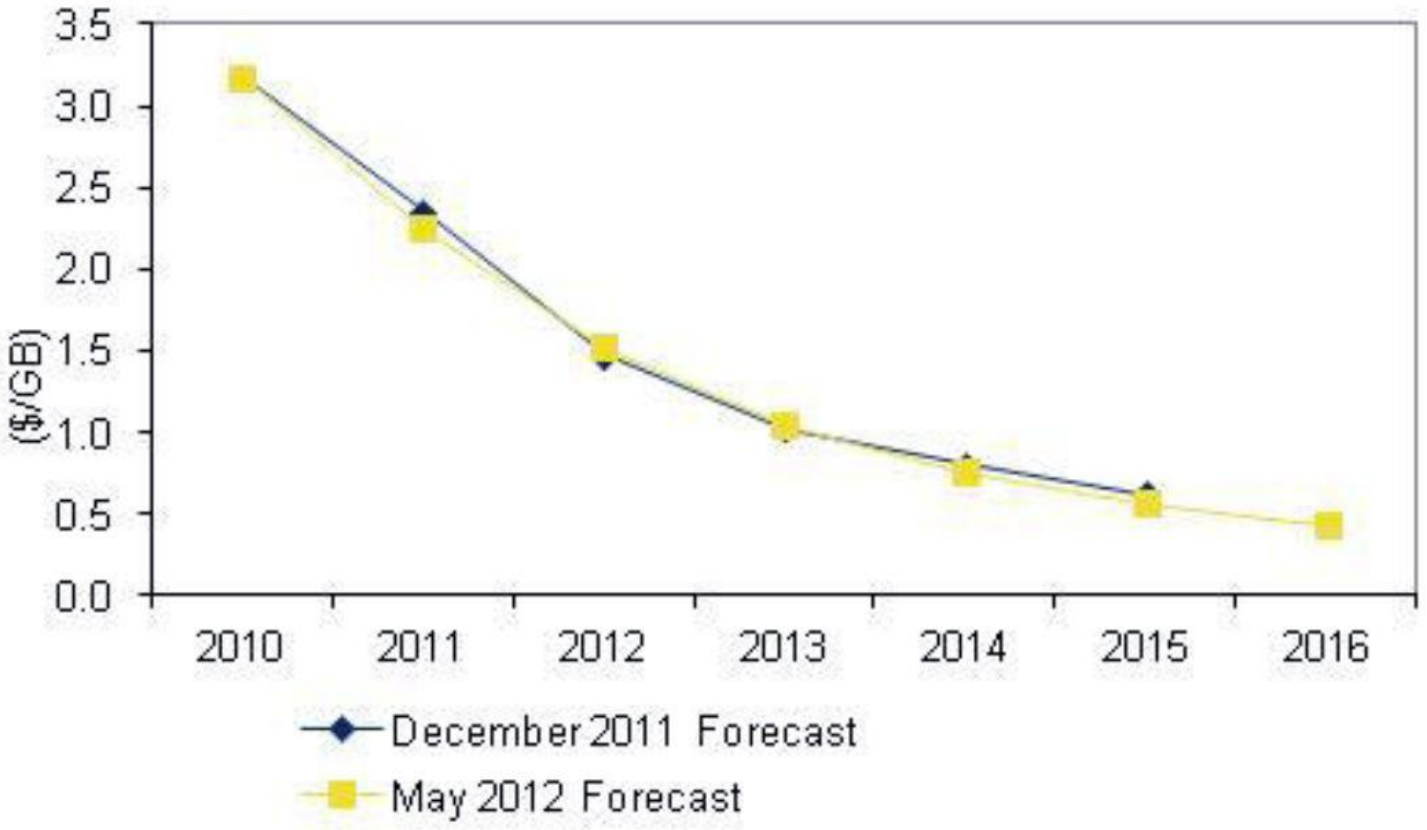

Source: IDC, 2012

"Worldwide Solid State Storage 2012 – 2016 Forecast and Analysis", June 2012, IDC #235155, Volume: 1

Trends In Non-Volatile Memory (NVM)

SNIA
Advancing storage & information technology

- We are just beginning to see the impact of NAND in enterprise storage
  - The percentage of NAND used in storage systems will continue to increase
- NAND going to get cheaper
  - "IDC expects NAND memory ASP to decline at a CAGR of 38% from 2011 to 2016"*
- NAND is going to get faster
  - Each semiconductor process shrink delivers higher speeds
- By the time the industry has absorbed the impact of NAND, it will likely be replaced by something 2 or 3 orders of magnitude even faster
  - Memistors? Phase-change?

*"Worldwide Solid State Storage 2012 – 2016 Forecast and Analysis", June 2012, IDC #235155, Volume: 1

What Impact Will NVM Have On Networking Requirements?

Fortunately, For Flash Storage, 10GbE Adoption Is Well Underway

SNIA
Advancing storage & information technology

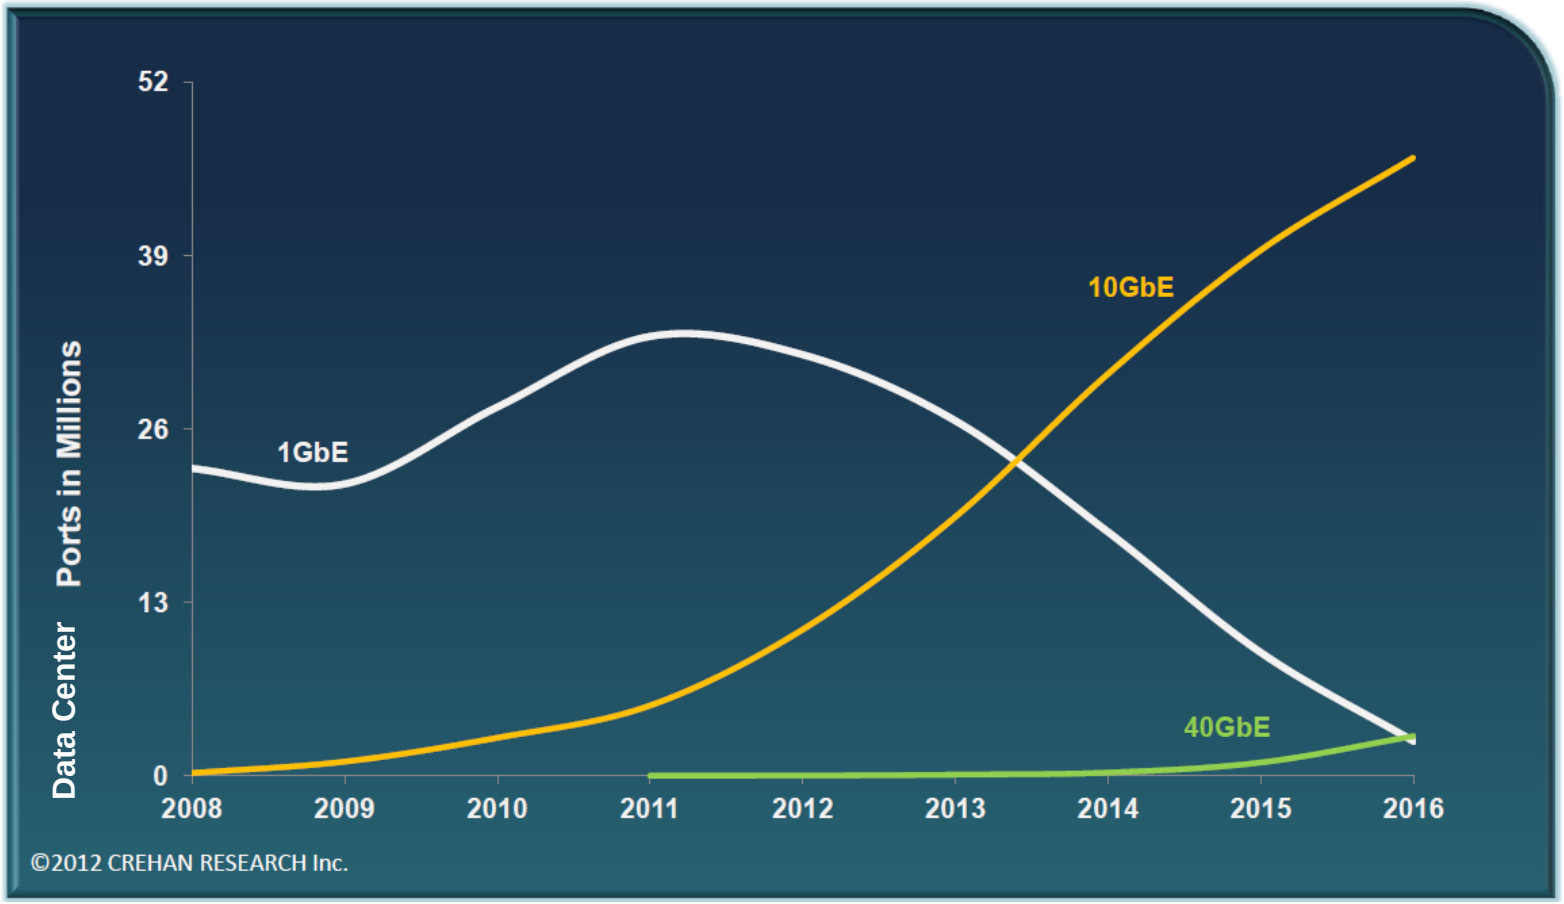

Reasons Driving 10GbE Adoption
(From "10GbE – Key Trends, Drivers and Predictions" SNIA Webcast)

SNIA
Advancing storage & information technology

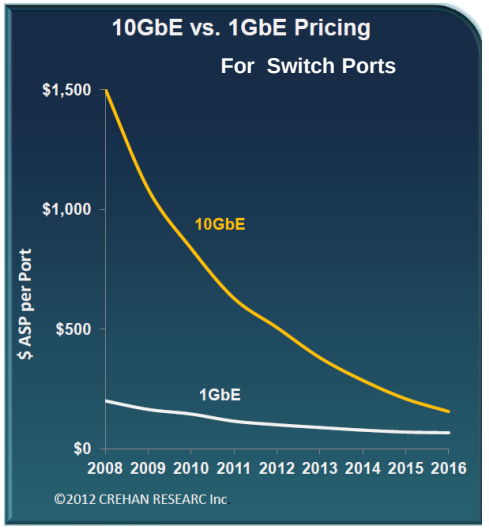

- Virtualization and network consolidation are driving requirements for faster Ethernet
  - Data Center Bridging and FCoE require 10GbE
- Prices for 10GbE switches and adapters are declining sharply
- 10GbE offers lower data center power
- 2012 server platforms can handle significantly more bandwidth
- 10GBASE-T and "flexible LOM" are further driving down deployment costs
- And, of course, the increasing prevalence of flash storage

Is 10GbE Enough To Support Flash Storage?

- For IDF 2012, Intel built a SAN with six 32GB prototype SSDs and commercially available switches and Ethernet adapters (NICs) as a technology demo proof of concept
- First key learning was that
  - Fast Storage + Fast Networking ≠ Fast Storage Networking
- Intel chose to work with the Open FCoE and NVM Express stacks because they were available for modifying
  - The Open FCoE stack had been optimized for rotating media
  - The NVMe stack had been optimized for DAS
  - Significant performance enhancements were achieved by harmonizing these two stacks to work together

Intel Achieved 3M IOPS In This Proof Of Concept

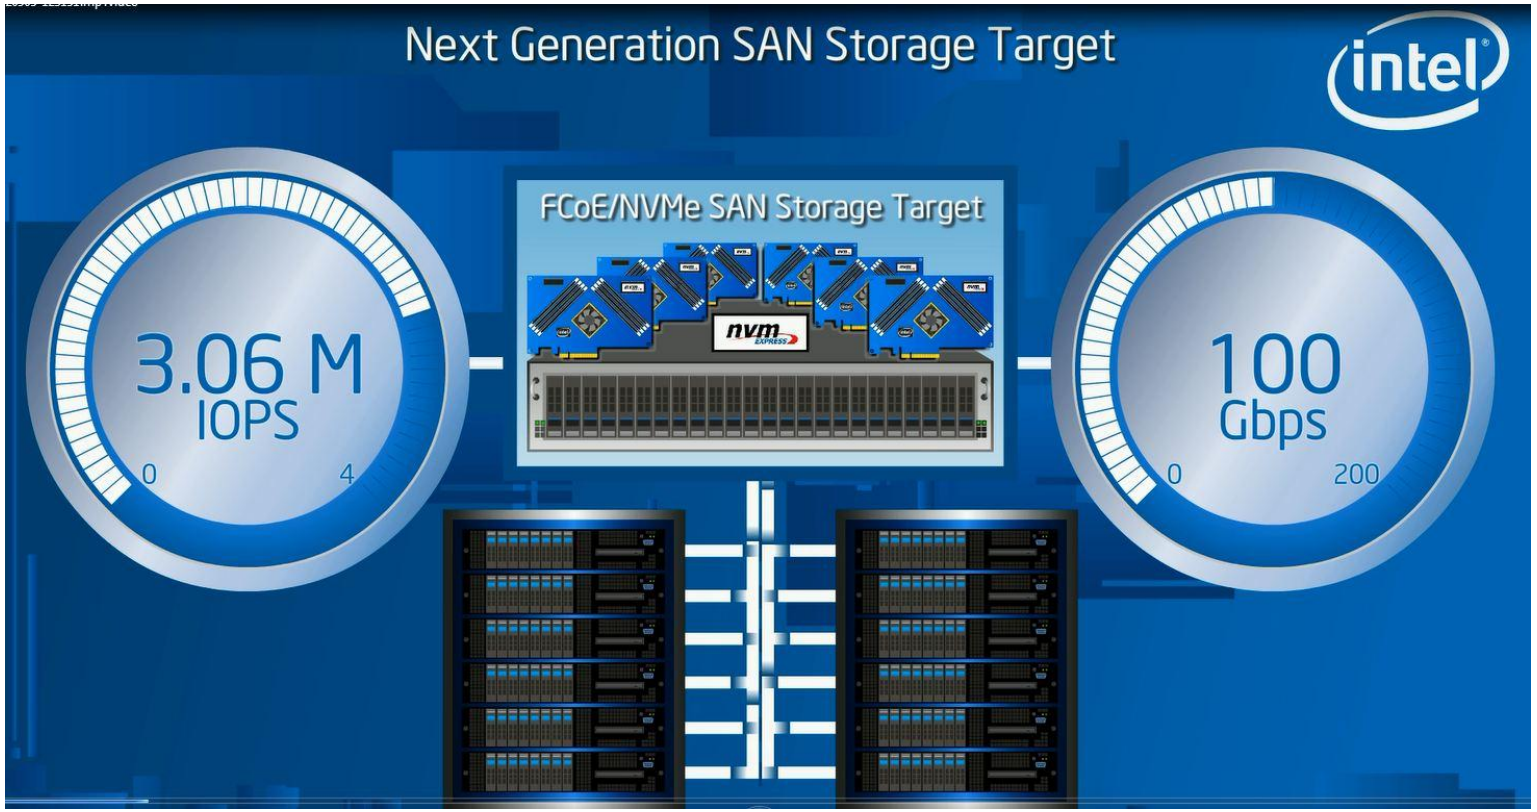

- 4k random read IOPS using fio workload
- 12.7 GBps of NVMe SSD throughput
- FCoE traffic peaking at 120 Gbps

3M IOPS SAN Consumes 12 Ports Of 10GbE

- This NVMe SAN POC is an indicator of where full NVM as well as hybrid storage is going in terms of peak requirements
- Three ports of 40GbE could support this SAN
- But NVM will penetrate greater percentages of storage
- NVM will get faster
- NVM storage will be a key force driving the transitions of 40GbE and 100GbE

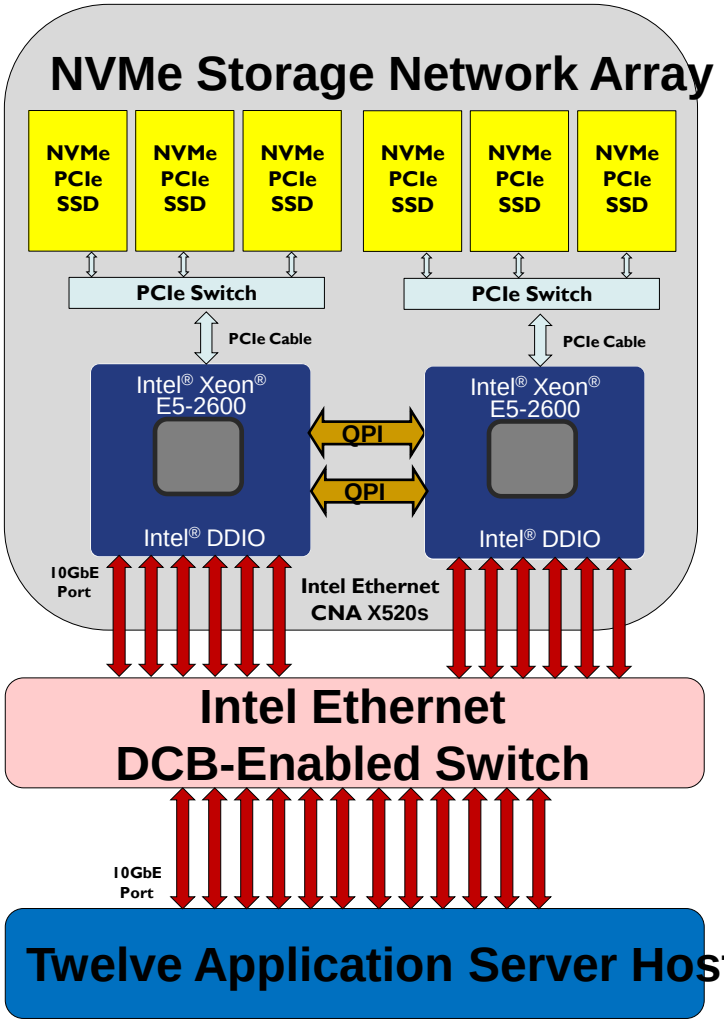

Summary

- Over the next five years solid state technologies will have a profound impact on enterprise storage
- It's not just about replacing mechanical media with solid state media
- The architectural balance of memory, cache and persistent storage will change
- Today's solid state implementations in enterprise storage demonstrate these changes
- It's only the beginning…

Questions?

SNIA
Advancing storage & information technology

- How to contact us:
  - David Fair – david.l.fair@intel.com
  - Paul Feresten – paul.feresten@netapp.com
- Full Q&A session from this Webcast will be posted on the SNIA-ESF Blog