G2M
RESEARCH

G2M Research Multi-Vendor Webinar: Advanced SSDs – PCIe Gen4, U.3, and New Form Factors

Tuesday August 25, 2020 (v1.0)

KIOXIA

intel

Webinar Agenda

9:00-9:05	Ground Rules and Webinar Topic Introduction (G2M Research)
9:06-9:35	Sponsoring Vendor presentations on topic (15 minute each)
9:36-9:36	Audience Survey 1 (1 minutes)
9:37-9:43	Key Question 1 (1-minute question; 3 minutes response per vendor)
9:44-9:44	Audience Survey 2 (2 minutes)
9:45-9:51	Key Question 2 (1-minute question; 3 minutes response per vendor)
9:52-9:52	Audience Survey 3 (2 minutes)
9:53-9:59	Key Question 3 (1-minute question; 3 minutes response per vendor)
10:00-10:09	Audience Q&A (10 minutes)
10:10-10:10	Wrap-Up

G2M RESEARCH

G2M Research Introduction and Ground Rules

Mike Heumann

Managing Partner, G2M Research

Panelists

KIOXIA

Matt Hallberg
Sr. Product Marketing Manager
www.kioxia.com

intel

Jonmichael Hands
Sr. Strategic Planner
www.intel.com

G2M RESEARCH

Mike Heumann
Managing Partner
www.g2minc.com

SSDs Get Faster, Smaller, and Smarter

- PCIe Gen4 doubles the max BW of SSDs
  - From 3.94GB/s to 7.88GB/s
  - Also reduces minimum latency by ~50%
- U.3 promises a standard interface between NVMe SSDs, SAS/SATA SSDs, and HDDs
- The new EDSFF form factor enables smaller and denser SSDs, with higher capacity
  - 32 drives in a 1U server chassis
  - 32TB per drive

PCI EXPRESS®

U.3 = NVMe + SAS/SATA + HDDs (?)

G2M RESEARCH

What Are The Impacts of This on Applications?

- How much faster do we need?
  - You can never be fast enough for Big Data or HPC…
  - Latency reductions help SSDs keep up with/reduce the gap between non-volatile storage and DRAM
- How does compatibility across media types simplify application deployment?
  - Do enterprise datacenter users really mix different flash media in the same servers?
- How much smaller/denser do we need?
  - A PB in 1U sounds awfully attractive IF you can get all of the data in and out fast enough..
  - SSD blast radius remains a concern

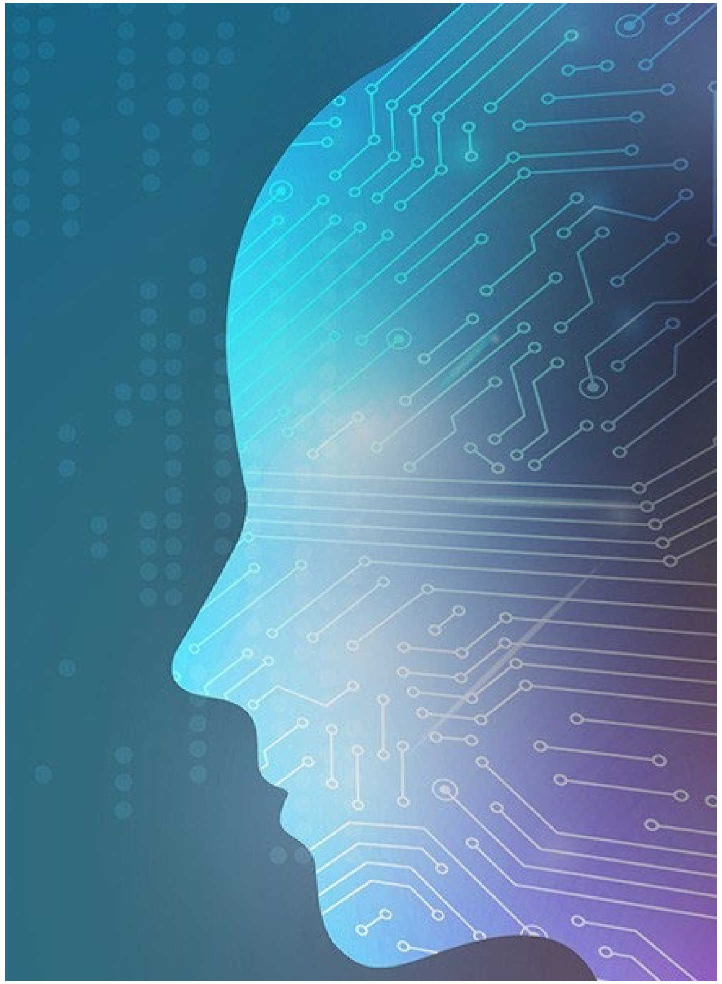

G2M RESEARCH

Kioxia

- Matt Hallberg
  Sr. Product Marketing Manager
  www.kioxia.com

Why PCIe Gen 4.0 for your Datacenter?

- Infrastructure
- Performance
- Disaggregation

How do PCIe Gen 4.0 SSDs fit into this?

Why PCIe Gen 4.0: Infrastructure

- PCIe Gen 4.0 Chipsets and systems bring more lanes to the table
  - More lanes → more devices→ higher system versatility
  - Add more storage, more NICs, more HBAs, more GPUs, etc.
- PCIe Gen 4.0 devices can take advantage of faster data pipes
  - GPUs are first to take advantage like the nVidia A100
  - Dual 200GbE NICs
  - HBAs
  - KIOXIA (and other manufacturer) SSDs
- Pick your preference: More storage or more performance? Both?
  - Can attach PCIe Gen 4.0 SSDs at 1x2 to PCIe Gen 4.0 Switch / HBA and get same/better performance than a Gen 3.0 deployment with twice the devices or half of the lanes
  - Can attach the same amount of PCIe Gen 4.0 SSDs and get twice the performance

Why PCIe Gen 4: Performance

- GPUs – nVidia announces A100, great improvements over the T100, uses PCIe Gen 4.0 1x16
- NICs – availability of dual port 200GbE PCIe Gen 4.0 NICs
  - Faster interconnect between systems, faster storage network
- HBAs – provide more bandwidth for CPU and SSDs
- SSDs – Faster performance, reduced latency
  - Example: Microsoft / KIOXIA / Broadcom / QCT were able to demonstrate double the performance on an Azure Stack HCI system running at PCIe Gen 4.0 versus a similarly configured PCIe Gen 3.0 system

G2M RESEARCH

Why PCIe Gen 4: Disaggregation

- Disaggregating the datacenter has been a long time vision
  - With the performance and latency improvements offered within PCIe Gen 4.0 infrastructure, composable disaggregation is becoming a reachable reality

Today's converged infrastructure model

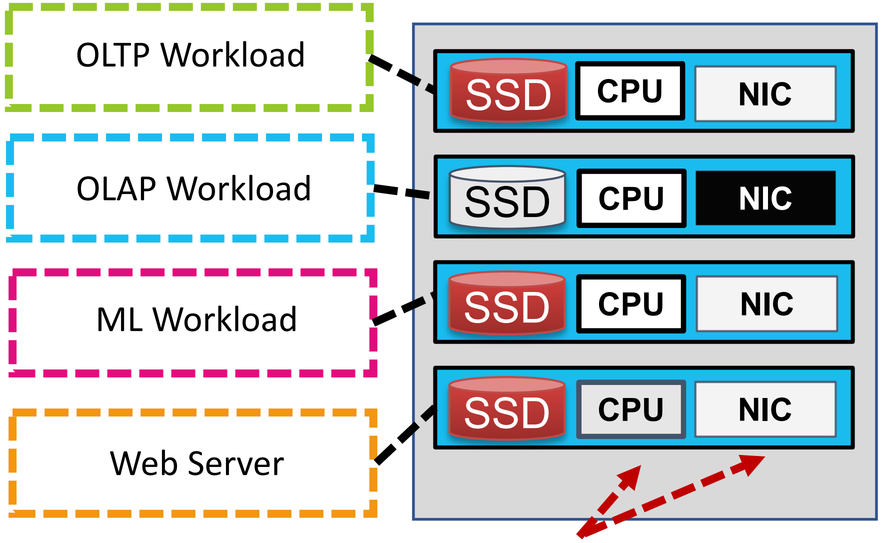

Stranded underutilized resources

Tomorrow's Composable Disaggregated Infrastructure

PCIe Gen 4.0 enables high speed connectivity between compute, storage, network, and GPU resources

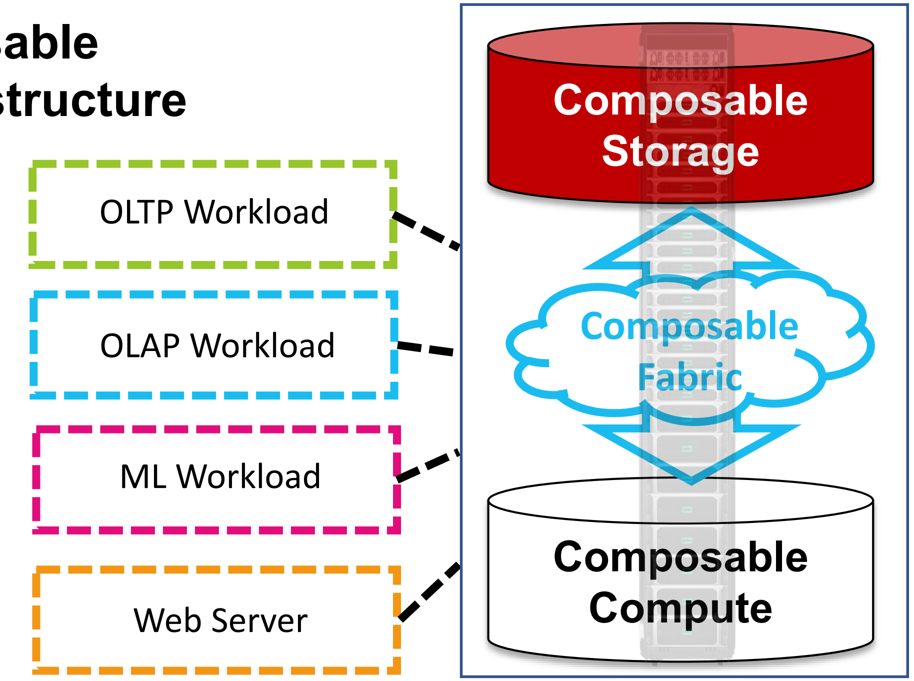

Efficient resource allocation per workload

Why PCIe Gen 4.0 SSDs?

- Compute Nodes
  - GPUs require large amounts of DRAM and very fast local storage to minimize idle time
    - KIOXIA's PCIe 4.0 Enterprise CM6 and Datacenter CD6 NVMe SSDs are **FAST!**
- Storage Nodes and General Purpose Servers
  - Faster reads / writes
    - KIOXIA's Enterprise CM6 and Datacenter CD6 NVMe SSDs can reach a million+ IOPS
  - Improved response times
  - Density
    - Up to 30TB 2.5" SSDs

KIOXIA’s Enterprise PCIe Gen 4.0 SSD: CM6

- Enterprise PCIe® 4.0, NVMe™ 1.4 SSDs
- Form factors: 2.5-inch, 15mm Z-height
- Proprietary KIOXIA architecture: controller, firmware and BiCS FLASH™ 96-layer 3D TLC memory
- SFF-TA-1001 conformant (U.3) works with Tri-mode controllers and backplanes
- Dual-port design for high availability applications
- 6th generation die failure recovery and double parity protection
- High performance with lower power consumption
- Power loss protection (PLP) and end-to-end data protection
- Suited for 24x7 enterprise workloads
- Data security options: SIE, SED (TCG Opal / Ruby), FIPS 140-2 (Planned)
- Five power mode settings

			CM6 (Mixed-Use)					CM6 (Read-Intensive)					
Endurance		DWPD	3					1					
User Capacity*		GB	800	1600	3200	6400	12800	960	1920	3840	7680	15360	30720
Sequential Read	128KB(QD32)	MB/s	6900	6900	6900	6900	6900	6900	6900	6900	6900	6900	TBD
Sequential Write	128KB(QD32)	MB/s	1400	2800	4200	4000	4000	1400	2800	4200	4000	4000	TBD
Random Read	4KB(QD256)	KIOPS	800	1300	1400	1400	1400	800	1200	1400	1300	1400	TBD
Random Write	4KB(QD32)	KIOPS	100	215	350	325	330	50	100	170	170	170	TBD

* KIOXIA Corporation definition of capacity: 1 GB = 1,000,000,000 (10^9) bytes (see end of presentation for full capacity disclaimer).

KIOXIA’s Datacenter PCIe Gen 4.0 SSD: CD6

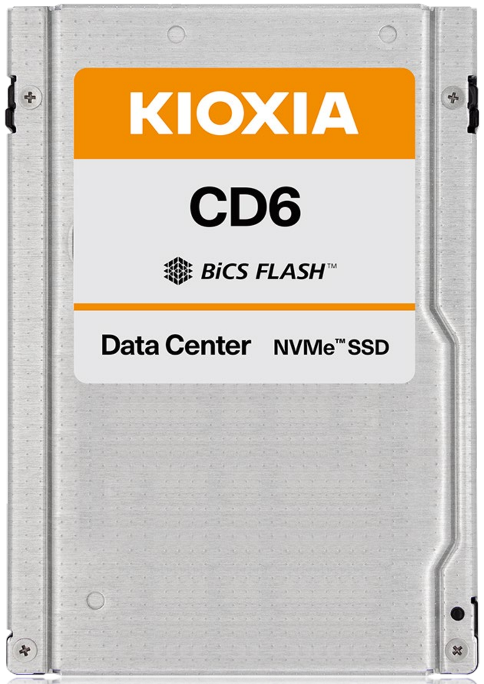

- Data Center PCIe® 4.0, NVMe™ 1.4 SSDs
- Form factors: 2.5-inch, 15mm Z-height
- Proprietary KIOXIA architecture: controller, firmware and BiCS FLASH™ 96-layer 3D TLC memory
- SFF-TA-1001 conformant (U.3) works with Tri-mode controllers and backplanes
- Single-port design, optimized for data center class workloads
- 6th generation die failure recovery and double parity protection
- Consistent performance and reliability in demanding 24x7 environments
- Designed for high density storage deployments
- Power loss protection (PLP) and end-to-end data correction
- Data security options: SIE, SED, FIPS 140-2
- Five power mode settings

			CD6 (Mixed-Use)					CD6 (Read-Intensive)				
Endurance		DWPD	3					1				
User Capacity*		GB	800	1600	3200	6400	12800	960	1920	3840	7680	15360
Sequential Read	128KB(QD32)	MB/s	5800	5800	6200	6200	5500	5800	5800	6200	6200	5500
Sequential Write	128KB(QD32)	MB/s	1300	1150	2350	4000	4000	1300	1150	2350	4000	4000
Random Read	4KB(QD256)	KIOPS	700	700	1000	1000	750	700	700	1000	1000	750
Random Write	4KB(QD32)	KIOPS	90	85	160	250	110	30	30	60	85	30

Advanced SSD: What is U.3?

- What: Universal backplane enables the ability to plug in any SAS, SATA or NVMe™ drive
- Why: Allows deployment of an NVMe SSD into any U.3-enabled socket with minimal investment
- How: Switchable architecture within the SSD for U.2 and U.3 connection types
- Vision: Storage tiering within the same server for SMB IT deployments
  - HDD for bulk storage
  - SATA/SAS for Warm Data
  - NVMe for high performance applications

G2M RESEARCH

Form Factors: What’s coming in the future?

EDSFF 3”: Designed for Server and Storage

- Higher Density Deployments
  - 1X or “Thin” SSDs are 7.5mm thick (2.5” is 15mm)
  - Can fit twice as many drives in a 1U or 2U system
    - (44) drives in a 2U system is more than 660TB!
- More Performance
  - Takes advantage of wider PCIe lane configurations
    - Connectors for 1x4, 1x8, 1x16 (and dual port!)
  - Allows for up to 40W power use (E3L.2X)
    - More power = more performance
- Unified Connectors for more flexibility
  - SFF-TA-1002 can be used for SSDs, NICs, GPUs, etc
- Better System Airflow / Cooling
  - Can bring critical devices to front of box for cool air
  - Flexible drive height and system configurations

Intel

- Jonmichael Hands
  Sr. Strategic Planner, SSDs
  www.intel.com

Intel EDSFF products in production today

Intel® SSD DC P4510, TLC NVMe
15.36TB, E1.L 9.5 & 18mm @ 25W

Intel® SSD D5-P4326, QLC NVMe
15.36TB, E1.L 9.5 & 18mm @ 25W

Intel® SSD DC P4511, TLC NVMe
4TB, E1.S @ 12.5W

E1.l. storage reimagined.

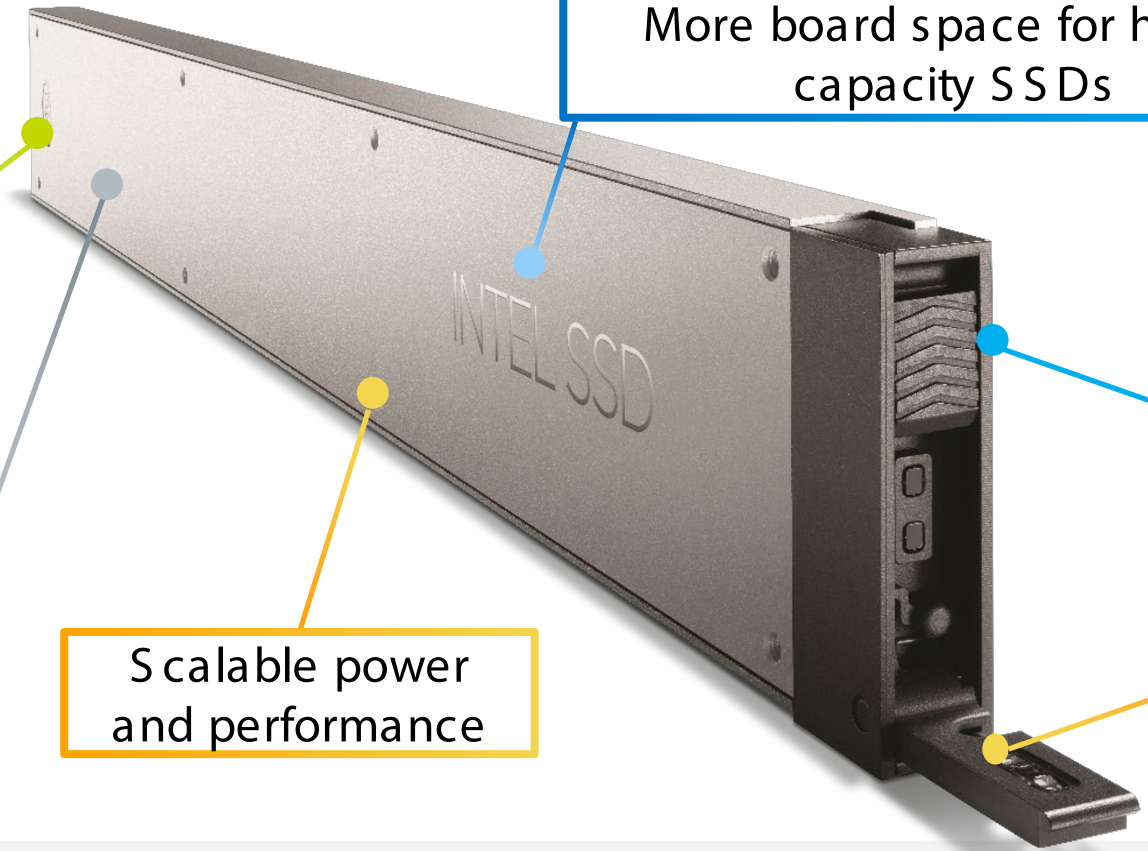

More board space for higher capacity SSDs

Optimized PCIe routing allows for better signal integrity and lower cost attach per drive

Thermally efficient design with enclosure as heatsink

Designed to fit in 1U Server

Scalable power and performance

Carrier-less design with integrated latch and LEDs

Scalable, thermal efficient, and dense, E1.L is a building block for high-volume storage. E1.L allows increased storage density, scaling, improved serviceability, and more efficient cooling optimized for 1U servers.

E1.L Systems Available Today

SuperStorage SSG-1029P-NEL32R

(Angled View – System)

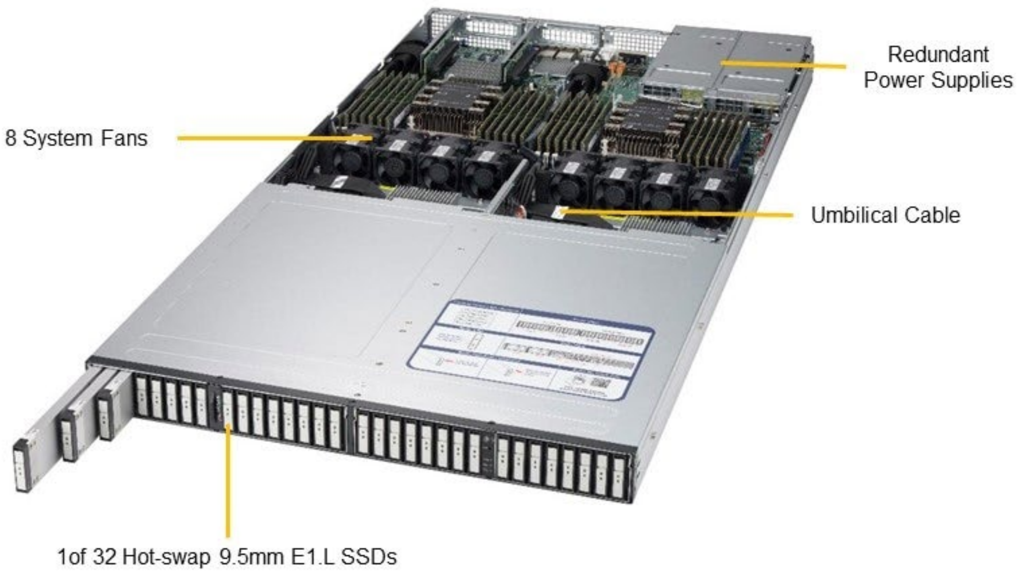

E1.L optimized for TB/rack unit & performance

PCIe JBOF - E1.L optimized for lowest TCO on QLC NVMe

Example - Dense NVMe U.2 Server

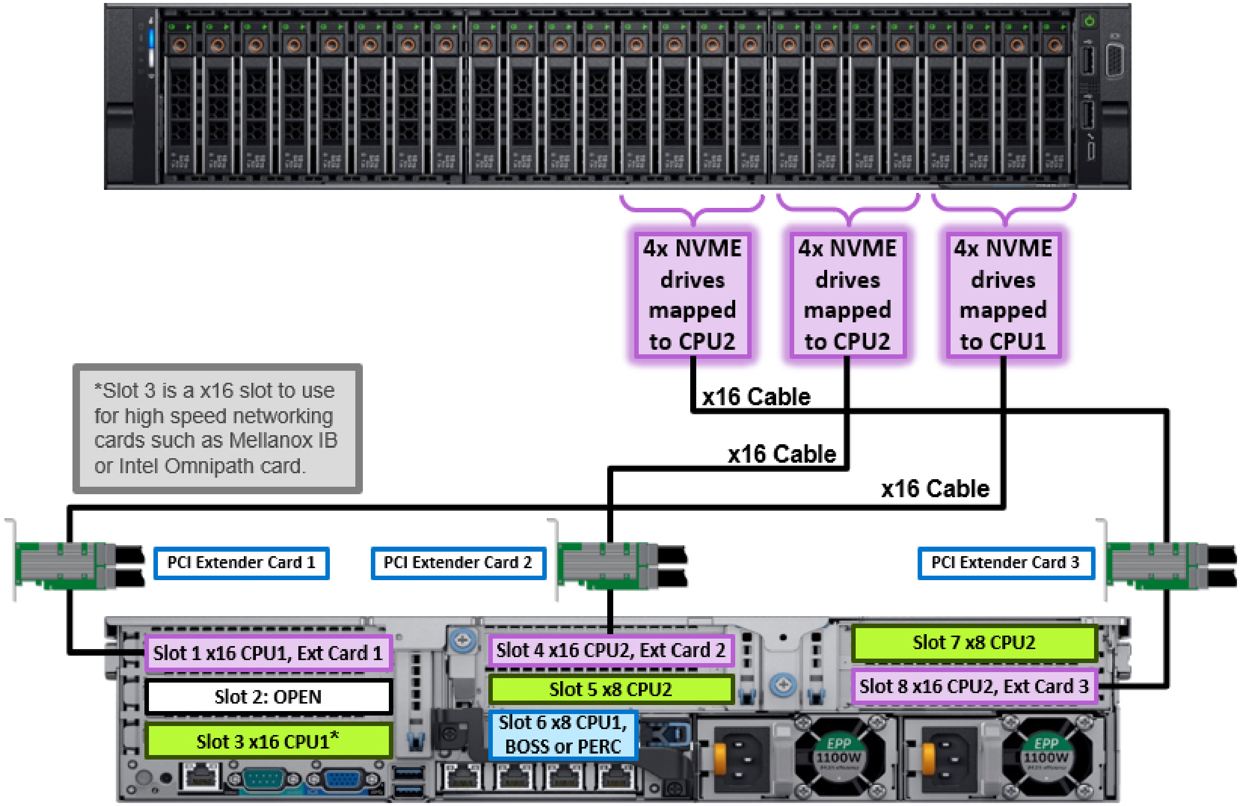

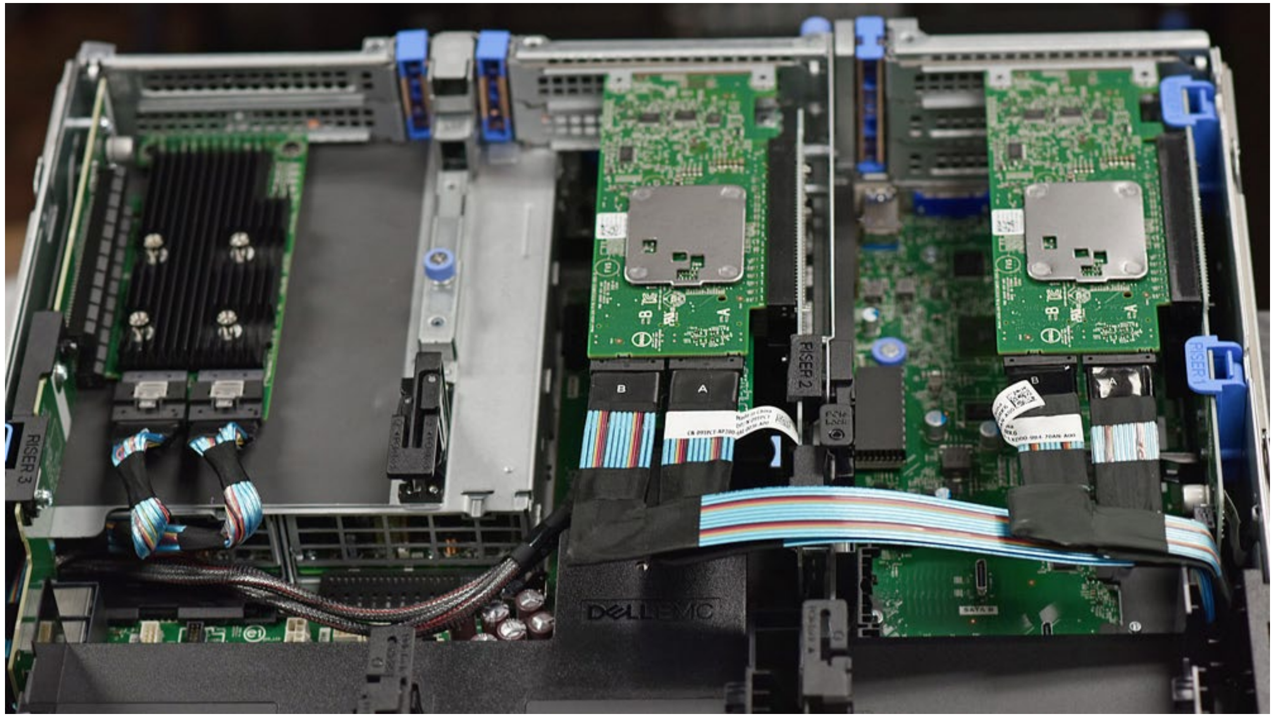

Intel E1.S Products Optimized for 1U Server

Intel® SSD DC P4511, 4TB, E1.S
5.9mm @ 12.5W
Production in Q3'2019

E1.S 15mm @ 15-20W. U.2 performance in smaller form factor
PCIe 4.0…coming soon

EDSFF Platforms

FACEBOOK

E1.S 25mm

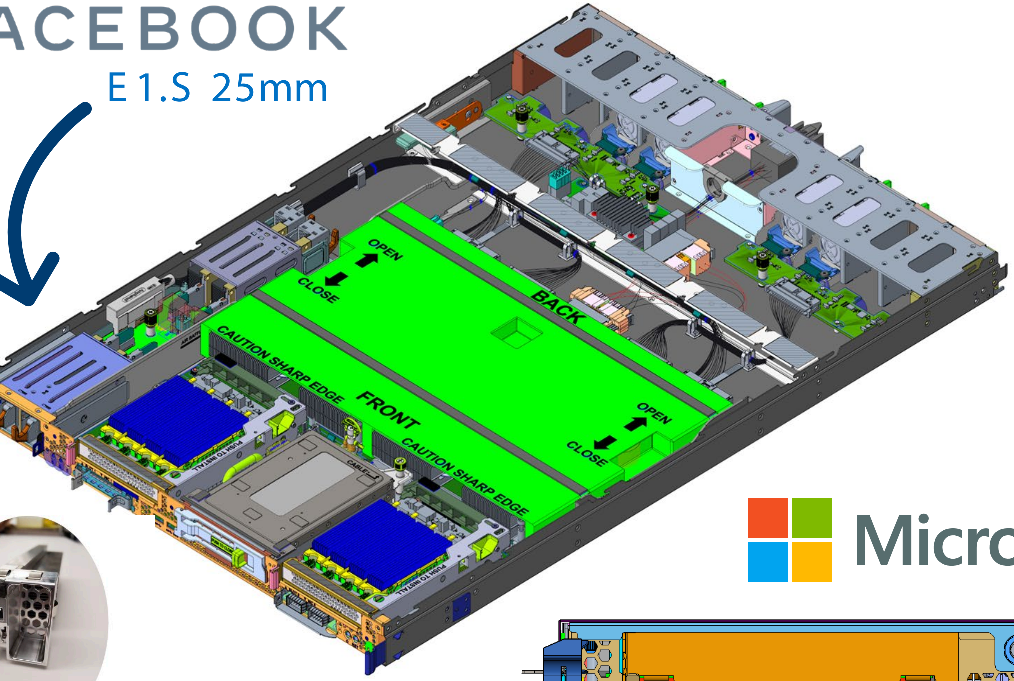

Yosemite V3

25mm E1.S Latch

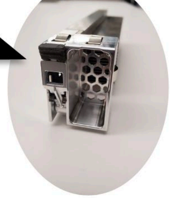

Blade

Chassis

Mainstream compute

Microsoft Concept Project Olympus Design

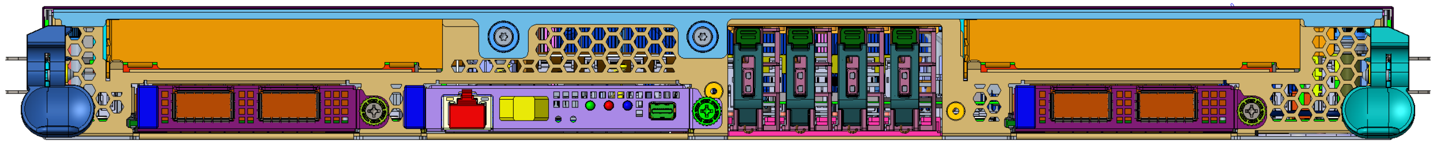

Dense Storage

E1.L 18mm

E1.L 9.5mm

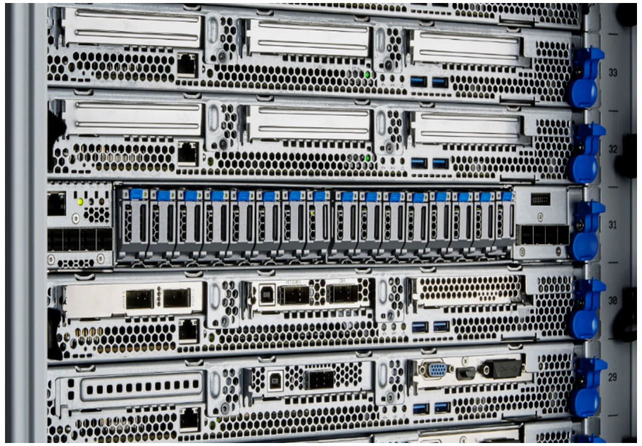

E1.S Optimal for 1U Performance Scalability

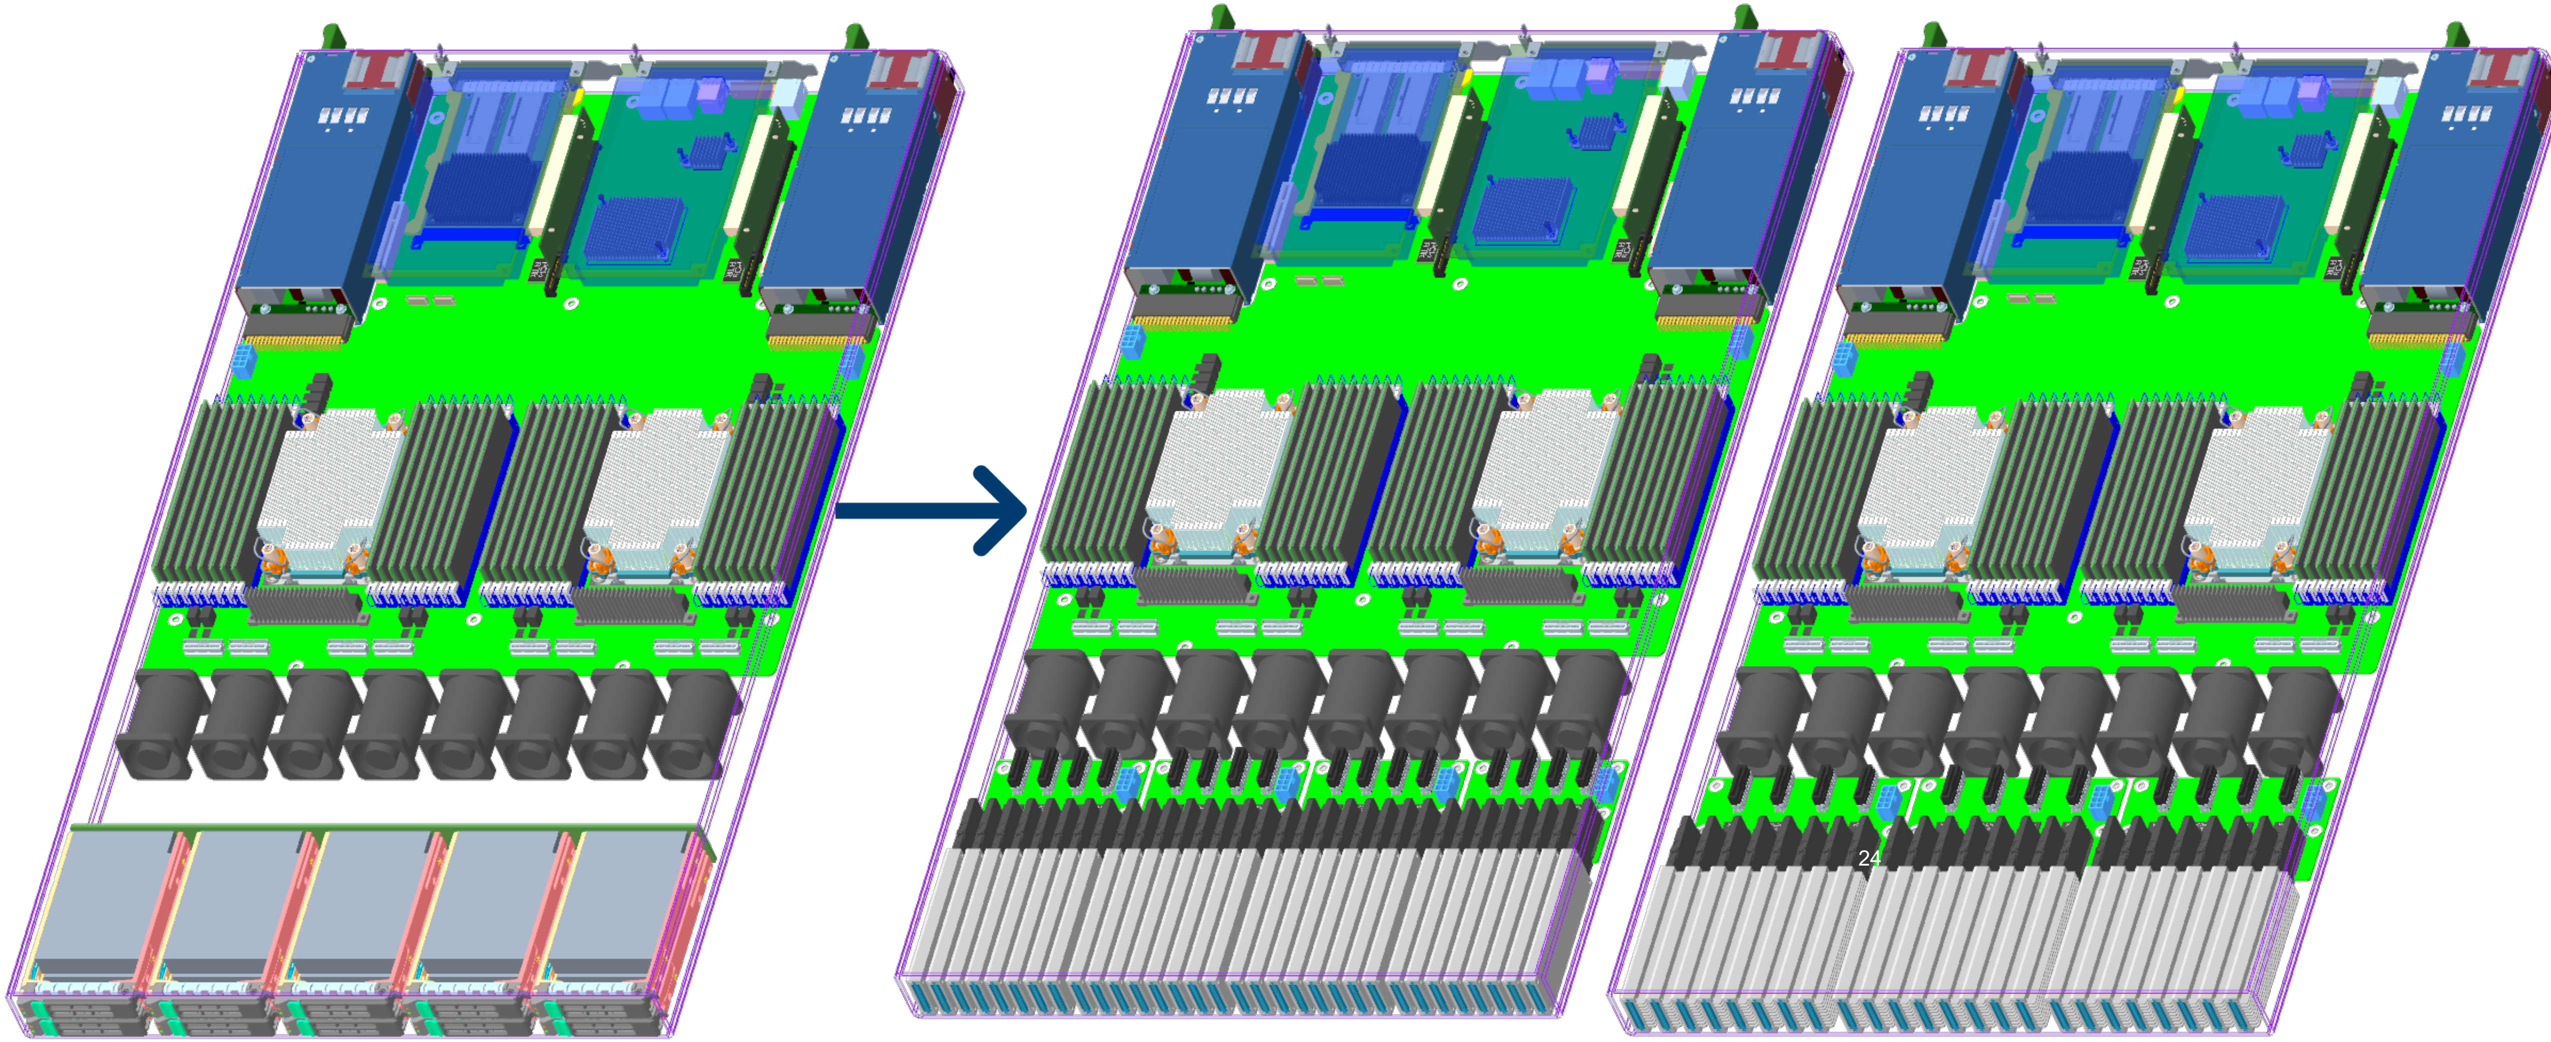

10 x U.2 SSD
Baseline

32 x E1.S 9.5 mm (Horizontal BP)
Max Fit

24 x E1.S 15 mm (Horizontal BP)
Max Fit

Intel® SSD D7-P5500/D7-P5600 Key Specifications

Predictably fast, high performance [10]

	D7-P5500	D7-P5600
4K Rand. Read/ Write	1M/130K IOPS	1M/260K IOPS
128K Seq. Read/Write	7000/4100 MB/s	7000/4100 MB/s
Endurance	1 DWPD	3 DWPD
4K Rand. Read/ Write Latency	78/15µs	78/14µs

NVMe Drive Serviceability

Hot plug support	Yes
In-band monitoring	NVMe SMART/health
Out-of-band monitoring	NVMe-MI

Range of Capacities

D7-P5500	U.2 15mm: 1.92TB, 3.84TB, 7.68TB
D7-P5600	U.2 15mm: 1.6TB, 3.2TB, 6.4TB

Data Integrity and Availability

Projected Annualized Failure Rate	0.44% [11]
End-to-End data protection Features	Industry-leading protection from silent data corruption [12]
Power Loss Imminent	Yes [13] + daily self test

See appendix for footnotes 10-13.

Intel® SSD D7-P5500 & D7-P5600

Enterprise inspired. Performance optimized

Extending NAND Technology Leadership

- **All-new firmware** optimizes performance for all-flash arrays
- New Intel **PCIe 4.0** ready controller
- Built on proven, industry-leading **Intel® 3D NAND Technology**[1]

Unlock the Value of Stored Data

- Predictably **fast**, **high performance**
- Up to **2.2x sequential read performance** vs. previous-gen[2]
- Up to **1,000,000 IOPS**[3]
- Up to **80% lower tail latencies** vs. previous-gen[4]

Advanced Feature Set

- **NVMe 1.3,** QoS and tail latency improvements for cloud & dbms, enhanced telemetry & debug
- **All-new TRIM architecture** optimizes QoS and endurance
- **Dynamic Multiple Namespaces** support enhanced runtime provisioning
- **NVMe active power states** for precision thermal and rack level TCO
- **TCG Opal 2.0**, **NVMe Sanitize**, and built-in **AES-256 Hardware Encryption** help to secure data at rest

See appendix for footnotes 1-4.
Performance results are based on testing as of the date set forth in the appendix. For more complete information on performance and test results, visit www.intel.com/benchmarks

Intel® SSD D7-P5500 and D7-P5600

New features improve IT efficiency, drive performance, and security

Telemetry 2.0

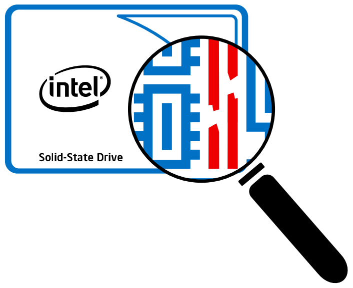

Enhanced drive debug helps improve customer experience through better issue tracking and logging, enabling accelerated qual cycles

New TRIM Architecture

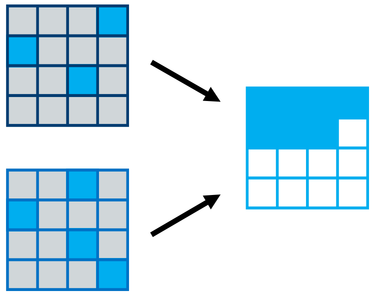

Optimized TRIM deallocate process now runs as a background operation to improve performance, reduce write amplification, and increase endurance

Dynamic Multiple Namespaces

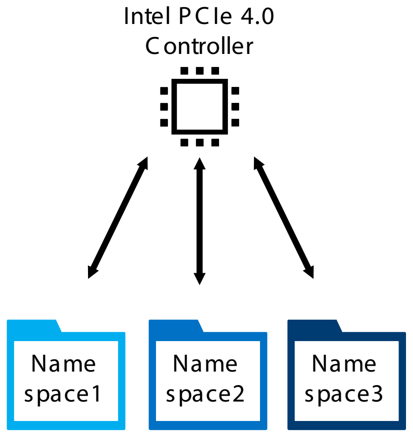

Supports create, delete, attach, and resize of namespaces during runtime

Robust Data Security Suite

TCG Opal 2.0, NVMe Sanitize, and built-in AES-256 Hardware Encryption help to secure data at rest

Appendix A

All performance claims: Test and System Configuration: CPU: Intel® Xeon® Gold 6139 @ 2.30GHz, BIOS: SE5C620.86B.00.01.0014.070920180847, RAM: 96GB, RAM Model: DDR4 2666MHz, Chipset: Intel® C624 Chipset, Hyper-Threading: Disabled, CPU Governor (through OS): Performance Mode, OS: CentOS 7.2, Kernel: 4.8.6. Measurements are performed 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. Power mode set at PM0.

1. **Industry-leading Intel® 3D NAND technology**: "Industry-leading" refers to the combination of industry-leading areal density and data retention. See footnotes 5 & 6 for more information or visit https://www.intel.com/content/www/us/en/architecture-and-technology/3d-nand-technology.html. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.
2. **2.2x sequential read performance**: Source – Intel. Comparing datasheet performance for 128KB (131,072 bytes) 100% Sequential Read QD128 (8 workers) performance between the Intel® SSD D7-P5500 7.68TB vs. Intel® SSD DC P4510 8TB. D7-P5500 and DC P4510 were measured on each drive's latest production firmware, "10019" and "1047E", respectively. Test and System Configuration as described at top of slide. Measured performance was 7000 MB/s and 3200 MB/s for the D7-P5500 and DC P4510, respectively.
3. **1,000,000 IOPS**: Source – Intel. Performance figures from D7-P5500/D7-P5600 Product Specification. Test and System Configuration as described at top of slide.
4. **80% lower tail latencies**: Source – Intel. Comparing measured performance for 4KB (4096 bytes) 100% Random Read QD1 (1 worker) latencies at 99.99999% (7 9s) availability between the Intel® SSD D7-P5500 7.68TB vs. Intel® SSD DC P4510 8TB. D7-P5500 and DC P4510 were measured on each drive's latest production firmware, "10019" and "1047E", respectively. Test and System Configuration as described at top of slide. Measured latency was 0.52ms and 3.29ms for the D7-P5500 and DC P4510, respectively.
5. **Industry-leading capacity scaling**: Source – IEEE International Solid State Circuits. ISSCC 2015, J. Im; ISSCC 2017 R Yamashita; ISSCC 2017 C Kim; ISSCC 2018; H. Maejima; ISSCC 2019 C. Siau. https://ieeexplore.ieee.org/Xplore/desktopReportingPrompt.jsp?tp=&arnumber=9063154&htmlRequest=true
6. **Industry-leading NAND reliability**: Source – Intel. Measurements were performed on components from SSDs using floating gate and charge trap flash technology. Measurement platform used was Teradyne Magnum 2 Memory test systems, and programming using random patterns and margins were quantified using customer commands. Data measured in 08/2019.
7. **56% Higher Random Read Performance**: Source – Intel. Comparing measured performance for 4KB (4096 Bytes) 100% Random Read QD256 (32x8w) performance between the Intel® SSD D7-P5500 7.68TB and Intel® SSD DC P4510 8TB. Test and System Configuration as described at top of slide. Measured performance was 1.0M IOPS and 641K IOPS for the D7-P5500 and DC P4510, respectively.
8. **44% Higher Mixed Workload Performance**: Source – Intel. Comparing measured performance for 4KB (4096 Bytes) 70/30 Random Read/Write QD256 (32x8w) performance between the Intel® SSD D7-P5600 6.4TB and Intel® SSD DC P4610 6.4TB. Test and System Configuration as described at top of slide. Measured performance was 520K IOPS and 360K IOPS for the D7-P5600 and DC P4610, respectively.
9. **80% lower tail latencies**: Source – Intel. Comparing measured performance for 4KB (4096 bytes) 100% Random Read QD1 (1 worker) latencies at 99.99999% (7 9s) availability between the Intel® SSD D7-P5600 6.4TB vs. Intel® SSD DC P4610 6.4TB. D7-P5600 and DC P4610 were measured on each drive's latest production firmware, "10019" and "1047E", respectively. Test and System Configuration as described at top of slide. Measured latency was 0.48ms and 2.57ms for the D7-P5600 and DC P4610, respectively.
10. **Predictably Fast, High Performance**: Source – Intel. Performance figures from D7-P5500/D7-P5600 Product Specification. Test and System Configuration as described at top of slide.
11. **Projected Annualized Failure Rate**: Source – Intel. Intel SSD Annualized Fail Rate Report for 2019 for products shipping >1Mu. Projected value for Intel® SSDs in high volume manufacturing.
12. **Industry-leading Silent Data Corruption**: Source – Intel. End-to-end data protection refers to the set of methods used to detect and correct the integrity of data across the full path as it is read or written between the host and the SSD controller and media. Test performed on single port SSDs, Intel® SSD DC S3520, Intel® SSD DC P3520, Intel® SSD DC P3510, Intel® SSD DC P4500, Intel® SSD DC S4500, Intel® Optane™ SSD DC P4800X, Samsung PM953, Samsung PM1725, Samsung PM961, Samsung PM863, Samsung PM963, Micron 7100, Micron 510DC, Micron 9100, Micron 5100, HGST SN100, Seagate 1200.2, SanDisk CS ECO, Toshiba Z6000 drives, Samsung 860DCT, Samsung PM883, Samsung PM983, Micron 5200, Toshiba XGS. Neutron radiation is used to determine silent data corruption rates and as a measure of overall end-to-end data protection effectiveness. Among the causes of data corruption in an SSD controller are ionizing radiation, signal noise and crosstalk, and SRAM instability. Silent errors were measured at run-time and at post-reboot after a drive "hang" by comparing expected data vs actual data returned by drive. The annual rate of data corruption was projected from the rate during accelerated testing divided by the acceleration of the beam (see JEDEC standard JESD89B/C). Projected performance of Intel® SSD D7-P5500 & Intel® SSD D7-P5600 based on these results.
13. **Power Loss Imminent**: Source – Intel. Intel Datacenter SSDs provide robust Power Loss Imminent (PLI) circuitry that helps to protect inflight data in the event of power loss. Intel SSDs monitor the health of the PLI circuitry via a Self Cap Test during power-up and operation and report it via log pages/SMART attributes.

Panel Questions and Audience Surveys

Audience Survey Question #1

- When considering SSDs for your datacenter, what factors drive your selections? (check all that apply):
  - The recommendations of application software providers: 29%
  - The recommendations of/choices available from our server vendor(s): 31%
  - SSD performance: 69%
  - SSD advanced features: 24%
  - SSD capacity: 54%
  - Other: 10%

Panel Question # 1

- What are the lead use cases for PCIe Gen4 SSDs?
  - Kioxia
  - Intel

Audience Survey Question #2

- The upcoming U.3 specification will enable U.2 NVMe SSDs, SAS/SATA SSDs, and HDDs to fit into the same slot. How important is the ability interchange media types to your organization (check one)?:
  - Very Important – We adjust our server configurations regularly: 20%
  - Important – It will allow us to update our servers and extend their life: 20%
  - Somewhat Important – It will provide us with more purchasing flexibility: 24%
  - Not that important – We can work with our current configurations: 7%
  - Not that important – We will transition to all-U.2 SSDs in the near future: 4%
  - Not that important – We will transition to EDSFF SSDs when they are available: 7%
  - No opinion: 18%

G2M RESEARCH

Panel Question # 2

- Another technology that could replace U.2 is the upcoming U.3 specification. What are the primary drivers for the adoption of U.3 by enterprises?
  - Intel
  - Kioxia

G2M RESEARCH

Audience Survey Question #3

- What is the most important factor for you to consider new technologies such as EDSFF and/or U.3 for use in your datacenter? (check one):
  - Use cases that are relevant to our organization: 10%
  - Support from relevant application vendors: 10%
  - Support from our server vendor(s): 25%
  - Availability from our normal storage vendors: 21%
  - Availability in our normal storage channel: 4%
  - Mass adoption of the technology in the market: 23%
  - Other: 6%

G2M RESEARCH

Panel Question #3

- EDSFF was THE hot topic in SSDs in early 2019? What are the hurdles to EDSFF's adoption, and how long until it replaces U.2?
  - Kioxia
  - Intel

G2M RESEARCH

Audience Q&A

G2M RESEARCH

Thank You For Attending

G2M

RESEARCH