INTEL SSD SOLUTIONS FOR HPC

Darrin Lynch,

BDM NSG, Intel Corp.

Agenda

- Recent product changes and strategy update
- HPC ingredients and opportunities for SSDs
- One problem has many solutions
- NVMe over Fabric (NVMeOF) is here, solutions - on horizon
- Memory Extension – are you ready for this today?

Agenda

- Recent product changes and strategy update
- HPC ingredients and opportunities for SSDs
- One problem has many solutions
- NVMe over Fabric (NVMeOF) is here, solutions - on horizon
- Memory Extension – are you ready for this today?

Technology Driven: Two Breakthrough Technologies

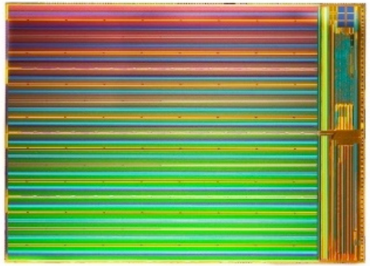

3D MLC and TLC NAND

Building block enabling expansion of SSD into HDD segments

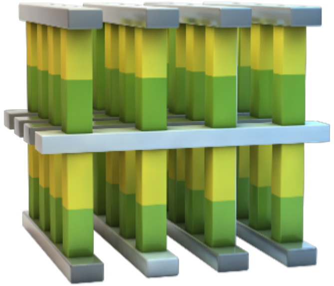

3D XPoint™ Technology

Building blocks for ultra high performance storage & memory

The Future is Here – NVM Express™

1

What is NVMe?

NVM Express is a standardized high performance software interface for PCI Express® Solid-State Drives

2

Built for SSDs

Architected from the ground up for SSDs to be efficient, scalable, and manageable

3

Developed to be lean

Streamlined protocol with new efficient queuing mechanism to scale for multi-core CPUs, low clock cycles per IO

4

Industry standard

software and drivers that work out of the box

5

Ready for next generation SSDs

New storage stack with low latency and small overhead to take full advantage of next generation NVM

NVMe* GETS TO THE RESULT FASTER

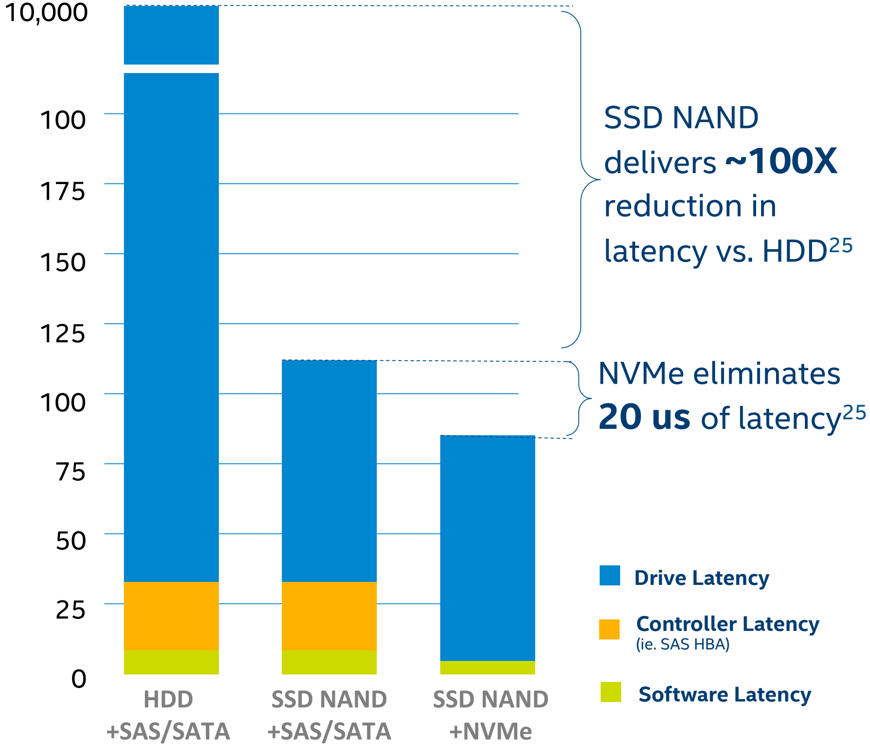

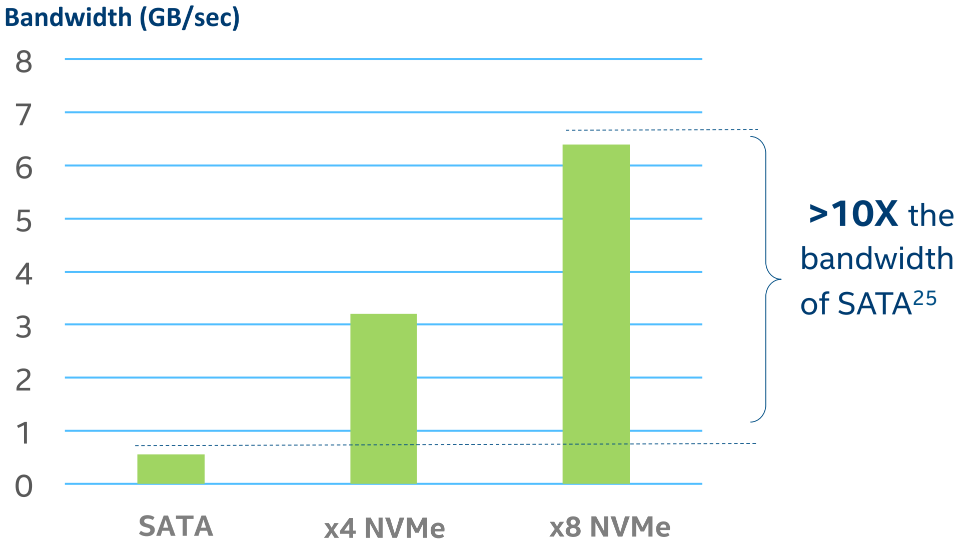

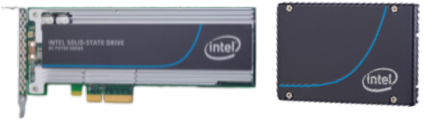

Intel® SSD DC P3700 Series

Intel® SSD DC P3600 Series

Intel® SSD DC P3500 Series

Capacity

Intel® SSD DC P3700 Series	Intel® SSD DC P3600 Series	Intel® SSD DC P3500 Series
400 GB, 800 GB, 1.6TB, 2TB	400 GB, 800 GB, 1.2TB, 1.6TB, 2TB	400 GB, 1.2TB, 2TB

Endurance

Intel® SSD DC P3700 Series	Intel® SSD DC P3600 Series	Intel® SSD DC P3500 Series
Up to 17 DWPD – High Endurance Technology	3 DWPD – Mixed use	0.3 DWPD – Read Intensive

Performance

	Intel® SSD DC P3700 Series	Intel® SSD DC P3600 Series	Intel® SSD DC P3500 Series
Random 4k Read	450k IOPS	450k IOPS	450k IOPS
Random 4k Write	175k IOPS	56k IOPS	35k IOPS
Random 4k 70/30 R/W	265k IOPS	160k IOPS	85k IOPS
Sequential Read	2800 MB/s	2600 MB/s	2500 MB/s
Sequential Write	2000 MB/s	1700 MB/s	1700 MB/s

Sequential latency of 20µs

Intel® SSD DC P3608 Series

NVM Express™ PCI Express® 3.0, x8

New dual controller and PCIe 3.0, x8 architecture delivers increased capacity and read performance compared to the Intel SSD DC P3600 Series

1.6TB 3.2TB 4TB

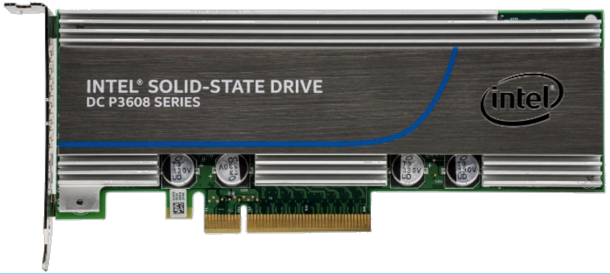

Performance

Random 4k Read	Up to 850k IOPS
Random 4k Write	Up to 150k IOPS
Random 4k 70/30 R/W	Up to 300k IOPS
Sequential Read	Up to 5000 MB/s
Sequential Write	Up to 3000 MB/s
Avg Active R/W Power	40W, 25-50W configurable
Idle Power	8-10 W

Uses

- HPC
- Database

Features

- One drive volume with Intel RSTe software
- Power loss protection
- End-to-end data protection UBER 10^{-17}, 2M hours MTBF
- 20nm HET NAND for mixed workload performance at 3 DWPD
- Consistently amazing performance Low latency of NVMe*

NEW ADDITIONS TO…

…the most widely adopted PCIe* SSD family[27]

Intel's **first active/active dual-port NVMe SSD** delivering breakthrough performance[19] to mission critical data

Intel® SSD DC D3700

Intel® SSD DC D3600

Break the SAS architecture bottleneck

Intel's best IOPS value built on the **highest density 3D NAND**[28] enabling solutions that are all NVMe

Intel® SSD DC P3520

Intel® SSD DC P3320

bringing **NVMe drives to the mainstream**

*Names and brands of others may be claimed as property of others.

INTEL® SSD DC P3520 AND P3320 SERIES

Bringing storage out of the cold.

Performance.[16]

	P3320 vs. Intel® SSD DC S3510		P3520 vs. P3320
Random 4K Read	Up to 365K IOPS	Up to **5X faster**	**Performance data available at later date**
Random 4K Write	Up to 22K IOPS	Up to **1.4X faster**	
Random 4K 70/30 R/W	Up to 65K IOPS	Up to **1.9X faster**	
Sequential Read	Up to 1600 MB/s	Up to **3.2X faster**	
Sequential Write	Up to 1400 MB/s	Up to **3.1X faster**	

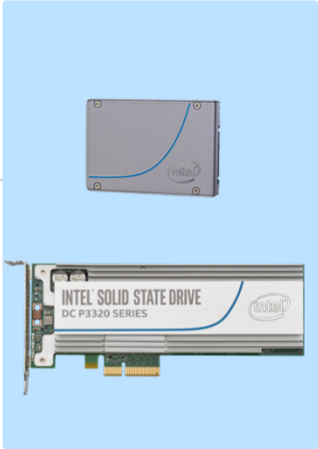

Capacity and Form Factor.

2TB

Capacities	450GB (2.5" only)
	1.2TB
	2TB
Form Factors	2.5" x 15mm
	X4 HHHL Add-in-card

Performance consistency.

IOPS consistency	Up to 90%[9]
Performance degradation	≤5% over product life[10]

Data integrity and reliability.

End-to-end data protection	Demonstrated 10^{-17} UBER[11]
	100X more reliable preventing SDC[11]
PLI	Self-test validated on 2M cycles[14]
AFR	Actual AFR consistently better than Intel goal of ≤0.44%[15]

Intel® SSD DC D3700 and D3600 series

Always available. Always fast.

Performance.[18]

		vs. SAS
Random 4K Read	Up to 470K IOPS	Up to **3.9X faster**
Random 4K Write	Up to 95K IOPS	Up to **1.3X faster**
Random 4K 70/30 R/W	Up to 213K IOPS	Up to **2.9X faster**
Sequential Read	Up to 2100 MB/s	Up to **1.8X faster**
Sequential Write	Up to 1500 MB/s	Up to **2.0X faster**

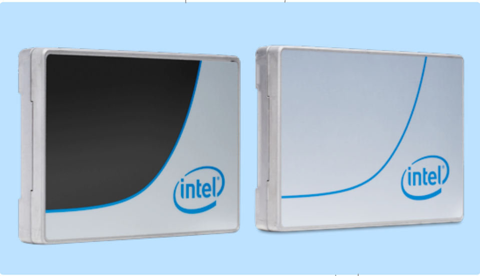

Capacity and Form Factor.

2TB

D3700 Capacities	800GB
	1.6TB
D3600 Capacities	1TB
	2TB
Form Factors	2.5" x 15mm

Performance consistency.

IOPS consistency	Up to 90%[9]
Performance degradation	≤5% over product life[10]

High Availability.

24

Dual port	Active-Active

Data integrity and reliability

End-to-end data protection	Demonstrated 10^{-17} UBER[11]
	100X more reliable preventing SDC[11]
PLI	Self-test validated on 2M cycles[14]
AFR	Actual AFR consistently better than Intel goal of ≤0.44%[15]

Refer to appendix for footnote 9, 10, 11, 14, 15, 18.

DRIVES WITH INTEGRITY. THE INTEL PCIe* SSD DC ADVANTAGE

Solution

Lasting integrity

≥100X more reliable
preventing Silent Data Corruption[11]

>2M PLI cycles and self-test
delivers trusted protection from data loss[14]

Reliably effective performance

Validated on a **massive scale**

≤5% performance degradation
through product life[10]

Platform confidence

Complete in-house design
of data center solutions

Up to **1 more year**
in extended platform-level validation

Foundation

Uncompromising supply chain quality

95% delivery
to commit[21]

<100 DPM
factory quality[22]

#4 overall,
#1 semiconductor
Gartner* 2015 Supply Chain Ranking[23]

Complete product life cycle support

>2000 engineers optimizing top ISV solutions[24]

>500 dedicated OEM design-in resources[24]

Platform-expert
post-sales support

Innovation leadership

~1000 researchers in Intel labs

Memory technology **innovation leader**

3D Xpoint™ on NVMe is the next quantum leap

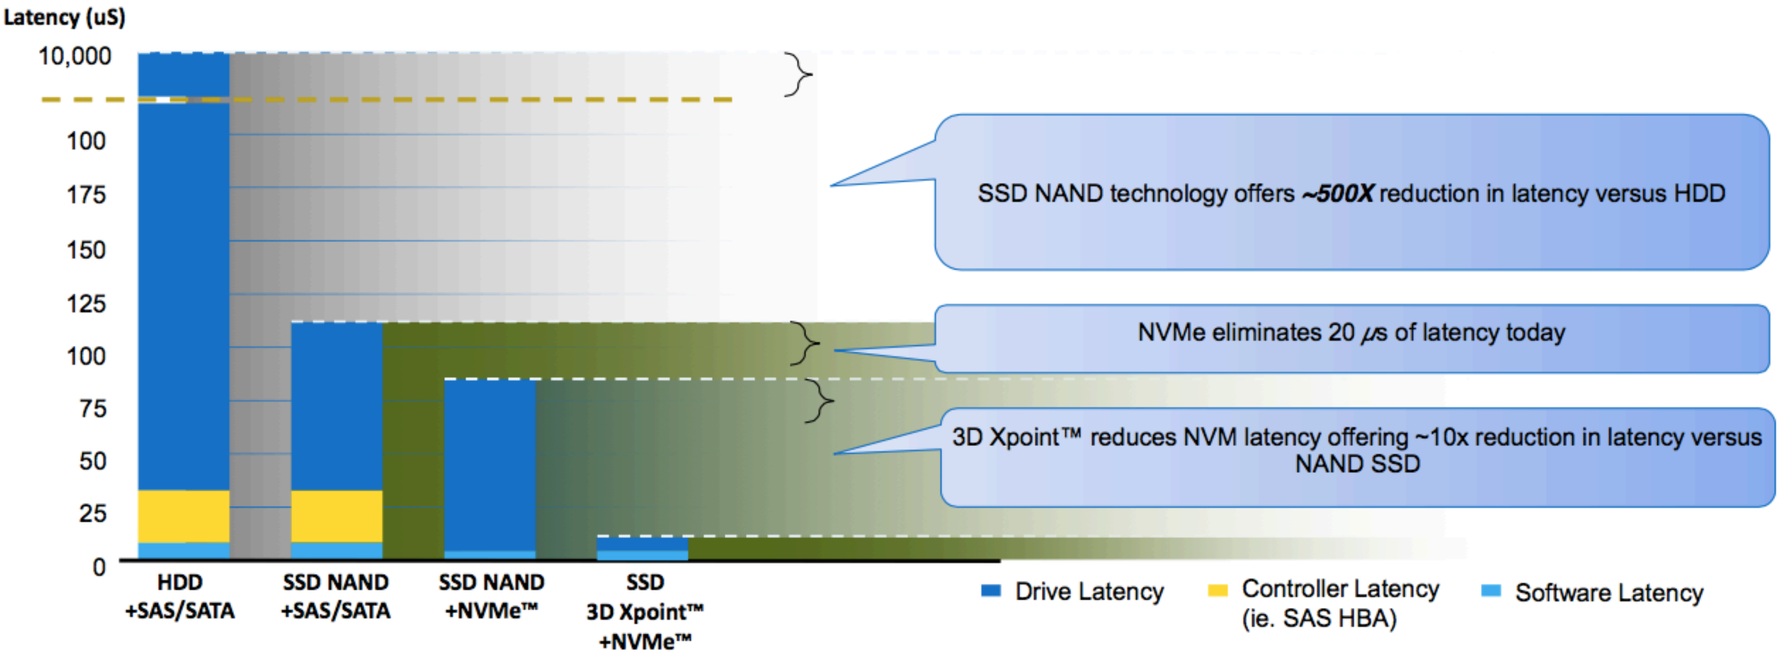

Source: Storage Technologies Group, Intel

NAND vs 3D XPoint™ technology: SSD IOPS vs Queue Depth

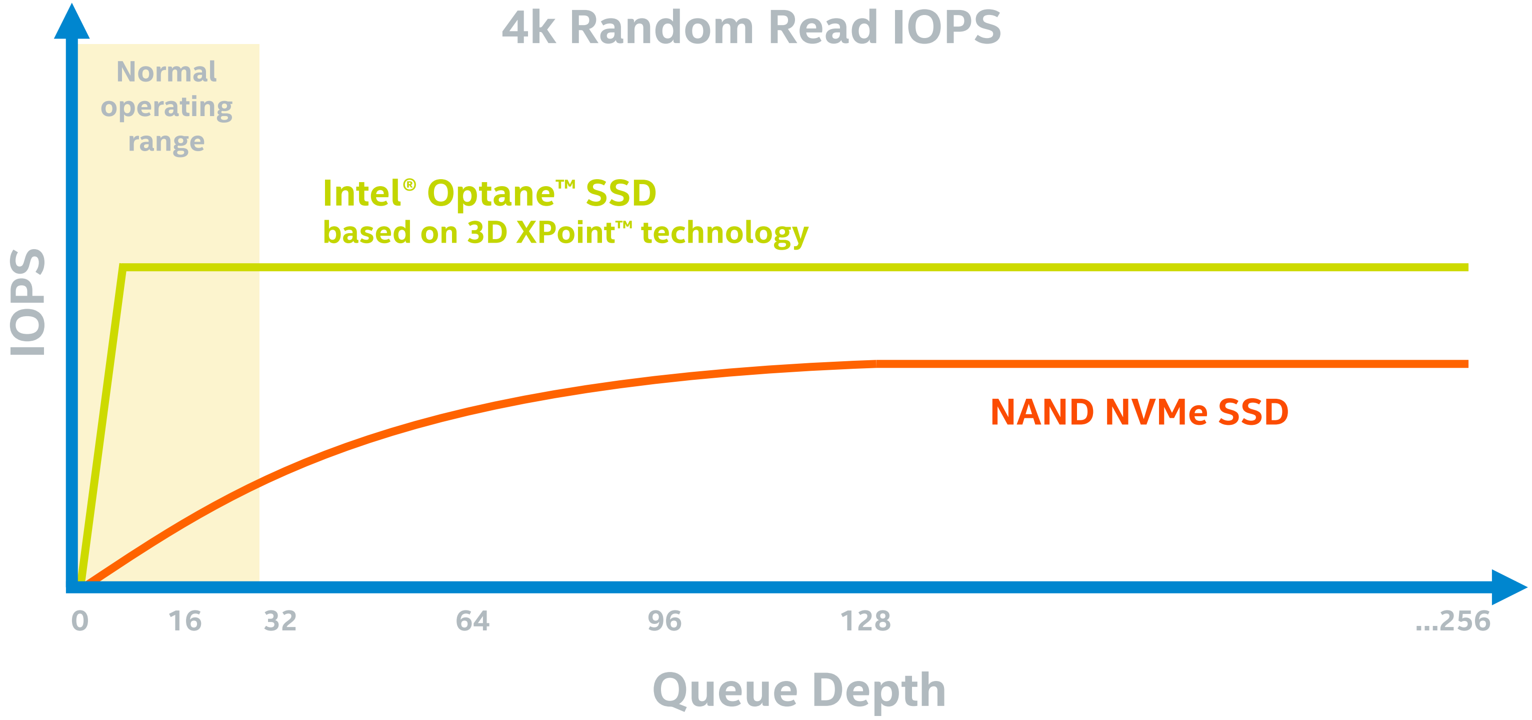

Agenda

- Recent product changes and strategy update
- HPC ingredients and opportunities for SSDs
- One problem has many solutions
- NVMe over Fabric (NVMeOF) is here, solutions - on horizon
- Memory Extension – are you ready for this today?

HPC ingredients and opportunities for SSDs

SSDs in local compute nodes (Server side)

Purpose:
- Scratch space to accelerate I/O intensive host workloads.
- Checkpoint restart / Memory snapshot.
- Decrease network utilization
- Improve KNL/GPU performance.
- Far memory tier for DRAM, Swap.
- Boot drive.

Technology ingredients:
- NVMe designs, all product range 3D NAND + 3D XPoint
- SATA / NVMe low power boot drives
- Alignment with Xeon/KNC/OPA feature set

I/O node and Burst Buffer designs

Purpose:
- Accelerating data transfer in-and-out of the compute by scaling I/O nodes in regards to compute nodes
- Improve small random I/O
- Cost saving to enable diskless compute noted maintaining close to local I/O latencies

Technology ingredients:
- All NVMe designs performance balanced with high speed fabric
- Great play for 3D NAND NVMe, less often for 3D XPoint
- NVMeOF designs, especially NVMe-over-OPA changes the storage paradigm

Improve HPC Storage performance

Purpose:
- Bring NVMe into typical HPC storage ecosystem to improve small I/O performance (HSM).
- Better together story with IEEL (Intel Enterprise Edition for Lustre) for MDS and Lustre DSS.
- Improve cold storage CEPH and parallel FS such as BeeGFS.

Technology ingredients:
- Dual Port NVMe designs for Lustre* in MDS and OSS
- Single port NVMe / SATA drives for scratch file systems (no HA)
- Intel CAS DSS for intellectual caching in Lustre*

SSDs in Technical Compute (Workstation)

Purpose:
- Improve overall system performance and critical workloads (CAD, simulations)
- Software development / compilation tasks
- Enable local cache SSD drive to decrease network utilization
- Boot drive / Swap

Technology ingredients:
- NVMe designs, all product range 3D NAND + 3D XPoint based on application requirements
- Intel CAS caching solutions

Compute Node optimized architecture

3D XPoint™ Technology

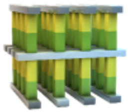

3D NAND

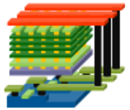

Purpose

Accelerate HPC applications by providing close to memory local storage capabilities.

Block Diagram

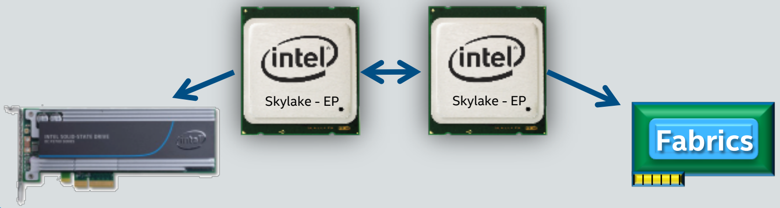

Architecture Ingredients:

- Attach 1-2 local SSDs to the compute node directly
- Improve workloads characterized by large dataset or swap drive requirements *in* ***Genomics, Fluid Dynamics, Computational chemistry, Structural analysis, Finance, HPDA or generic app development.***
- Implement Checkpoint restart/Memory snapshot for critical workloads
- Minimize I/O requirements to an external storage
- Opportunity for low power SATA/PCIe boot SSDs

Supercharge Xeon Phi and GPU processing

3D XPoint™ Technology

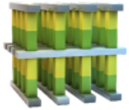

3D NAND

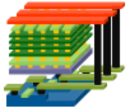

Purpose

Drive mass parallelism evolution in Many-Core processing by supporting KNL architecture with NVMe tier and close to memory performance.

Block Diagram

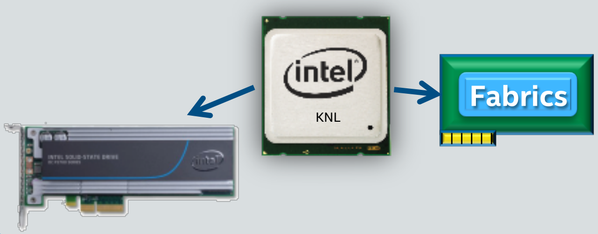

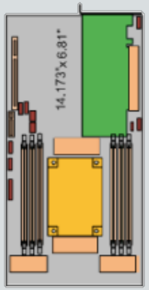

Architecture Ingredients:

- Direct attach of NVMe SSDs to the KNL standalone node
- Host sharing NVMe SSD between GPUs
- Improve multithreaded workloads by offering direct access to NVMe many queue approach
- Minimize I/O requirements to an external storage

Use Cases

- Cray's Compute cluster architecture with SSD at Chevron

Improving performance of HPC storage

Purpose

Accelerating traditional HPC parallel storage by introducing new features to improve small I/O

Architecture Ingredients:

- Bringing NVMe into typical HPC storage ecosystem to improve small I/O
- Intel SSDs in IEEL (Intel Enterprise Edition), BeeGFS, CEPH
- Enabling NVMe Dual Port drives topologies

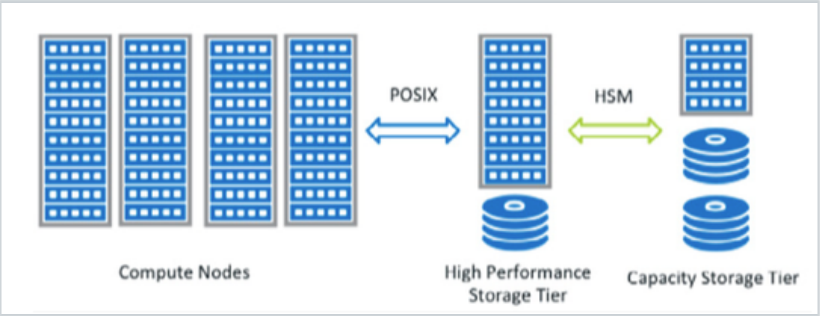

Use cases

Lustre:

- Metadata server (MDS)
- HSM storage tier
- DSS with CAS (in progress)
- All flash scratch or HA solutions

Other alternative solutions:
BeeGFS, CEPH for cold storage

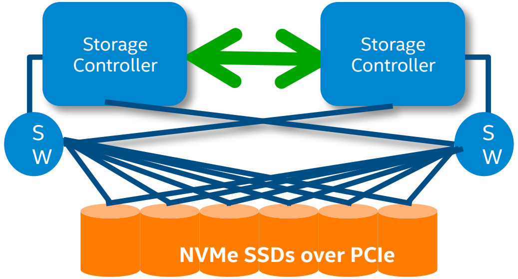

High Availability Storage Systems

Must do. Storage redundancy.

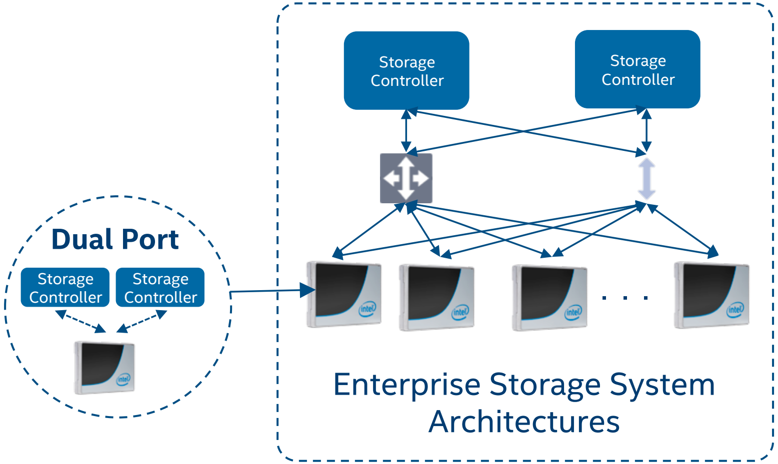

Do it better. Faster results with NVMe*

Accelerate mission critical database workloads — up to **3.9x** faster 4K RR vs. SAS SSD[18]

Improve data access for better QoS — up to **3.3x** Lower latency vs. SAS SSD[18]

INTEL® SSD DATA CENTER D3700 AND D3600

New features supported for mission critical storage

Reservations

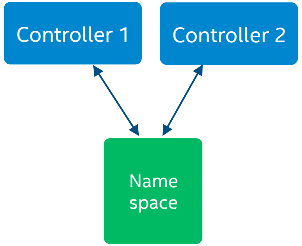

Manages who can write to the device and when

Scatter Gather List

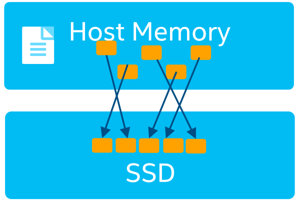

Connects scattered data in host memory to reduce IO operations

Multiple Namespace

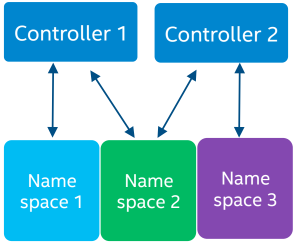

Smaller granularity in managing data across drives

In Controller Memory Buffer

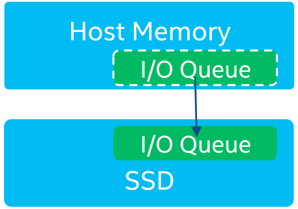

Moves IO queue to the SSD to reduce latency

Caching – Differentiated Storage Services & CAS

Intel® Cache Acceleration Software (CAS) 3.0 introduced DSS support

- Use CAS with SSD storage to provide a complete DSS-enabled read cache for Lustre*

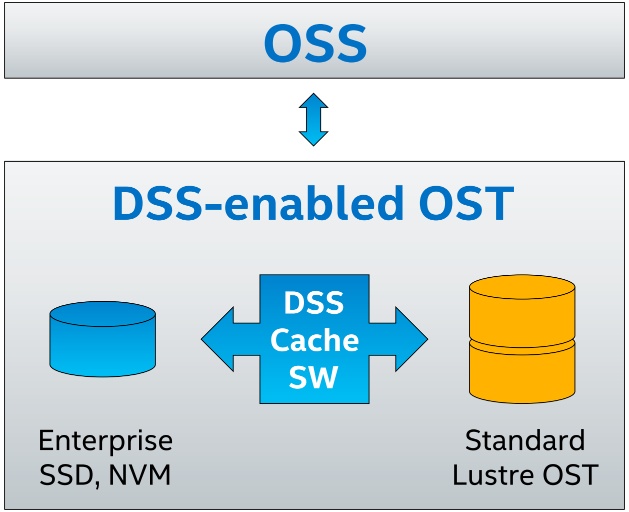

Lustre "tags" data to identify priority and passes to storage → Intel CAS interprets tag & directs IO to cache or sends to HDD → Data is written to HDD and/or flash, depending on priority

I/O node, Burst Buffer, NVMe over Fabric (NVMeF)

Purpose

Accelerating data transfer in-and-out of the compute by scaling I/O nodes in regards to compute nodes with close to local I/O latencies

Block Diagram

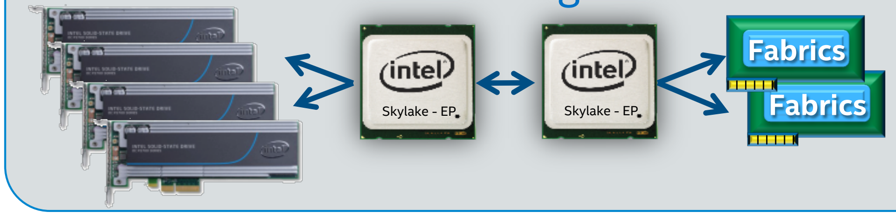

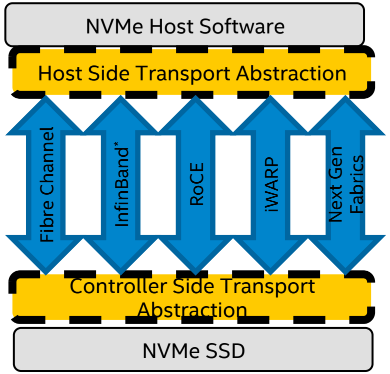

3D XPoint™ Technology

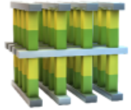

3D NAND

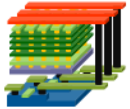

Architecture Ingredients:

- PCIe IO rich nodes with balanced input-output configuration, so, the internal IO processing capabilities can scale externally.
- Time to market to OmniPath and Ethernet products.
- Typical attach rate is 30 compute nodes to 1 burst buffer node
- Optimal NVMe SSD count depends on the used fabrics solution and available bandwidth.

Data Transfer Node “DTN”

Purpose

Enabling the rapid transfer of large datasets (>1TB) across long distances, largely avoiding latency issues.

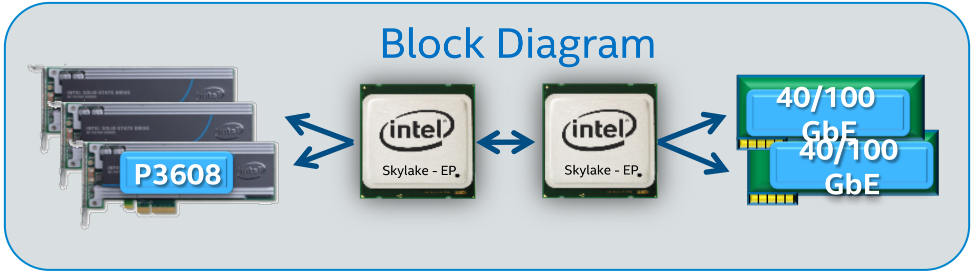

Architecture Ingredients:

- Nodes with >72 PCIe lanes to support 4 SSDs and 1-3 Ethernet NICs
- Time to market to Ethernet products (100GbE).
- Sold as matched pairs of DTNs installed at each end of file transfer paths

Software:

- Aspera* FASP data transfers
- Globus
- Zettar zx* – Ultra Performance Parallel Data Transfer

Agenda

- Recent product changes and strategy update
- HPC ingredients and opportunities for SSDs
- One problem has many solutions
- NVMe over Fabric (NVMeOF) is here, solutions - on horizon
- Memory Extension – are you ready for this today?

Improving Genome Sequencing at German Cancer Research Center

Goal:
- The genome comparison for individual patients of cancerous tissue with healthy tissue
- Apply the right treatment, new way of "personalized medicine".

Challenges:
- 70GB of raw data per sequencing run (1.3 PB per year)
- Requires fast processing for compute as well as for storage.
- Data exchange between pipeline steps is done by IO of millions of small files

Solution:

Enable every compute node with 1.6 TB Intel DC P3700 NVMe scratch space

- >= 4 drive writes per day requirements (1.3PB in a year overall)
- Minimize CPU idle time due to the low-latency, high-performance NVMe technology
- IO from and to NAS is limited to a few large files, streams at the start and at the end of the pipeline.

"Intel's NVMe technology boosts IOPS by avoiding SCSI and SAS interfaces of classic SSDs and opens the full PCIe performance and bandwidth for small file IO."

Full details: http://www.govirtual.se/2015/08/26/go-virtual-attended-isc-in-june-2015/

nvm EXPRESS™

GOVIRTUAL

dkfz. GERMAN CANCER RESEARCH CENTER IN THE HELMHOLTZ ASSOCIATION

50 Years – Research for A Life Without Cancer

The German Cancer Research Center (known as DKFZ), is a national cancer research center based in Heidelberg, Germany. It is a member of the Helmholtz Association of German Research Centers, the largest scientific organization in Germany.

System configuration: 2 Xeon 2660v2, 512GB DDR3, P3700 -1.6TB, RHEL 6.5, Intel GATK

Qiagen: Analyze Your $1,000 Genome for as Little as $22

- Reference Architecture to improve genome sequencing time by hosting data sets at optimized Lustre (IEEL)
- 700GB data per genome, more than 30TB of data analysis per day using 32-node HPC cluster
- Cost optimized 165TB storage solution is driven by target capacity and required performance. Combination of HDD for the storage and SSDs for metadata.
- Intel DC SSD for NVMe P3600 series meets performance requirements for MDS at data center reliability (2M hours MTBF)

https://communities.intel.com/community/itpeernetwork/healthcare/blog/2015/10/22/analyze-your-1000-genome-for-as-little-as-22

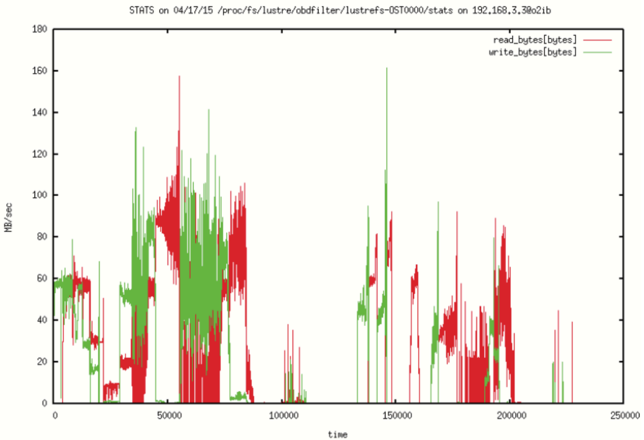

Heavy I/O read and write demands are heavy on the object storage servers

System configuration:
HPC cluster: 2 x Intel® Xeon® processor E5 v3 (14 core, 2.60 GHz), 128GB DDR4, CLC Genomic Server, version 7.1,
Storage: 165 TB useable disk capacity (256 TB total) , Intel® Enterprise Edition for Lustre* 2.3, 2 x Intel Xeon processor E5 v3 (14 core, 2.60 GHz), 128 GB DDR4 memory, RAID Controller (RAID 10), 4 x 800 GB Intel® Solid State Drive Data Center P3700 Series

Agenda

- Recent product changes and strategy update
- HPC ingredients and opportunities for SSDs
- One problem has many solutions
- NVMe over Fabric (NVMeOF) is here, solutions - on horizon
- Memory Extension – are you ready for this today?

The Need to Extend NVM Express™ over Fabrics

PCI Express® ideal for in-server and in-rack, but difficult to scale beyond 100's of nodes:

- Address routing rather than endpoint routing
- Want to converge storage + networking at scale
- Want to leverage standard switch infrastructure

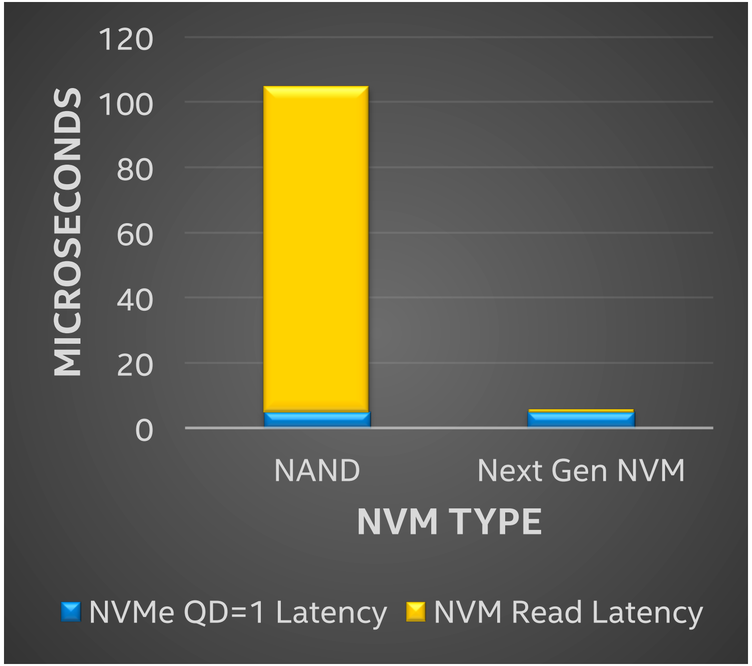

Existing Fabric interface (e.g., iSER / SRP) ecosystem is not well suited for this:

- Inconsistent adoption across OS/VMs
- Protocol is overly complex, adding latency
- Issues even worse when we move to NG-NVM

Delivering < 20 µs across Fabric requires new, simple, efficient protocol

NVMe over Fabrics

Industry standard definition of NVMe over Datacenter Network Fabrics

- **Shares same base architecture and NVMe Host Software as PCIe**
- **Specification defines interface supporting multiple network fabrics**

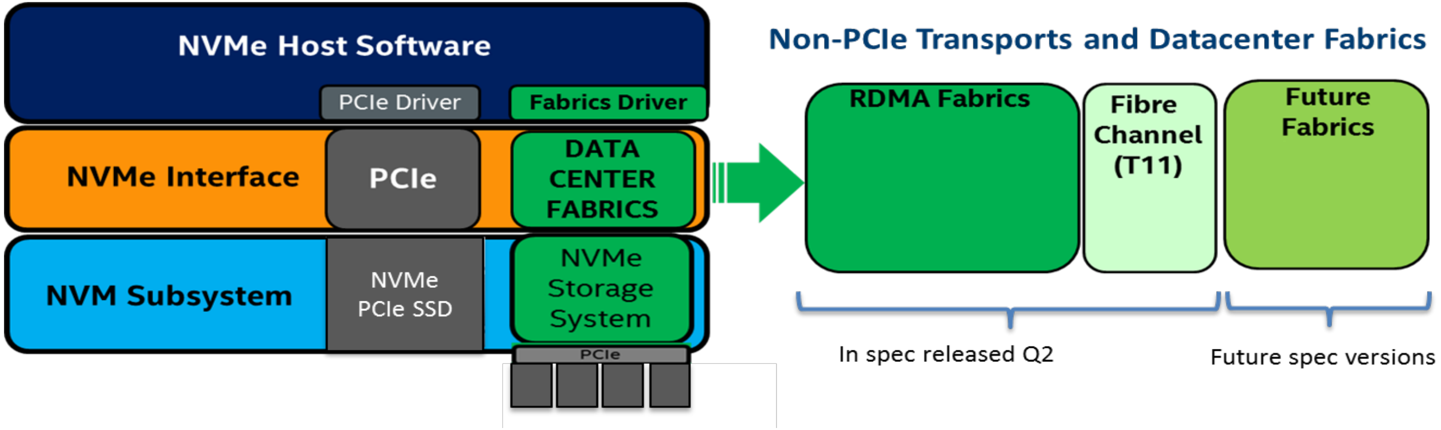

NVMe over Fabrics specification defines how NVMe is extended to

support new network fabrics

NVMe over Fabrics

Technology Overview

The back-end of many deployments is PCIe Express® based NVM Express™ SSDs

With 10-100Gb reliable RDMA fabric and NVMe SSDs, the remaining issue is the software necessary to execute the protocol

Use NVMe end-to-end to get the simplicity, efficiency, and low latency

- Simple protocol => Simple host and SSD software
- No translation to/from another protocol like SCSI

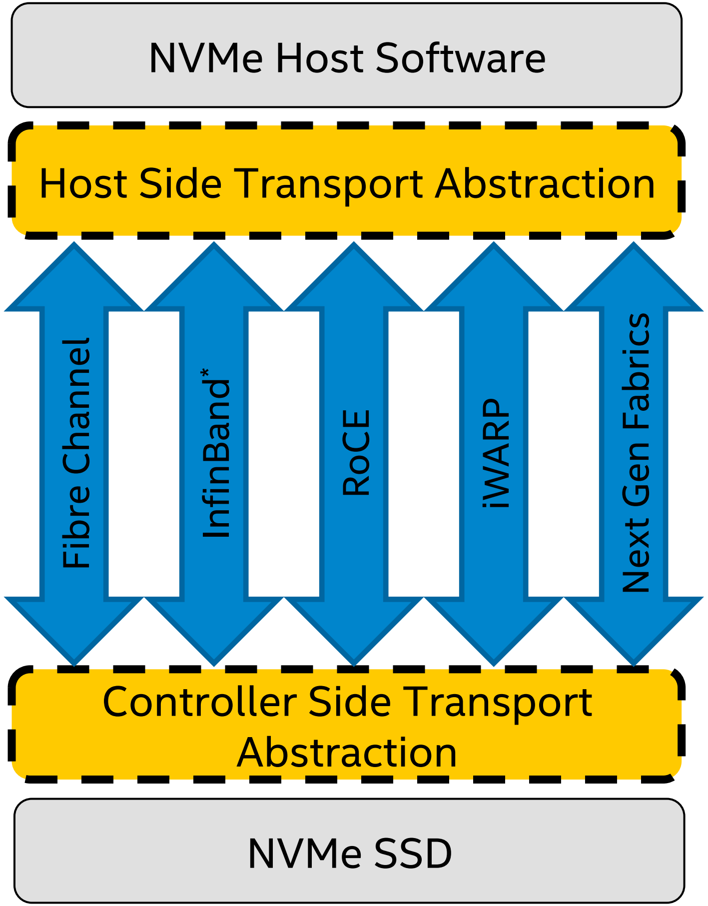

Standard abstraction layer enables NVMe across range of Fabrics

NVMe over Fabrics

Host Components

Architecture

- Separation of NVMe Core Logic and Transport
- Common NVMe Core across all transports: OS and Block interface
- Common NVMe Fabric agnostic functions: Capsules, Properties, Connect, and Discovery
- Thin transports with limited NVMe awareness that can interoperate with multiple targets

Implementation

- Separate Driver Components
- Components interface with Multi-Q Block layer
- Pluggable Transports

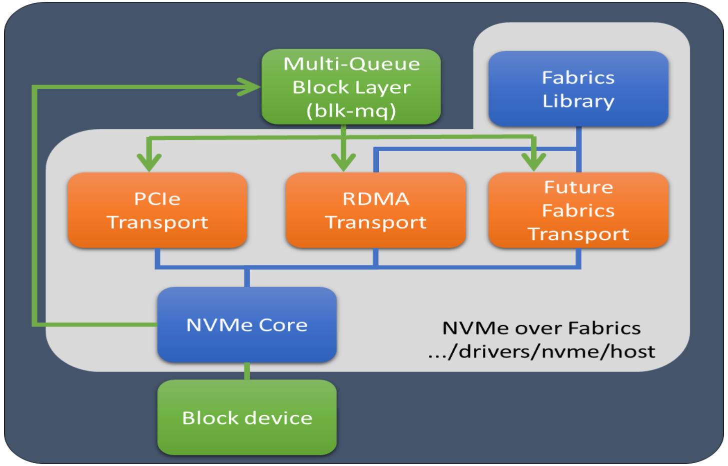

Retains efficiency of PCIe NVMe access over fabrics

NVMe over Fabrics

Target Components

Architecture

- Virtual NVMe subsystem representation of PCIe SSD Subsystem
- Separate NVMe subsystem core components and transports
- Discovery Subsystem and controllers

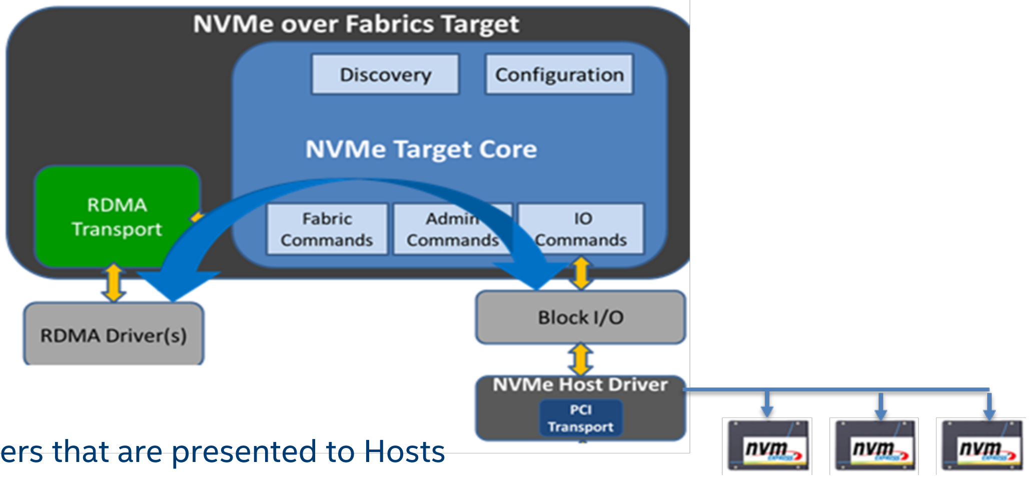

Implementation

- Creates logical NVMe Subsystems and Controllers that are presented to Hosts
- NVMe Namespaces are logically mapped to physical block device

NVMe over Fabrics

Host and Target Stack Status

Linux Host and Target Kernel drivers operational

- Target configured with NVMe PCIe SSD
- To ensure fabrics transport is fabric agnostic it has been tested on InfiniBand™*, iWARP, RoCE RDMA adapters
- Host has been tested with multiple target implementations
- Drivers are still being performance tuned

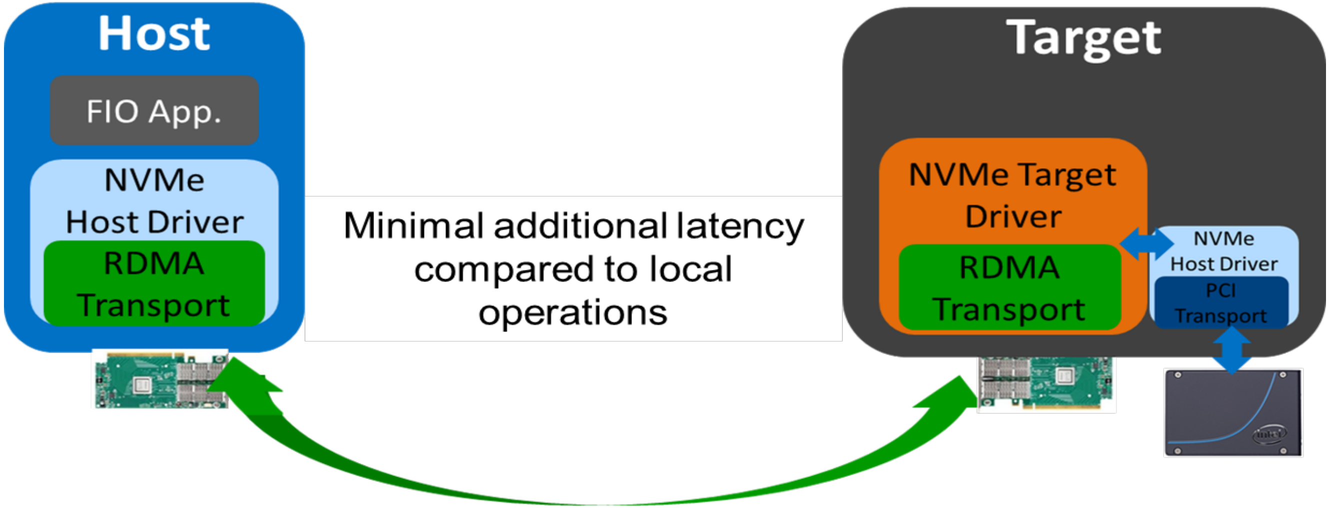

NVMe over Fabrics

NVMe over Fabrics Standardization and Enabling

- **NVMe over Fabrics Specification complete**
  - In release candidate final review
  - Available publicly on nvmexpress.org: 05/16
- **Host and Target drivers will available in upstream Linux kernel**
  - Upstream Kernel under drivers/nvme/{host/target}
  - Available: shortly after specification released in ~4.8
  - Fabrics supported:
    - Local: PCIe
    - RDMA: Fabric agnostic kernel verbs transport
    - Fibre Channel (under development)
    - Other fabrics transports under evaluation / pathfinding (e.g., kFabric)

When the specification and Host and Target code is available

Please download, test with your fabric hardware and environments

Agenda

- Recent product changes and strategy update
- HPC ingredients and opportunities for SSDs
- One problem has many solutions
- NVMe over Fabric (NVMeOF) is here, solutions - on horizon
- Memory Extension – are you ready for this today?

NVME SSD AS A SYSTEM MEMORY BY SCALEMP VSMP

Transparent: no need to change OS or applications
- –Acts as DRAM, no persistency
- –P3700 performance close to 80% of DRAM

Lower TCO: NVMe SSD to replace DRAM
- –Lower CAPEX: significantly lower cost vs. DRAM
- –Lower OPEX: up to 33% lower power vs. DRAM

Increased Capability: NVMe for DRAM expansion
- –Up to 8TB on 2-socket system, and over 20TB on 4-socket system
- –Target in-memory databases and multi-tenant deployments

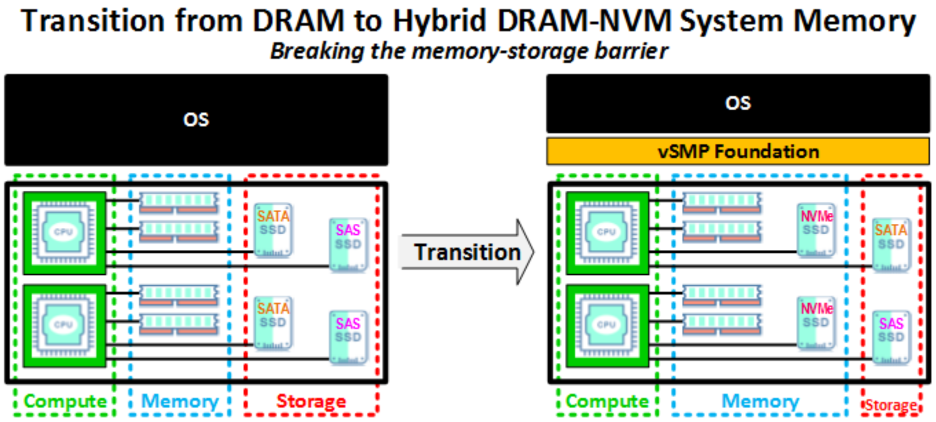

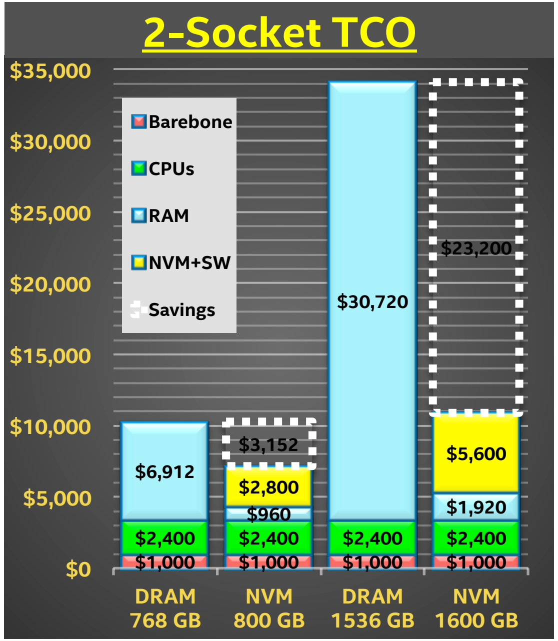

https://communities.intel.com/community/itpeernetwork/blog/2016/02/01/ssd-as-a-system-memory-yes-with-scalemp-s-technology

System configuration: Motherboard S2600WT, 2x Intel Xeon E5 2690v3, 768GB DDR4 or 256GB + 4x Intel SSD P3700 Series 800GB, ScaleMP vSMP

UNIFIED PROGRAMMING MODEL TO ACCESS ANY MEMORY TIER

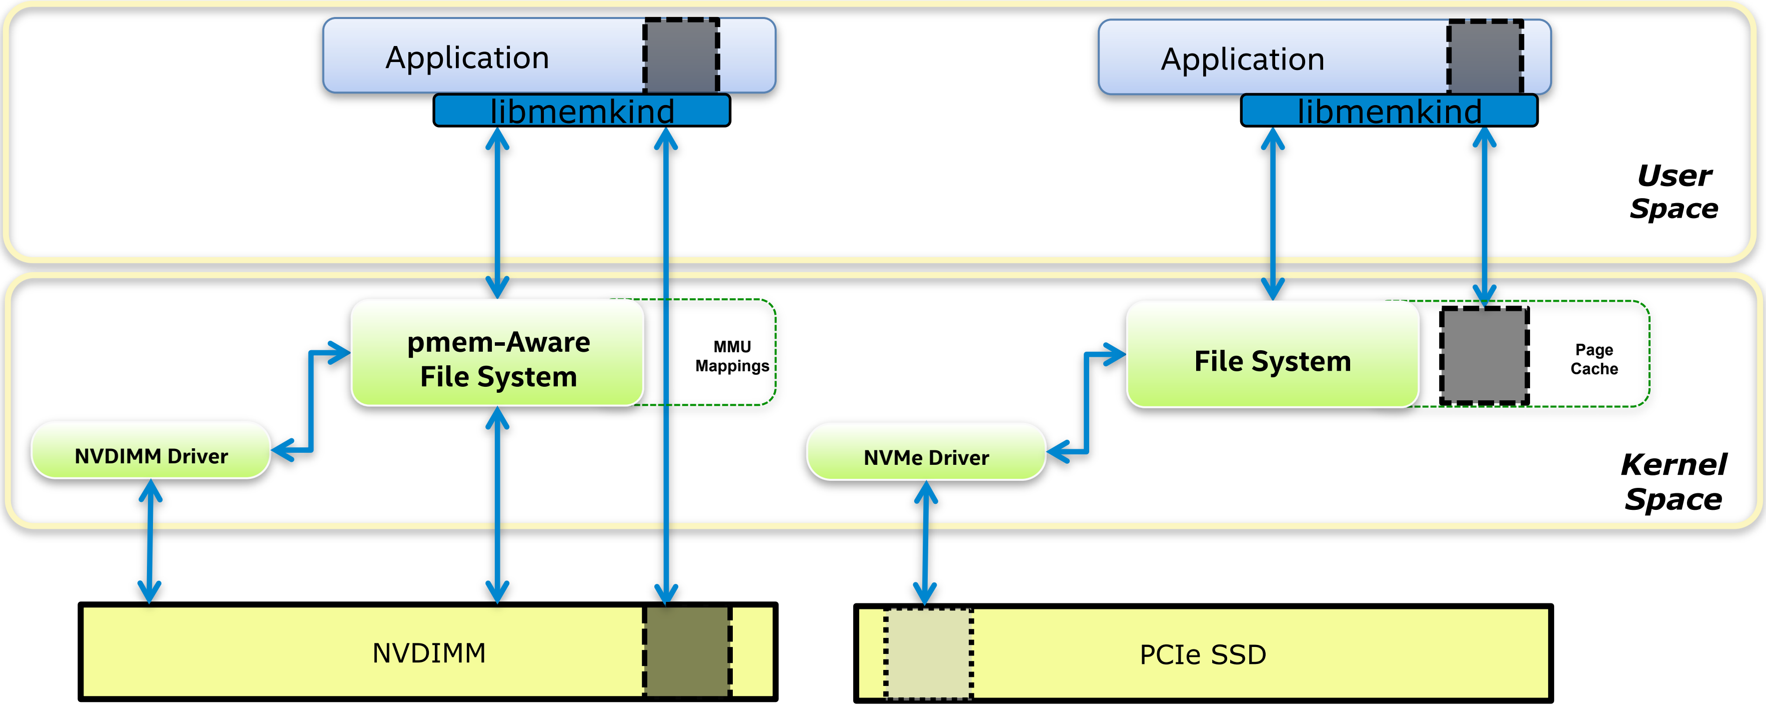

https://github.com/memkind/memkind

Intel® SSD DC D3700, D3600, P3320 press briefing

FOOTNOTES AND DISCLAIMERS

9. IOPS consistency. Source – Intel. Measured performance of Intel® SSD DC S3710 and DC P3700 on 4K Mixed (70/30) workload. Device measured using Iometer. Quality of Service measured using 4 KB (4,096 bytes) transfer size on a random workload on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability. Based on Random 4KB QD=1, 32 workloads, measured as the time taken for 99.9(or 99.9999) percentile of commands to finish the round-trip from host to drive and back to host.

10 .Performance degradation. Source – Intel. Data collected on Intel SSD for Data Center for NVMe family on a standard endurance offering. Performance data collected on cycled drives using short stroke approach adhering to JESD218 method.

Intel® SSD DC D3700, D3600, P3320 press briefing

FOOTNOTES AND DISCLAIMERS

11. **Proven end-to-end protection.** Source – Intel. Demonstrated 10^{-17} UBER as compared to JEDEC UBER specification of 10^{16}. https://www.jedec.org/standards-documents/focus/flash/solid-state-drives. 100X more reliable in preventing Silent Data Corruption. Test performed on Intel® SSD S3x00 drives, Samsung PM853T and SM843T, Micron P400e, Seagate 600 Pro and SanDisk Lightening drives. Drives were exposed to increasing amounts of radiation. After a drive "hang", a power cycle was performed to determine whether the drive would re-boot. If a drive re-booted it was read, and data was compared to the tester's master copy of the up-to-date data that the drive was expected to contain based on writes the drive had acknowledged as completed prior to the "hang" event. If the drive returned data that differed from the expected data, it was recorded as failing for silent errors. The annual rate of silent errors was projected from the rate during accelerated testing divided by the acceleration of the beam (see JEDEC standard JESD89A).
12. **Massive validation.** Source – Intel. Intel validates its data center SSDs across 1000s of drives, 1000s configurations, >5000 workloads and over 1M power cycles.
13. **Value-time curve.** Source – Decision Management Solutions. http://jtonedm.com/2012/11/21/decision-latency-revisited
14. **PLI**. Source – Intel. Intel® Datacenter Drives provide robust Power Loss Imminent (PLI) circuitry that helps to protect inflight data in the event of power loss. Intel drives monitor the health of the PLI circuitry via a Self Cap Test using SMART attributes. Samsung PM853T and SM843T drives were checked for capabilities and flags. No PLI monitoring capabilities (e.g. SMART Attributes) were listed in the Samsung drive specification sheet. Additionally, the drives were tested by powering off a drive and removing one electrolytic (or any other type) capacitor. The drives were then powered up to recollect SMART attribute data to determine is the cap test detected the removal of the capacitor. The Samsung drives did not detect capacitor removal.
15. **Annual Failure Rate**. Source - Intel. Intel SSD Annualized Fail Rate Report for all of 2015. Intel® SSD DC S3500, S3700, P3700.
16. **Intel SSD DC P3320 performance.** Source – Intel. Results for Intel® SSD DC P3320 2TB capacity have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Simulated configuration: Intel Core i7-3770K CPU @ 3.50GHz, 8GB of system memory, Windows1 Server 2012, IOMeter. Random performance simulated for 4 workers with 32 QD. Any differences in your system hardware, software or configuration may affect your actual performance. Comparison to Intel® SSD DC3510 based on that product's spec datasheet for 1.2TB capacity.
17. **Intel SSD DC P3320 benchmarks**. Source – Intel. Any differences in your system hardware, software or configuration may affect your actual performance. System Configuration: Supermicro 2U SuperServer® 2028U-TNR4T+, Dual Intel® Xeon® Processor E5-2699 V3 (45M Cache, 2.30 GHz), 192 GB DDR4 DRAM, Boot Drive – Intel® SSD Data Center S3710 Series (200GB), NVMe – Intel® SSD Data Center P3320 Series (2TB), SATA – Intel® SSD Data Center S3510 Series (1.6TB). Running Vmware ESXi 6.0 with two identical Windows Server VMs running various workloads housed on two different SSD-backed Datastores.
18. **Intel SSD DC D3700 and D3600 performance.** Source – Intel. Performance comparison using Intel DC D3700 1.6TB spec datasheet vs HGST ME SAS 1.6TB drive spec datasheet: Intel DC D3700 specs: http://www.intel.com/content/www/us/en/solid-state-drives/solid-state-drives-dc-d3700-series.html, HGST specs: https://www.hgst.com/products/solid-state-drives/ultrastar-ssd1600mm

Intel® SSD DC D3700, D3600, P3320 press briefing

FOOTNOTES AND DISCLAIMERS

19. **Intel SSD DC D3700 vs. SAS SSD performance comparison.** Source – XIO. Configuration - External Host with windows server 2008 running. External host specifications: HP DL360, G7 with dual intel E5-2620 and 25GB ram. Storage array system using E52699v3 with 40*Intel DC D3700 10 DWPD 800GB & Storage array system using E52699v3 with 40* SAS 10 DWPD 400GB . Test - 8K transfer with 80/20 Read/Write workload on QD 1,2,4 accessing 1 volume on the shared storage array. Measurements taken on IOMeter.
20. **Data Center SSD shipments.** Source – Gartner, Q4 2015
21. **Delivery to commit.** Source – Intel. Actual measurement of percent shipments for all Intel SSD products that arrive on committed delivery date.
22. **Factory quality**. Source – Intel. Average Functional Outgoing Quality Monitor measurement across all Intel SSD products.
23. **Best-in-class supply chain:** Source - Gartner. "The Gartner Supply Chain Top 25 for 2015." Intel #4 ranked overall and top ranked in the technology industry. http://www.gartner.com/technology/supply-chain/top25.jsp
24. **ISV enabling, OEM design-in support.** Source – Intel Software and Solutions Group, Intel Non-volatile Memory Solutions Group, Intel Influencer Sales Group, Intel Direct and Channel Sales.
25. **NVMe gets to results faster.** Source - Storage Technologies Group, Intel. Comparisons between memory technologies based on in-market product specifications and internal Intel specifications.
26. **Performance degradation range for all data center drives** - Source – Intel. Data collected on Intel SSD for Data Center for PCIe family on a standard endurance offering. Performance data collected on cycled drives using short stroke approach adhering to JESD218 method. Data collected on Intel SSD for Data Center for SATA family on a standard endurance offering. Configuration – Windows 2012 Server, DDR4-32GB, Xeon DP Haswell-EP E5-2699 v3 LGA2011 2.3GHz 45MG 145 W 18 core, G60T0045 firmware. Workload flow Sequence - Prefill seq WR, Seq WR highest QD BS -> lowest, Seq RD highest QD BS ->lowest, RND WR highest QD BS ->70WRhighest QD BS->30WR highest QD BS -> RND RD highest QD BS.
27. **Most widely adopted NVMe family**. Source - Forward Insights. Inside Solid State Drives. Q4 2015.