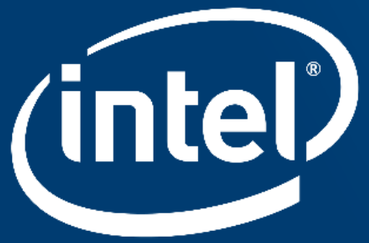

INTEL® OPTANE™ TECHNOLOGY BASICS

March 2017

NVM Solutions Group

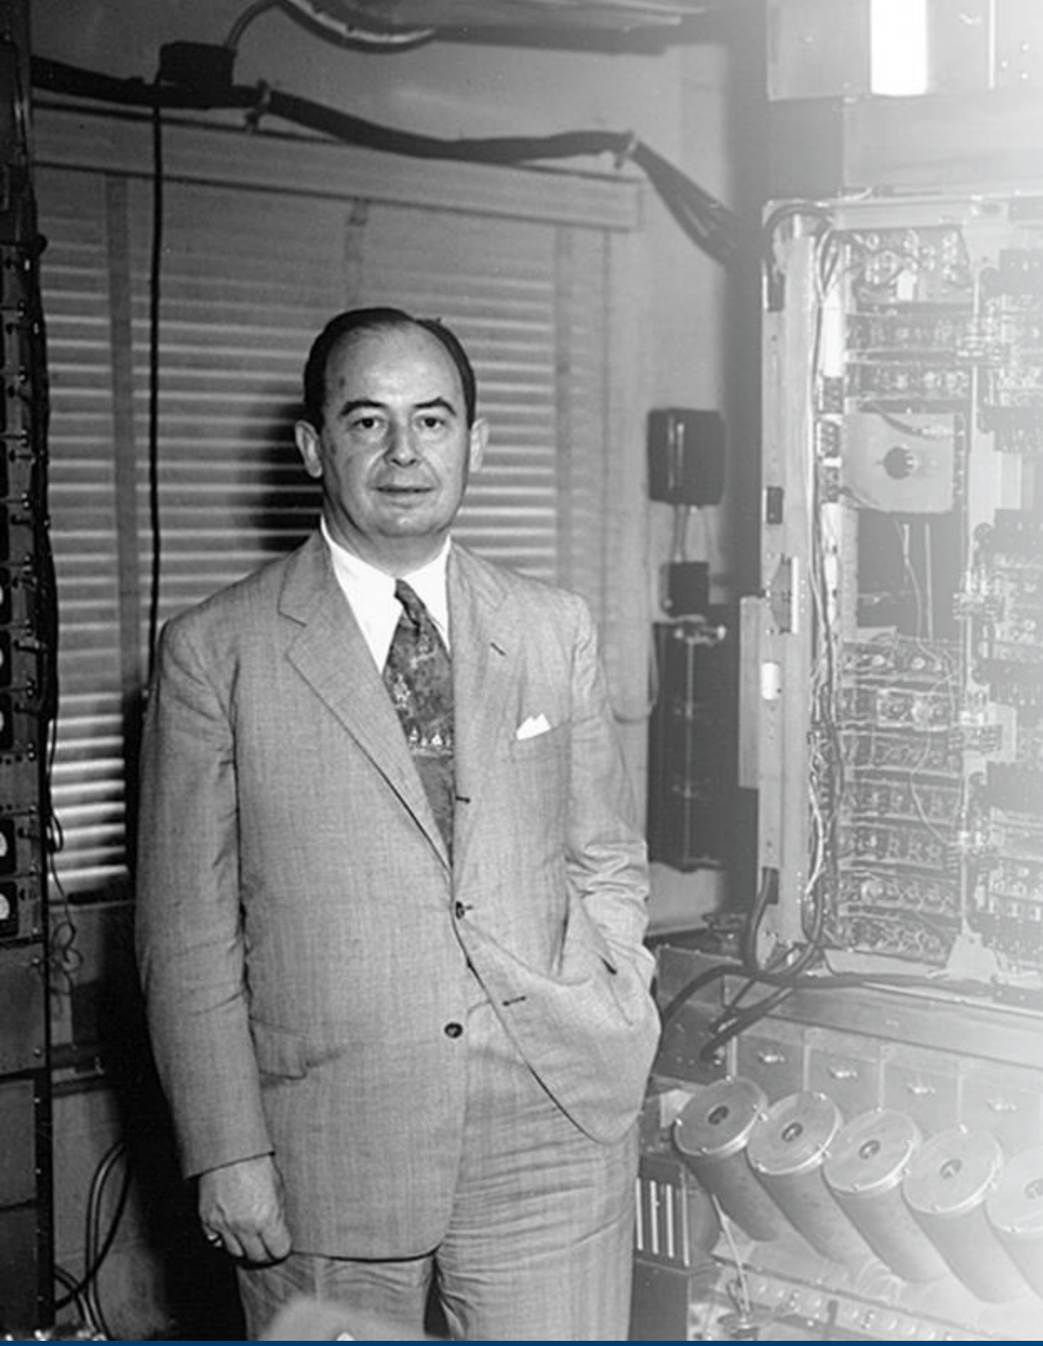

“ Ideally one would desire an indefinitely large memory capacity such that any particular … word would be immediately available. … It does not seem possible physically to achieve such a capacity. We are therefore forced to recognize the possibility of constructing a hierarchy of memories, each of which has greater capacity than the preceding but which is less quickly accessible.”

Preliminary Discussion of the Logical Design of an Electronic Computing Instrument
Arthur Burks, Herman Goldstine and John von Neumann, 1946

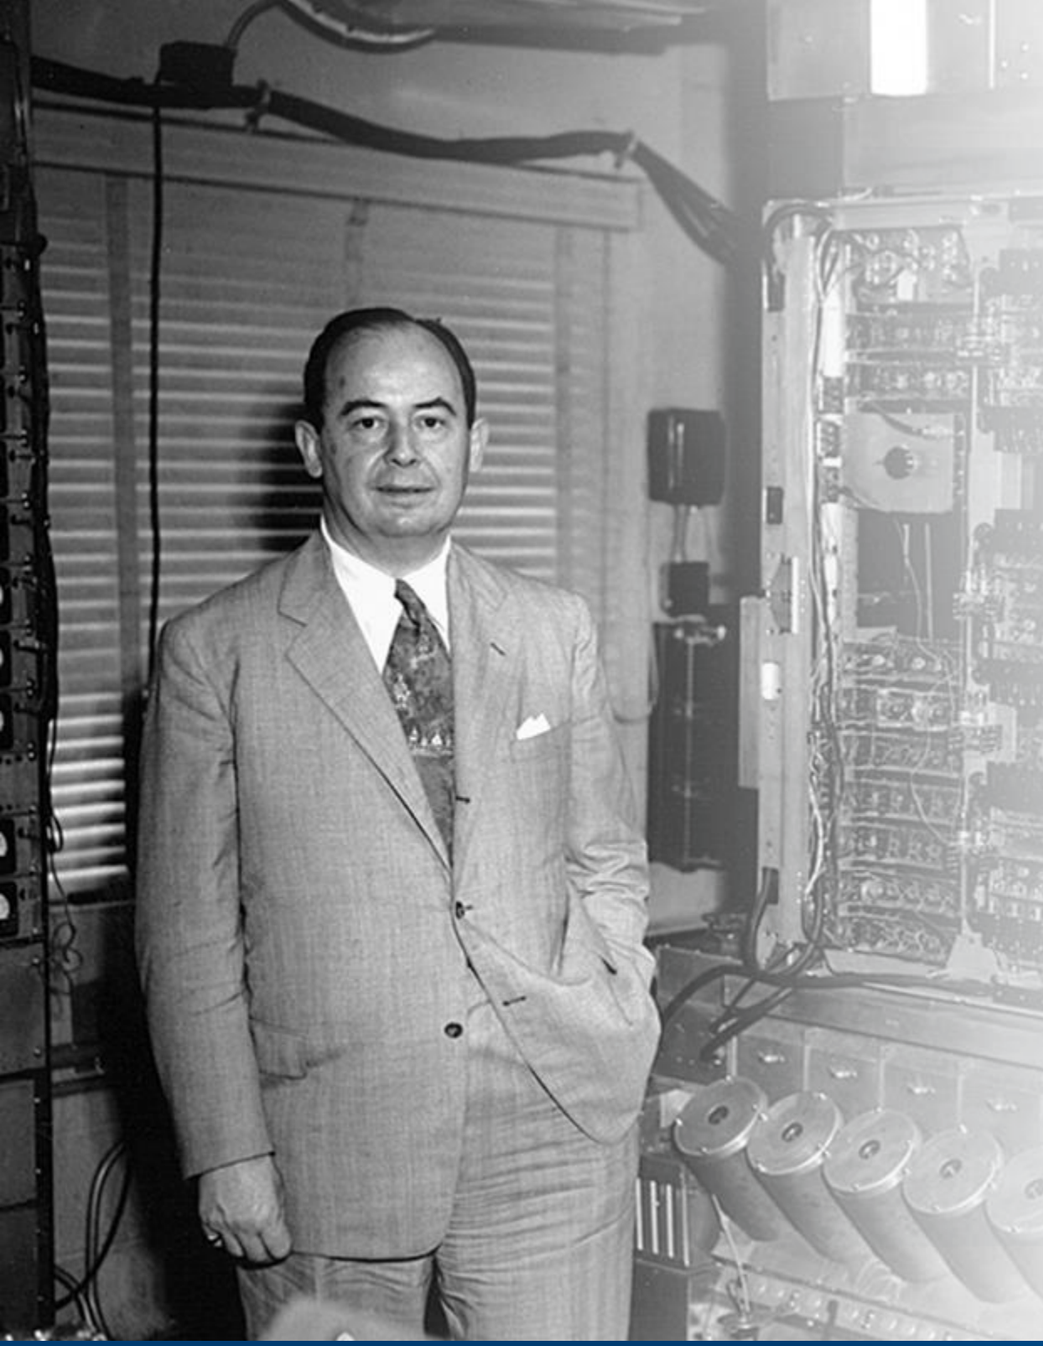

“ Ideally one would desire an **indefinitely large memory capacity** such that any particular … word would be **immediately available**. … It **does not seem possible physically** to achieve such a capacity. We are therefore forced to recognize the possibility of **constructing a hierarchy of memories**, each of which has **greater capacity than the preceding** but which is **less quickly accessible**.”

Preliminary Discussion of the Logical Design of an Electronic Computing Instrument
Arthur Burks, Herman Goldstine and John von Neumann, 1946

MEMORY AND STORAGE HIERARCHY

The Performance Problem Grows Bigger Over Time

MEMORY

SRAM
Latency: 1X
Size of Data: 1X

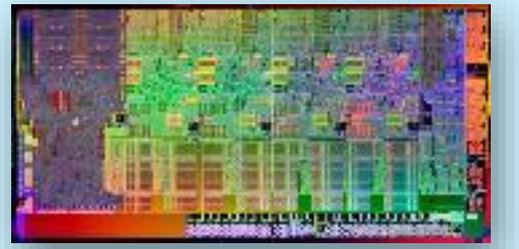

DRAM
Latency: ~10X
Size of Data: ~100X

STORAGE

CPU Perf.
Normalized Media
Access Time for 20K Read

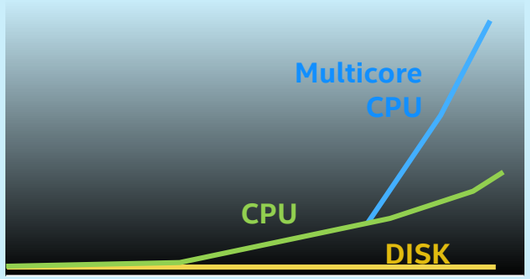

Source: Intel Measurements

HDD
Latency: ~10 Million X
Size of Data: ~10,000X

Technology claims are based on comparisons of latency, density and write cycling metrics amongst memory technologies recorded on published specifications of in-market memory products against internal Intel specifications.

MEMORY AND STORAGE HIERARCHY

NAND SSDs Help Alleviate The Gap In The Hierarchy

MEMORY

SRAM
Latency: 1X
Size of Data: 1X

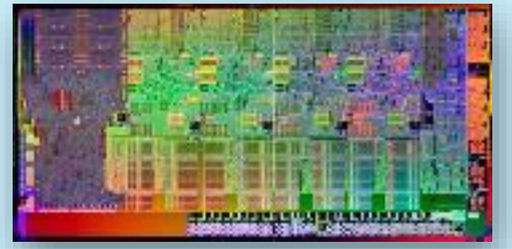

DRAM
Latency: ~10X
Size of Data: ~100X

STORAGE

NAND SSD
Latency: ~100,000X
Size of Data: ~1,000X

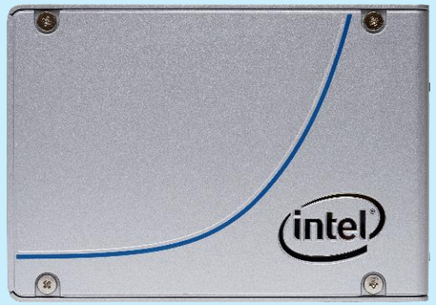

HDD
Latency: ~10 Million X
Size of Data: ~10,000X

Technology claims are based on comparisons of latency, density and write cycling metrics amongst memory technologies recorded on published specifications of in-market memory products against internal Intel specifications.

EVERY DAY DATA IS EXPLODING

1.5 GB[1]
AVERAGE INTERNET USER

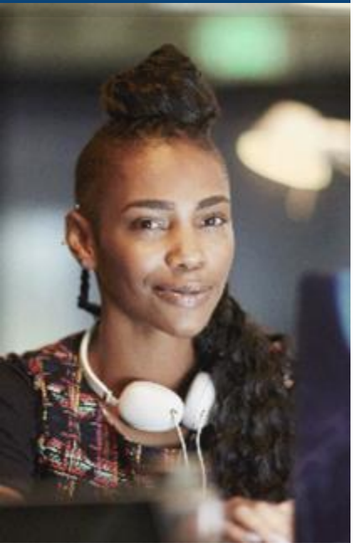

3,000 GB[2]
SMART HOSPITAL

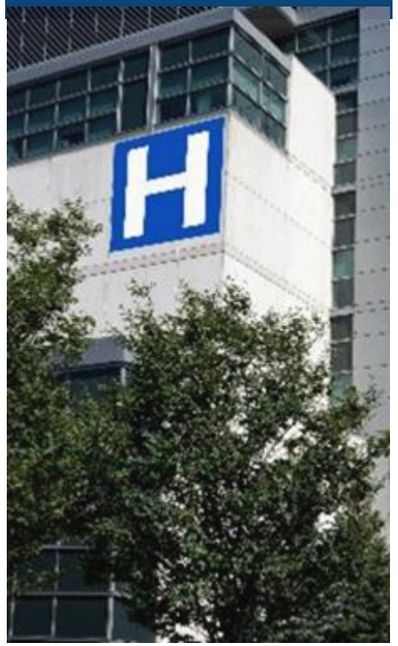

4,000 GB[3]
AUTONOMOUS DRIVING

40 K GB[4]
AIRPLANE DATA

1 M GB[5]
SMART FACTORY

1. Source: http://www.cisco.com/c/en/us/solutions/service-provider/vni-network-traffic-forecast/infographic.html
2. Source: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html
3. Source: https://datafloq.com/read/self-driving-cars-create-2-petabytes-data-annually/172
4. Source: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html
5. Source: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html

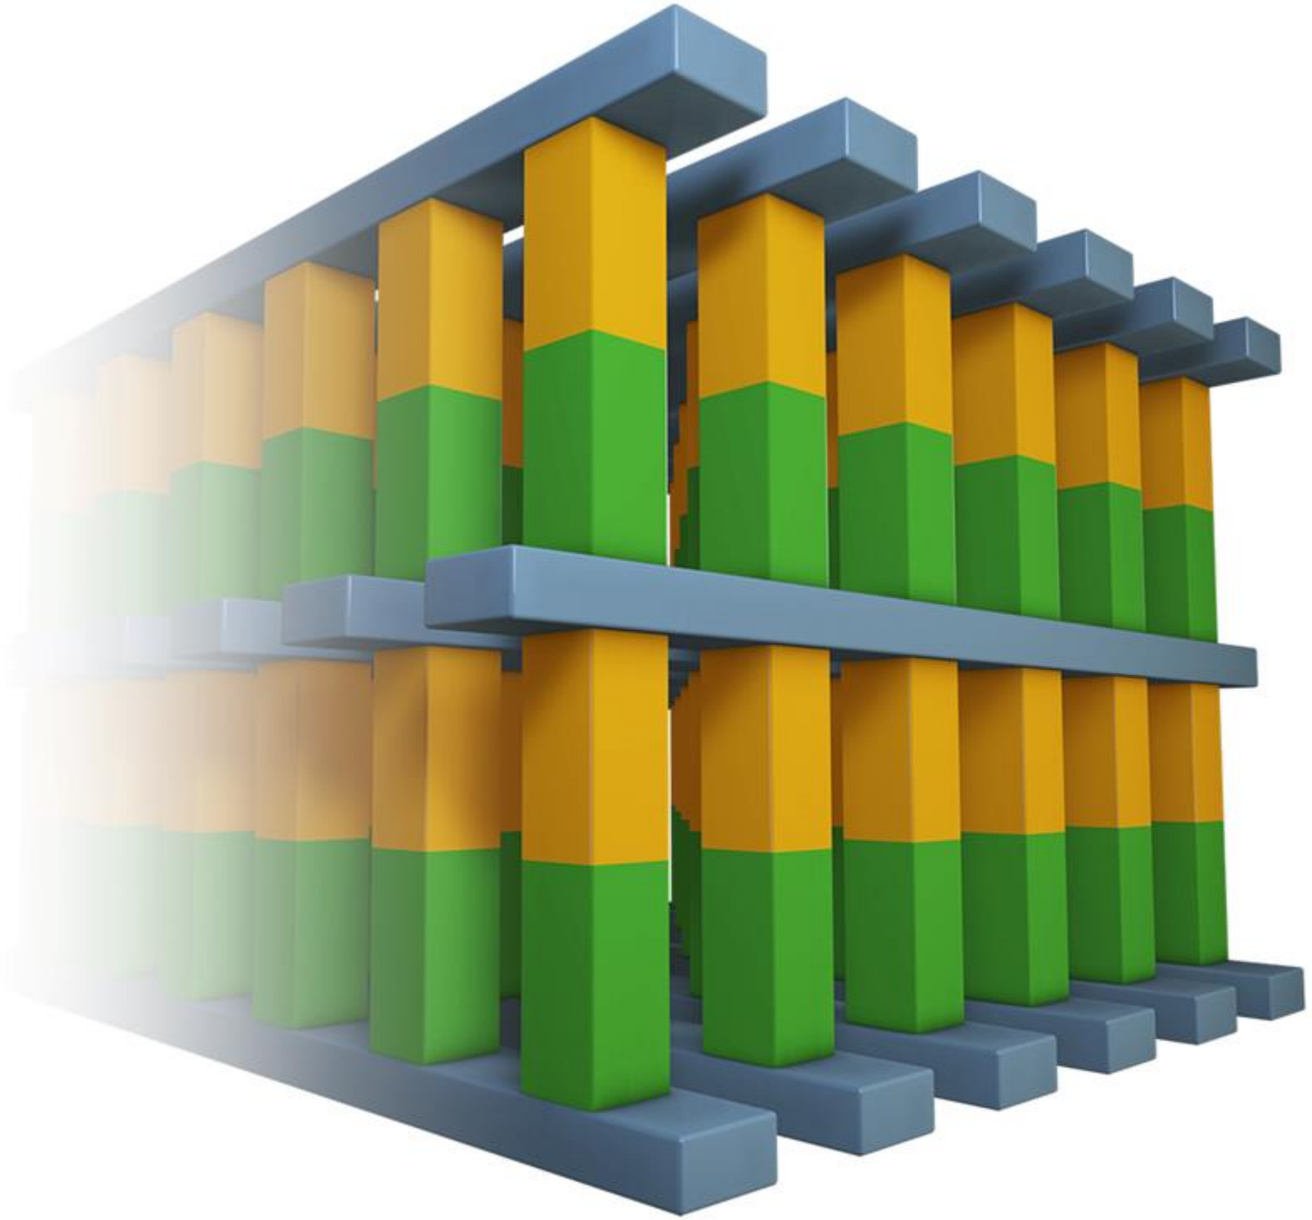

A NEW CLASS OF STORAGE AND MEMORY

3D XPOINT™ MEMORY MEDIA

3D XPOINT™ MEMORY MEDIA

In Pursuit of Large Memory Capacity … Word Access … Immediately Available …

Word (Cache Line)

Crosspoint Structure

Selectors allow dense packing and individual access to bits

Large Memory Capacity

Crosspoint & Scalable

Memory layers can be stacked in a 3D manner

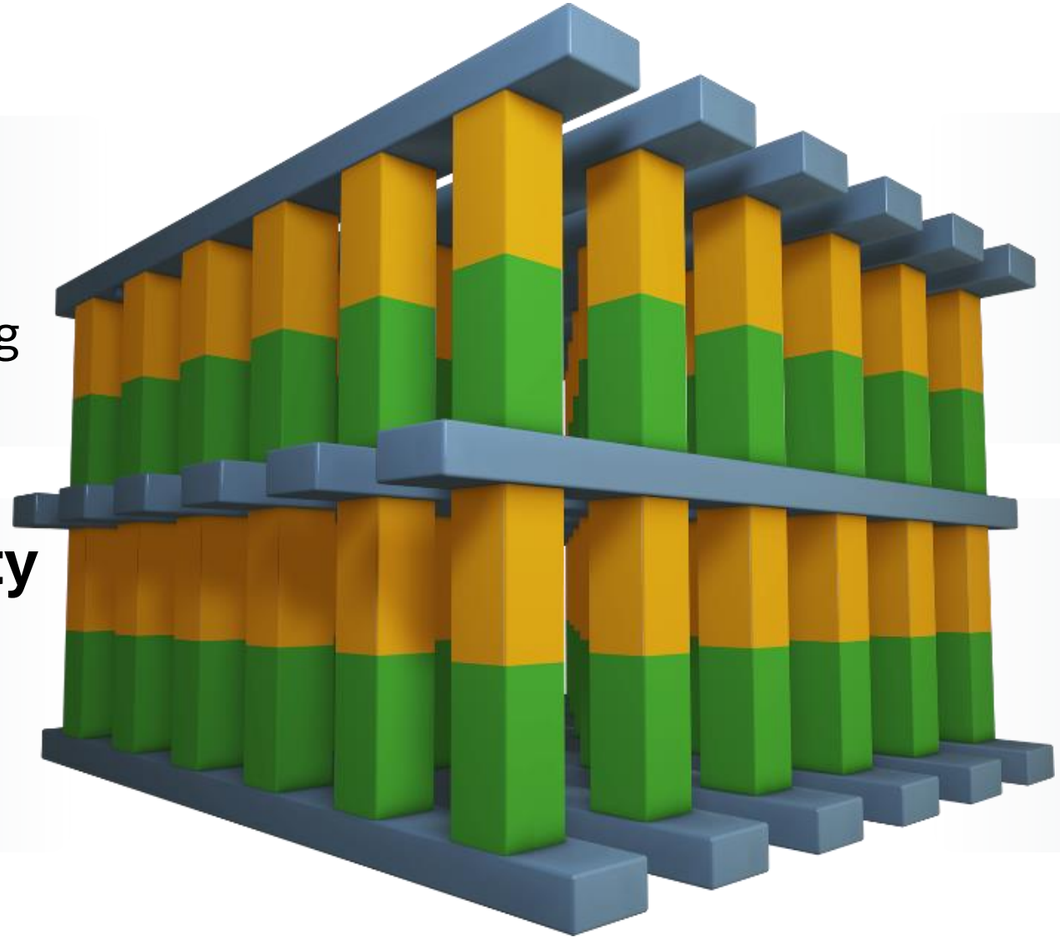

NVM Breakthrough

Material Advances

Compatible switch and memory cell materials

Immediately Available

High Performance

Cell and array architecture that can switch states 1000x faster than NAND

3D XPOINT™ MEMORY MEDIA

Breaks the memory/storage barrier

MEMORY + STORAGE

SRAM
Latency: 1X
Size of Data: 1X

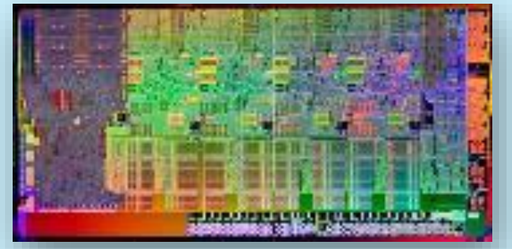

DRAM
Latency: ~10X
Size of Data: ~100X

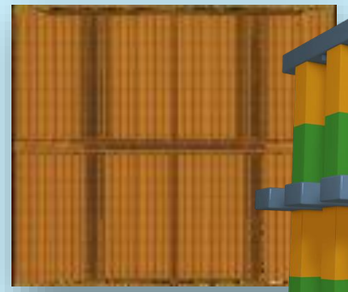

3D XPoint™
Latency: ~100X
Size of Data: ~1,000X

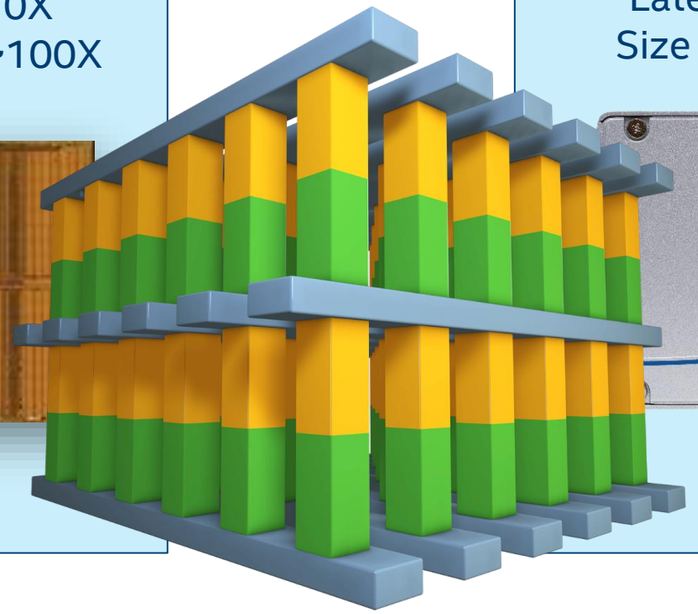

NAND SSD
Latency: ~100,000X
Size of Data: ~1,000X

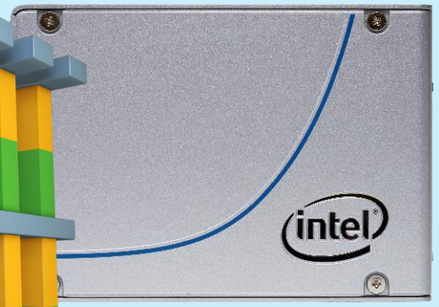

HDD
Latency: ~10 Million X
Size of Data: ~10,000X

Technology claims are based on comparisons of latency, density and write cycling metrics amongst memory technologies recorded on published specifications of in-market memory products against internal Intel specifications.

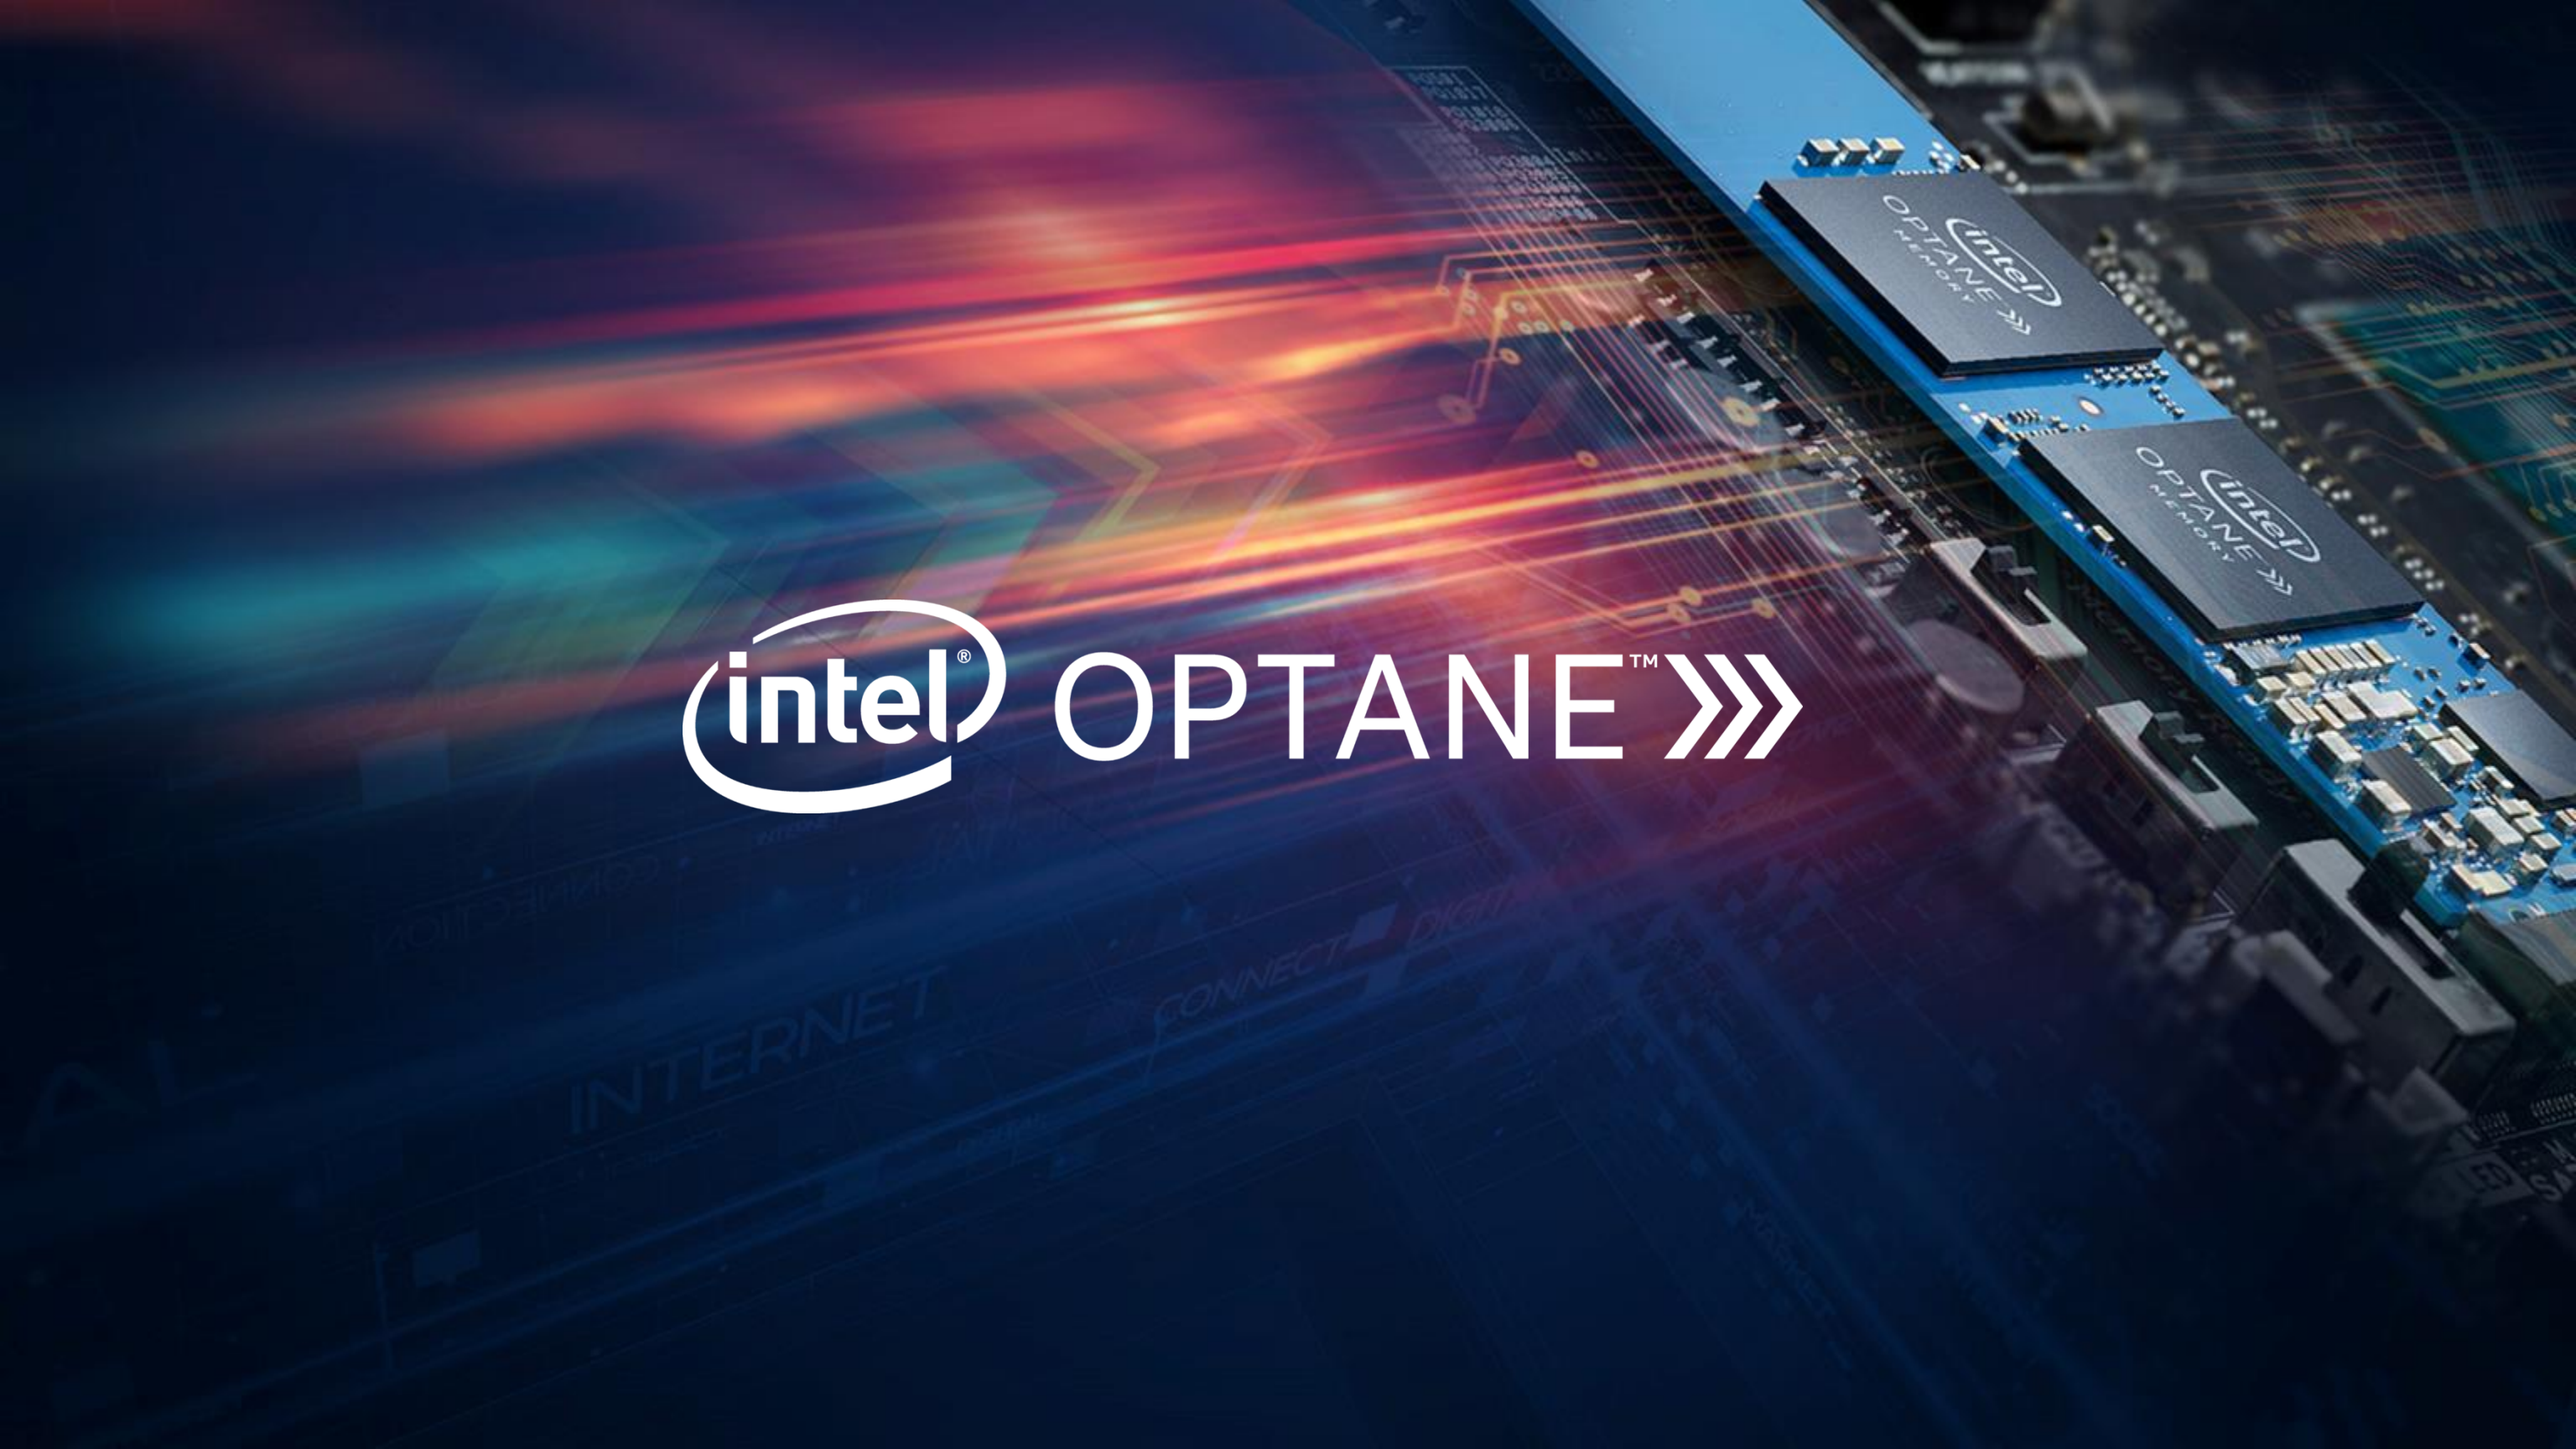
intel
OPTANE™

INTEL® OPTANE™ TECHNOLOGY: BUILDING BLOCKS

Unleashing Breakthrough Performance for a New Generation of Computing

3D XPoint™ Memory Media

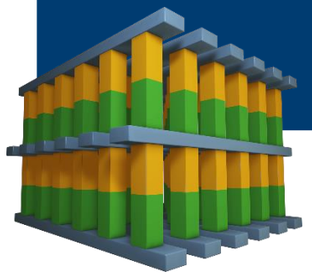

+

Intel Memory and Storage Controllers

+

Intel Interconnect IP

+

Intel Software

OPTIMIZED AT EVERY LEVEL TO DELIVER 3D XPOINT™ MEMORY MEDIA ADVANTAGES TO THE PLATFORM

INTEL® OPTANE™ SSD: ALL NEW DESIGN

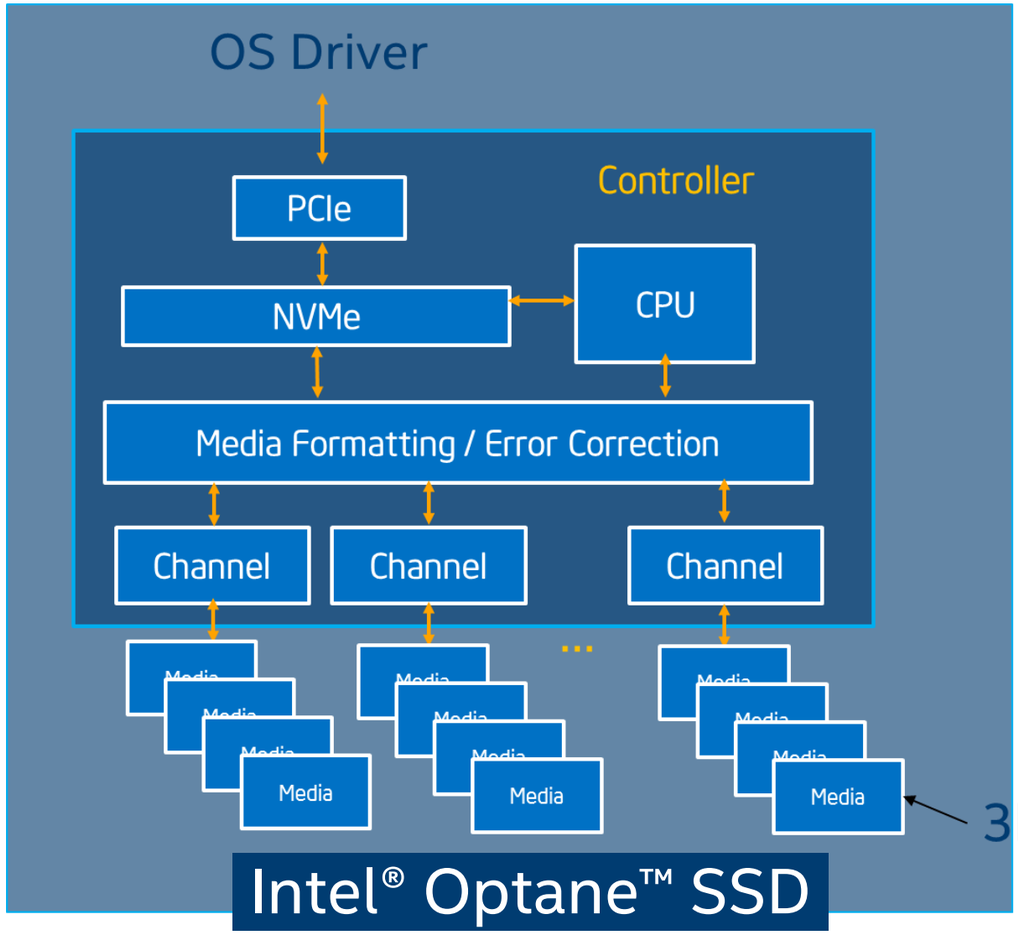

- Optimized storage interface PCIe*/NVMe*
- Hardware-only read/write path controller
- Highly parallel media access
- Write-in-place design
- Completely new media management
- Co-architected, co-designed, and co-optimized with 3D XPoint™ memory media

CO-ARCHITECTED, CO-DESIGNED, CO-OPTIMIZED WITH 3D XPOINT™ MEMORY MEDIA

*Other names and brands may be claimed as the property of others.

UNPRECEDENTED PERFORMANCE

10x latency reduction†

- < 10usec latency

100x QoS improvement†

- < 200usec 99.999th r/w

Uniform behavior for reads and writes

High Endurance

- 30 Drive Writes Per Day

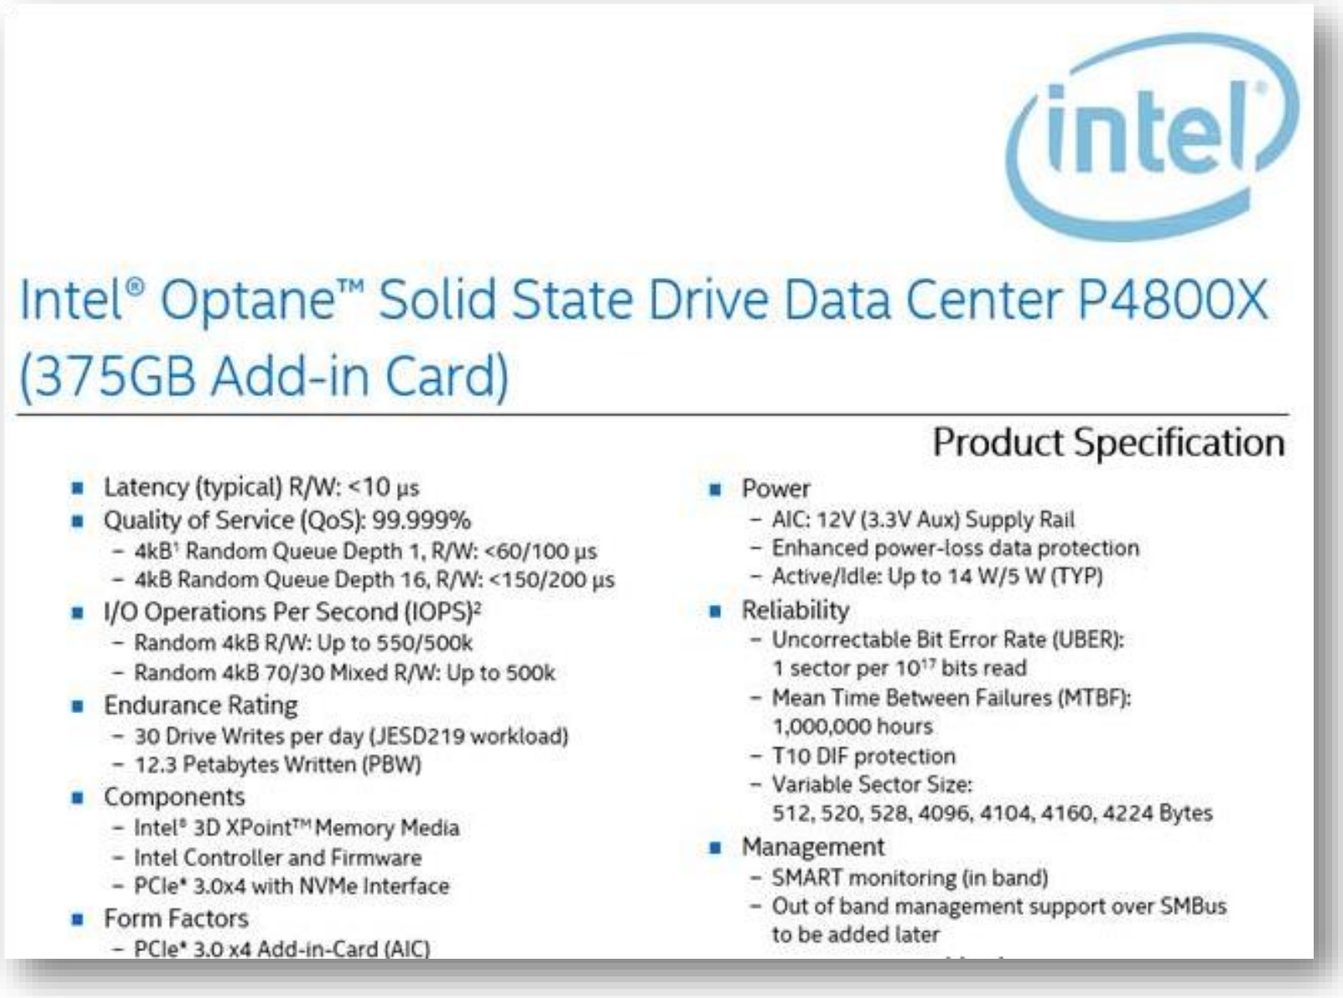

† vs. NAND based SSD as shown on succeeding slides

SSD PERFORMANCE: AT VARYING QUEUE DEPTHS

4K RANDOM READ

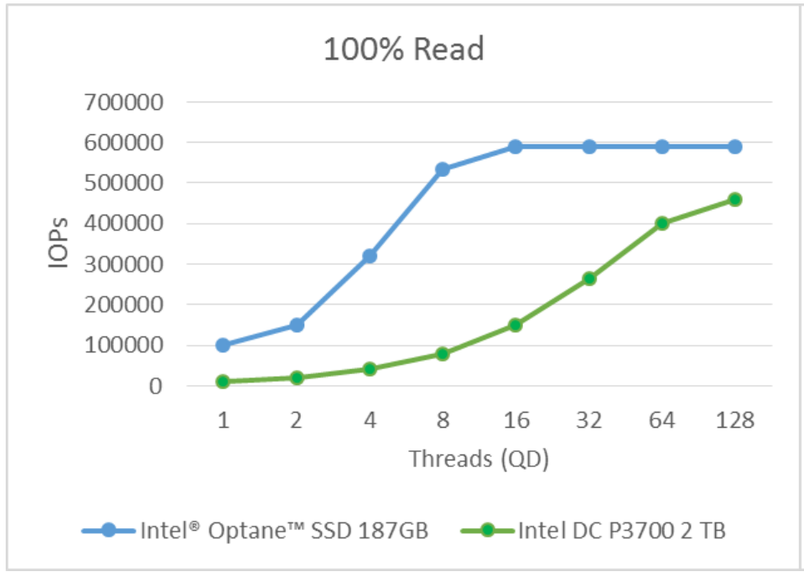

4K RANDOM WRITE

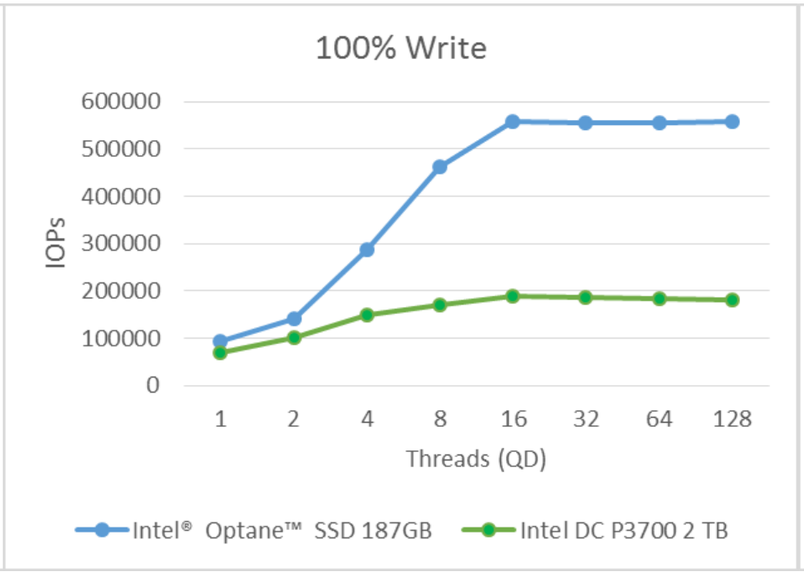

4K RANDOM 70/30 MIX

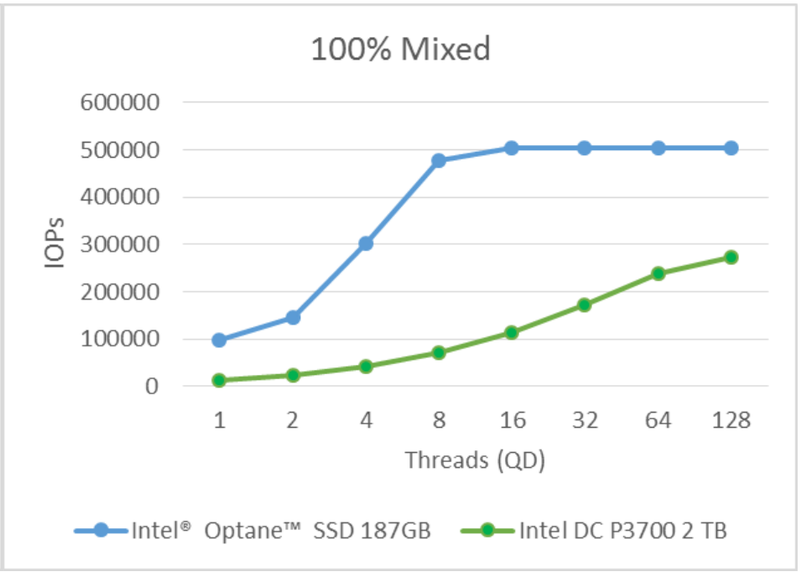

INTEL® OPTANE™ SSDs DELIVER HIGH IOPS FOR A SMALL # OF THREADS BUT THIS MEASURE IGNORES TIME PER I/O

Results measured by Intel based on the following configurations. Ubuntu 16.04.2 LTS (GNU/Linux 4.4.0-21-generic x86_64); Intel S2600WT motherboard with 2x Xeon E5-2699v4 @ 2.20GHz, Turbo @ 3.6GHz, 256GB RAM, fio-2.2.10, irqbalance off, smp affinity changed, cpu governor = performance; Prototype Intel Optane SSD: 187GB, FW: E2010211, Intel DC P3700: 2 TB, FW: 8DV101F0

SSD QUALITY OF SERVICE

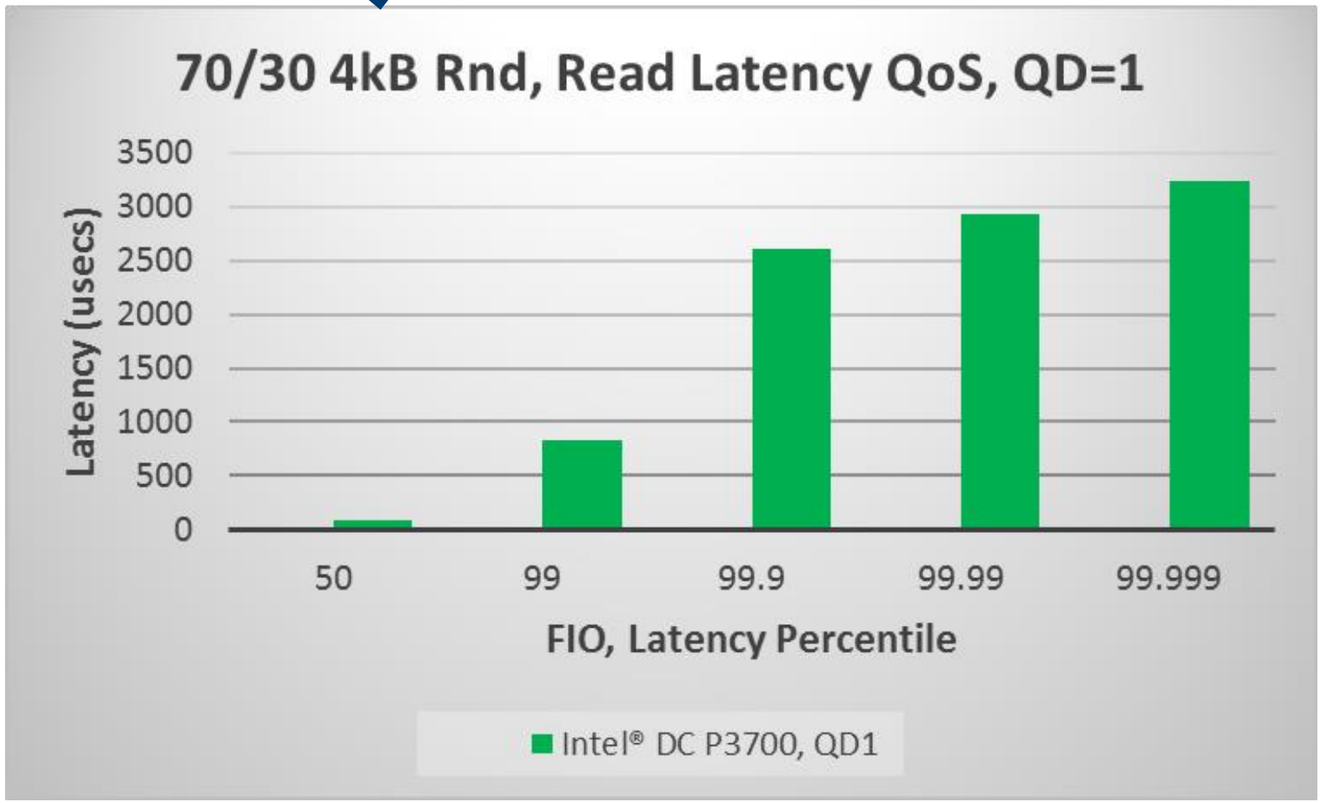

Config: I7-6700K Turbo to 4.3GHz, ASUS* Z170m-plus 4x4GB DDR-2133, Hyperthreading disabled CPU C-state disabled, Intel P3700 SSD 800GB, Ubuntu* 14.04 LTS 64 bit server, kernel 4.4 (polling enabled), FIO 2.1.11

SSD QUALITY OF SERVICE

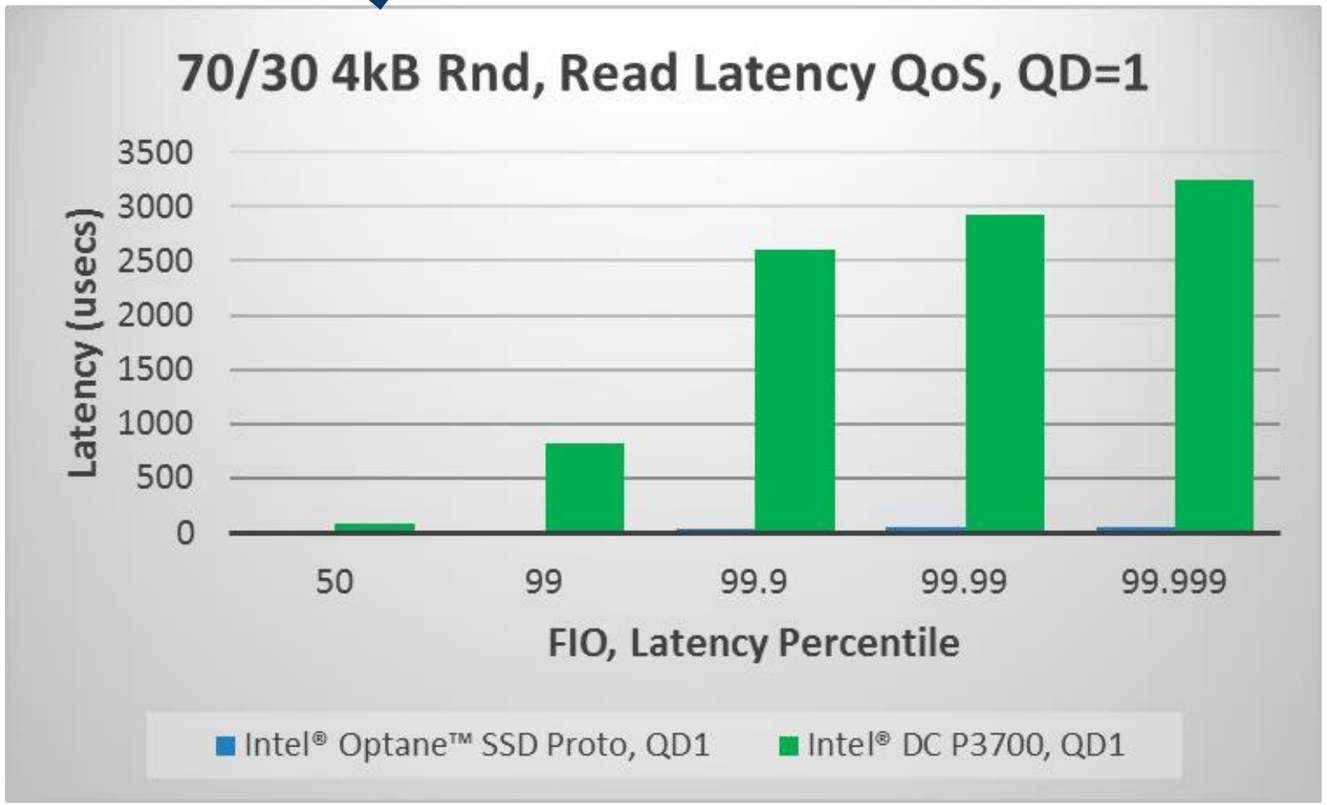

INTEL® OPTANE™ SSDs DELIVER HUGE RESPONSE LATENCY ADVANTAGES BUT THIS MEASURE IGNORES THROUGHPUT (IOPS)

Config: I7-6700K Turbo to 4.3GHz, ASUS* Z170m-plus 4x4GB DDR-2133, Hyperthreading disabled CPU C-state disabled, Intel P3700 SSD 800GB, Ubuntu* 14.04 LTS 64 bit server, kernel 4.4 (polling enabled), FIO 2.1.11

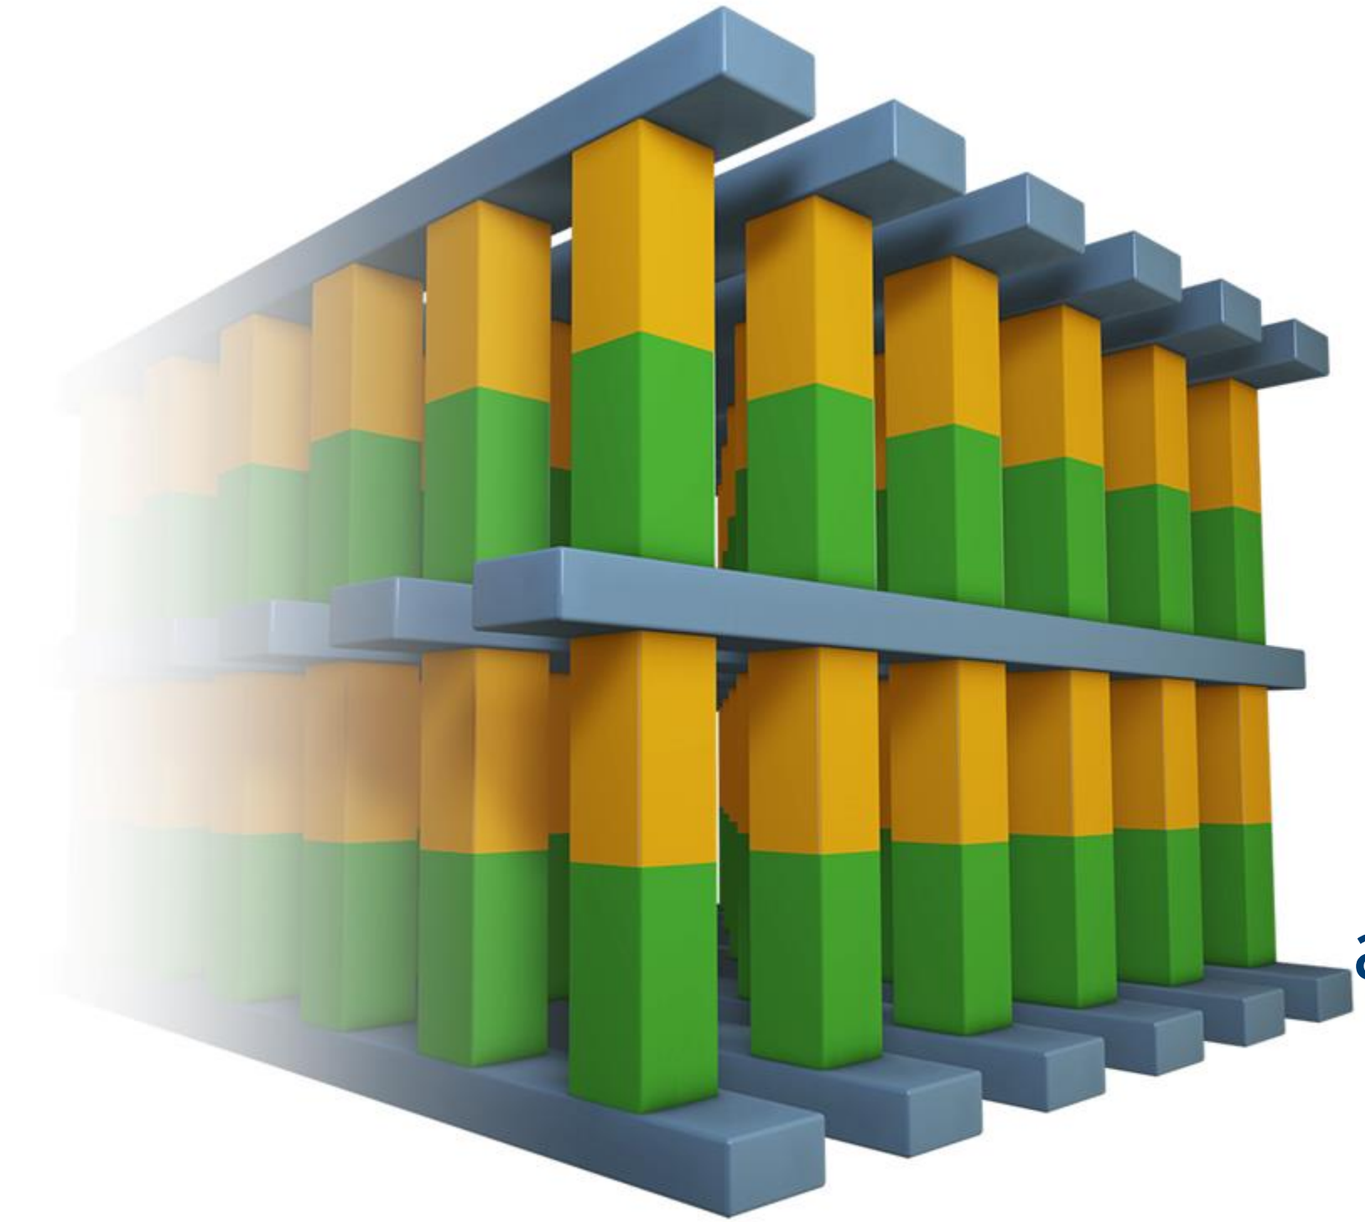

Storage this fast demands a new measurement

MEASURE IT LIKE DRAM

MEASURE
LATENCY AT LOAD

This measure shows up as application **Responsiveness**

Responsiveness is equivalent to read latency; Load represents a defined workload

STORAGE PERFORMANCE

Latency vs. Load: NAND SSD vs. Intel® Optane™ SSD (Intel® DC P3700 vs. Intel® P4800x)

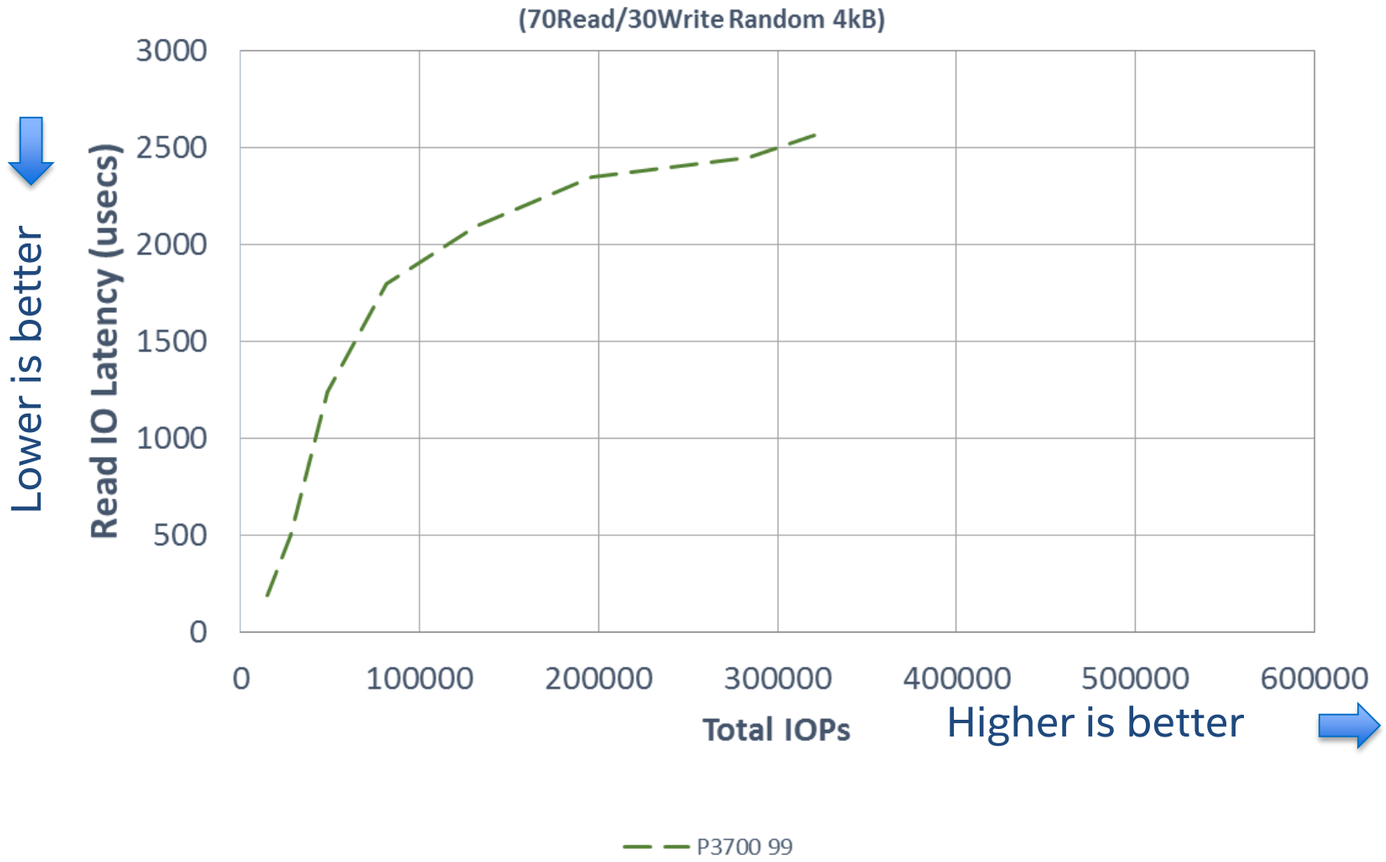

Results measured by Intel based on the following configurations. 375GB P4800X or 800GB P3700, Intel(R) Xeon(R) CPU E5-2695 v3 @ 2.30GHz, Wildcat Pass, 4 x 8GB DDR4 32GB total, Hyper-threading disabled, CPU C-state disabled, Ubuntu 15.04 LTS 64 bit server (v3.19), FIO 2.1.11. Performance on final samples is subject to change.

STORAGE PERFORMANCE

Latency vs. Load: NAND SSD vs. Intel® Optane™ SSD (Intel® DC P3700 vs. Intel® P4800x)

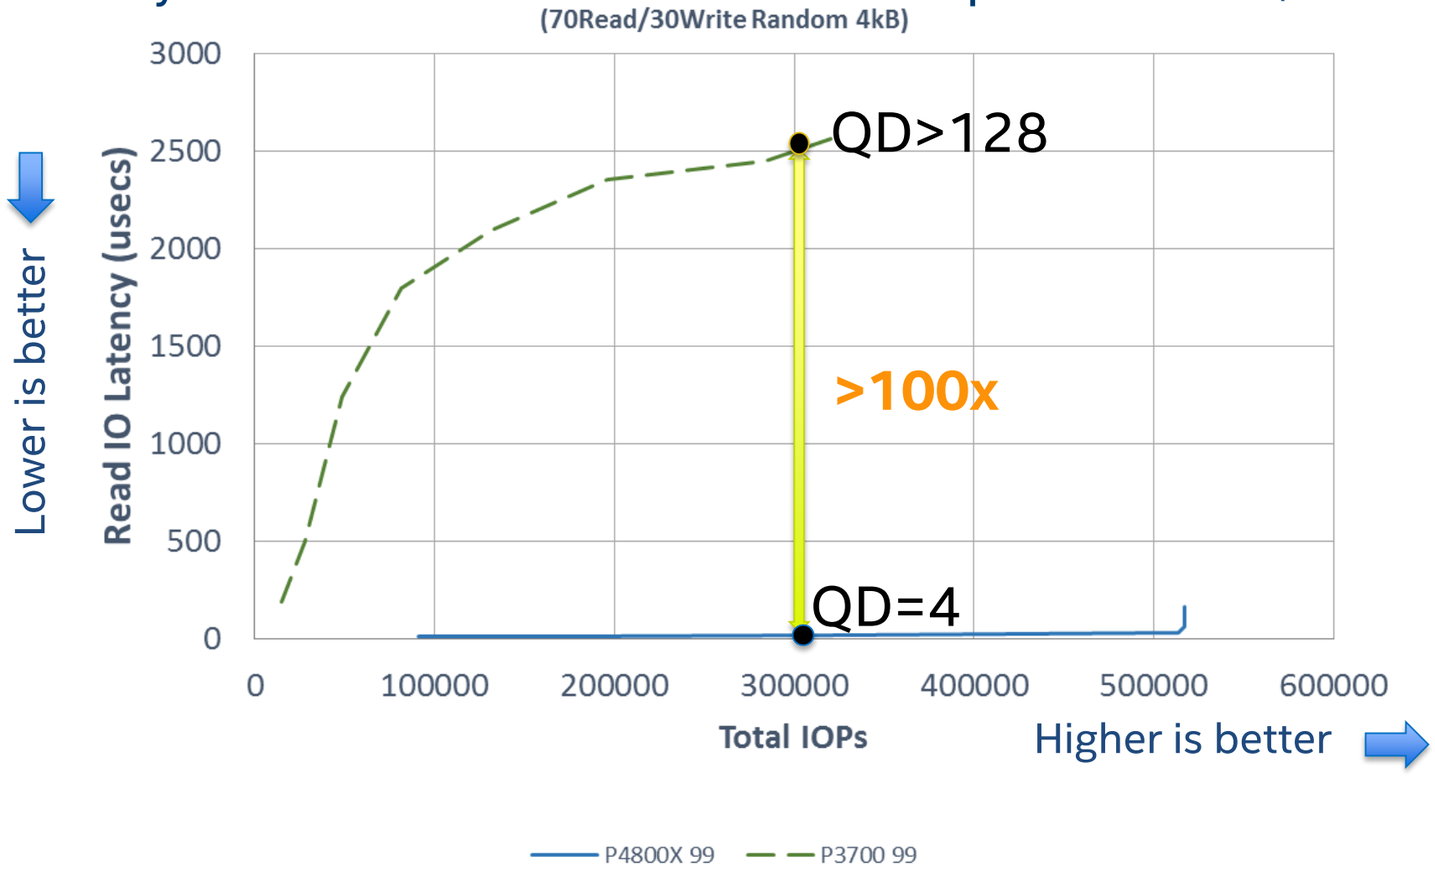

A MORE COMPLETE MEASURE OF SSD PERFORMANCE: RESPONSIVENESS UNDER LOAD

STORAGE PERFORMANCE CHARACTERIZATION

Latency vs. Load: NAND SSD vs. Intel® Optane™ SSD (Intel® DC P3700 vs. Intel® P4800x)

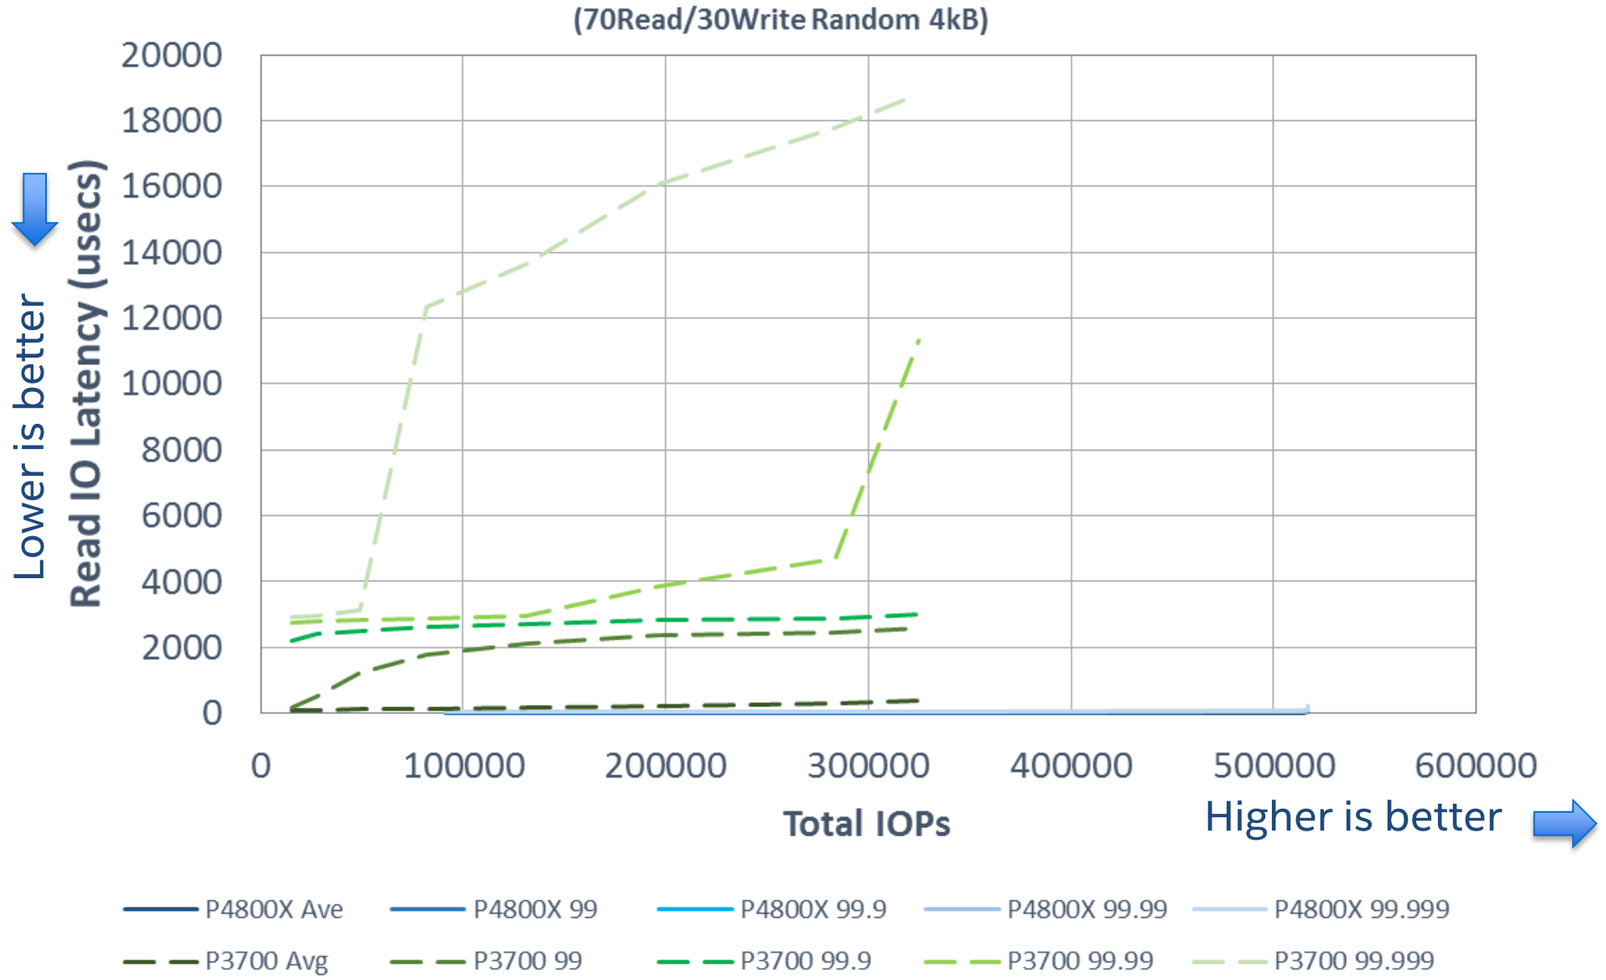

STORAGE PERFORMANCE CHARACTERIZATION

Latency vs. Load: NAND SSD vs. Intel® Optane™ SSD (Intel® DC P3700 vs. Intel® P4800x)

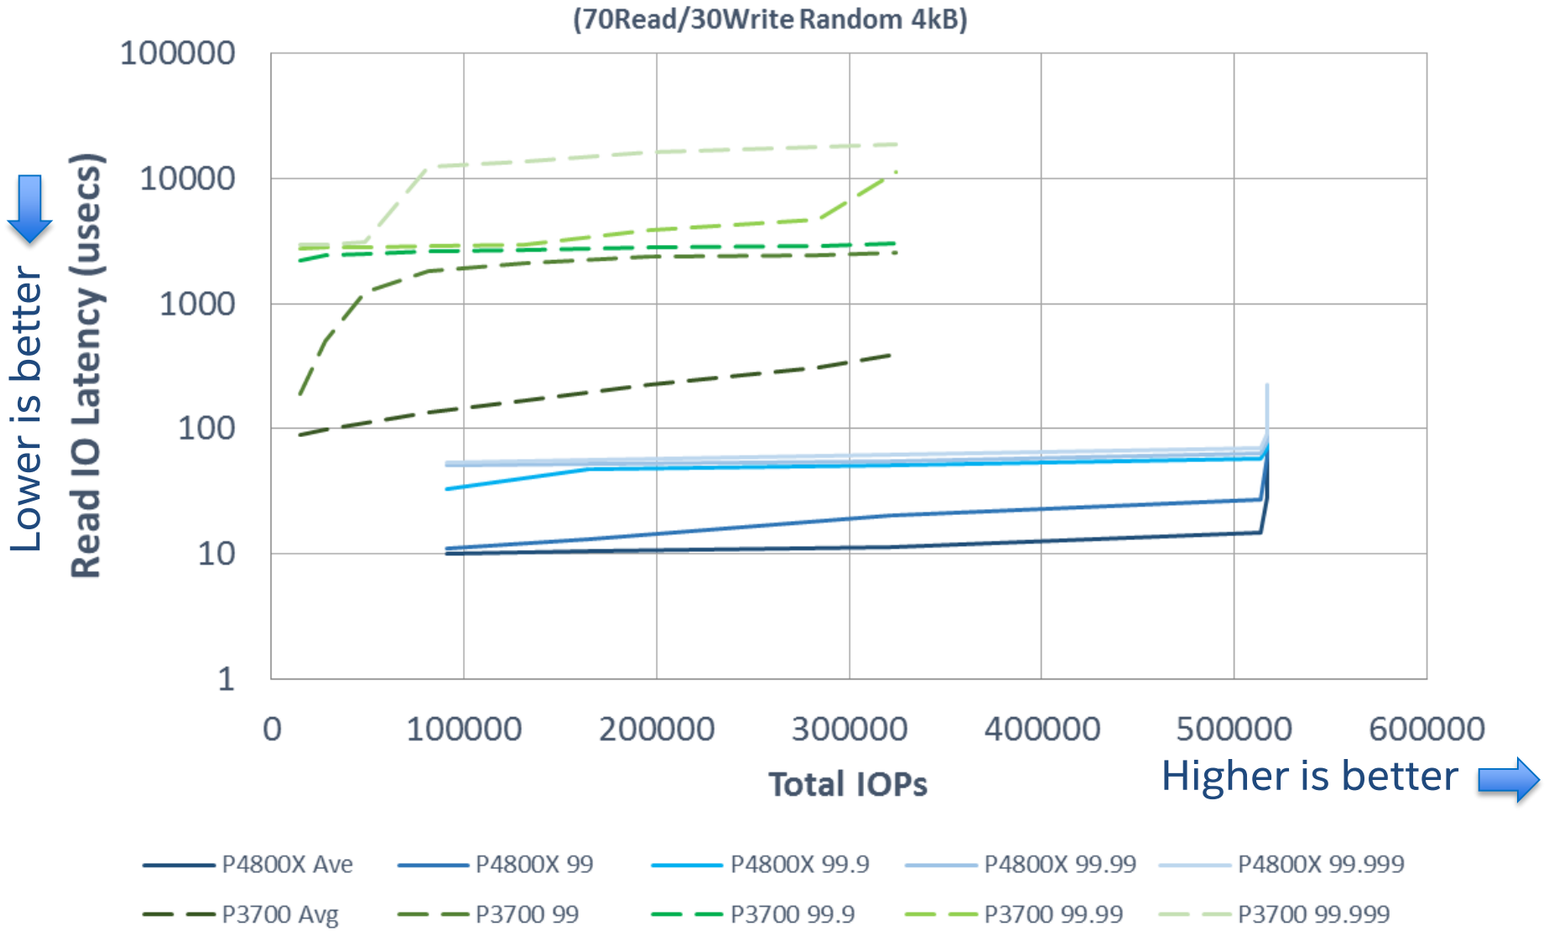

10x latency reduction

- < 10usec latency †

100x QoS improvement

- < 200usec 99.999th r/w †

† vs. NAND based SSD

STORAGE PERFORMANCE CHARACTERIZATION

Latency vs. Load: NAND SSD vs. Intel® Optane™ SSD (Intel® DC P3700 vs. Intel® P4800x)

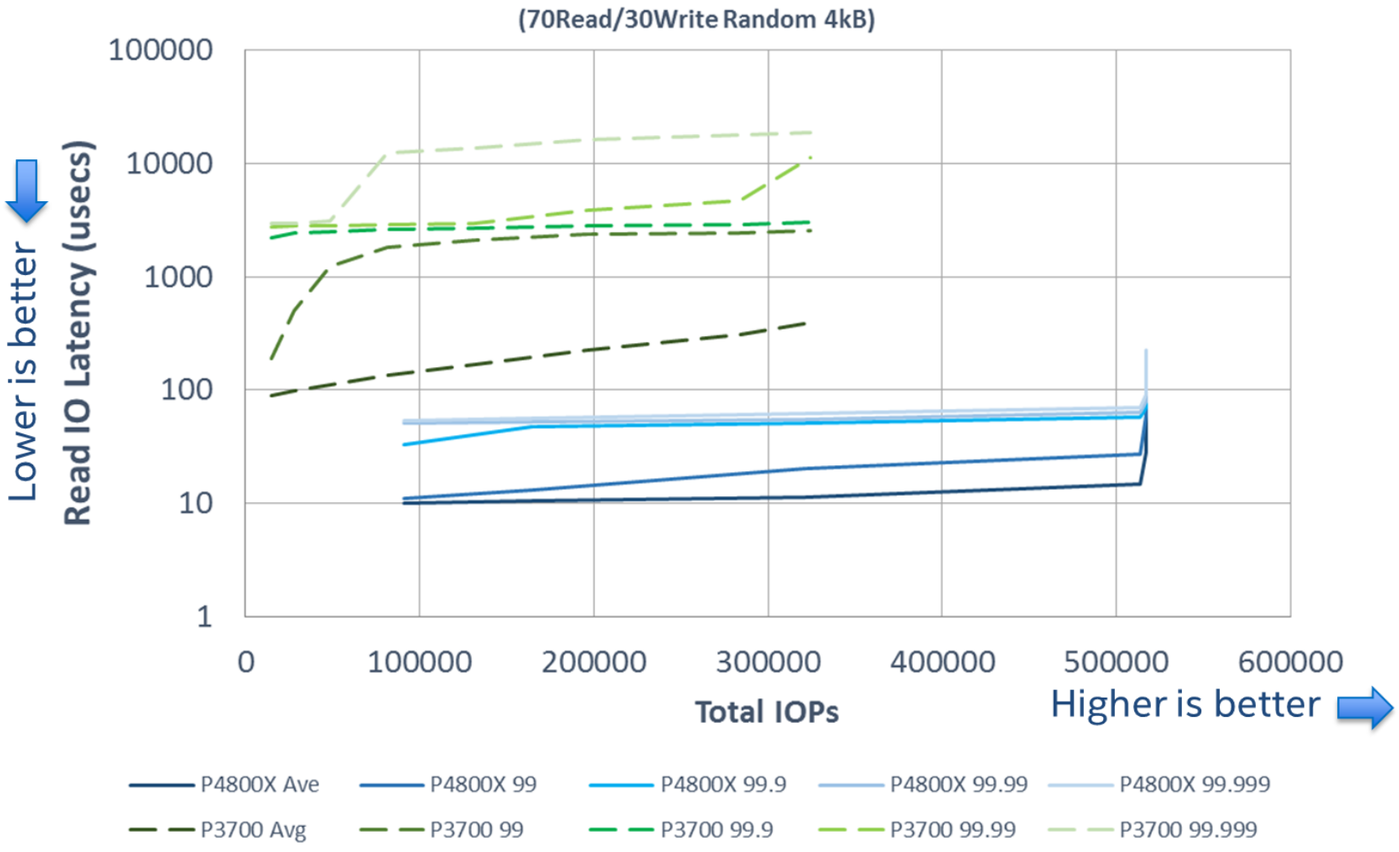

Get work done orders of magnitude faster

- Low latency at high IOPs

Always faster

- Low latency until saturation
- <10x change from ave to 99.999

Easier software development

- Low QD is OK
- Reads and writes the same

Endurance supporting high write rate

WICKED FAST STORAGE - FUNDAMENTALLY DIFFERENT

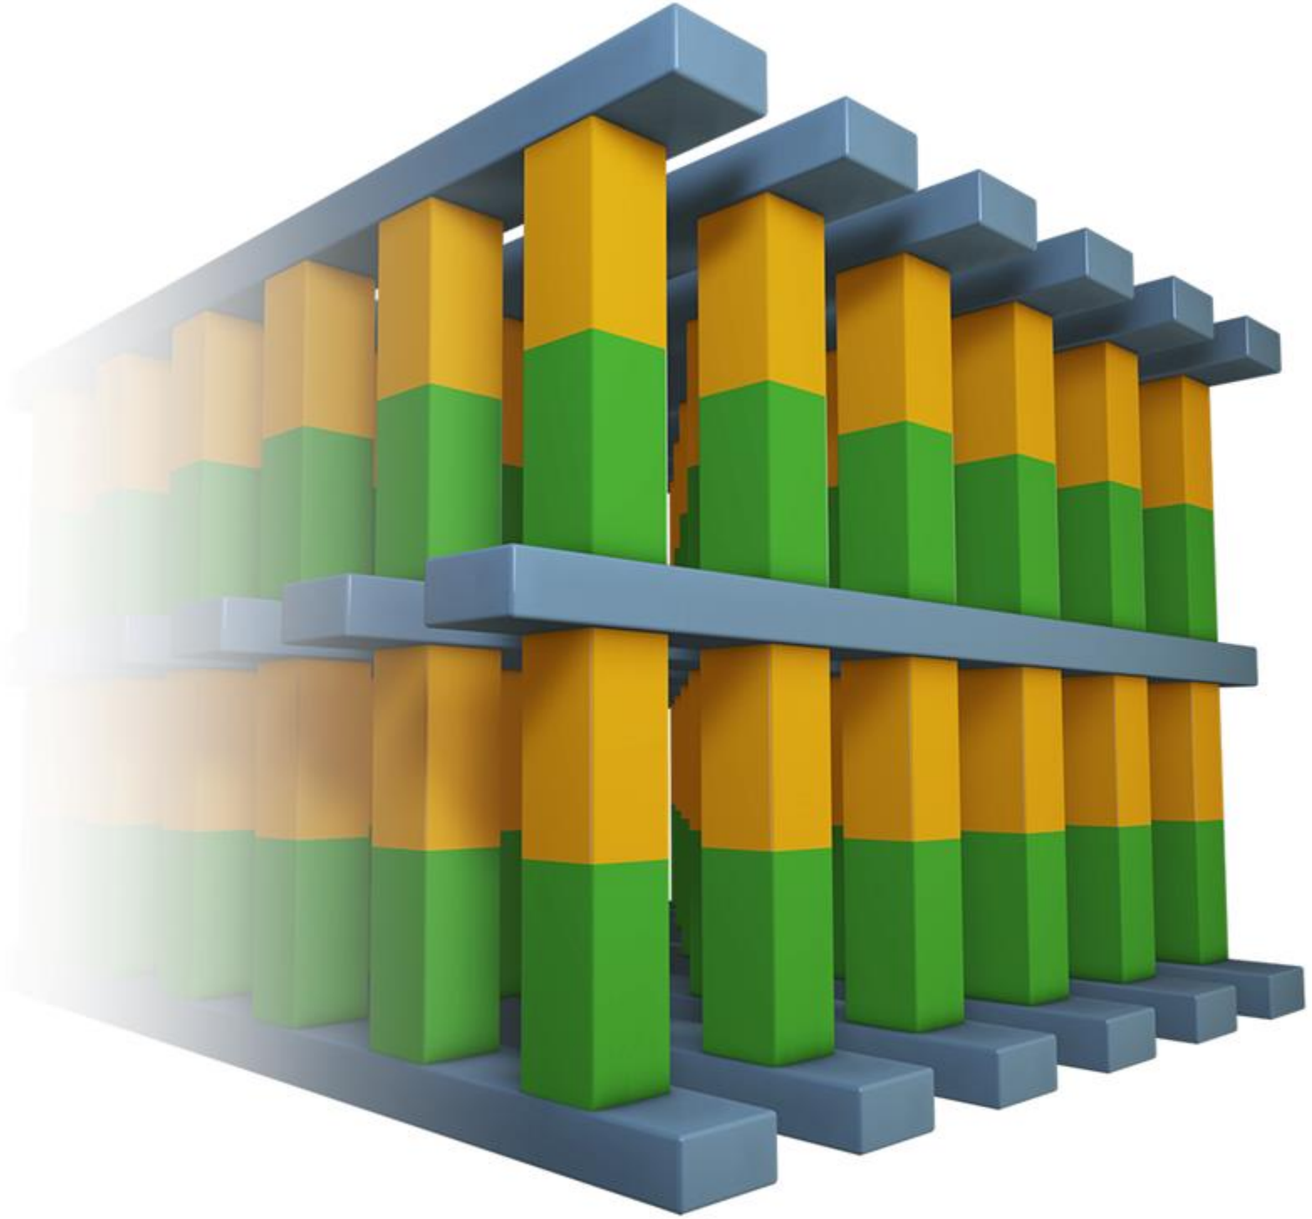

But…….can the OS and applications access this performance?

YES!

INTEL® OPTANE™ SSD LOW LATENCY IS ACCESSIBLE

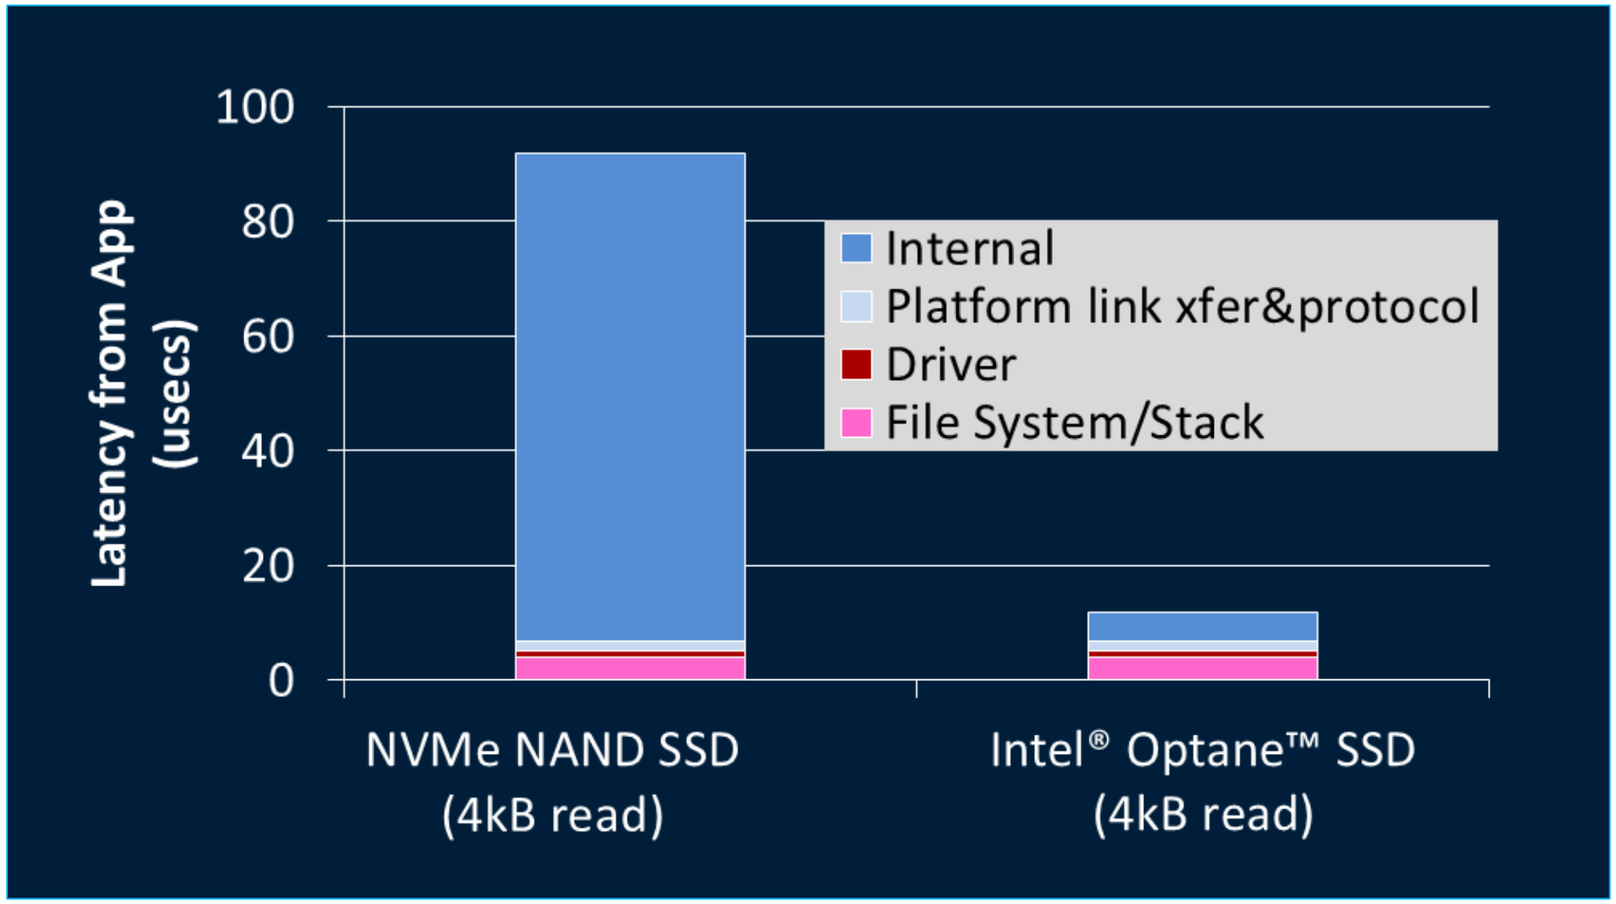

INTEL® OPTANE™ SSD LATENCY IS ACCESSIBLE TO APPLICATIONS

Sources: "Storage as Fast as the rest of the system" 2016 IEEE 8th International Memory Workshop and measurement, Intel® Optane™ SSD measurements, Intel P3700 measurements with FIO as detailed in paper.

PCIe*/NVMe* DELIVERS SUPERIOR LATENCY AND THROUGHPUT

Platform HW/SW Average Latency Excluding Media, 4KB Reads

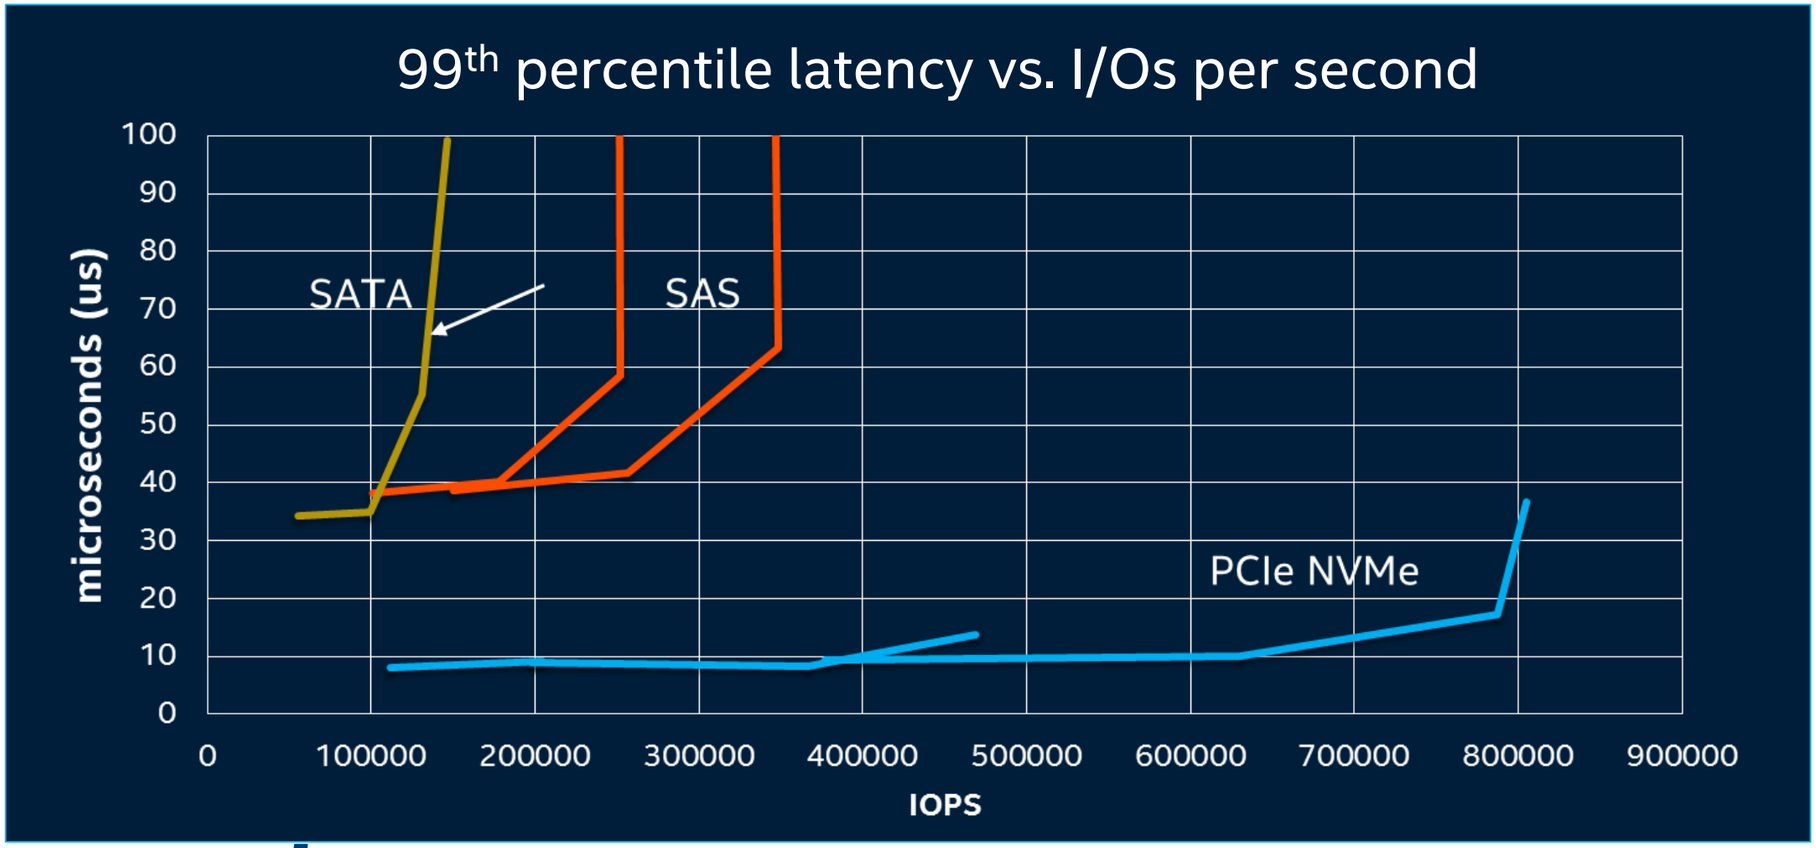

PCIe/NVMe ENABLES HIGH IOPS AT LOW LATENCY

Results measured by Intel based on the following configurations: Wildcat Pass Haswell Server Platform with 28 CPUs, 2 sockets, 2.3 GHz clock speed per CPU, Ubuntu 14.04.1 LTS (GNU/Linux 3.16.0-rc7tickles x86_64), idle=poll kernel settings, SAS HBA is LSI SAS9207-4i4e with controller LSI SAS 2308. SATA SSDs are Wolfsville at XYZ GB. NVMe-based Intel® SSD DC P3700 at 2 TB. Drives tested empty to test interface only (no NVM access.) *Other names and brands may be claimed as the property of others.

MULTIPLE INTEL® OPTANE™ SSD I/O SCALING WITH INTEL® XEON® CORES

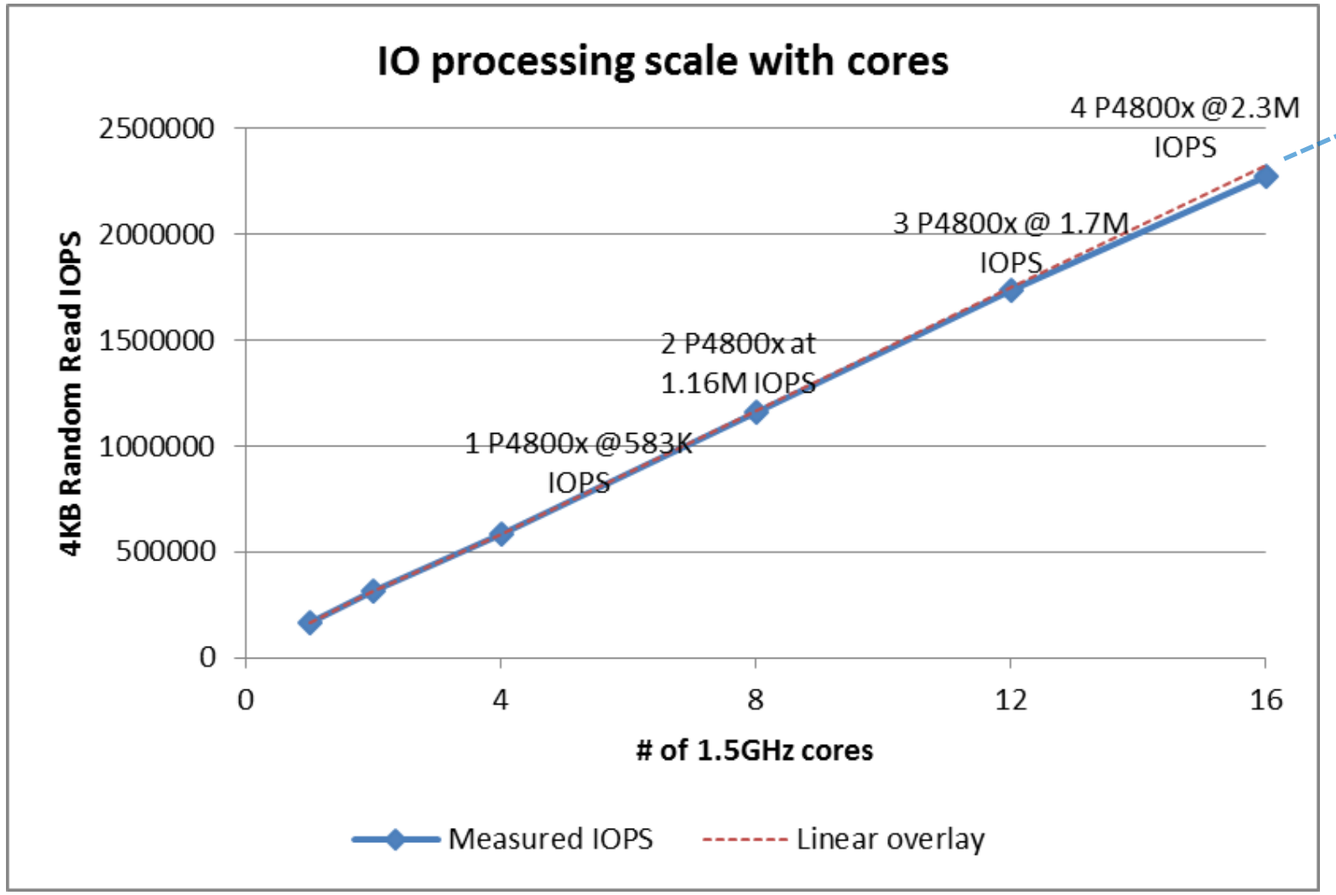

MULTIPLE INTEL® XEON® CORES SCALE IOPS EFFICIENTLY WITH MULTIPLE INTEL® OPTANE™ SSDs

1. Estimates based on Intel internal testing using 2x 16C Intel® Xeon® processor, Linux 4.6.7,256GB DRAM, P4800X 375GB, OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, using fio-2.15. Actual performance depends on system configuration.

PAGING PERFORMANCE

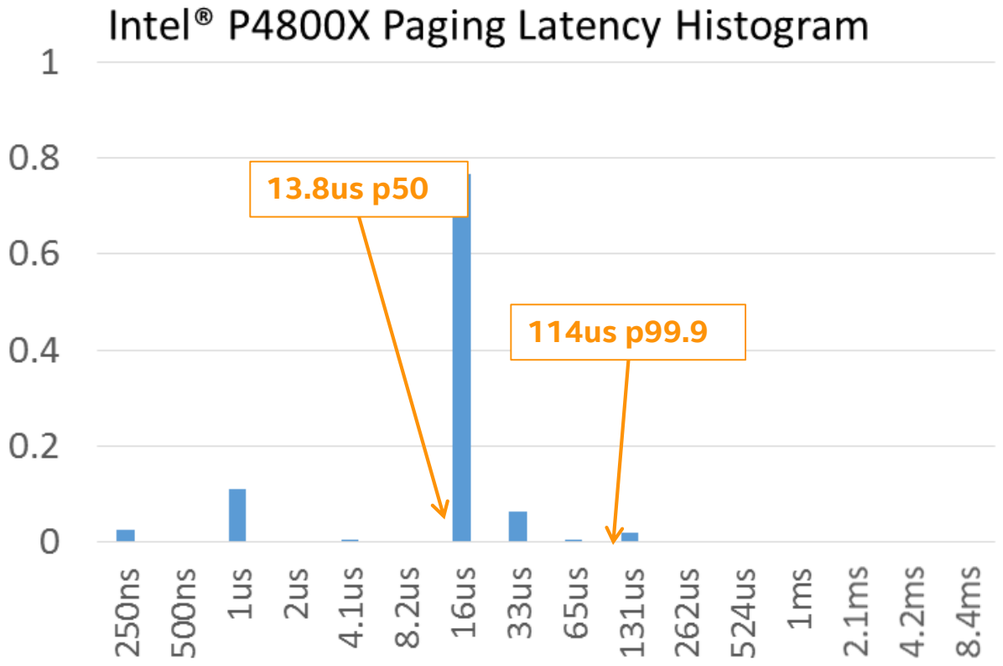

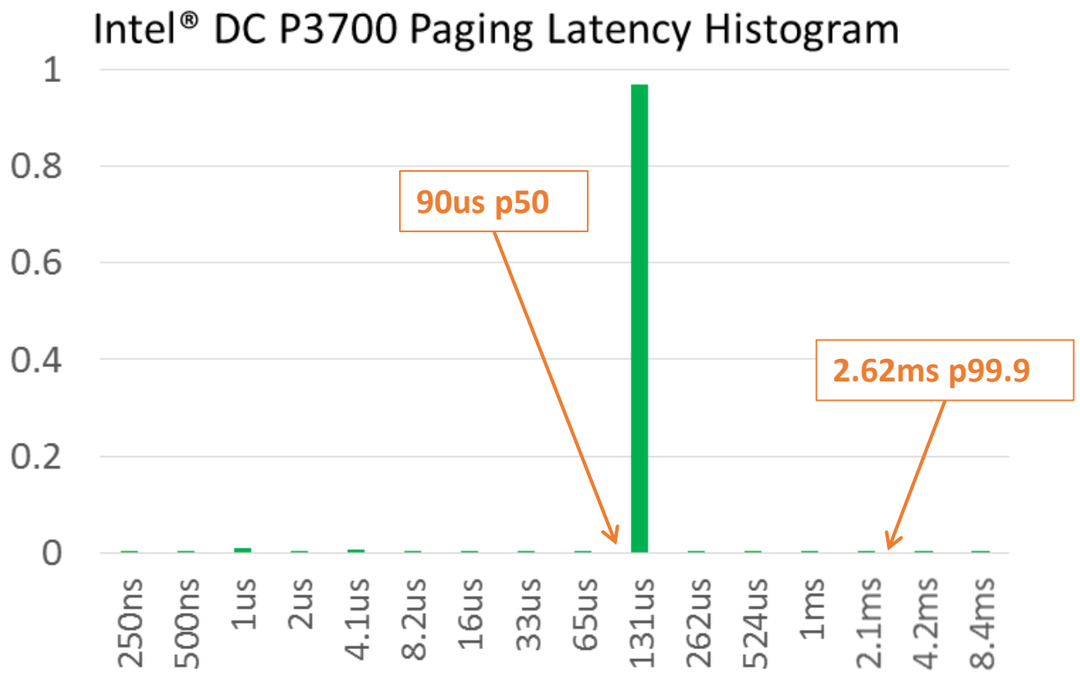

Average paging time = 14 usecs
99.9% paging time = 114usecs

~6x faster

~23x faster

Average paging time = 90 usecs
99.9% paging time = 2.6msecs

OS PAGING WITH INTEL® OPTANE™ SSDs EXTENDS MEMORY CAPACITY

1. Estimates based on Intel internal testing using 2x 16C Intel® Xeon processor, Linux 4.6.7,256GB DRAM, P4800X 375GB, OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, using fio-2.15. Actual performance depends on system configuration.

INTEL® OPTANE™ SSDs ON APPLICATIONS

Extreme performance

- Application waits less, completes faster

SideFX* Houdini*

Predictably fast service

- Improved application responsiveness

Rocks DB*

- Do more in a “click time”

Aerospike*

Fast enough for paging

- Bigger data set accessible

Quantum Sims*

*Other names and brands may be claimed as the property of others.

DATA CENTER: ROCKSDB* PERF ON TEST5 (FROM ROCKSDB.ORG)

Open source persistent key-value store

All threads randomly reads keys, one writer thread updates up to ~80K keys/second

Higher is better

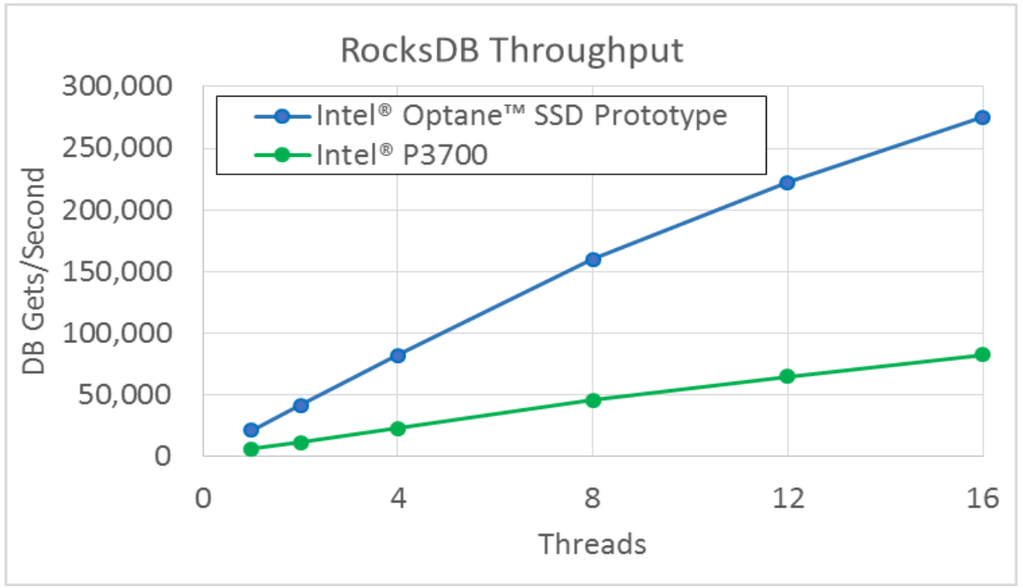

~3x Throughput advantage

Lower is better

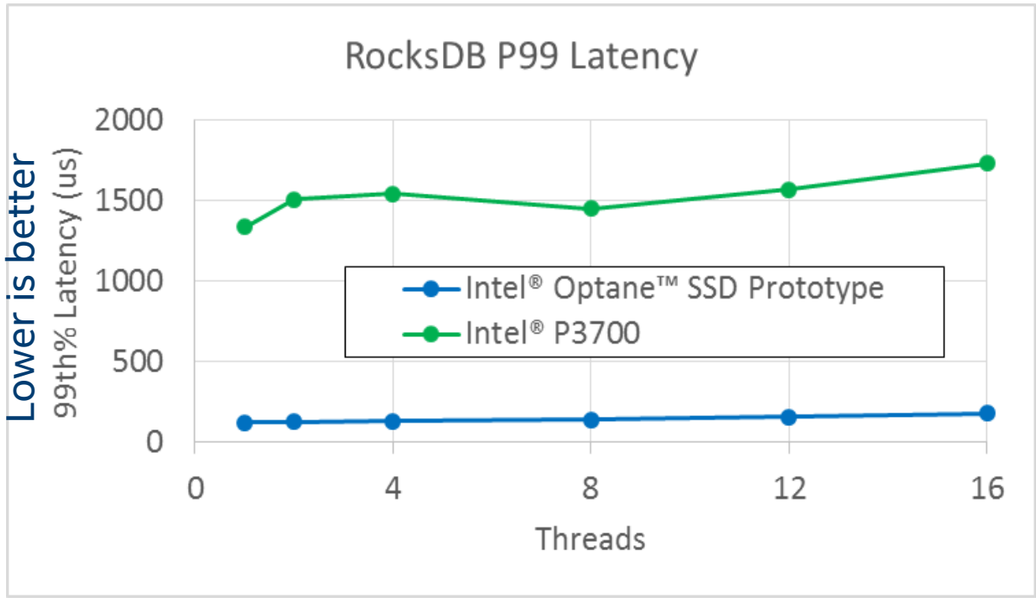

~10x Latency advantage (99th percentile)

INCREASED PERSISTENT KEY-VALUE STORE THROUGHPUT WITH BETTER QoS

RocksdB setup based on published tests at rocksdb.org: 1B Key Database used, 8 “Shards” of 25M Key/Values each, 20 byte keys, 800 byte values, 50% compression, ~100 GB on-disk. Read: All threads randomly read all keys. Read/Write: All threads randomly reads keys 1 writer thread updates up to ~80K keys/second. Quanta Leopard base board, 2x Intel Haswell CPUs (2.5 GHz, 12 core, HT Enabled, 8 DDR4 DIMMs, 256GB, 32GB Used, CentOS* 7.2, no OS changes XFS FS with FB build/mount opts, TRIM enabled, P3700 (50% capacity used) and Intel Optane Based Prototype (75% capacity used).

DATA CENTER: ACT* BENCHMARK

Aerospike* Certification Tool emulates the I/O pattern of a real-time database:

- 1.5kB random reads that meet Service Level Agreement
- 128kB background writes
- Measure multiplier while maintaining SLA

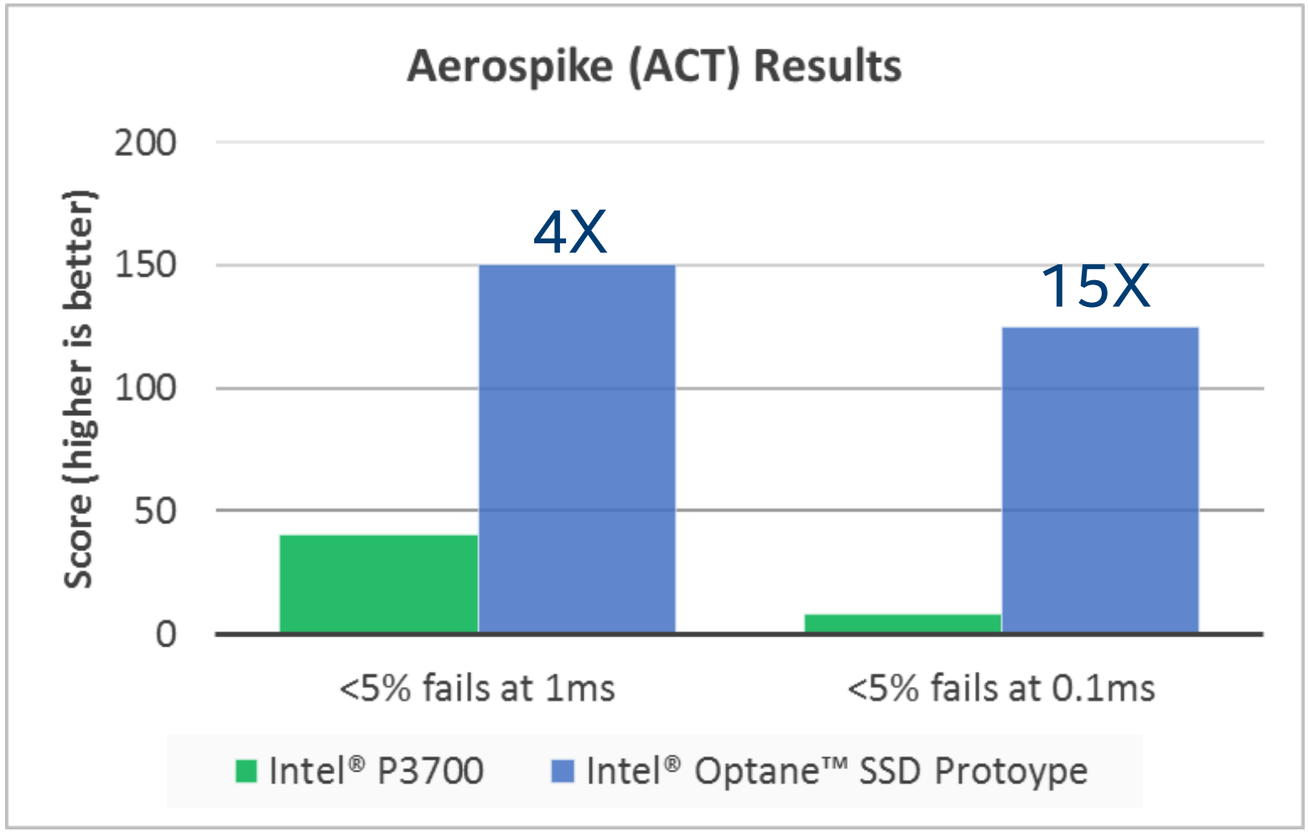

HIGHER REAL-TIME DATABASE THROUGHPUT AT MUCH TIGHTER DEADLINES

CentOS* Linux* release 7.1.1503, Intel® Core™ i7 4770, ASUSTeK COMPUTER INC., H87I-PLUS, Boot Drive: Intel SSD DC S3500 160GB SATA SSD, 4GB DDR3 Single Channel

*Other names and brands may be claimed as the property of others.

INTEL® OPTANE™ SSD USED TO EXTEND MEMORY

Easy scientific computing access in the Cloud

Say you want to experiment with your quantum computer algorithm idea:
quantum Fourier transform 35 Qubits (512GB) x 64 Instances

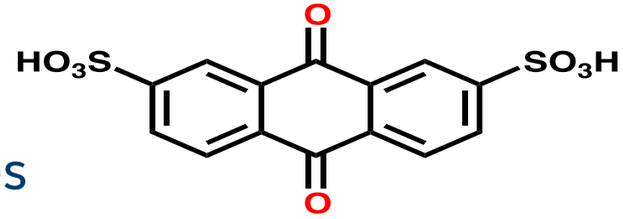

Intel® Corporation newly released Open source Quantum Compute simulator "qHiPSTER"

- TACC Supercomputer → 64 Processors, 16GiB DRAM → 179 min
- 1 Cloud Node → 1 Processor, 16GiB DRAM → ⊘
- 1 Cloud Node (Intel® Optane™ SSD) → 1 Processor, 16GiB DRAM, 896GiB SSD configured as memory → 7360 min → 64 Nodes → 115 min

"Intel® Optane™ SSD is truly a game changer for chemistry, simulating molecules and strongly correlated materials directly in RAM. With resources such as Intel Optane SSD, academic computing and sophisticated scientific jobs can be moved to the cloud."

- Prof. Alán Aspuru-Guzik – Dept of Chemistry and Chemical Biology, Harvard

IN THE FUTURE…

MOVING TOWARDS A PERSISTENCE HIERARCHY

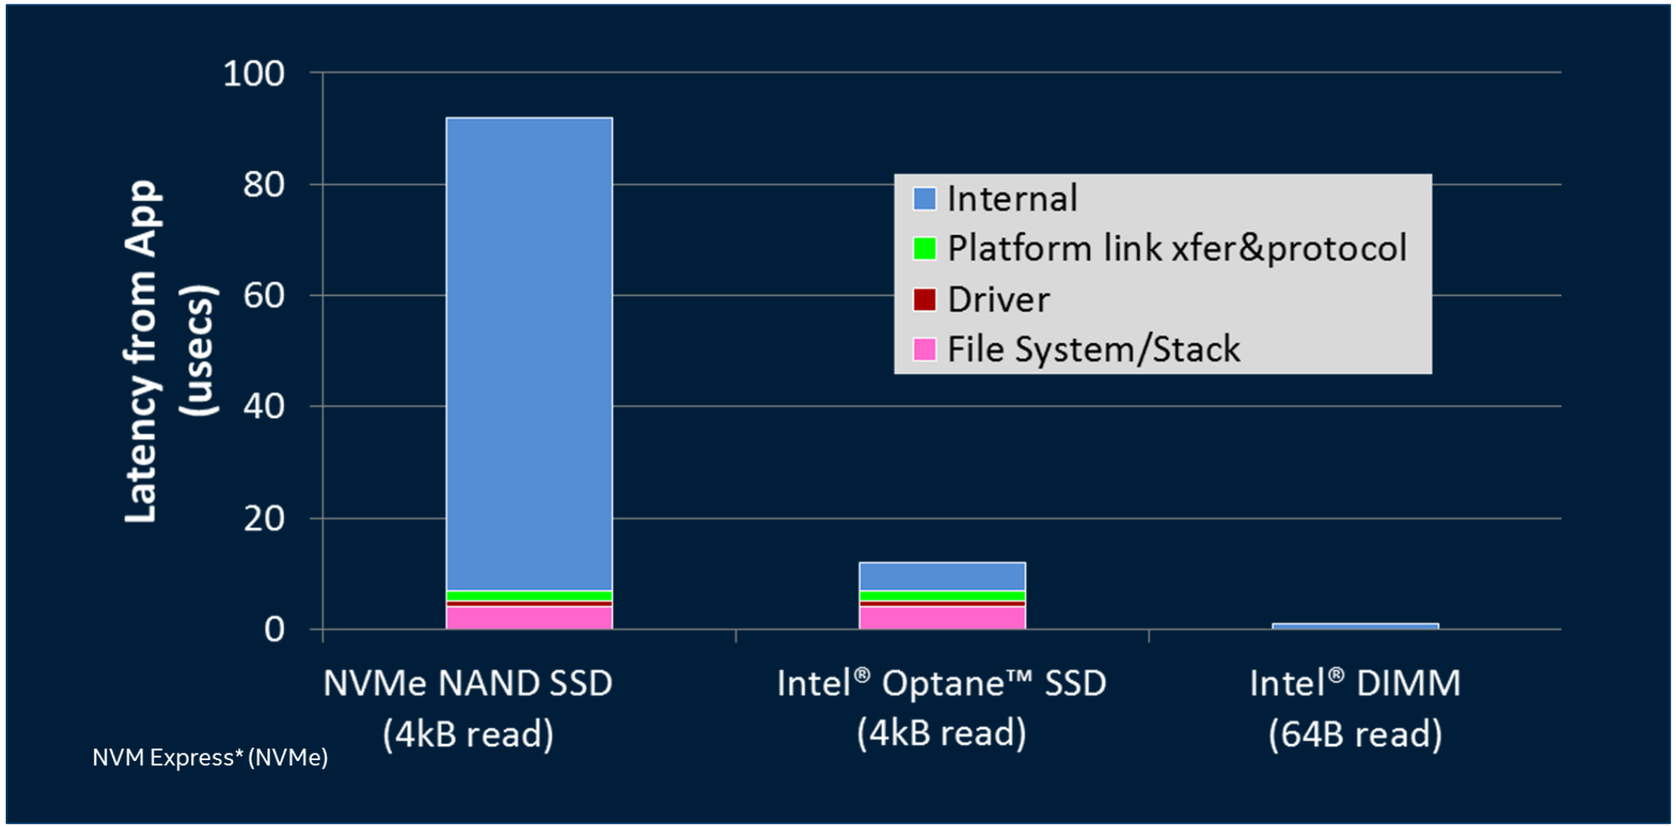

THE THREE DIFFERENT PERFORMANCE/CAPACITY SOLUTIONS WILL WORK TOGETHER IN FUTURE PLATFORMS

Sources: "Storage as Fast as the rest of the system" 2016 IEEE 8th International Memory Workshop and measurement, Intel® Optane™ SSD measurements and Intel P3700 measurements, and technology projections

INTEL® DIMMs BASED ON 3D XPOINT™ MEMORY MEDIA

- **Supported on next generation Intel® Xeon® processor-based platforms**
- **Attaches directly to the processor memory bus**
- **Enables both volatile and persistent memory capabilities**
- **Delivers up to 4x the capacity of today's DRAM DIMMs**

Intel® DIMM

(based on Intel® 3D XPoint™ memory media)

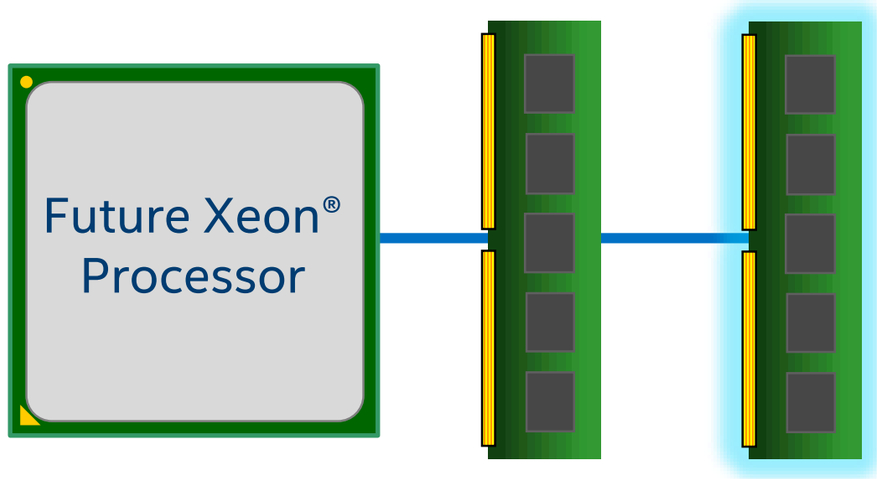

DIMM

(acts in conjunction with Intel DIMM)

INTEL® OPTANE™ TECHNOLOGY IS HERE

- **Intel® Optane™ SSD: 1st step in journey**
  - –3D XPoint™ memory media is here
  - –Purpose-built, all new
  - –Revolutionary latency and consistency
  - –So fast we measure it and use it like DRAM
- **Applications will run faster consistently**
  - –Responsive under any load

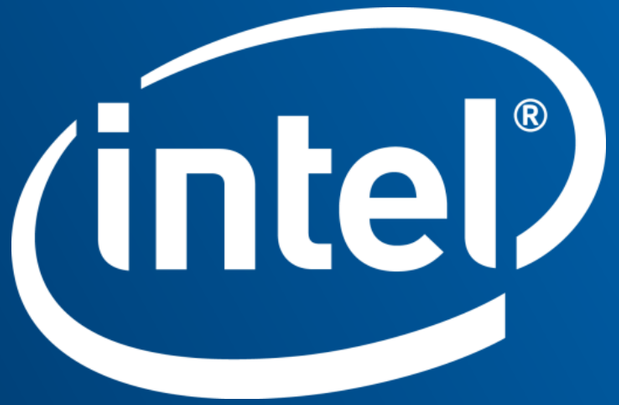
intel