State of Datacenter NVMe SSD – 2021
OCP Workshop

Jeff Janukowicz
Vice President

Spending on Enterprise SSDs now outpaces Enterprise HDDs

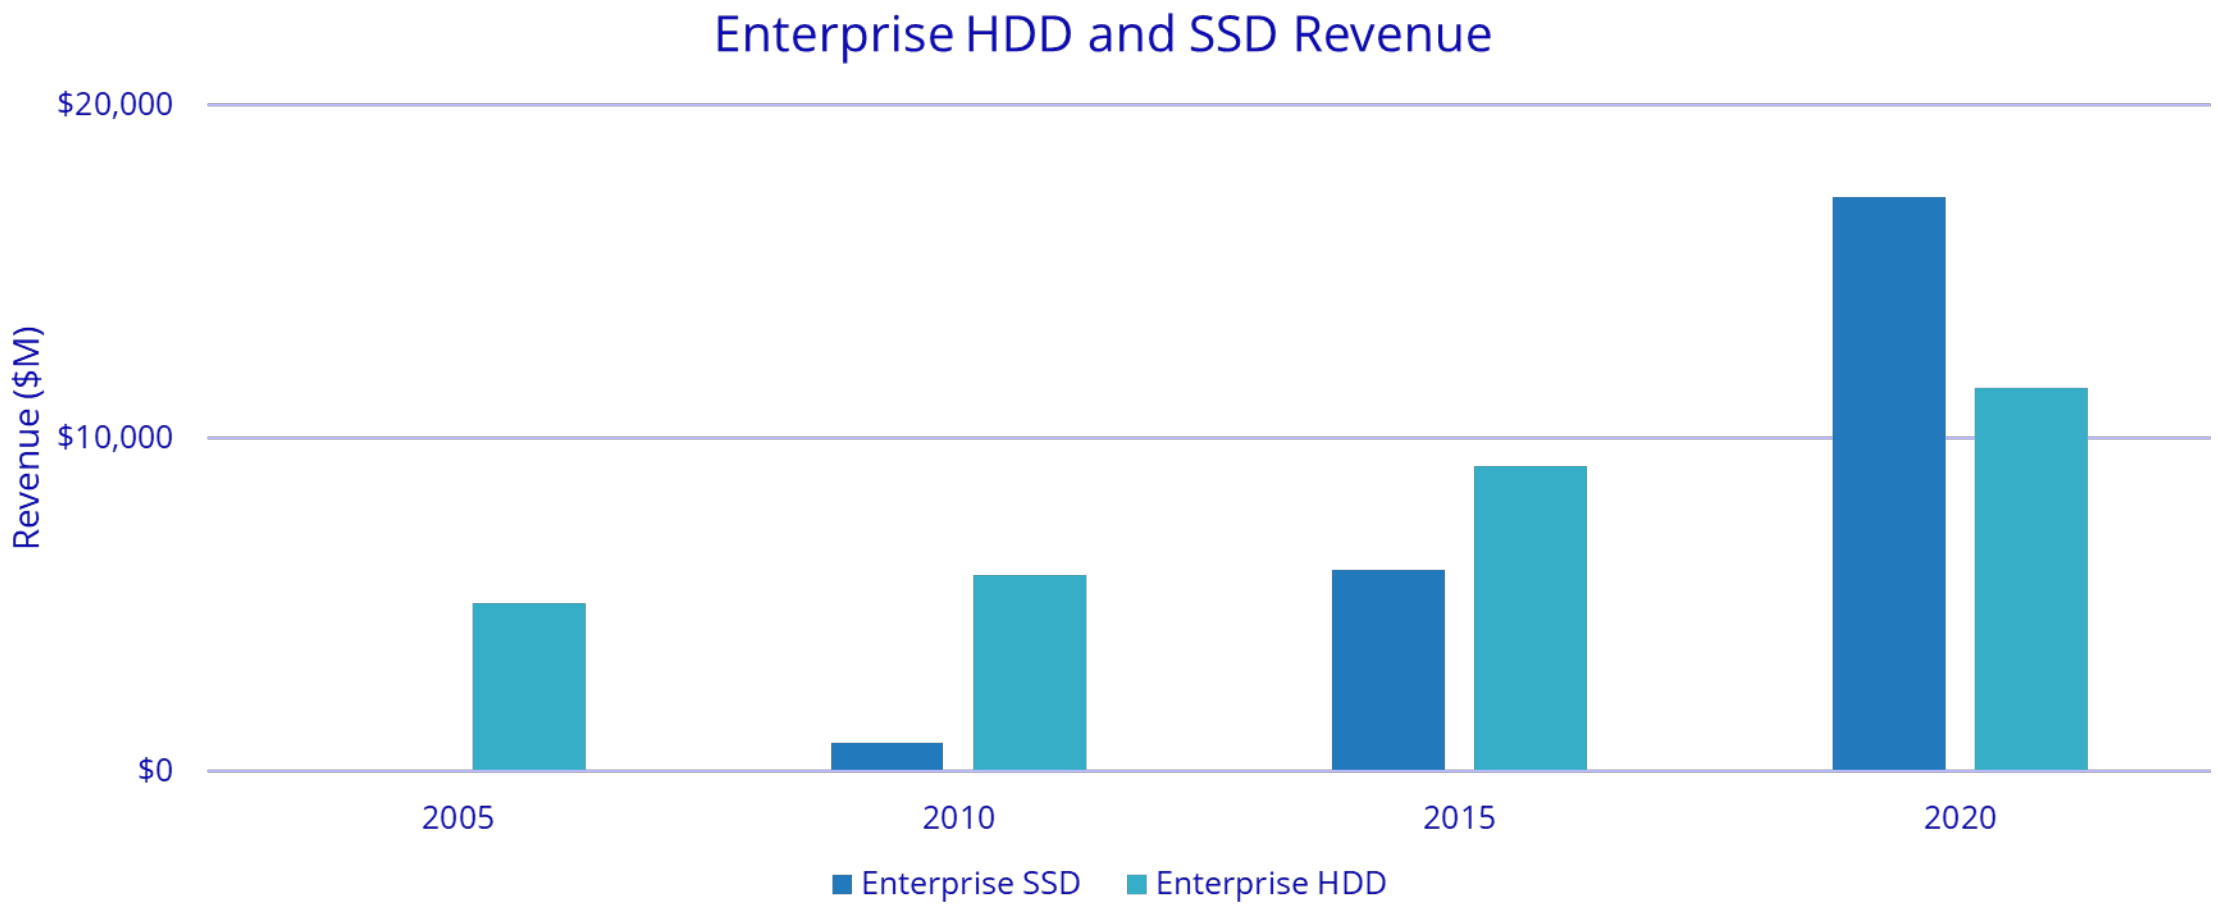

Sources: Worldwide Hard Disk Drive Forecast Update, 2020–2024, Doc # US47039720, Dec 2020

Worldwide Solid State Drive Forecast Update, 2019-2024 Doc # US45909420 , Dec 2020

SSDs Continue to Expand in the Datacenter

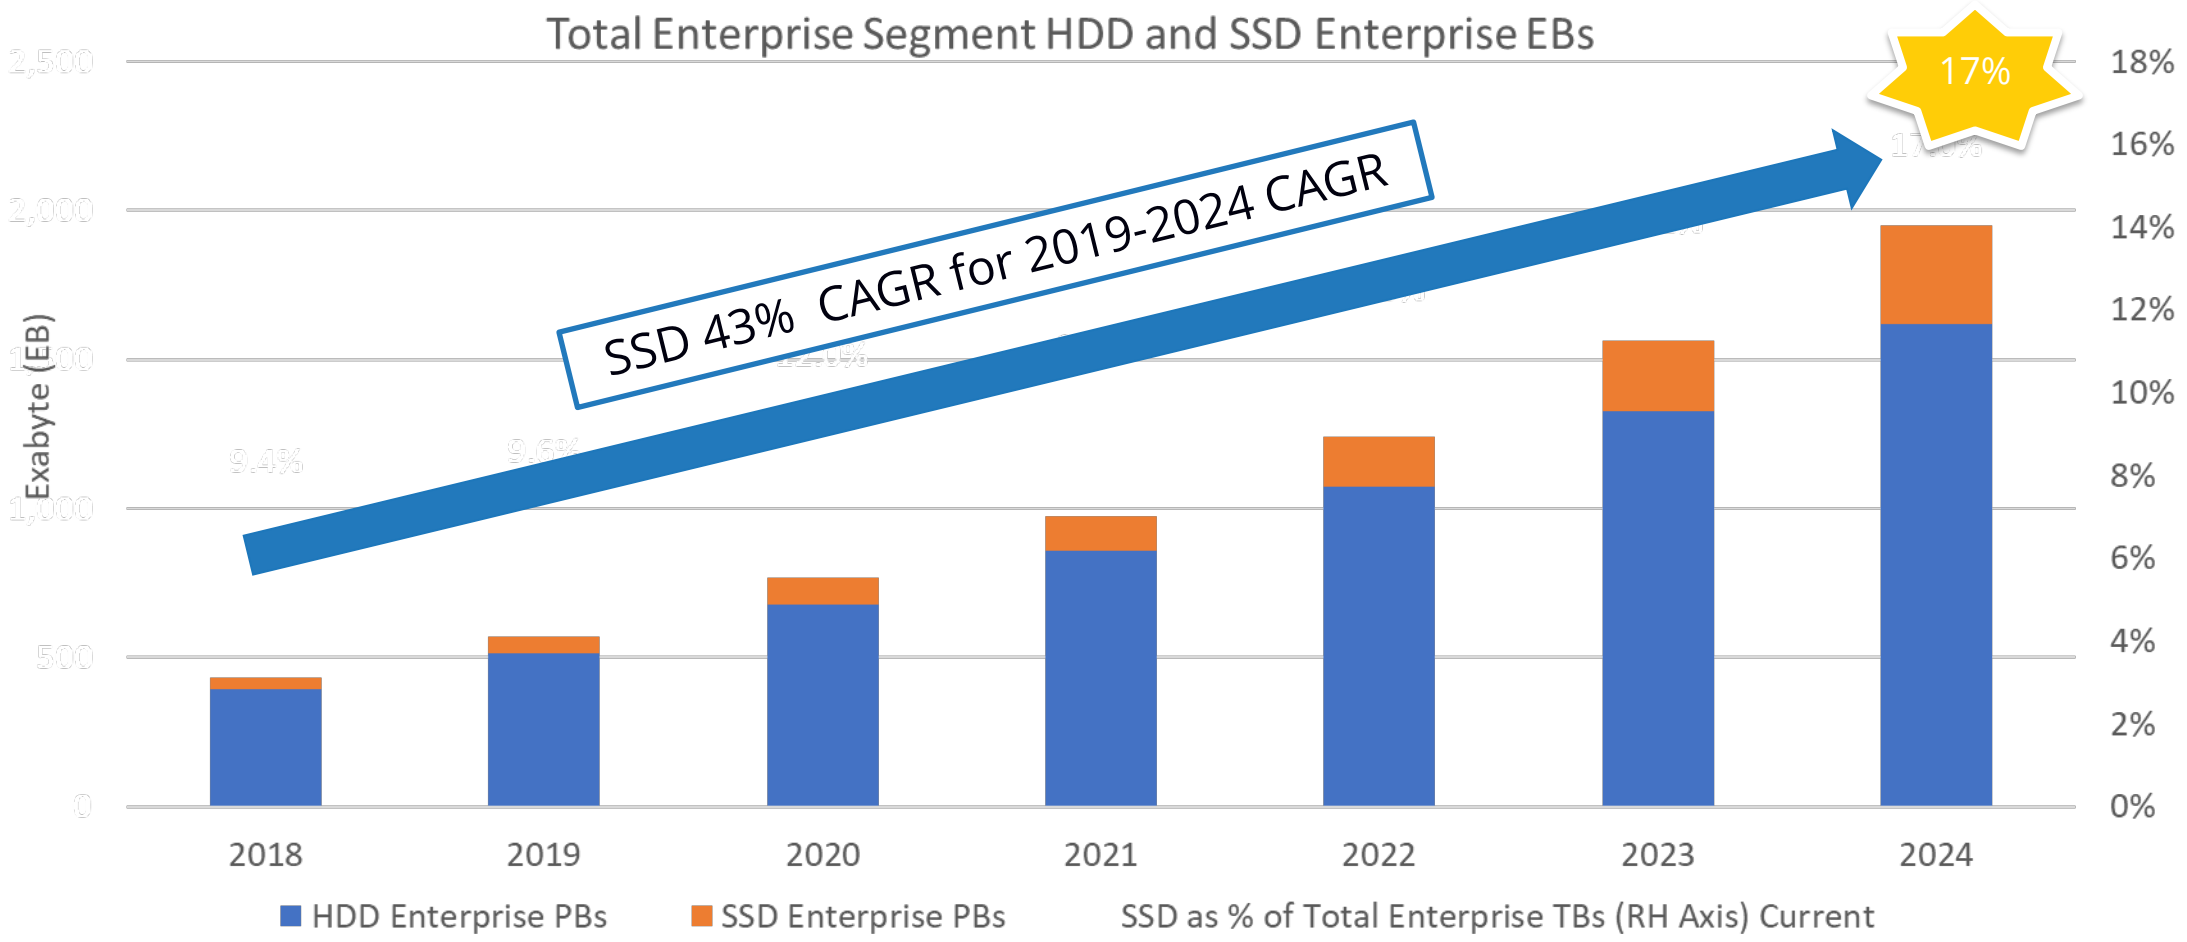

Sources: Worldwide Hard Disk Drive Forecast Update, 2020–2024, Doc # US47039720, Dec 2020

Worldwide Solid State Drive Forecast Update, 2019-2024 Doc # US45909420 , Dec 2020

IDC ANALYZE THE FUTURE

Optimizing the Ecosystem for Solid State Drives

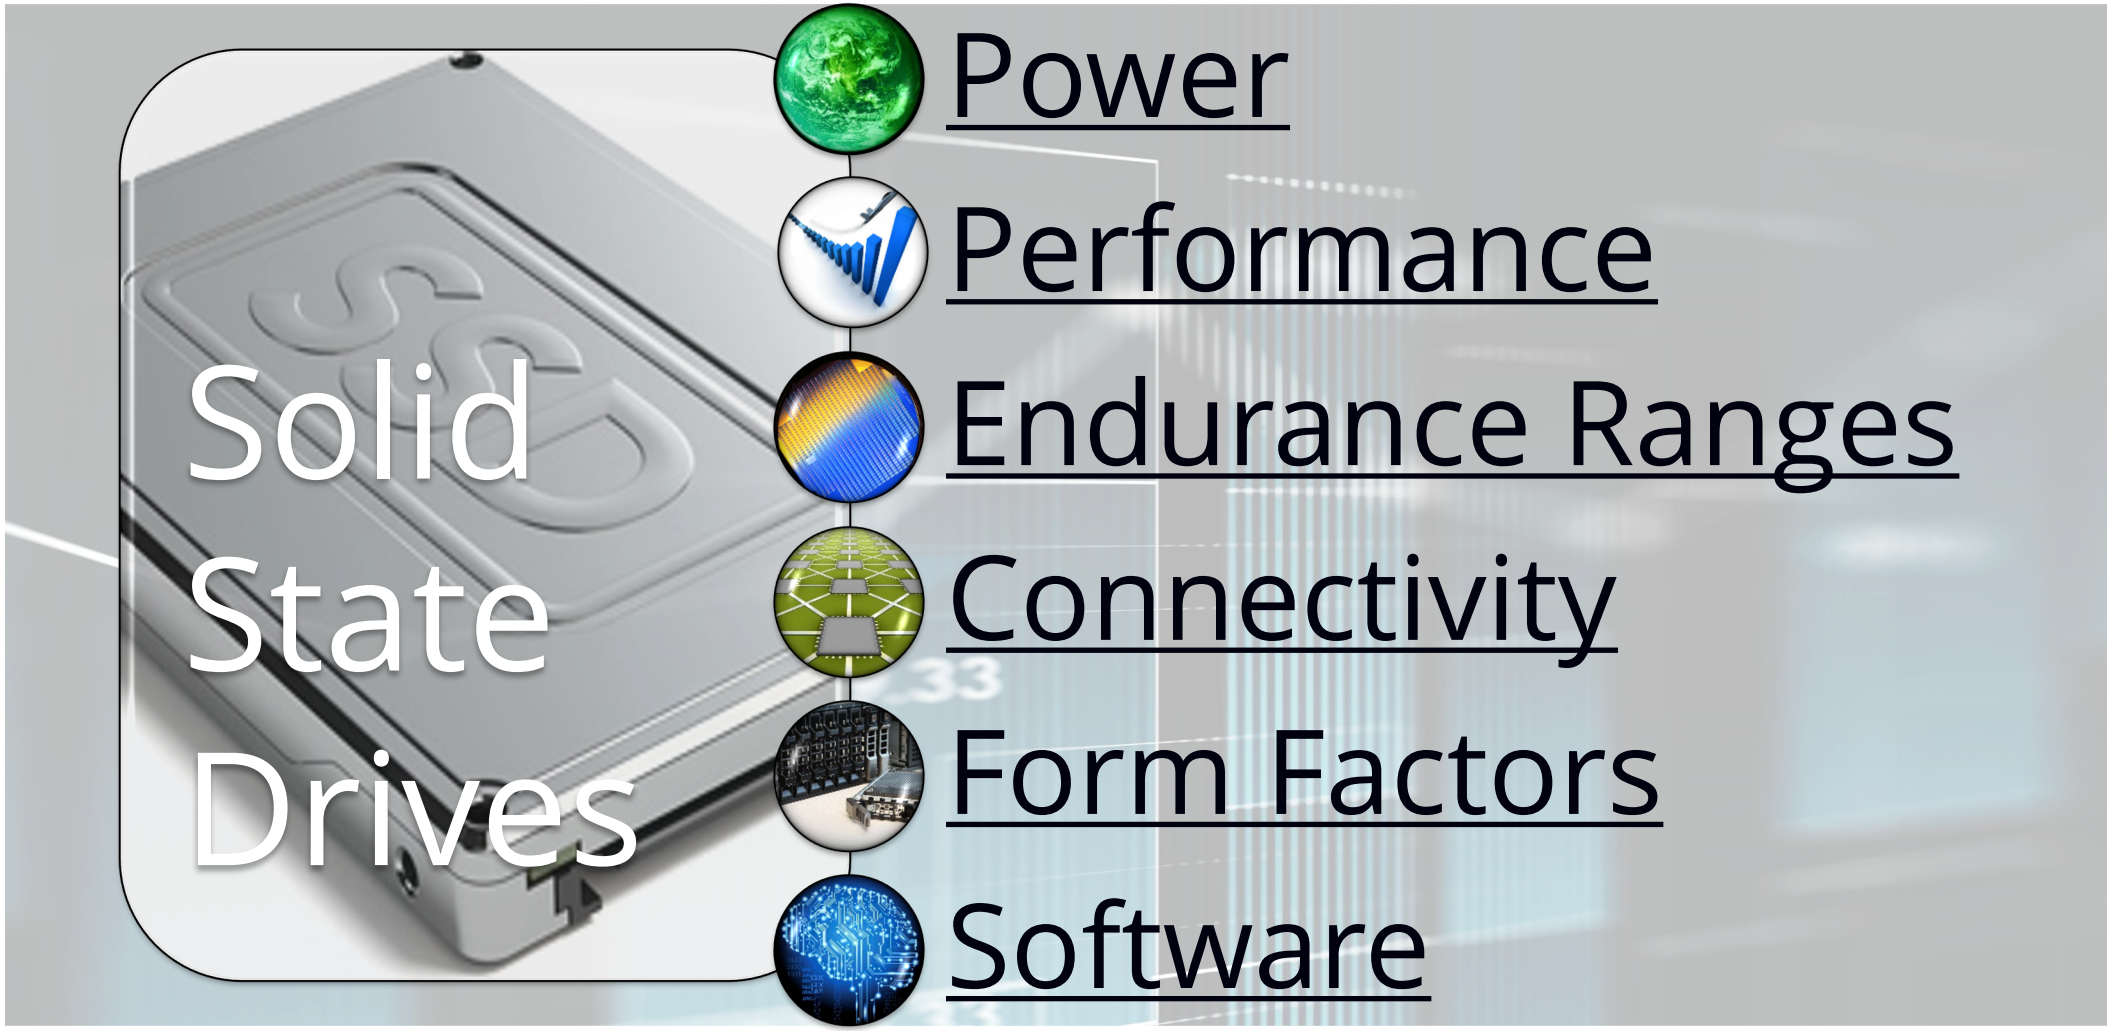

Enterprise SSD Capacity Shipment Forecast by Interface

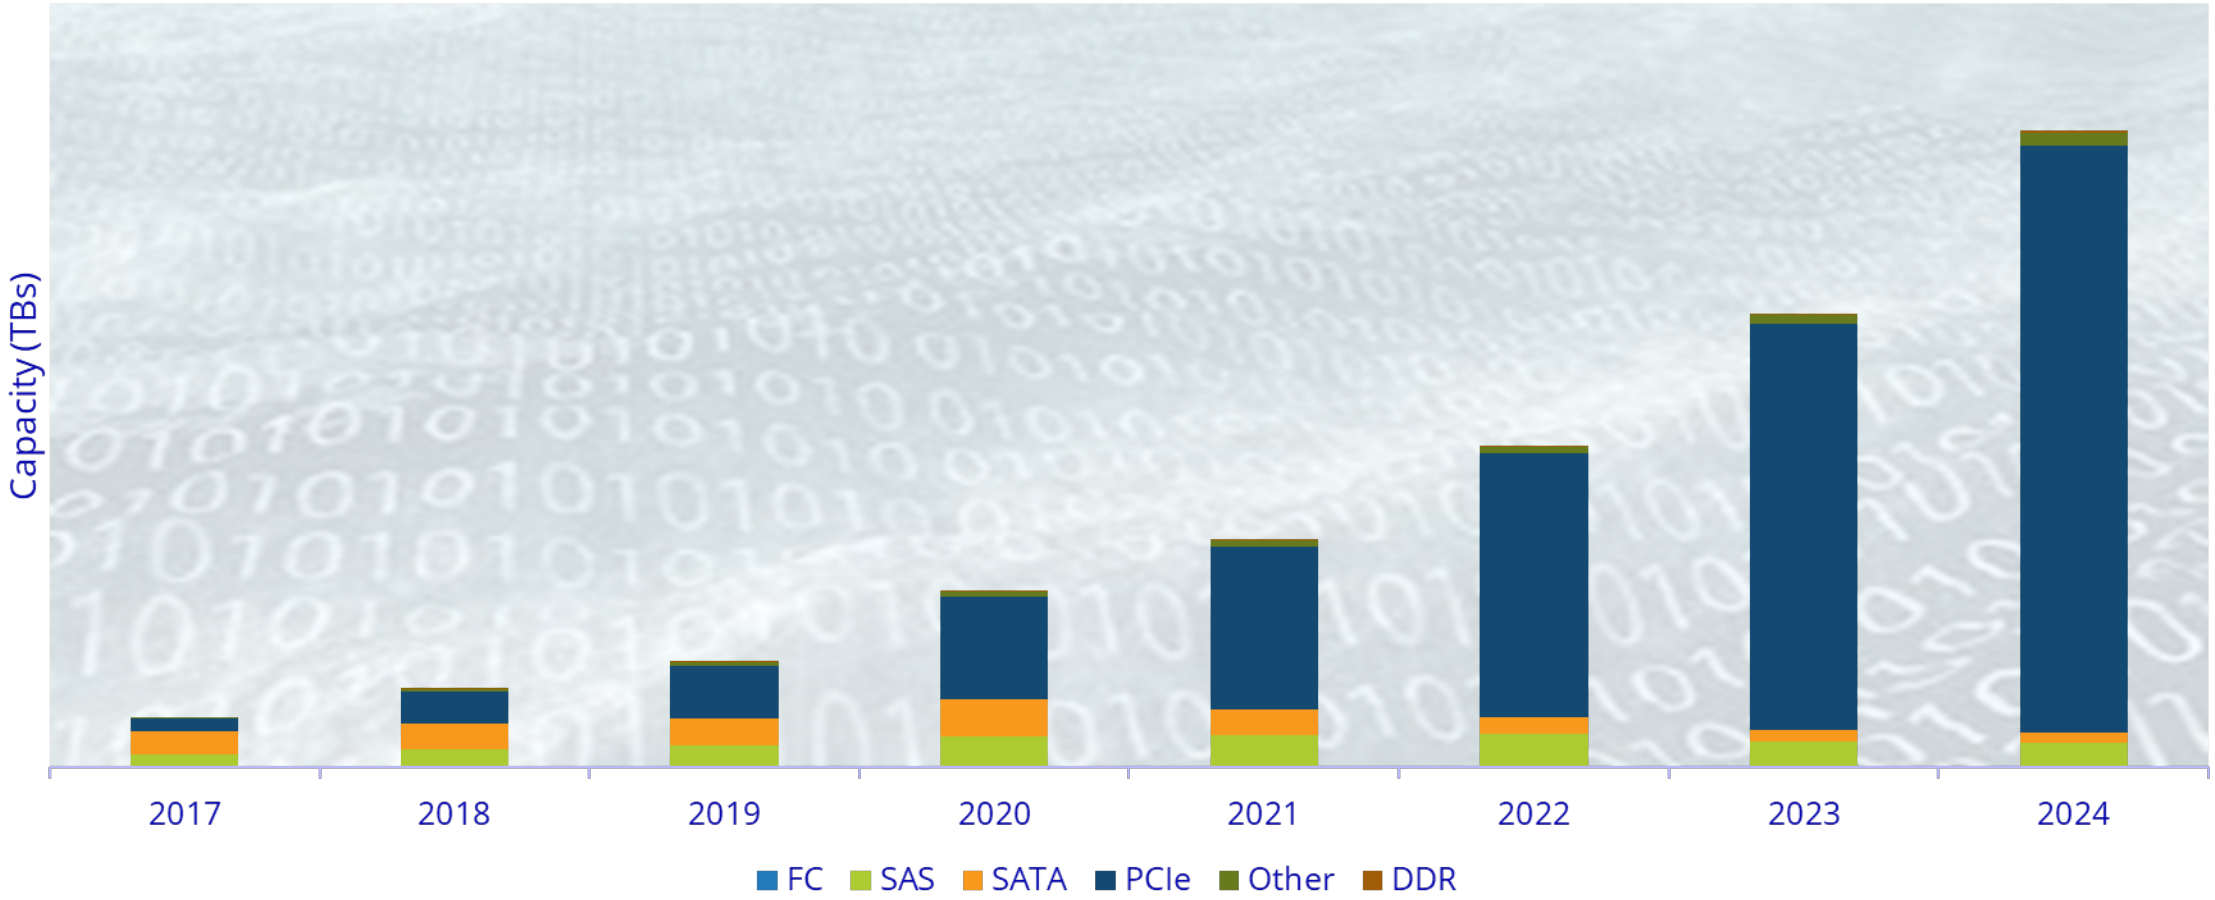

Worldwide Solid State Drive 2019-2024 Forecast Update (IDC#US45909420, December 2020)

Datacenter NVMe Growth Outpaces the Market

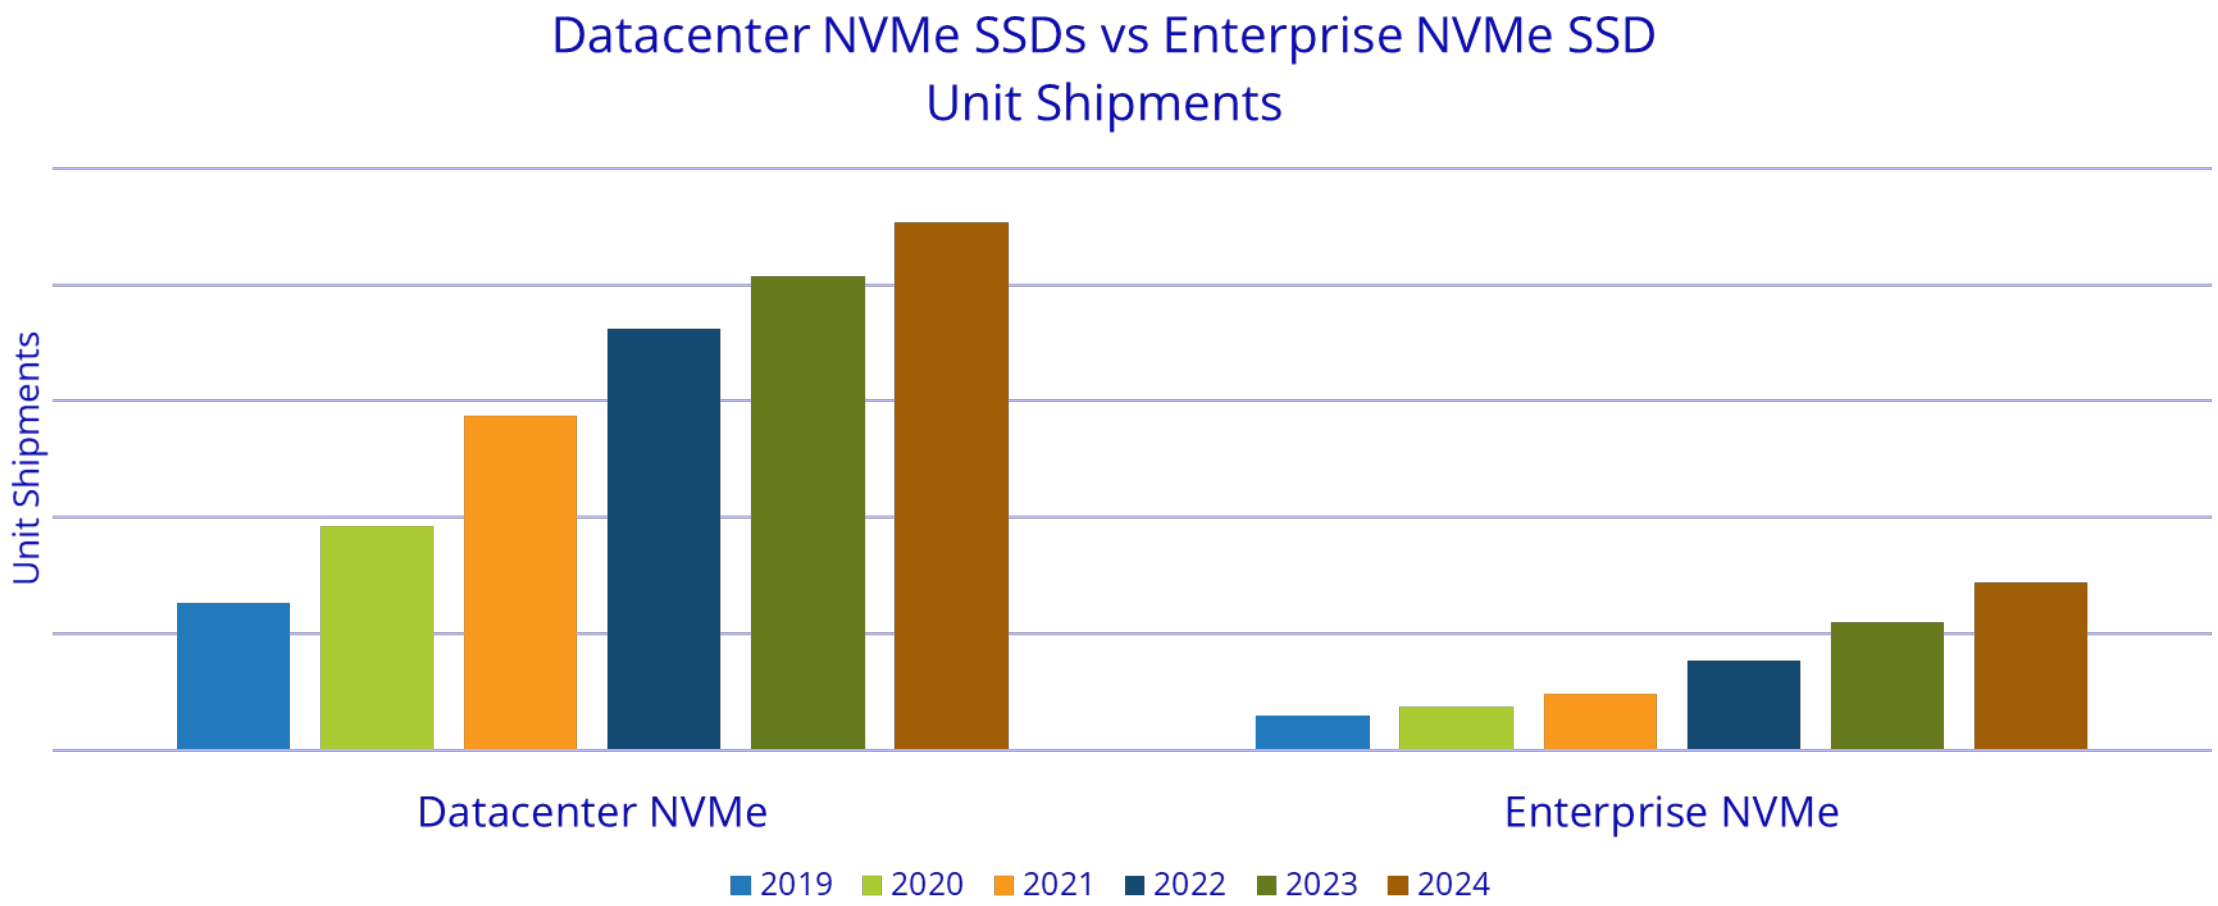

Worldwide Solid State Drive 2019-2024 Forecast Update (IDC#US45909420, December 2020)

Evolution of Enterprise SSD Form Factors

Flexibility to offer various options for length, width, and height with manageability, serviceability, and thermal consideration built-in

BGA

m.2

2.5” (U.2)

E3

AIC

E1.L

E1.S

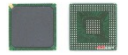

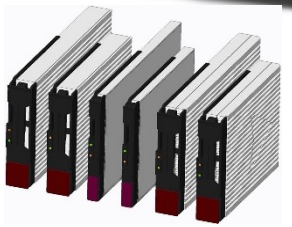

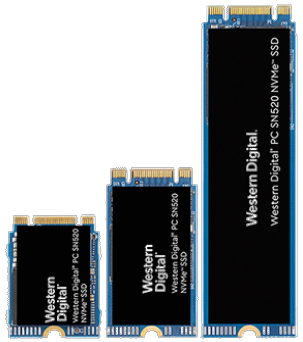

Key Takeaways

1. A New Era in the Enterprise Market with Solid State Storage in the Forefront of modern Datacenters Driving Digital Transformation
2. NVMe SSDs have moved in the mainstream and Datacenter NVMe SSD Growth Outpaces the Market
3. Optimizing IT Infrastructure for Solid State Storage to be More Efficient, Flexible, and Scalable will be Key to Drive Future of Digital Infrastructure

For More Information

Jeff Janukowicz
Research Vice President
jjanukowicz@idc.com

For more information, please visit, www.idc.com and search for:
NAND flash or Solid State Drive

OCP Workshop: Data Center NVMe SSD and EDSFF

A Media Perspective

www.ServeTheHome.com

Serve The …

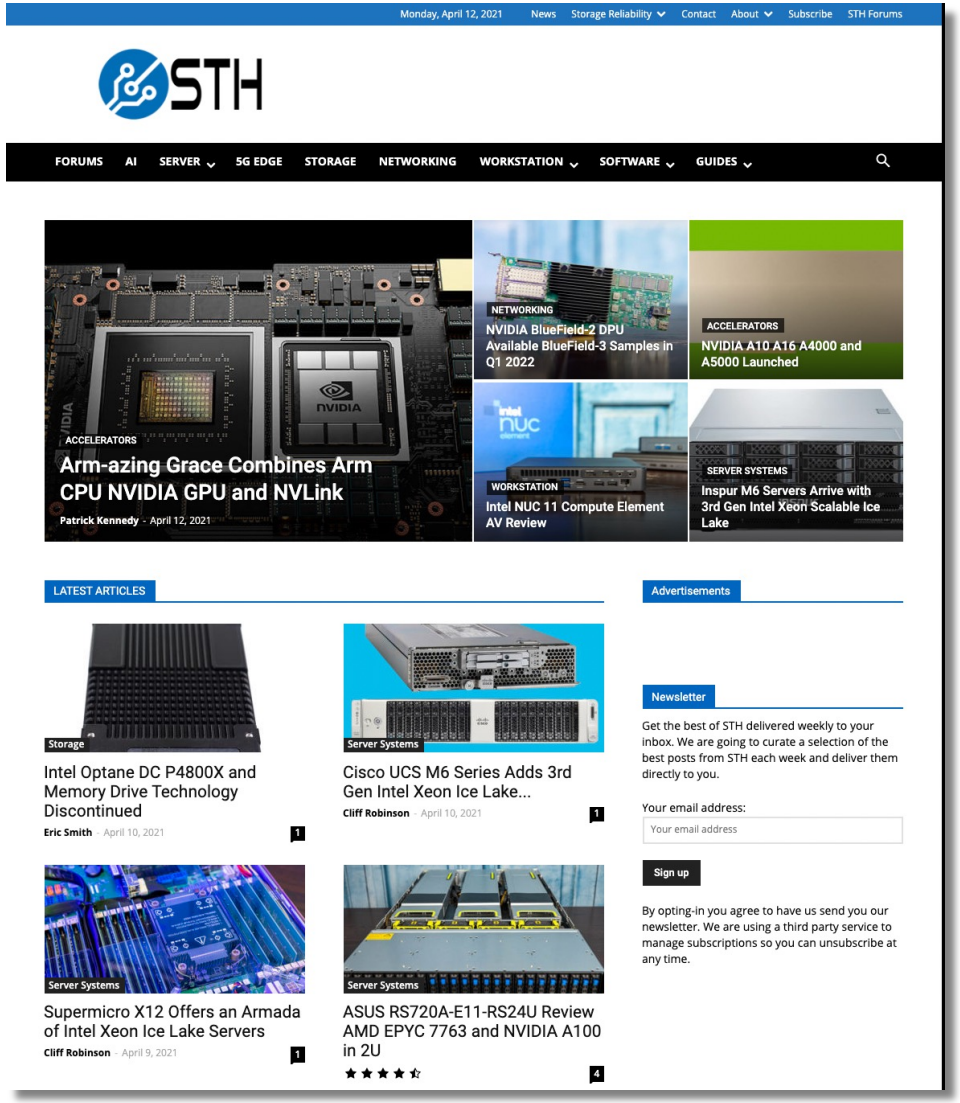

STH was founded in June 2009

- Focus on the data center market, specifically sever, storage, and networking
- Hardware reviews are in-data center in three Silicon Valley data centers
- Covered OCP since OCP Summit was in the Santa Clara Convention Center
- Readers come from different aspects including IT admins, programmers, procurement, “server enthusiasts”

	WEB	VIDEO
CONTENT CREATION	✓	✓
ORGANIC DISTRIBUTION	✓	✓

Key Topics

- Insights STH Coverage Thus Far
- Marketing Perspective on Overcoming Adoption Barriers
- Closing Thoughts

Sample “Ruler” Coverage

The Intel Ruler SSD: Already Moving Markets

Intel Ruler SSD Connector

Hands-on with the 1U Half-Petabyte Supermicro EDSFF Server

Where Cloud Servers Come From Visiting Wiwynn in Taipei

AIC FB127-AG Innovative NF1 AMD EPYC Storage Solution

Kioxia EDSFF SSD adding Momentum with New Development

Key Coverage Insights

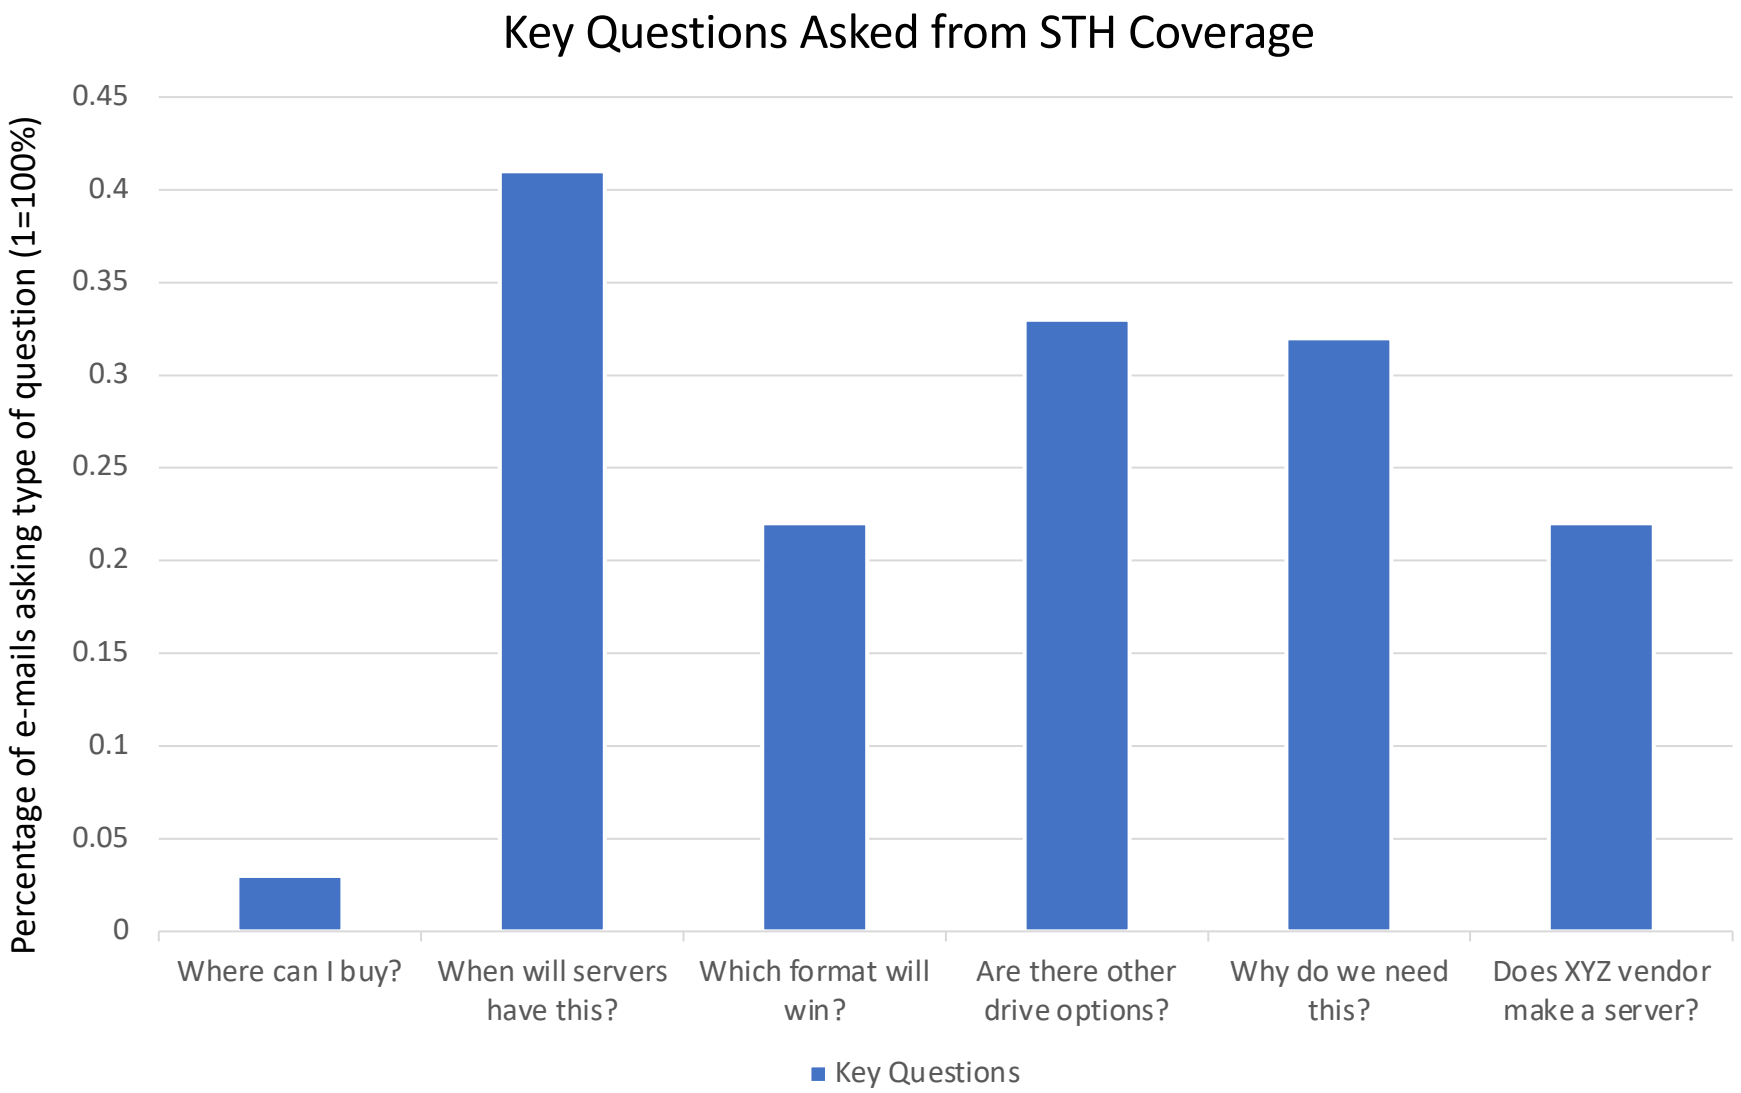

Patrick's Key Insights:

- Early in cycle
- Uncertainty on scope
- Uncertainty on "winner"

***Sample size**: 912 e-mails + article comments*
Counting instances so one email/ comment may Have multiple summarized questions

High Engagement Rate

Coverage Statistics

- Next-gen SSD Form Factor Articles (non-review)
  - 22% longer time on page than benchmark
  - 41% higher page views in first 72 hours
  - 9% better “long tail” page views after first quarter

- Hands-on piece with Supermicro
  - #1 Video at the time

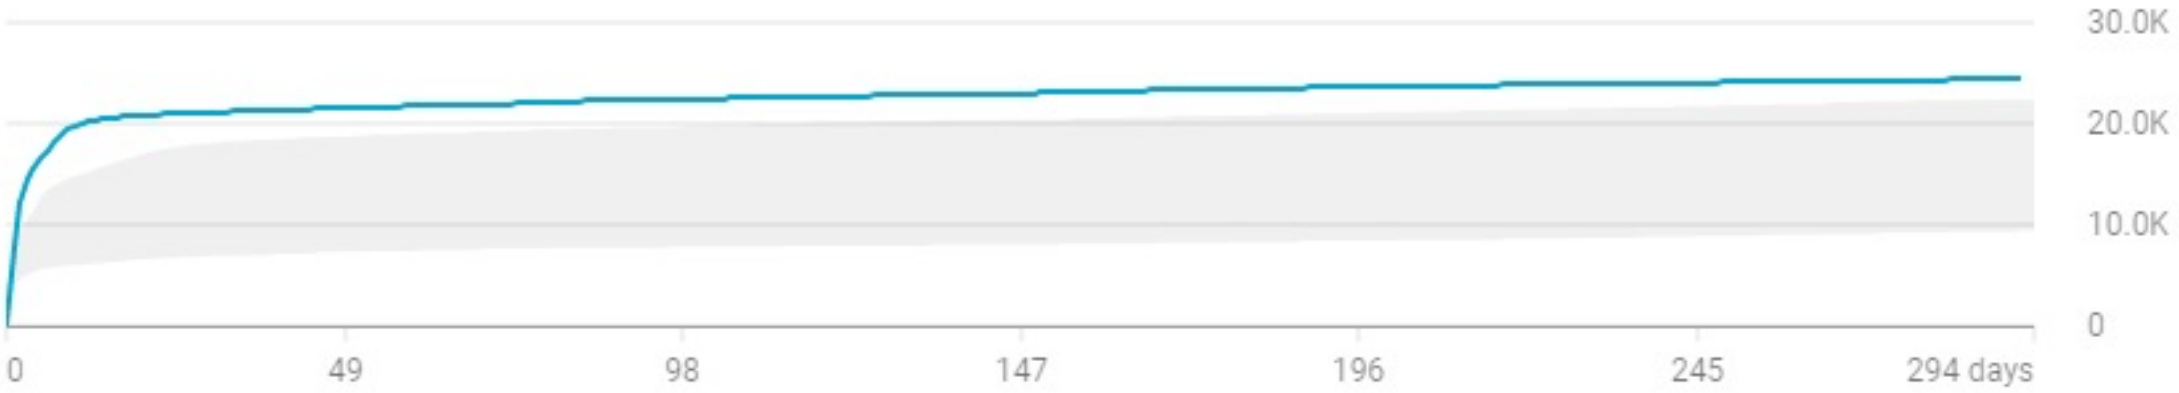

Market Perspective

2.5” seen as a “safe” choice in the industry

V.

“Do I want to be an infrastructure hero?”

Market Perspective

Need benefits beyond the “ultra-technical”

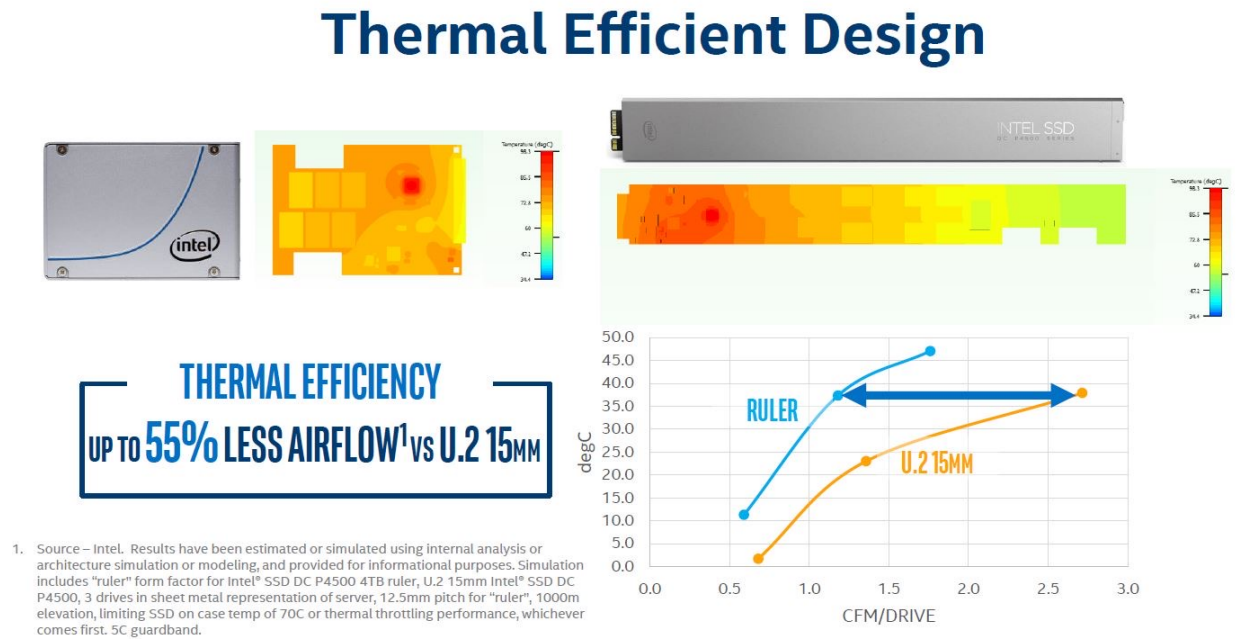

2.5” SSDs had benefits of:

- Higher reliability
- Lower latency
- Higher throughput

Why does a buyer care about thermal design?

Market Perspective

Diverging Data Center Densities

As power consumption rises, we are seeing two classes emerge. Those prepared for higher-power and those without.

Next-generation NVMe form factors such as the OCP NVMe SSD standard will be easier for hyper-scale to consume than many higher-margin enterprise customers

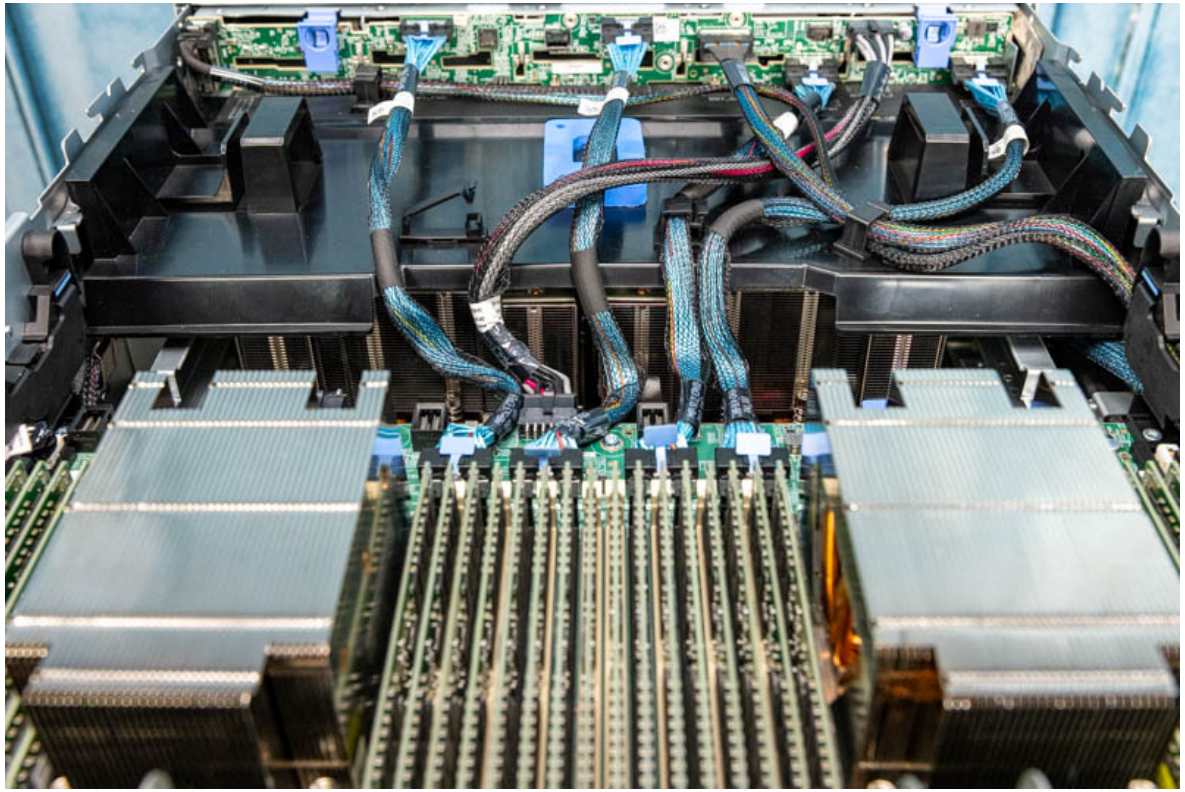

Closing Thoughts

- There is a surprising amount of market interest
- Now that standards are being driven, it is time to get messaging aligned
- Need a vision around the art of possible for next-gen drives

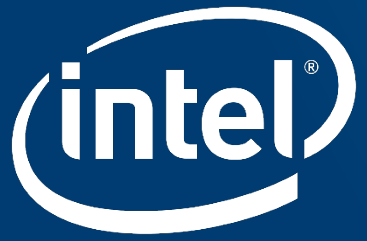

EDSFF: Meeting System Needs

Anthony Constantine

Principal Engineer, Intel

What is EDSFF?

- Multiple usages = Multiple form factors

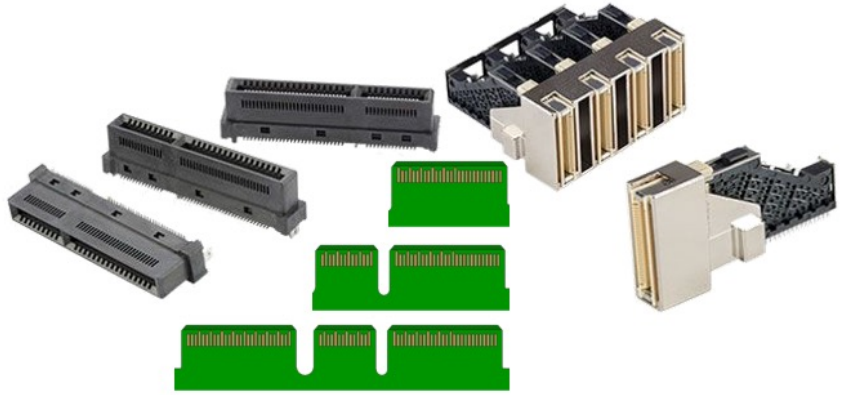

- One connector and interface

SNIA™
SFF TWG
Technology Affiliate

SFF-TA-1009

Specification for

Enterprise and Datacenter Standard Form Factor Pin and Signal Specification

- One common set of specifications

Updates in the Specifications

SFF-TA-1009 Rev 3.0 (pin/signal spec)

- Electrical Requirements for 16 and 32 GT/s
- Lots of cleanup

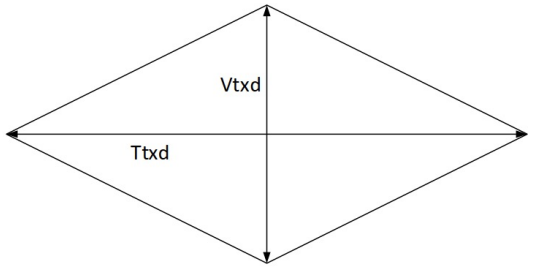

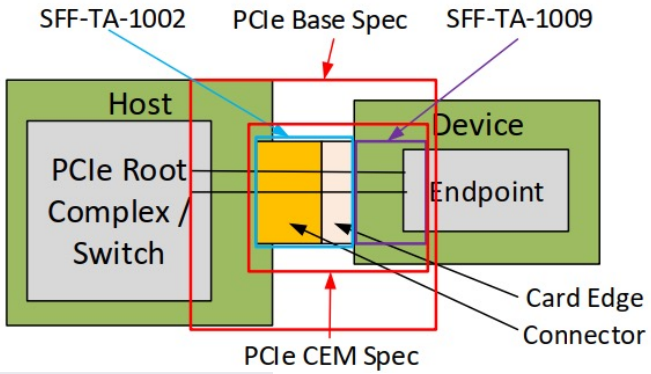

SFF-TA-1008 Rev 2.0 (E3 form factor)

- Length, connector position change
- Other changes

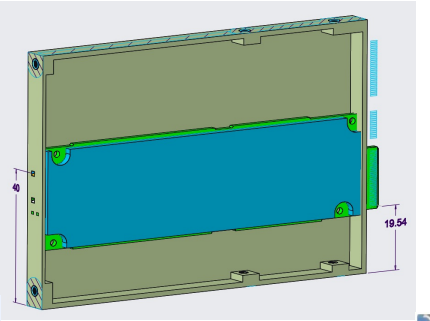

New: SFF-TA-1023 (Thermal Characterization)

- Currently a draft
- Device characterization methodology

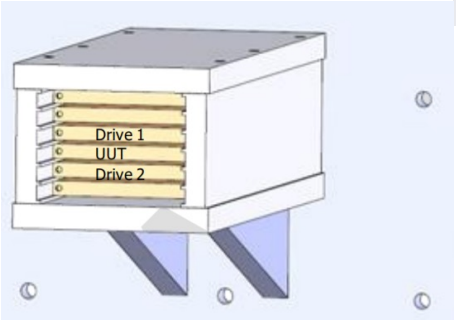

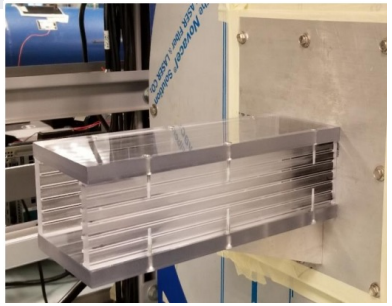

Expect more updates soon!

Want to know more?

EDSFF: Enterprise and Datacenter SSD Standard Form Factor

Visit: http://www.snia.org/sff/specifications

- SFF-TA-1002: Card Edge multilane protocol agnostic connector
- SFF-TA-1006: Enterprise and Datacenter 1U Short Standard Form Factor (E1.S)
- SFF-TA-1007: Enterprise and Datacenter 1U Long Standard Form Factor (E1.L)
- SFF-TA-1008: Enterprise and Datacenter Form Factor (E3)
- SFF-TA-1009: Enterprise and Datacenter Standard Pin and Signal Specification

Participate:

- SFF: https://www.snia.org/sff
- OCP: https://www.opencompute.org/projects/storage

Adopt EDSFF!

Microsoft

Azure M&S

Memory & Storage Center of Excellance

OCP v2.0 Major Changes

04/21/2021

Lee Prewitt - Microsoft

Major Additions

- Latency Monitor (Log Identifier C3h) and Latency Monitor (Feature Identifier C5h)
- Latency Monitoring Feature Set Theory of Operation
- Unsupported Requirements (Log Identifier C5h)
- PLP Health Check Interval (Feature Identifier C6h)
- Power Loss Protection requirements
- Additional Thermal Throttling requirements
- Additional NVMe Basic Management Command requirements
- Additional Security requirements
- Additional Label requirements

New Participants

HPE and Dell Add an Enterprise Perspective

- Renamed from “NVMe Cloud SSD Specification” to “Datacenter NVMe SSD Specification”
- Support for Multiple Namespaces
- Added requirements for several more commands: Sanitize, NVMe-MI Send/NVMe-MI Receive, Write Zeros, Compare, Write Uncorrectable
- Added Device Capabilities (Log Identifier C4h)
- Added NVMe-MI 1.1b Requirements
- Component Measurement and Authentication (CMA) requirements
- Device Profiles

Device Profiles

Surprisingly Few Differences

Requirement ID	Description	Configuration Setting	
		A	B
DP-CFG-1	Factory Default Sector Size.	4096-byte	512-byte
DP-CFG-2	Number of Namespaces Supported.	NSM-4 (16 Namespaces)	NSM-5 (16 Namespaces per TB)
DP-CFG-3	Retention Time based on RETC-1 (data retention time).	1 Month	3 Months
DP-CFG-4	NVMe Basic Supported.	R	P
DP-CFG-5	Max M.2 top side height.	FFM2-3 (top-side height)	TBD

Minor Additions

NVMe-OPT-4	The device shall only reset the Timestamp value on a power cycle.

SMART-29	199:192	PCIe Link Retraining Count	8	This is a count of the number of PCIe Link Retraining events. This count shall only increment during run time. It shall not increment during training or power fail. This shall be cleared to zero on factory exit.

FWUP-10	Firmware activation shall not cause user data to be lost or destroyed.
FWUP-11	Firmware activation without reset shall preserve the running state of the device (e.g., Opal locking state, Set Features, etc.).

SECTOR-4	The device shall support 512-byte and 4096-byte logical block sizes.

TTR-19	When the CC.SHN register is written to notify the device to shutdown it shall not be assumed that power will be lost even after CC.EN is cleared to 0. Under these conditions the device shall continue to function properly based on the NVMe and PCIe Specifications.

REL-5	The AFR targets in REL-1 (MTBF of 2.5 million hours) and REL-2 (MTBF of 2.0 million hours) shall be maintained up to a continuous reported composite temperature of 77˚C (170 ˚F) (WCTemp) with less than 1% of the device lifetime above WCTemp.
GFF-6	A CAD file of each supported form factor shall be provided.
PWR-8	A device shall not consume more power than the PCI-SIG Slot Capabilities register Slot Power Limit regardless of other settings (e.g., PSD in NVMe) and report Entry and Exit Latencies of FFFFFFFFh for all Power State Descriptors above the Slot Power Limit.
FFE1S-12	A x4 device shall only have a 1C connector.

DEVLMT-1	The device shall not have any restrictions on the number of times any of the following events can occur: Firmware downloads and activation supported. See FWUP-2 (number of firmware downloads); or Changing password when taking/changing ownership via TCG; or Crypto erase or block erase (format); or Power cycles; or Set/get features (including power state changes); or Log page or debug log retrievals.

FFE1L-12	A x4 device shall only have a 1C connector.

Next Generation Flash

Ross Stenfort, Facebook

Facebook @ Scale

3.30 Billion People Monthly

2.60 Billion People Daily

E1.S 25mm Flash Platforms

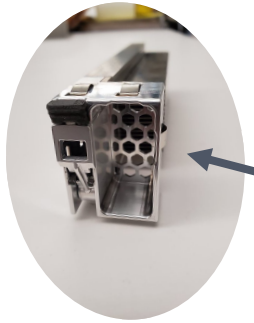

25mm E1.S Latch

1 OU Blade

Hyperscale Benefits:

- High Density
  - Enables up to ~6 PB per rack
- Low Airflow (CFM/W)
- Flexible CPU/Flash Ratios
- Excellent Serviceability

Chassis with 1 OU Blades

Chassis with 2 OU Blades

Rack with up to 96 blades

U41	Air Baffle			U41
U40	TOR Wedge100/100S			U40
U39	Air Baffle			U39
U38				U38
U37	Yosemite V3	Yosemite V3	Yosemite V3	U37
U36				U36
U35				U35
U34				U34
U33	Yosemite V3	Yosemite V3	Yosemite V3	U33
U32				U32
U31				U31
U30	Power Shelf			U30
U29				U29
U28				U28
U27				U27
U26	Yosemite V3	Yosemite V3	Yosemite V3	U26
U25				U25
U24				U24
U23				U23
U22	Yosemite V3	Yosemite V3	Yosemite V3	U22
U21				U21
U20				U20
U19				U19
U18	Yosemite V3	Yosemite V3	Yosemite V3	U18
U17				U17
U16				U16
U15				U15
U14	Yosemite V3	Yosemite V3	Yosemite V3	U14
U13				U13
U12				U12
U11	Power Shelf			U11
U10				U10
U9				U9
U8				U8
U7	Yosemite V3	Yosemite V3	Yosemite V3	U7
U6				U6
U5				U5
U4				U4
U3	Yosemite V3	Yosemite V3	Yosemite V3	U3
U2				U2
U1				U1

Datacenter NVMe SSD Specification Benefits

- Aligns SSD needs and requirements between Hyperscale/OEMs and SSD makers
- What is included in the specification?
  - NVM Express
  - PCI Express
  - SMART Logs
  - Reliability
  - Thermal
  - Power
  - Security
  - Form Factor
  - SMBUS
  - Tooling
- *Everything Needed to build a Datacenter NVMe SSD*

➢ ***Datacenter NVMe SSD and E1.S: Next generation technology ready to solve today's problems***

Link to 1.0 specification can be found under OCP Contributions:
https://www.opencompute.org/documents/nvme-cloud-ssd-specification-v1-0-3-pdf

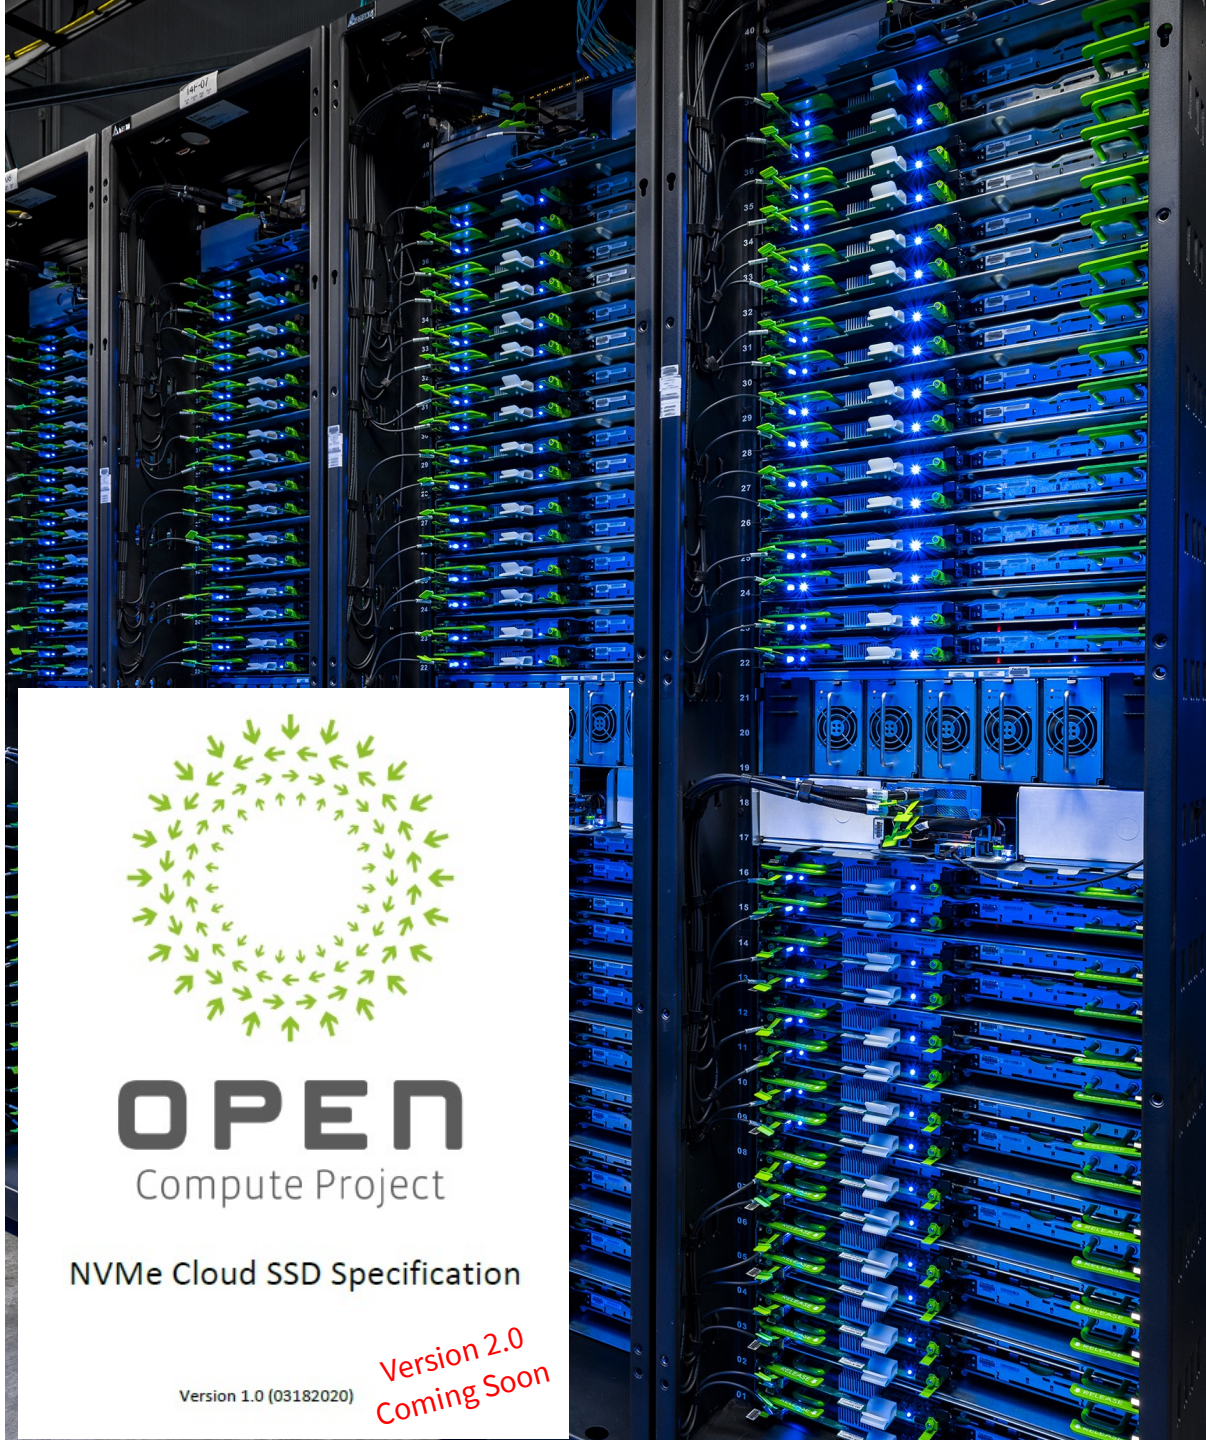

Thank You

E3.S ENABLES NEXT-GEN DEVICES AND OPEN NVME SSD SPECS

Paul Kaler, Future Storage Architect

E3.S ENABLES EASY TRANSITION TO NEXT-GEN DEVICES

- E3.S 2U designs can share a chassis with existing form factors
  - Support both E3.S and 2.5” drive cages for easy customer transition—mix SAS/SATA/NVMe
  - Swap two E3 thins for one E3.S 2T (thick)
  - Intermix NVMe and CXL devices
  - Shared bays increases flexibility and reduces cost
- Supports large FPGA and SoCs
  - Future devices types (e.g., NIC, TPU/GPU, CSD)
- E3.S better airflow and thermals than 2.5”
  - Enables higher TDP downstream components
  - Higher performance devices

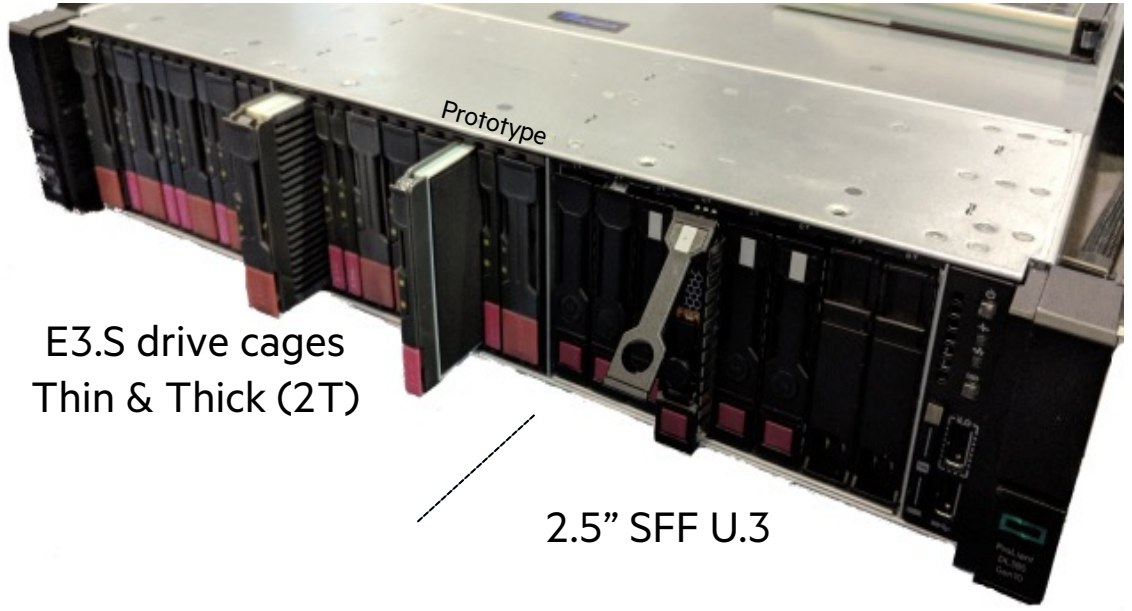

E3.S ENABLES EASY TRANSITION TO NEXT-GEN DEVICES

- Smaller connector enables smaller backplanes–reduces airflow impedance
- Better thermals enables up to 40W for E3.S 2T
  - Enables full saturation of PCIe Gen5 x4 NVMe and CXL devices
  - Provides thermal room to grow for PCIe Gen6 performance
- Cost effective performance scaling
  - Mix E3.S thin and thick to optimize performance without requiring PCIe switches
  - Higher MTBF & lower solution cost
- E3 thin enables excellent performance density for 1U as well
  - 20 drives for 2x the IOPS and bandwidth compared to 2.5”

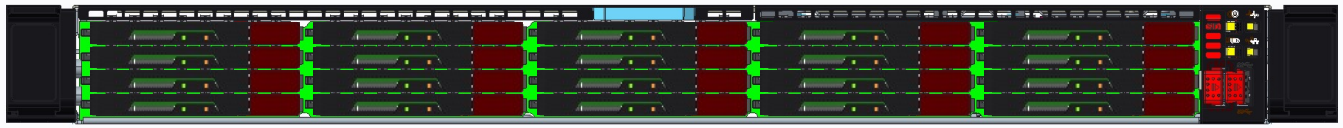

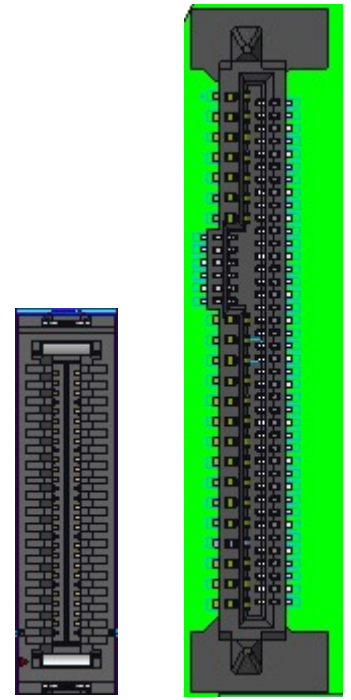

E3
(x4)

SFF-8639
(x4)

OCP DATACENTER NVME SSD SPECIFICATION BENEFITS

- HPE has traditionally developed custom firmware specifications for drives to achieve critical benefits:
  - Ensure consistent behavior
  - Assurance of Supply – consistent behavior enables multi-source
  - Faster issue resolution by specifying additional telemetry and metadata logs
  - Improved quality – spec out best practices from lessons learned
- HPE sees significant commonality with the features and requirements specified in the Cloud spec
- Leverage to drive economies of scale and improve quality
- Open requirements enable more complete 3rd party compliance tests
- HPE, Dell, Facebook, and Microsoft are working towards a common spec for both Enterprise and Cloud use cases which encompasses new form factors like E3.S – part of the EDSFF family
- Almost Complete! We hope to be releasing the final Datacenter NVMe SSD Specification very soon.


Hewlett Packard
Enterprise
THANK YOU
Paul Kaler, Future Storage Architect

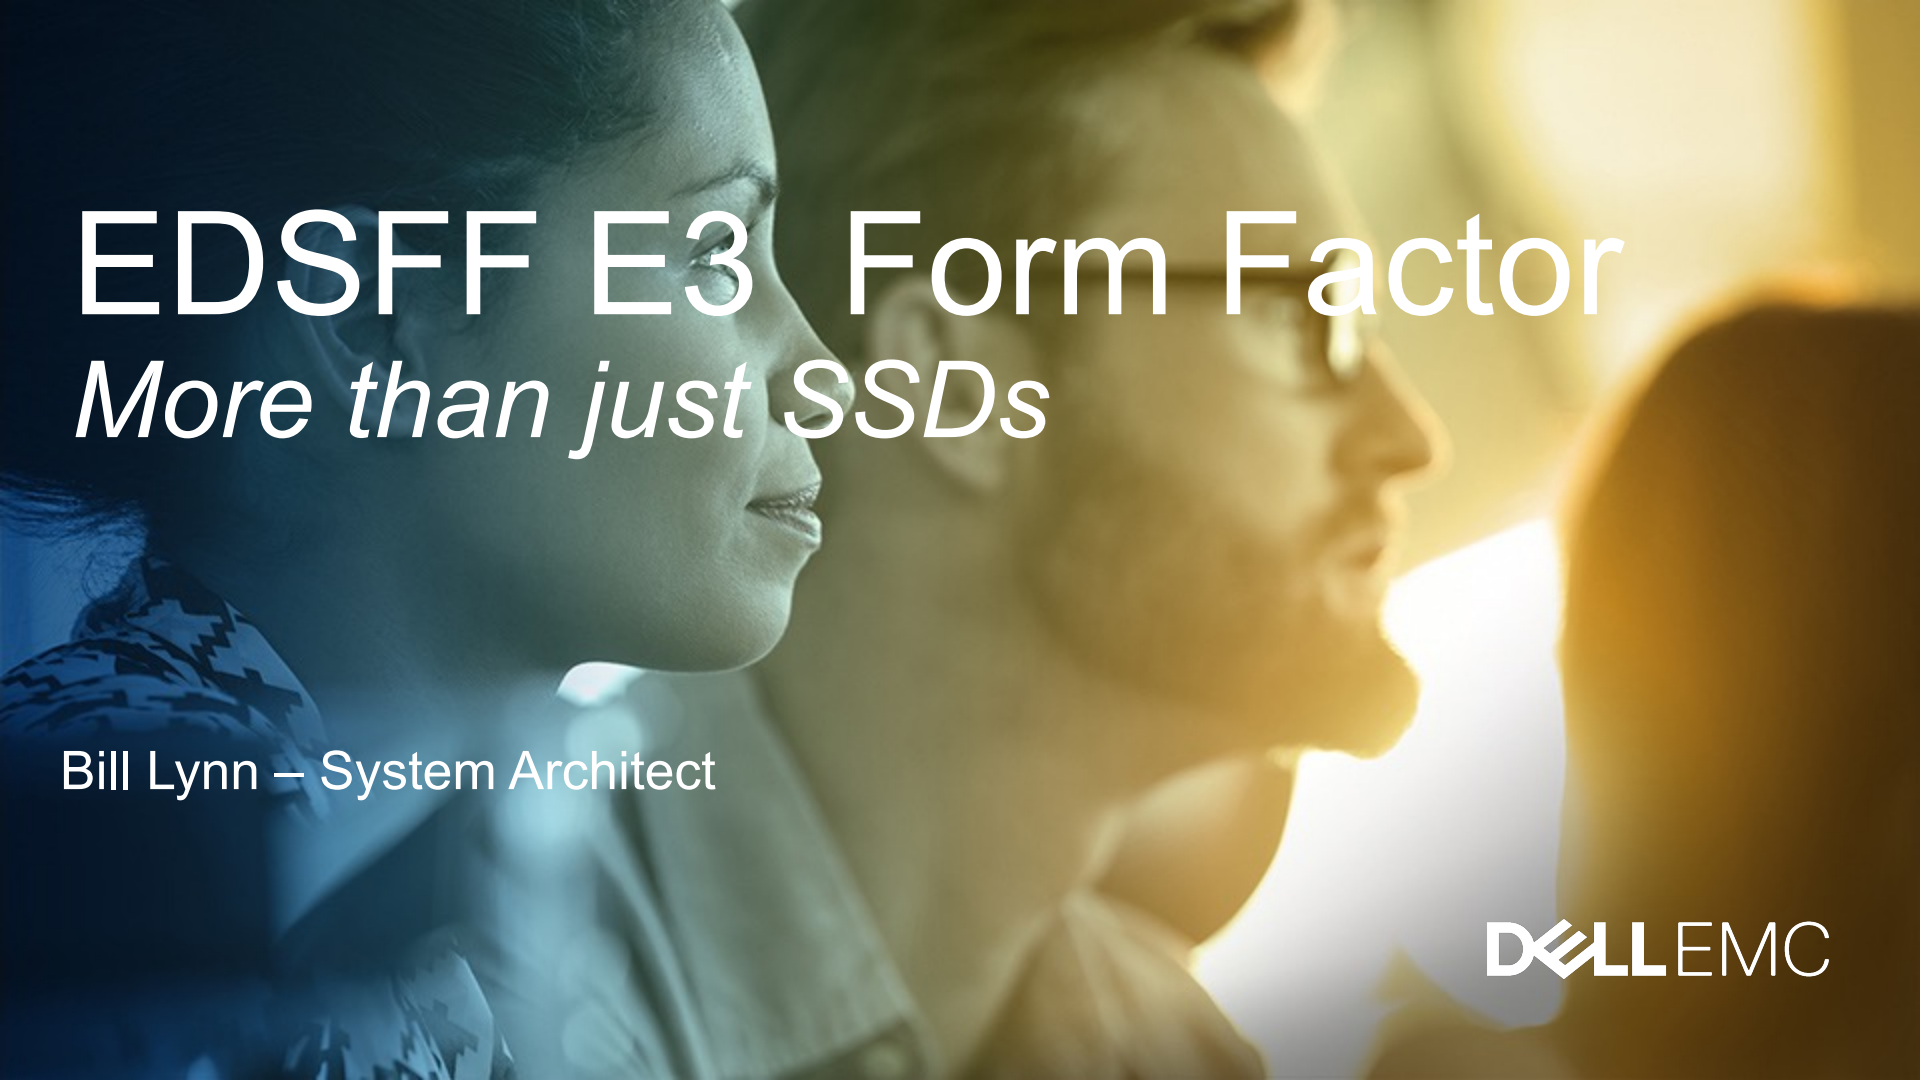

EDSFF E3 Form Factor

More than just SSDs

Bill Lynn – System Architect

DELL EMC

E3 - Flexible Form Factor

- Family of devices allows multiple device types with different power profiles
- Common device bay mechanics (smaller devices fit into larger device bays)
- Supports multiple link widths (x2, x4, x8, and x16)
- Requires a mechanical device carrier which allows for OEM ID customization

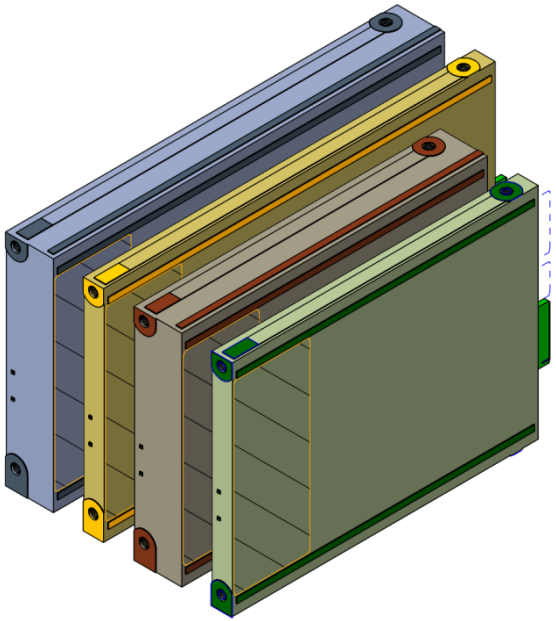

1U E3 Example Chassis Configuration

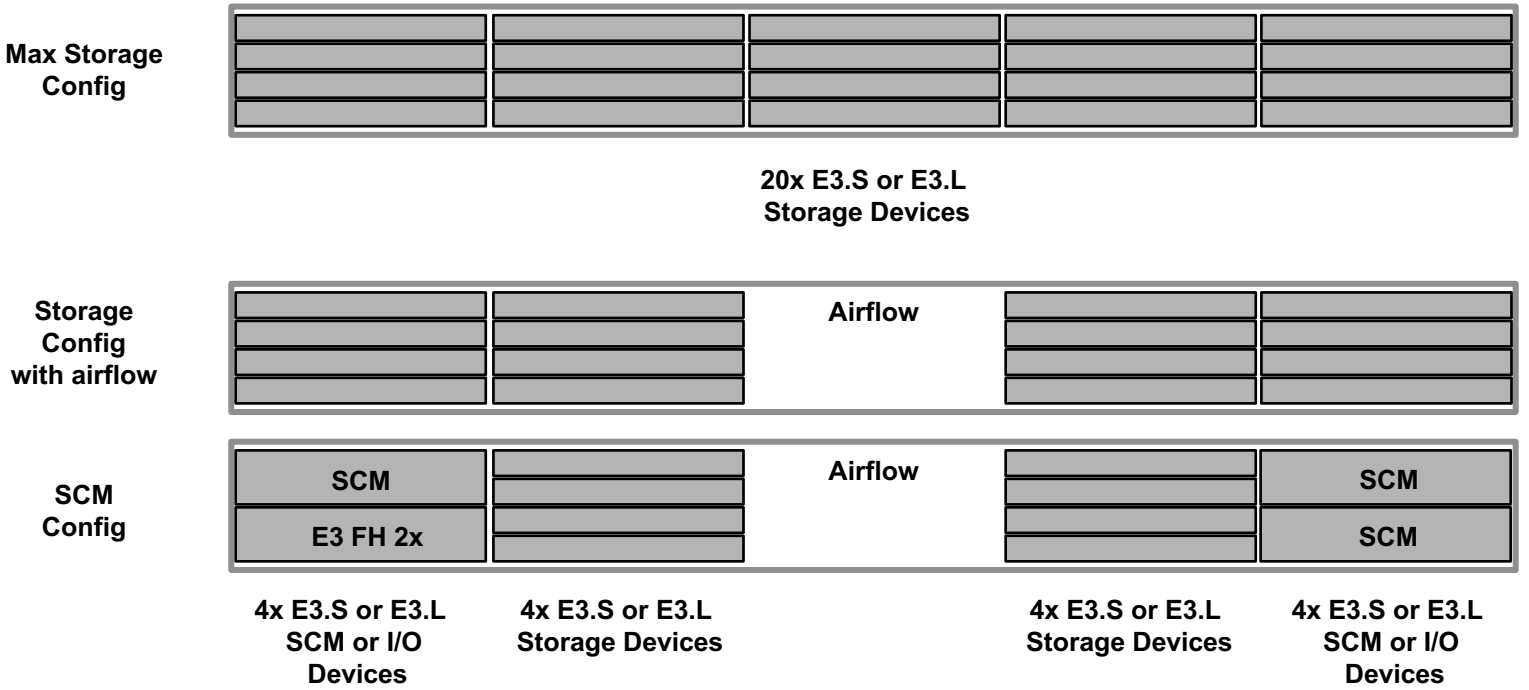

2U E3 Example Chassis Configuration

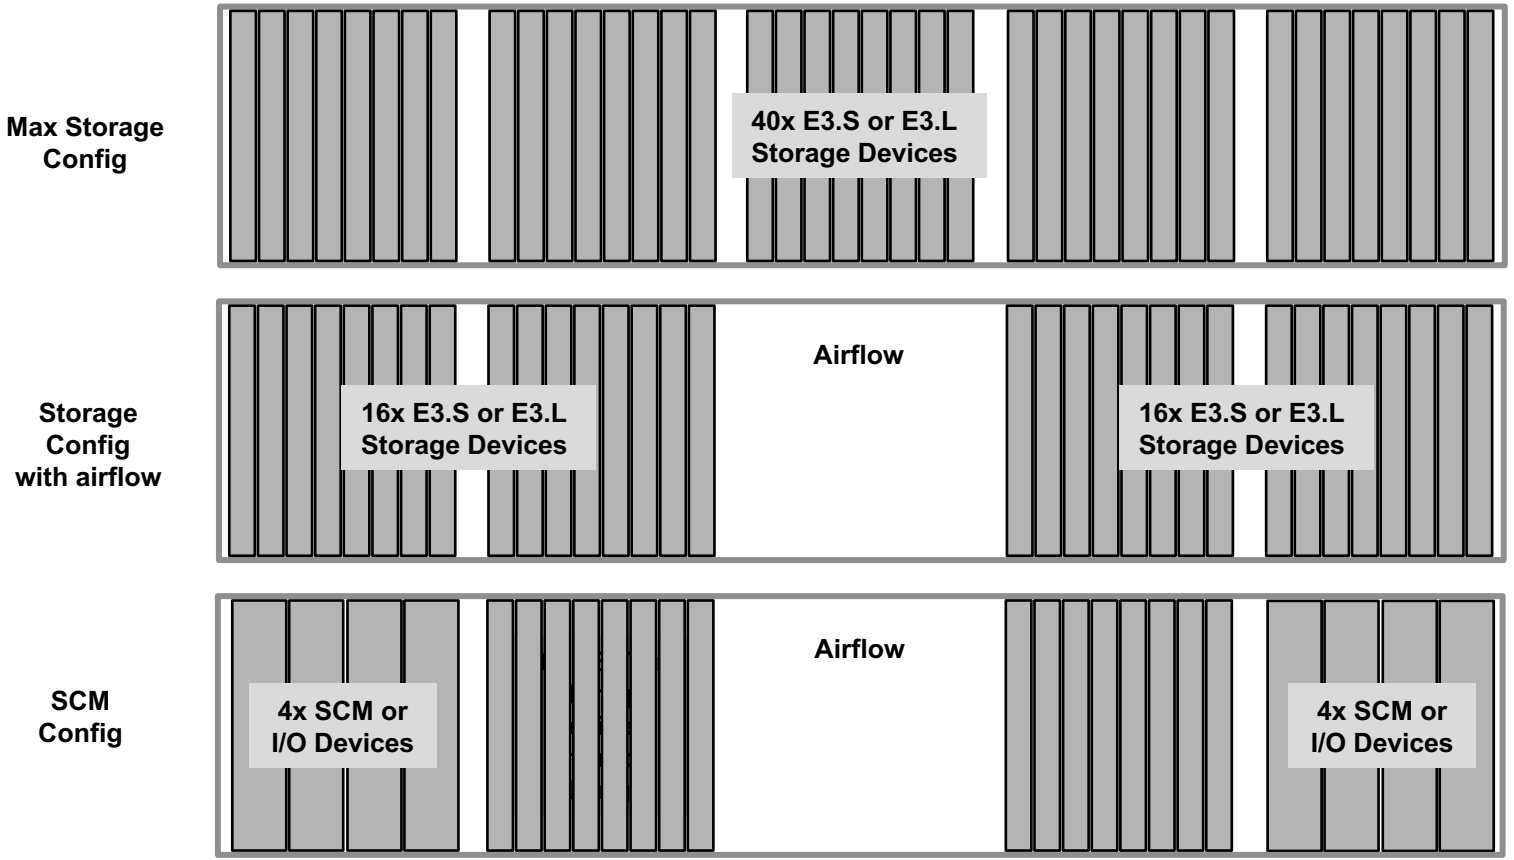

Future Device Types

- Moving the E3 connector to 19.54mm allows for the use of a 4C+ connector used by OCP 3 NIC
- Allows standard networking connectors in an E3 2T form factor
- Allows for potential future higher power devices

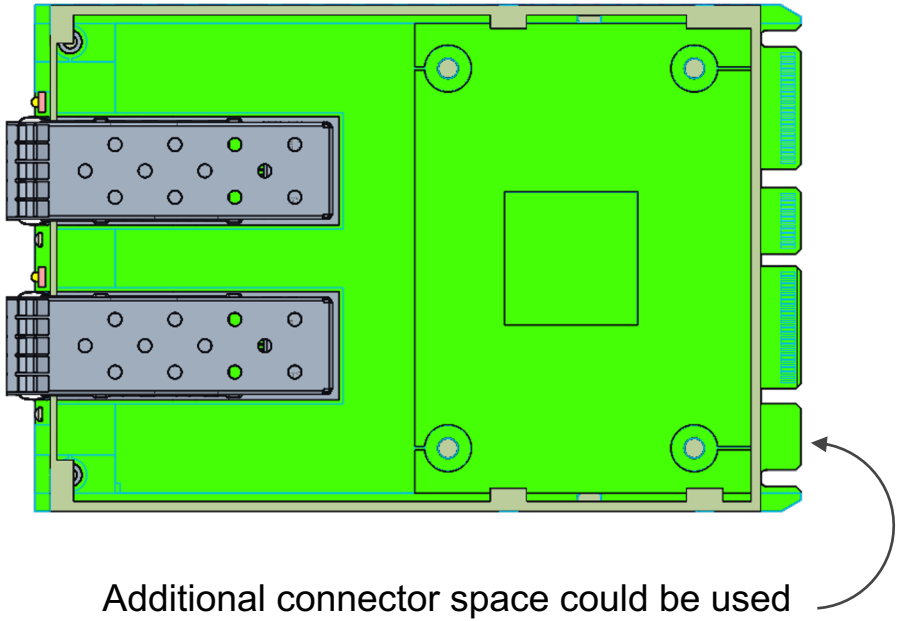

Additional connector space could be used for a 4C+ or a higher power connector tab

Smarter technology for all

Shipping now in a server near you: EDSFF

Jonathan Hinkle, Executive Director and Distinguished Researcher – System Architecture, Lenovo Enterprise and Cloud Research | April 2021

Promise of EDSFF (Enterprise and Datacenter Standard Form Factor)

Optimized for scaling solid-state devices in datacenter systems

- Family of cards leveraging same connector, pinout, behavior
- Better cooling, streamlined power delivery, modularity, density of drives
- Higher and scalable system capacity and performance to meet various workload needs
- Ready for future interfaces like CXL and new applications like memory, accelerators

Industry standard EDSFF connector

Improved airflow through connector

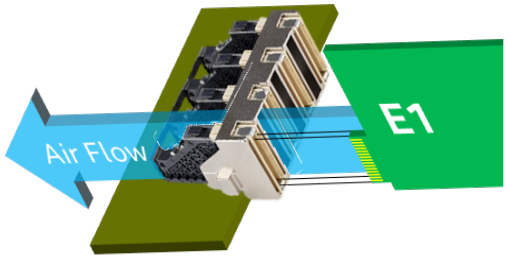

1U rack-space optimized E1.S

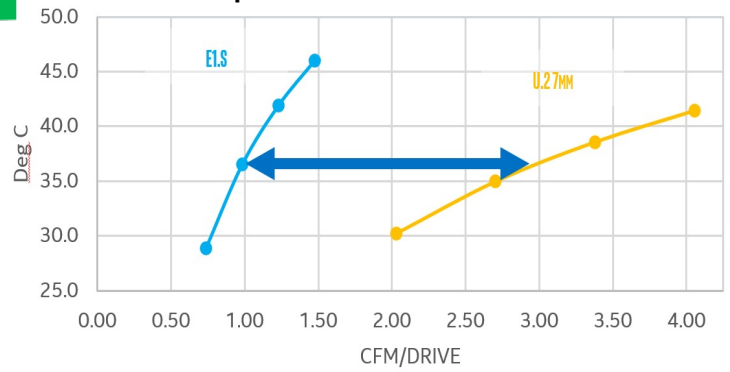

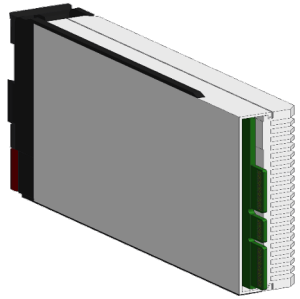

2U rack-space optimized E3.S thick

1U capacity-optimized E1.L

Promise of EDSFF – shipping now

4TB E1.S NVMe drive

Now Available: Lenovo's new Icelake servers leveraging the EDSFF E1.S

- 1U rack server with 16 x E1.S drives: higher system storage performance with lower power and lower cost mainstream NVMe drives
- Improved airflow and system cooling for overall better TCO

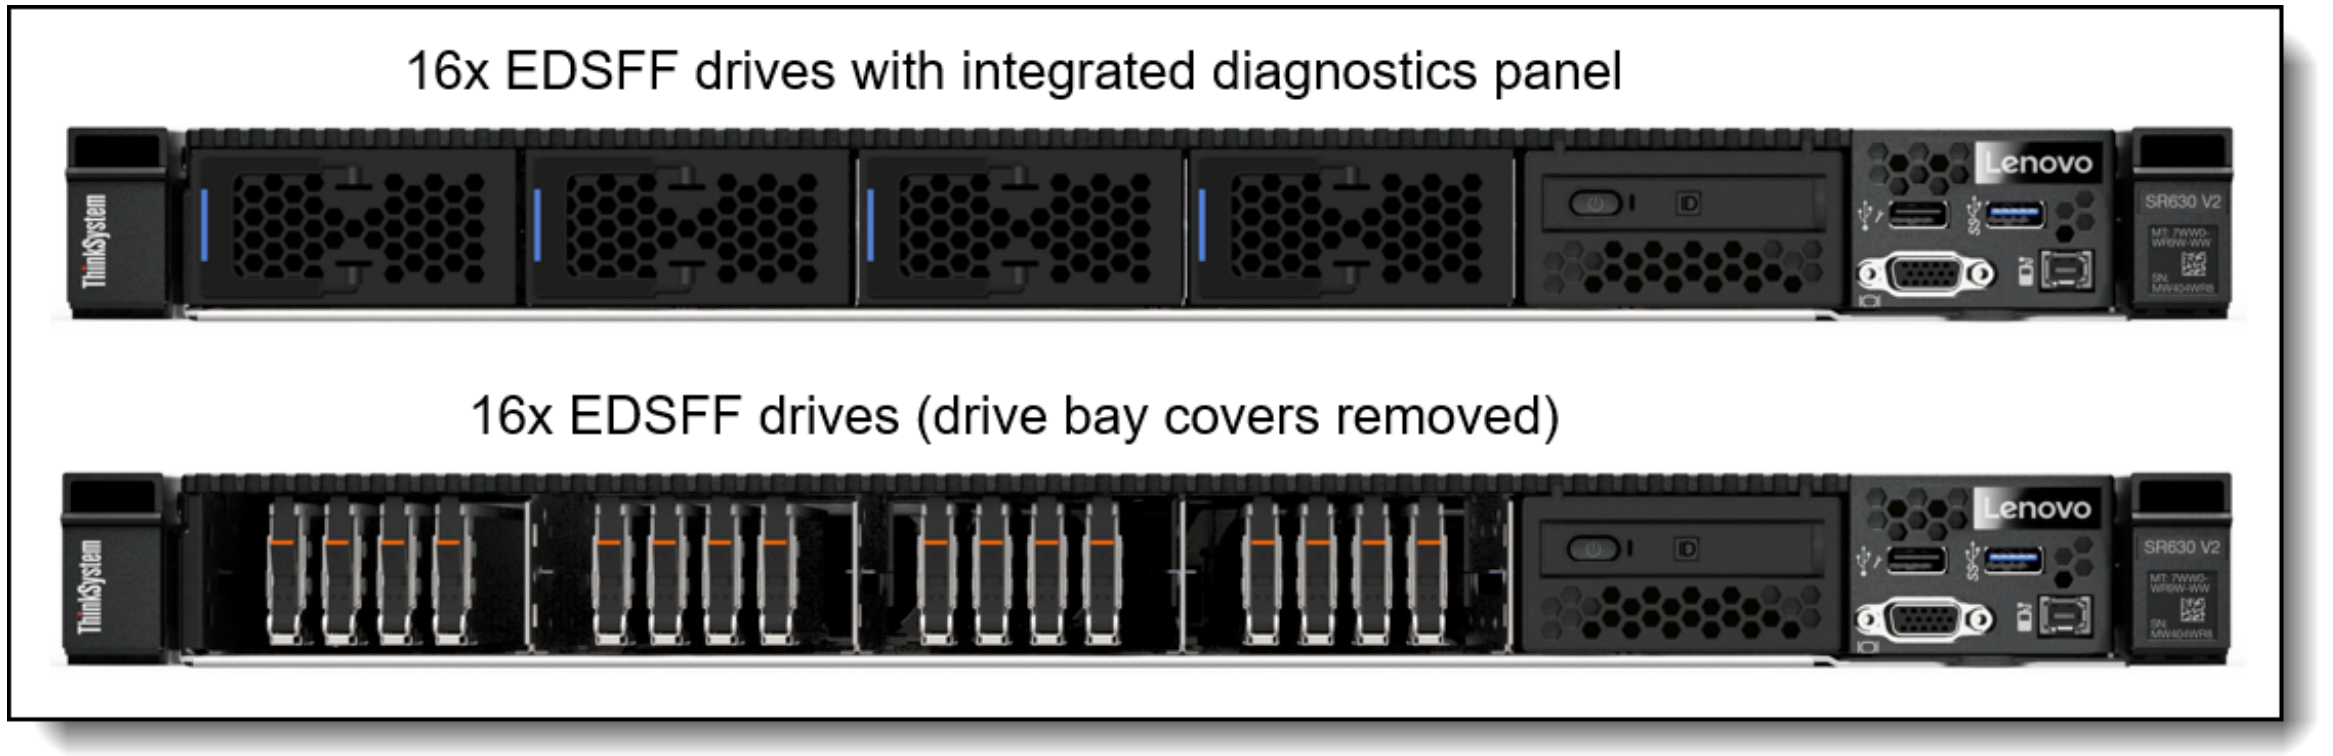

Lenovo SR630 V2: 2-socket high-volume 1U rack server

Promise of EDSFF – shipping now

EDSFF E1.S also enables high performance with 3X density, providing new value in computing systems

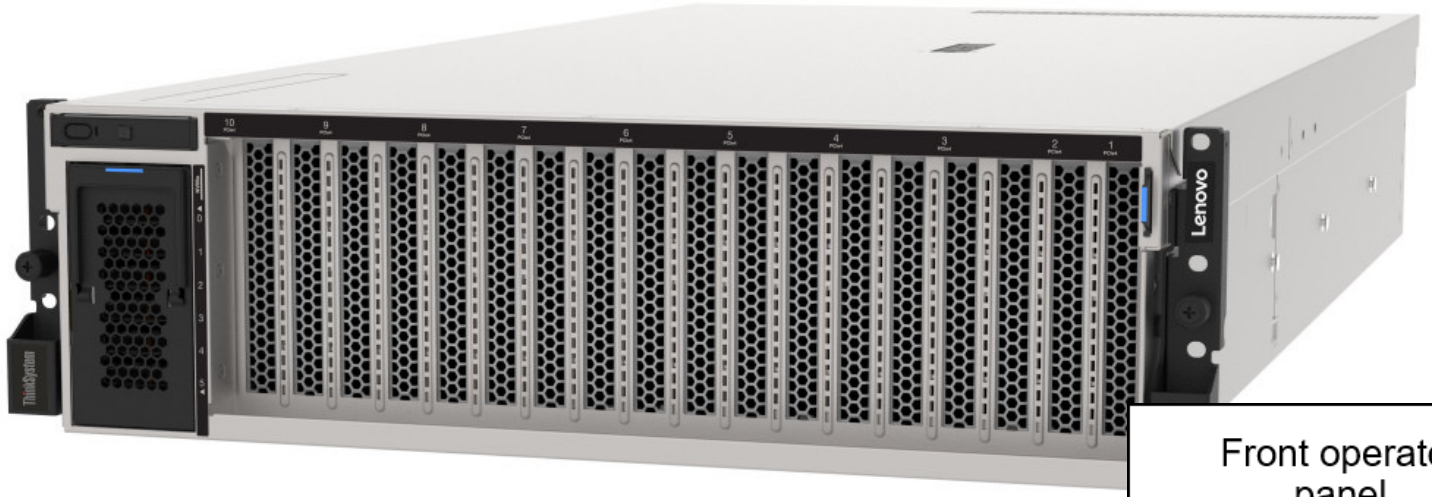

- Lenovo 3U SR670 V2 (dense) supports two CPUs and up to 8 double-wide adapter cards for applications like AI acceleration
- It can do this while *still* offering high performance storage with 6x E1.S NVMe drives.

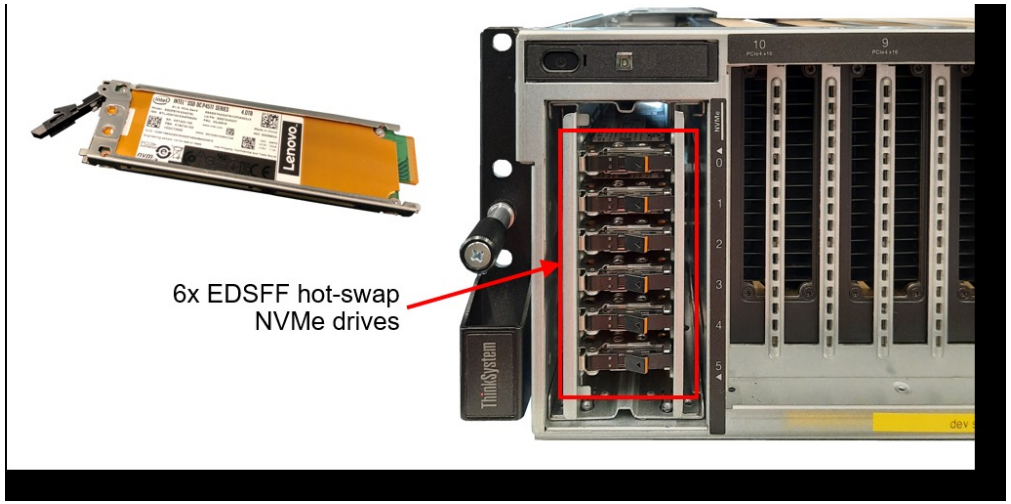

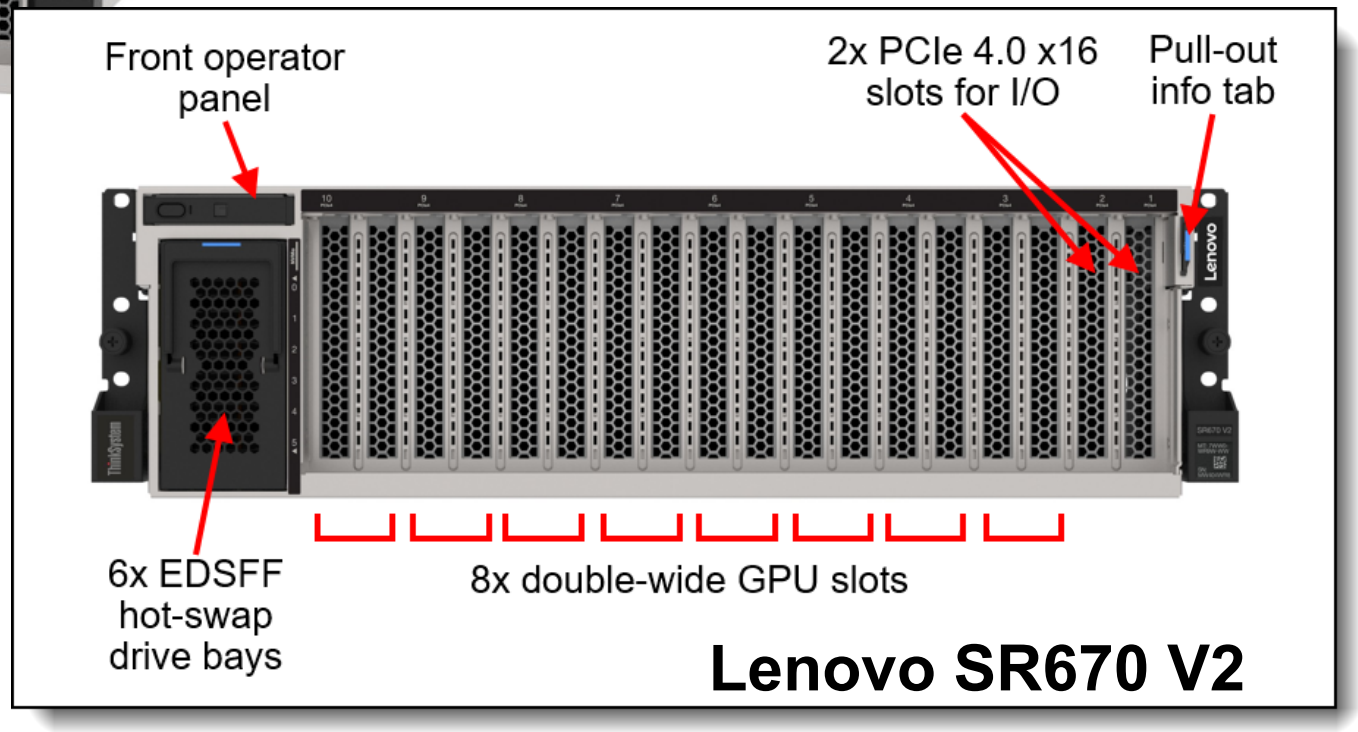

Promise of EDSFF – shipping now

Effective storage performance and capacity in every node

SN550 V2 Compute Sled: 2 x u.2 ➡ 6 x E1.S, enabling new density of integrated compute and storage

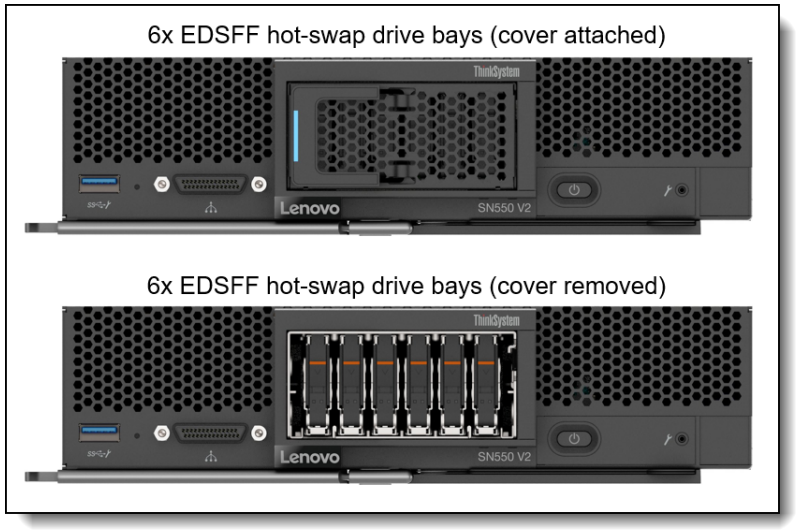

Lenovo SN550 V2

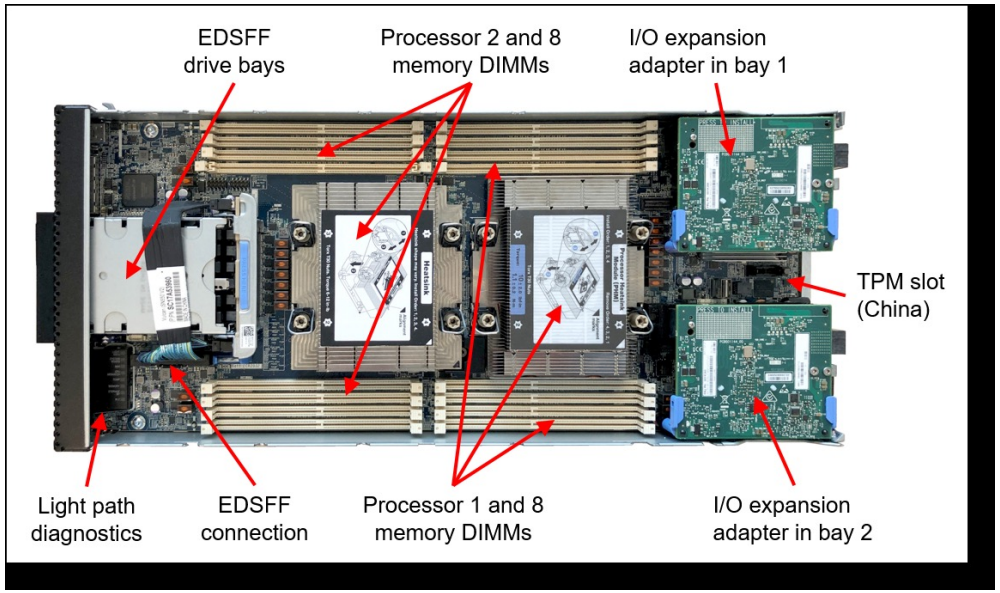

Lenovo Flex Chassis

thanks.

Smarter technology for all

Lenovo

EDSFF Overview

Jonmichael Hands, Sr. Strategic Planner

Intel Non-Volatile Memory Solutions Group

Legal Disclaimer

All product plans and roadmaps are subject to change without notice.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel technologies may require enabled hardware, software, or service activation.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

OC P NVMe SSD and EDSFF Workshop – April 27

- EDSFF overview
- Market data
- Hyperscale view
- OEM view
- SSD vendor view

OPEN
Compute Project ®

STORAGE

What is EDSFF?

- Enterprise and Data Center SSD Form Factor
- Improved thermals, power, and scalability
- High-speed common connector, pinout – scalable to faster speed PCIe
- Integrated serviceability, hot-plug support
- Built in LEDs, carrier-less design
- Customizable latch for toolless serviceability

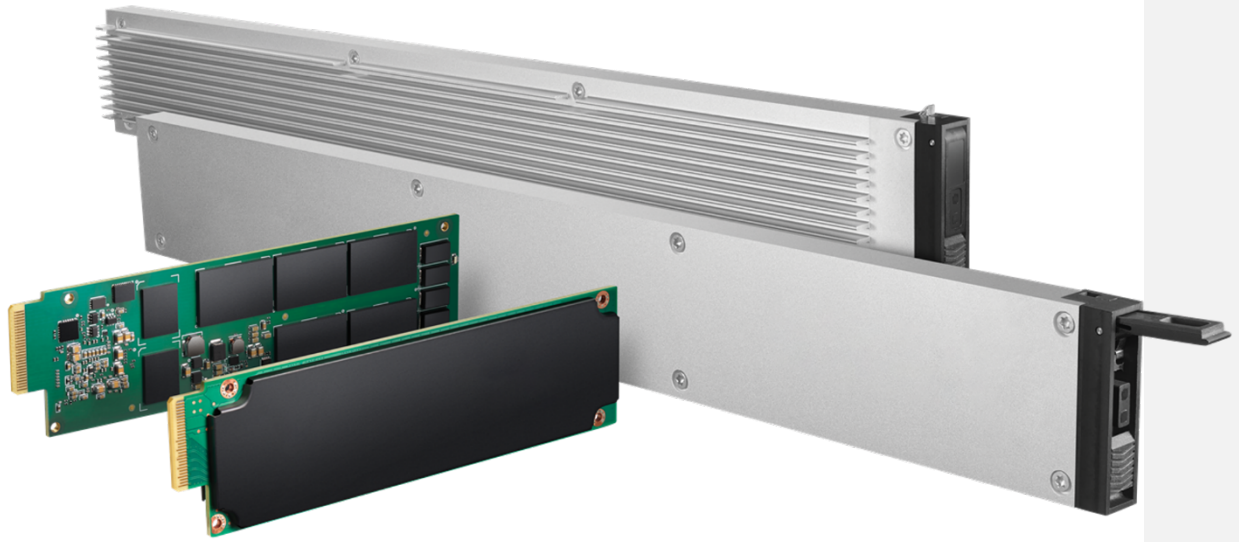

EDSFF History

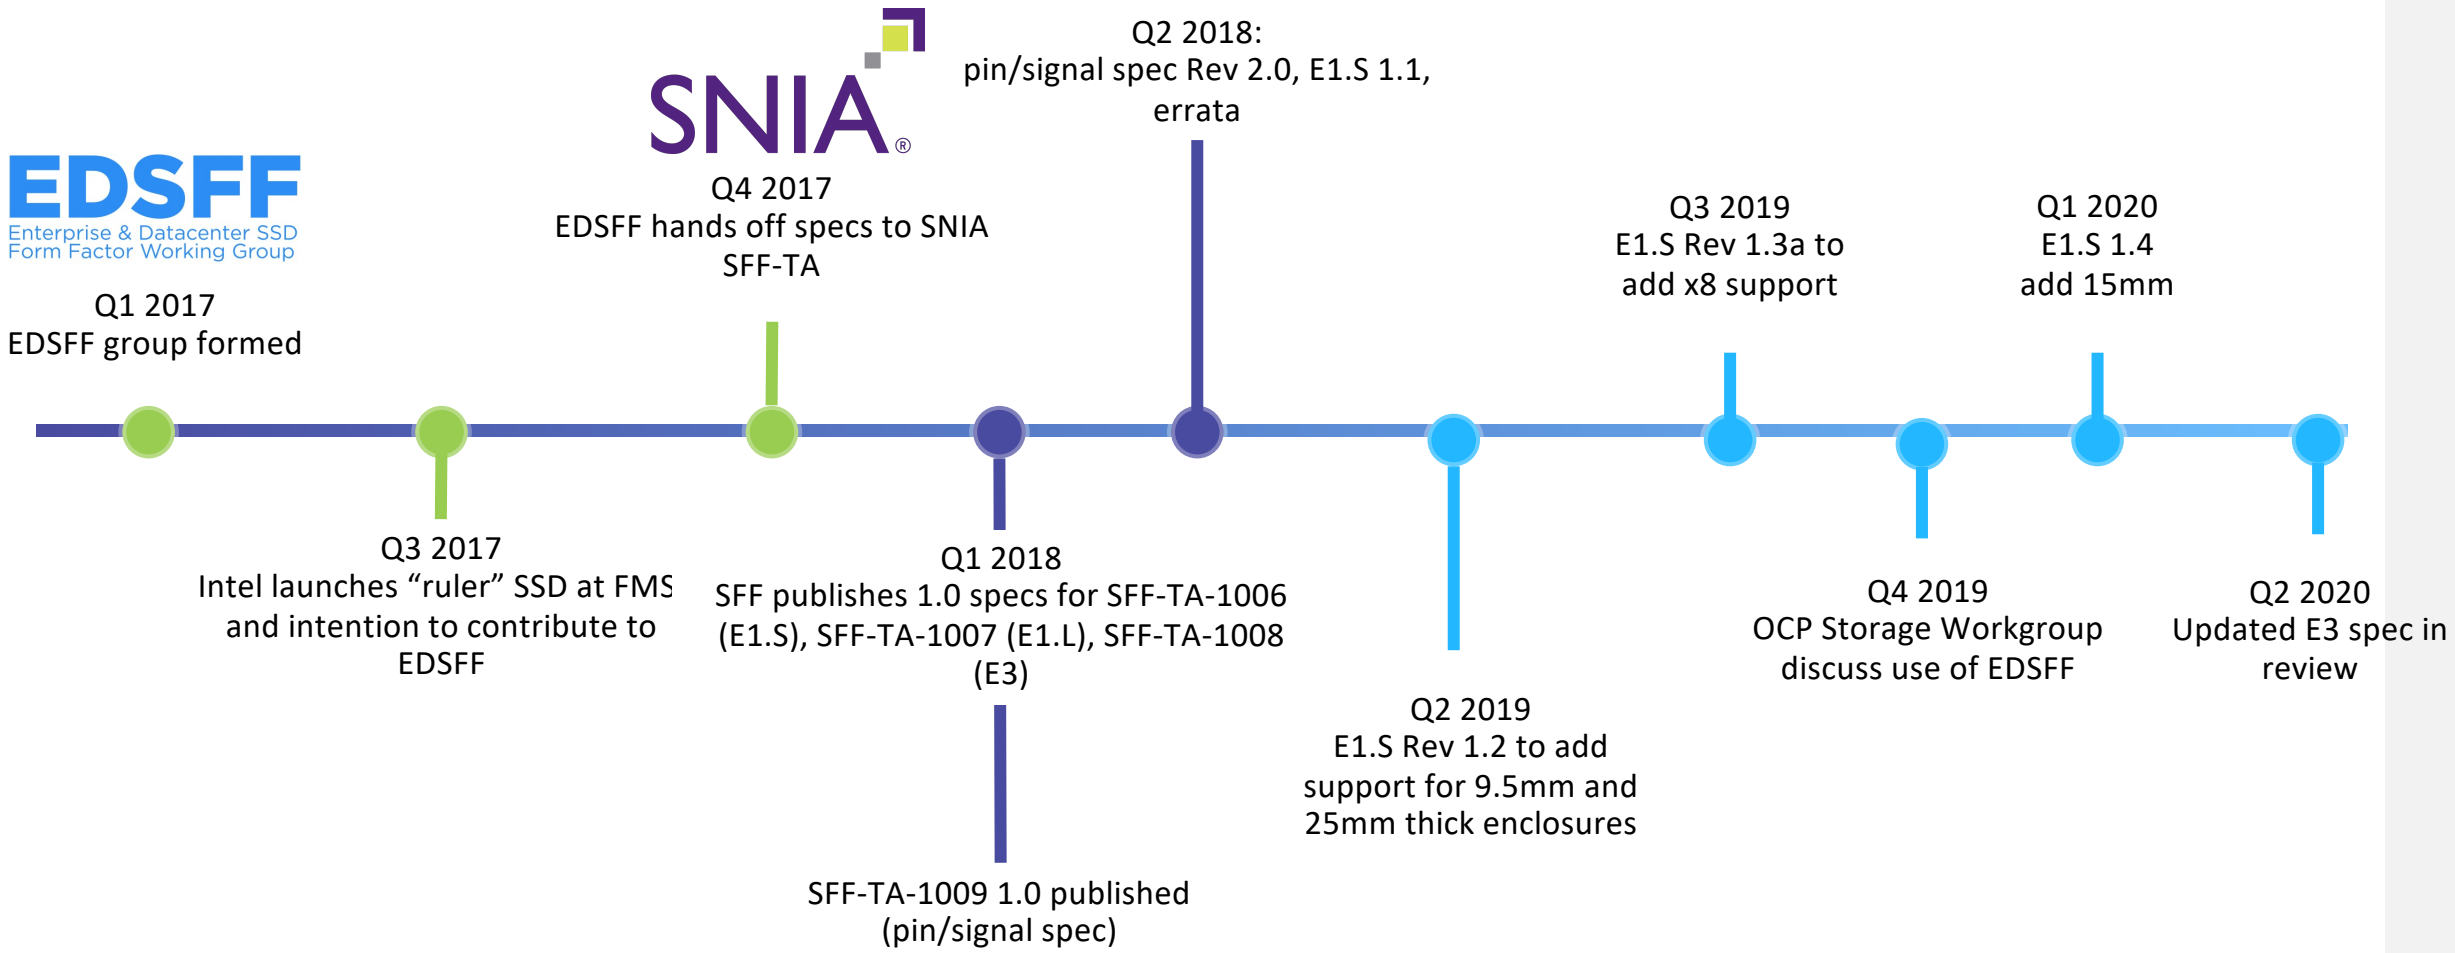

EDSFF Family

- Family of form factors and standards for data center NVMe SSDs
- E1.S for scalable & flexible performance storage
- E1.L for high capacity storage (e.g. QLC)
- E3 high performance SSD for 2U server / storage

E1.S

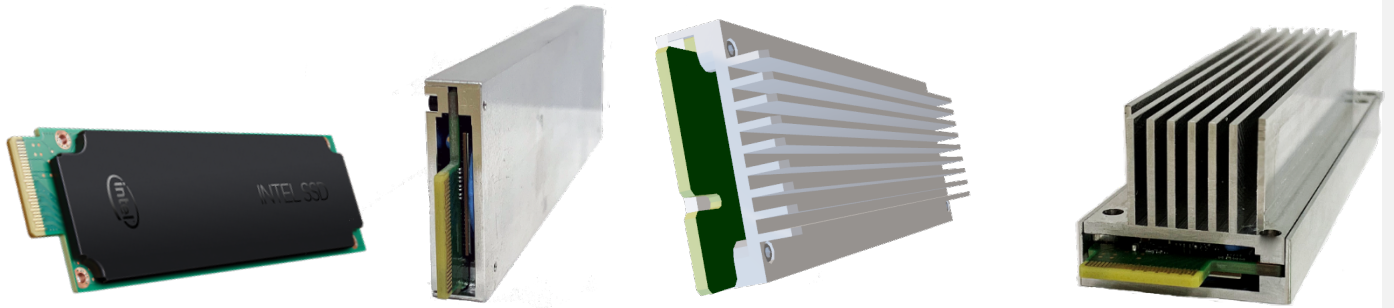

5.9mm 9.5mm 15mm 25mm

E1.L 9.5mm

E1.L 18mm

E3

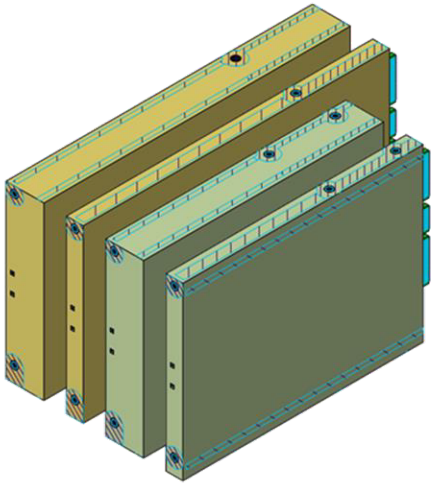

Intel Recommended Platform Design Guidance

	2U Server	1U Server	Storage/JBOF	Enterprise Storage Array	Boot
(PCIe 4.0)	U.2 15mm	OEM: U.2 Hyperscale: U2 or E1.S	E1.L or U.2	U.2 Dual Port	M.2
(PCIe 4.0 → 5.0)	U.2 & E3.S	OEM: U.2/E1.S/E3 Hyperscale: E1.S	E1.L	U.2 Dual Port	M.2
(PCIe 5.0)	E3.S	E1.S & E3.S	E1.L	E3.S or E3.L	M.2 → E1.S

PCIe 4.0 →5.0. U.2 & M.2 → EDSFF

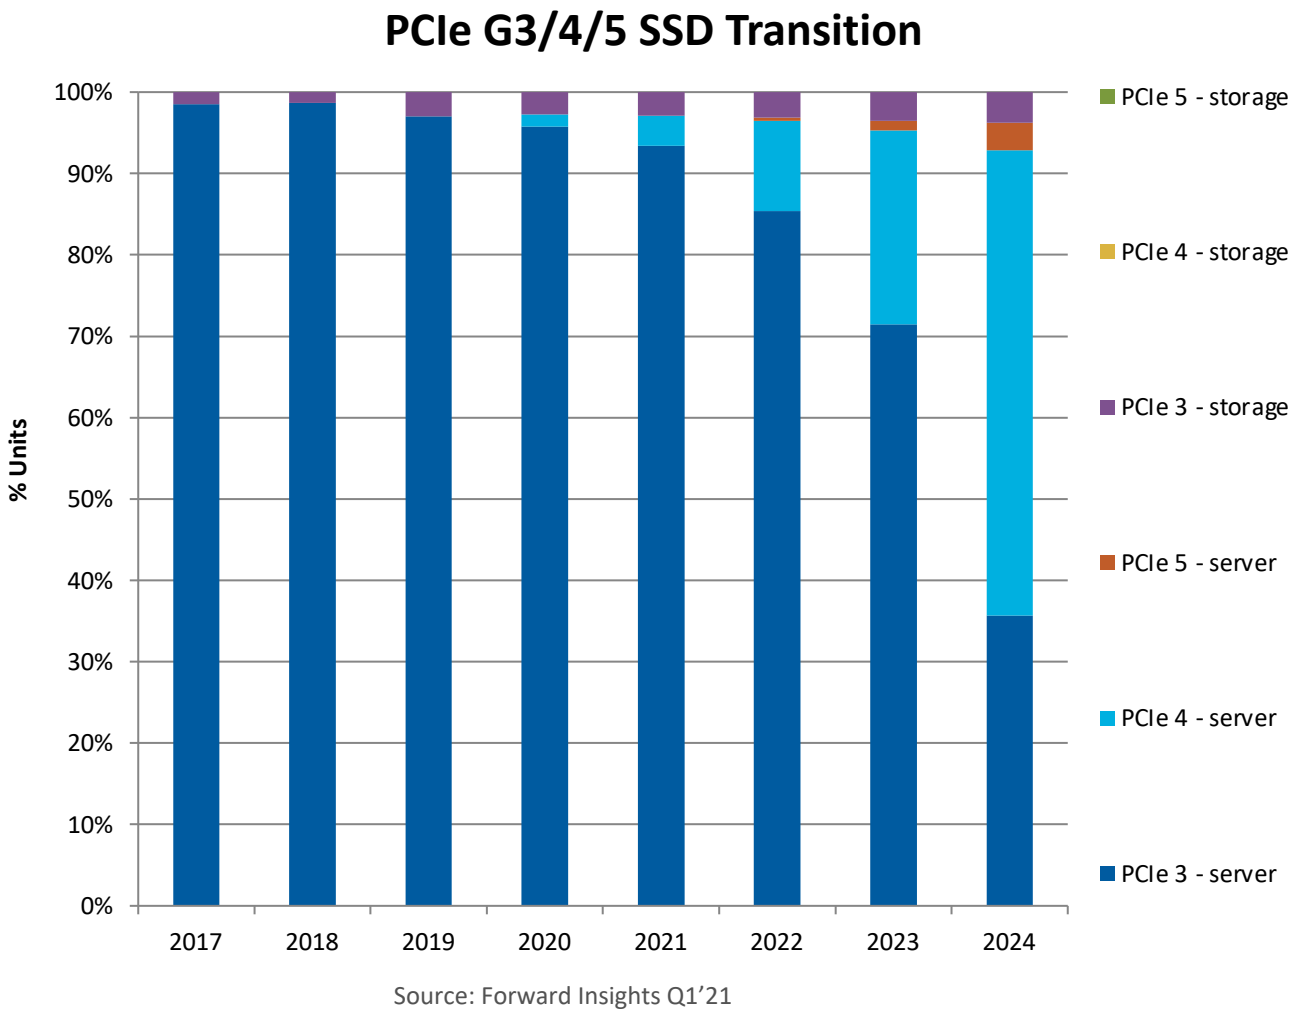

Source: Forward Insights Q1'21

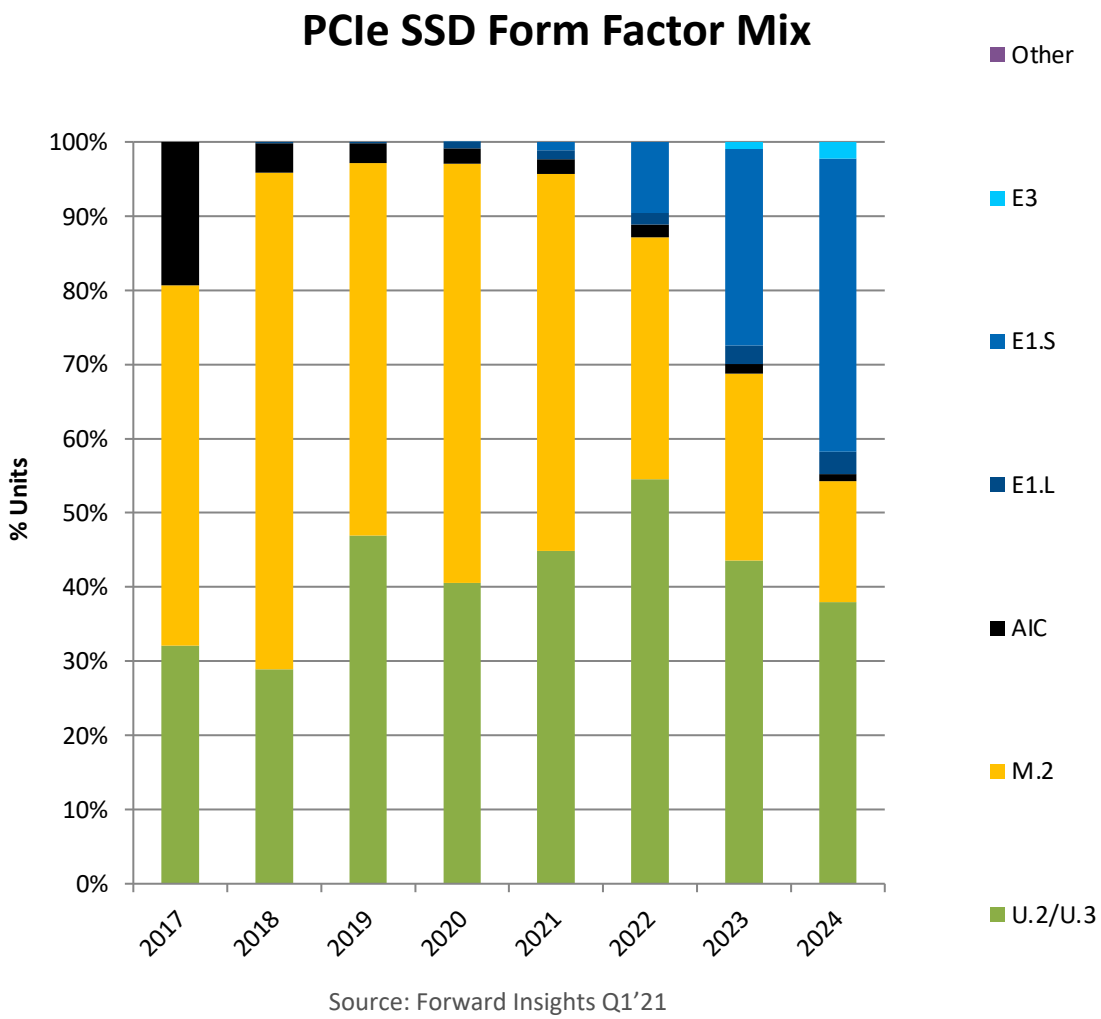

Source: Forward Insights Q1'21

Example - Dense NVMe U.2 Server

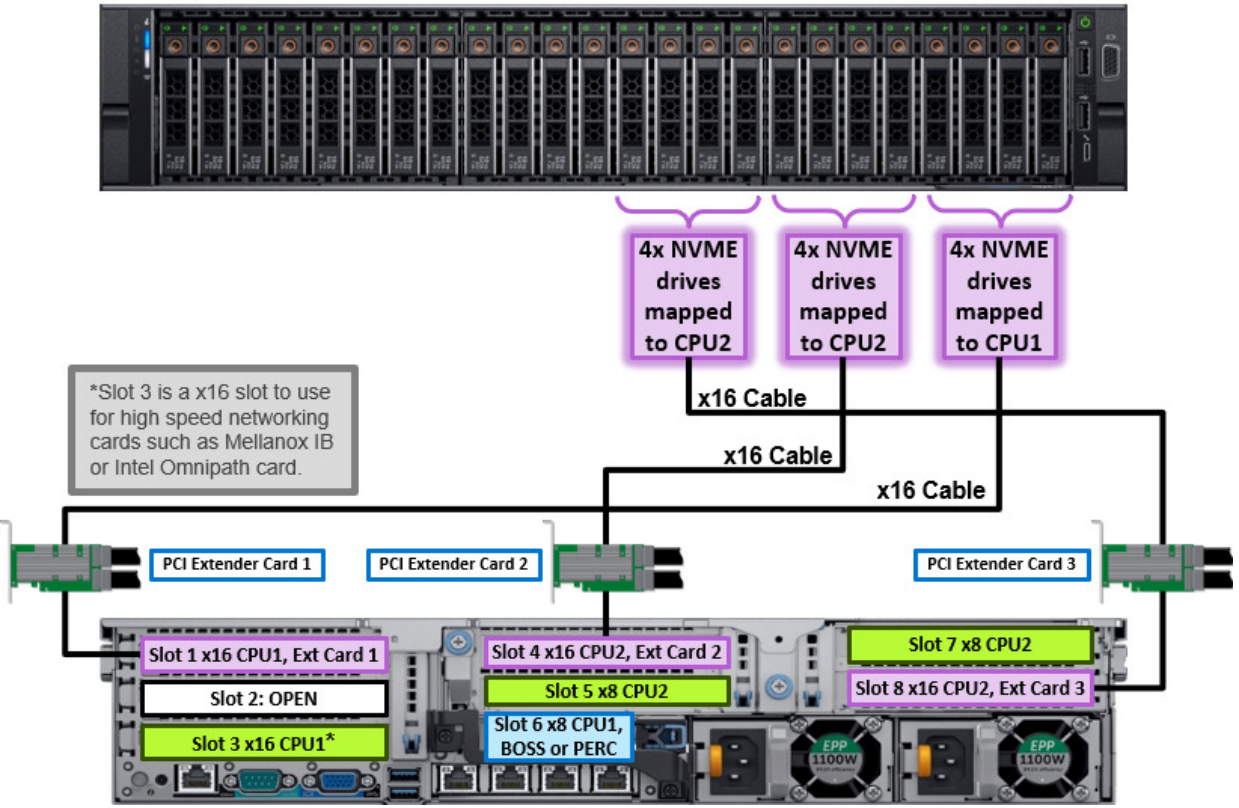

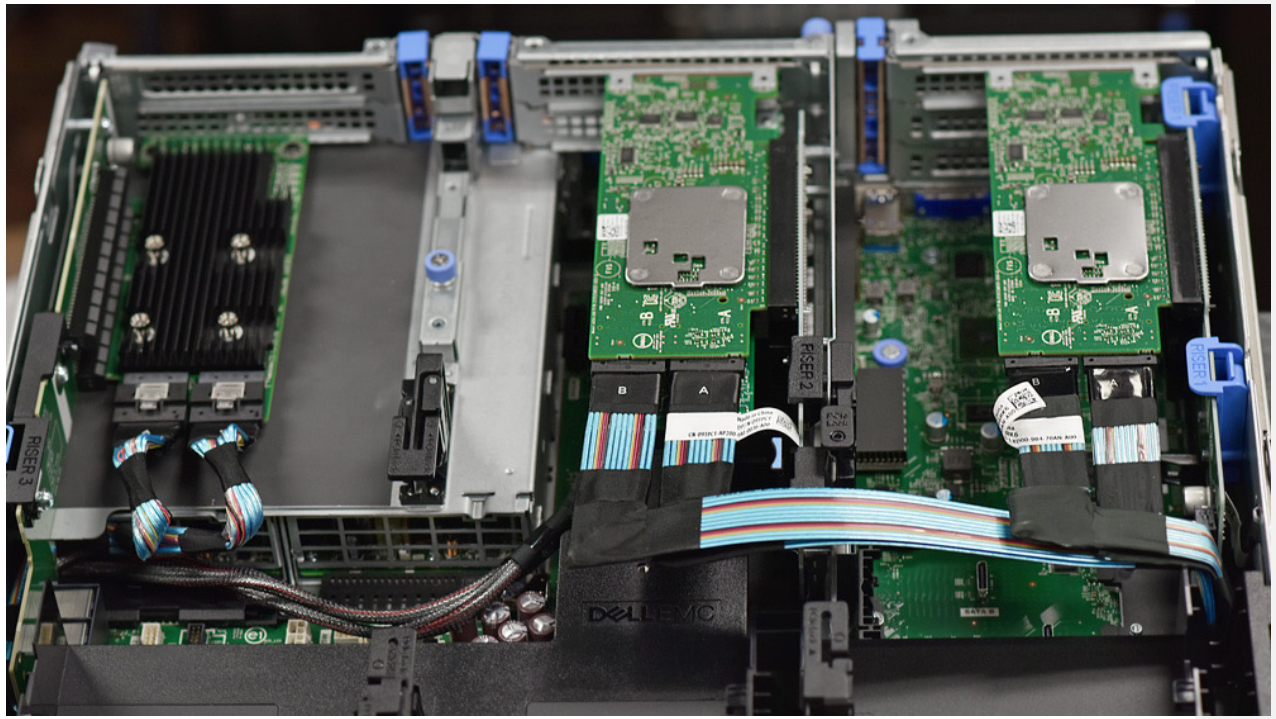

Source: https://www.storagereview.com/review/dell-emc-poweredge-r740xd-nvme-server-review

E1.L Storage Reimagined.

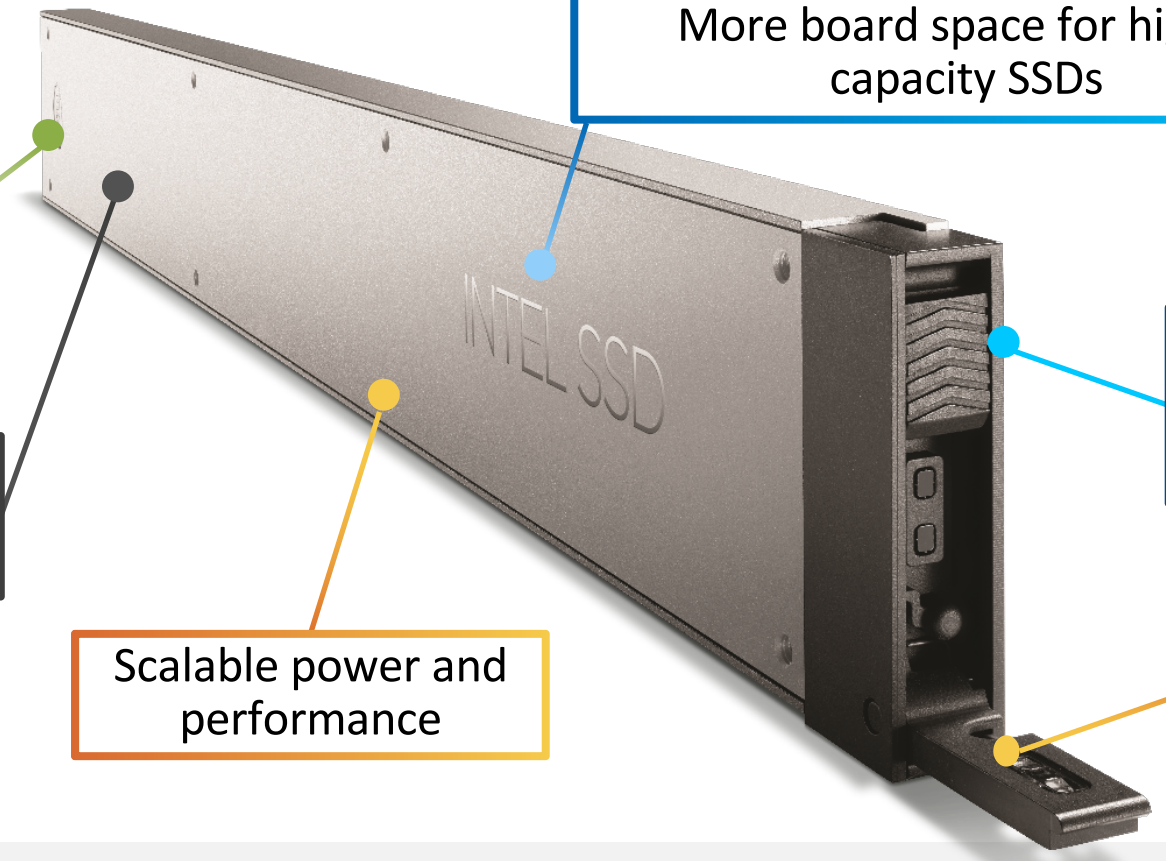

More board space for higher capacity SSDs

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

Carrier-less design with integrated latch and LEDs

Scalable power and performance

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 optimized for capacity storage

System Level Design Principles	Benefit
Integrated drive enclosure, customizable LEDs, and latch for servicing	Better serviceability, manageability, and removes cost from server
Fully passive backplane	Reduced complexity increases reliability, lowers cost
Optimized PCIe routing - No cables, no PCIe AICs to connect SSDs	Lower cost attach per drive, better signal integrity
Simplified enclosure management	For servicing drives and reducing downtime

E1.L Server and JBOF

E1.L optimized for TB/rack unit & performance

E1.L optimized for lowest TCO on QLC NVMe

E1.S – power and thermal options

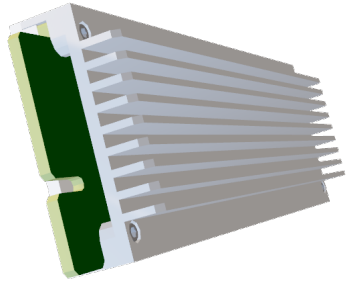

Enclosure Parameter	5.9mm Device	Device with Heat Spreader (8.01mm)	Device with Symmetric Enclosure (9.5mm)	Device with Asymmetric Enclosure (15mm)	Device with Asymmetric Enclosure (25mm)
Recommended sustained power (W)	12	16	20	20	25
Enclosure Max Inlet air temperature, 950 m to 3050 m (° C)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)
Add in card to add in card pitch (mm)	9	11	13	17	26
Recommended Fan Pressure loss across device (Pascal)	83	52	64	40	21
Airflow, average min per device (CFM). 1 CFM = 1.7 m3/h)	1.41 – (0.01 CFM for every 1° C below 35° C inlet temp)	1.71 – (0.06 CFM for every 1° C below 35° C inlet temp)	2.02 - (0.02 CFM for every 1° C below 35° C inlet temp)	1.5 - (0.02 CFM for every 1° C below 35° C inlet temp)	4.10 - (0.04 CFM for every 1° C below 35° C inlet

E1.S Optimal for 1U Performance Scalability

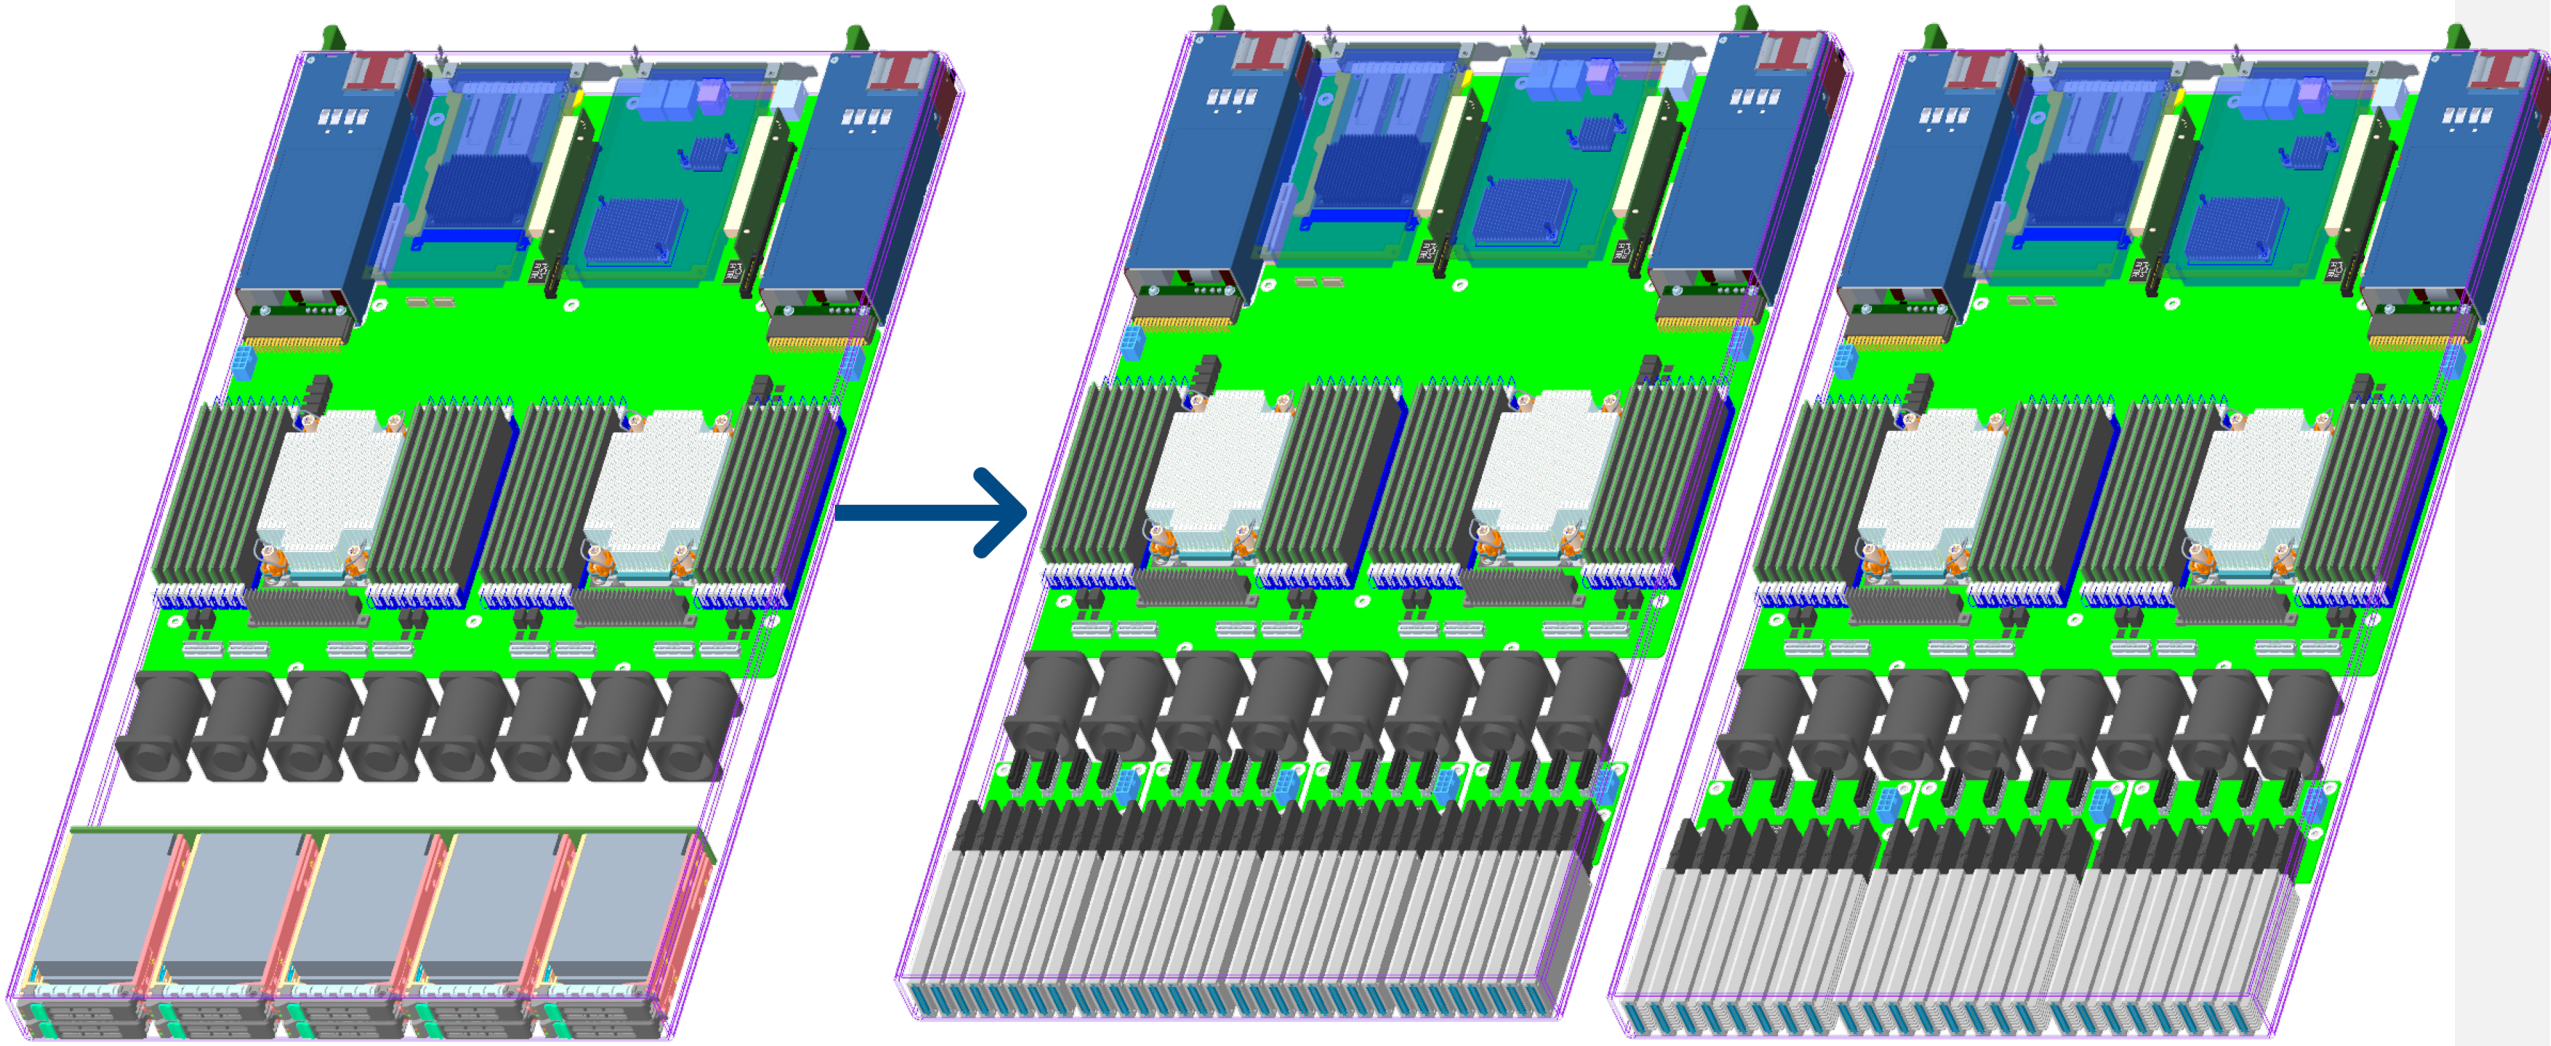

10 x U.2 SSD

32 x E1.S 9.5 mm (Horizontal BP)

24 x E1.S 15 mm (Horizontal BP)

Baseline

Max Fit

Max Fit

EDSFF E3 for Dummies

- E3 is a family of four form factors with a common 76mm height

  - E3.S
  - 76mm x 112.75mm x 7.5mm
  - Target to support from 20W to 25W
  - Optimized for primary NAND storage in Servers
  - E3.S, 2x
  - 76mm x 112.75mm x 16.8mm
  - Target to support from 35W to 40W
  - Support for higher power devices like CXL based SCM
  - E3.L
  - 76mm x 142.2mm x 7.5mm
  - Target to support up to 40W
  - Support for higher capacity NAND storage
  - E3.L, 2x
  - 76mm x 142.2mm x 16.8mm
  - Target to support up to 70W
  - Support for higher power devices like FPGAs and accelerators

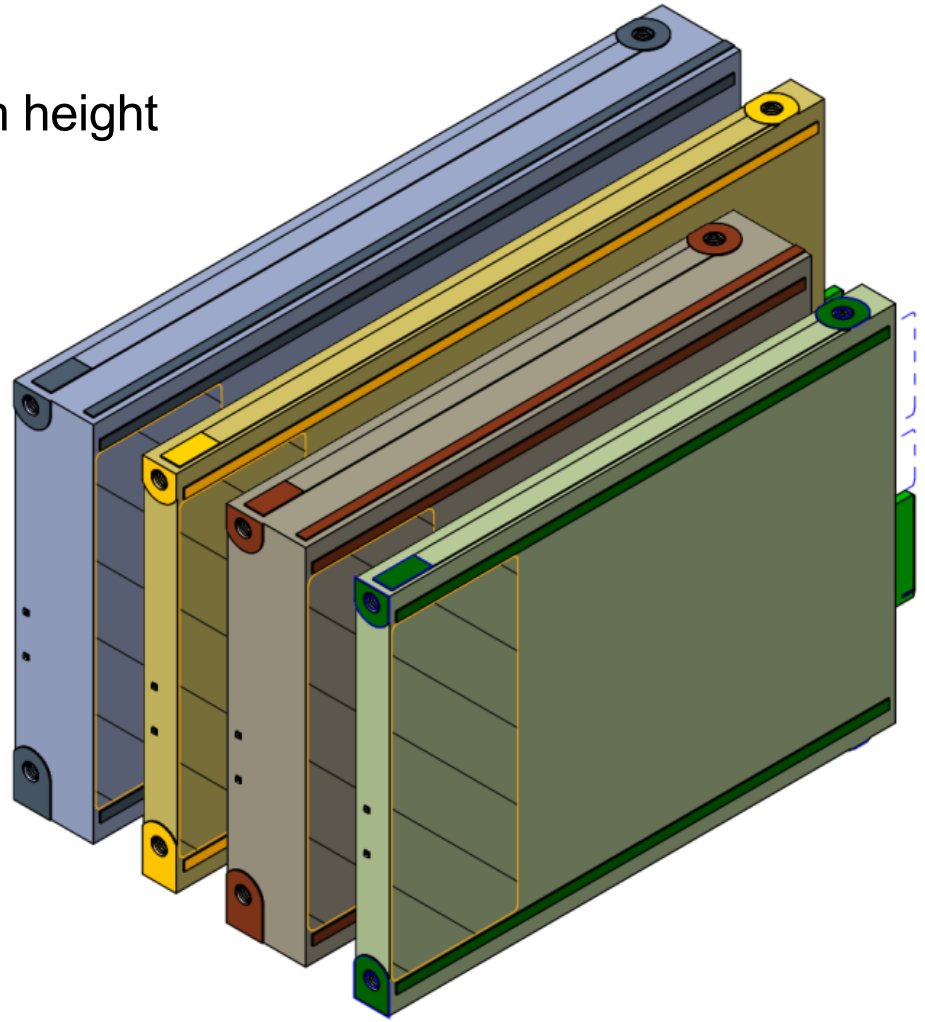

Note* - A thick device will fit into two thin slots

- A short device will fit into a long slot

2U 2S Spread – E3.S x 7.5 & 16.8

Mechanical Fit Study

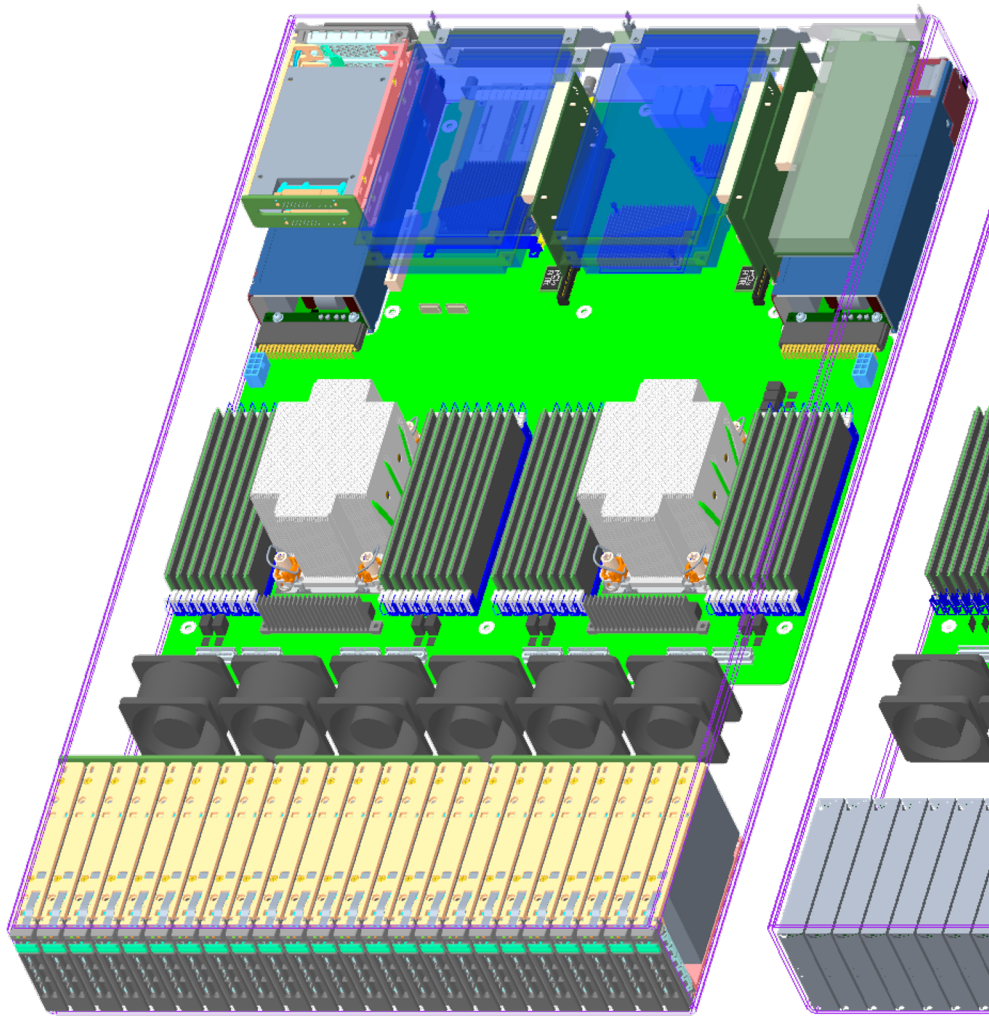

24 x U.2 SSD

Baseline

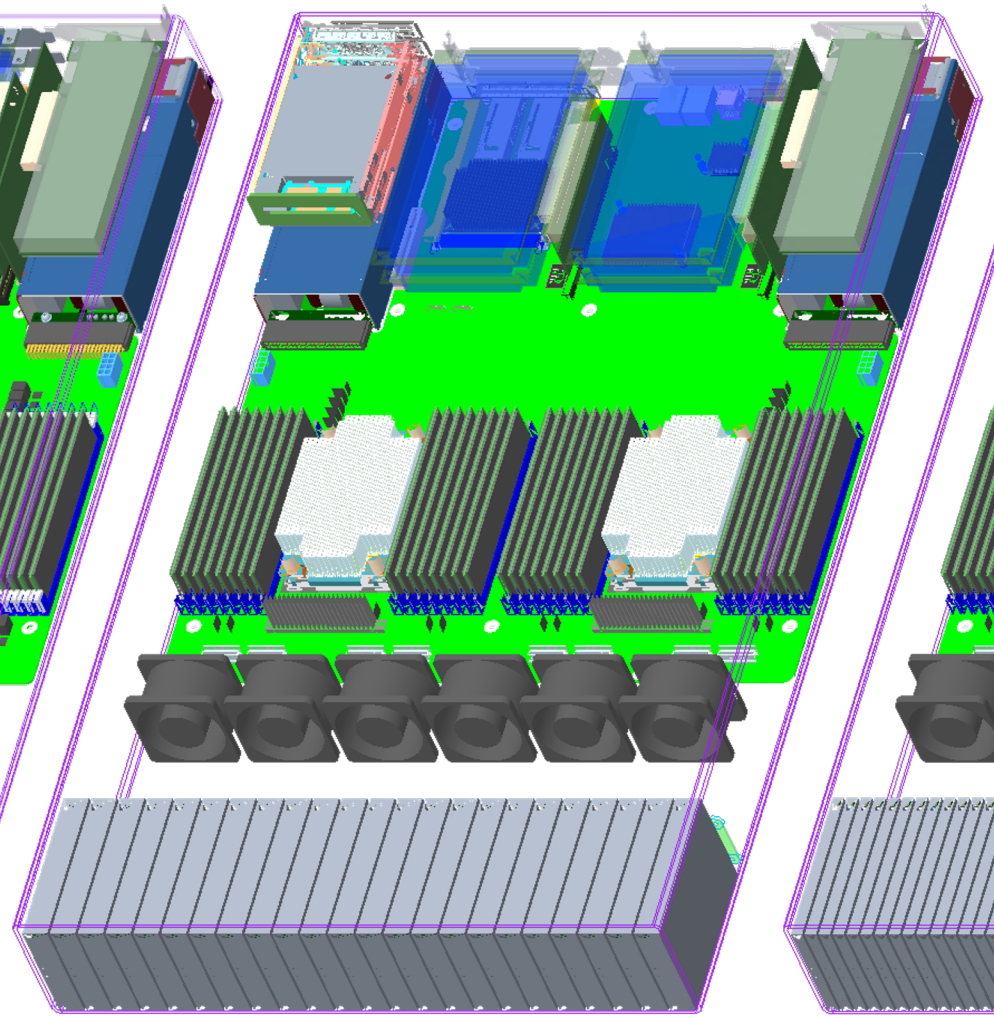

23 x E3 Short/Thick

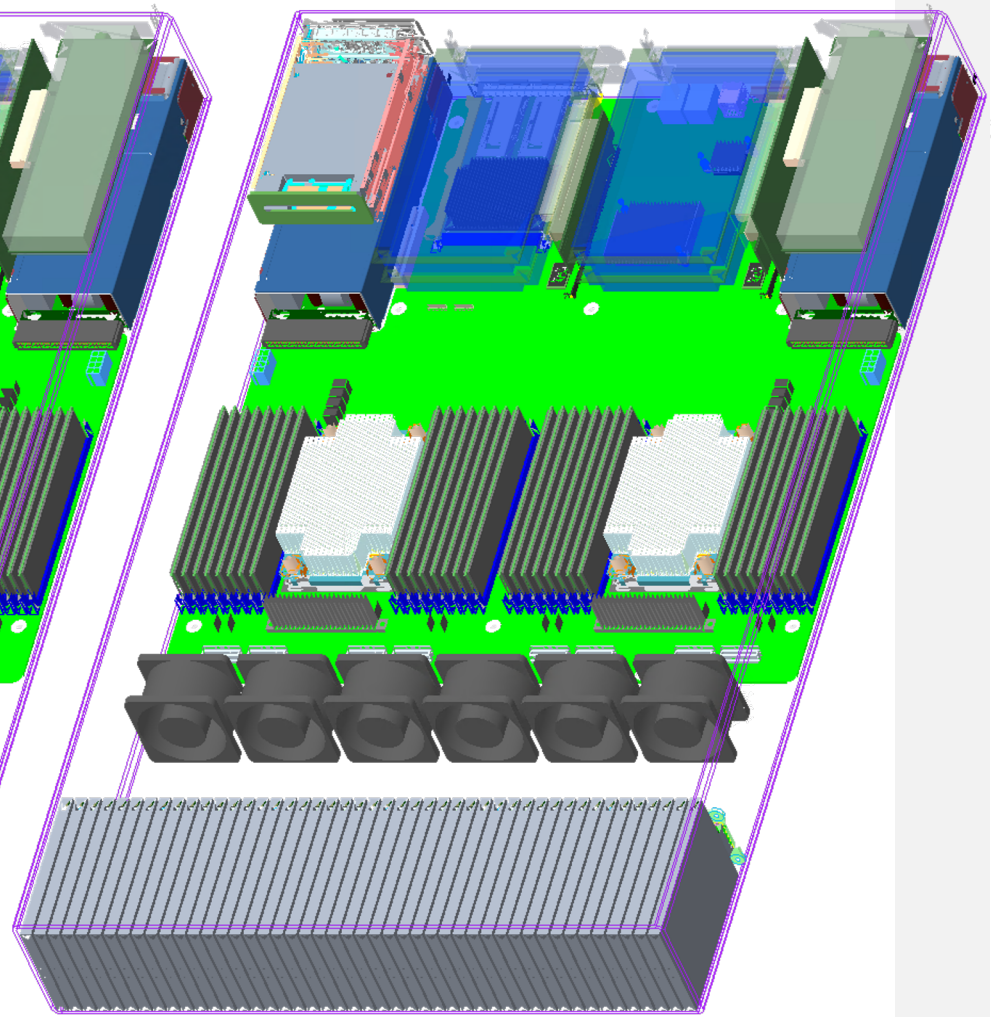

46 x E3 Short/Thin

1U 2S Spread – E3.s x 7.5 & 16.8

Mechanical Fit Study

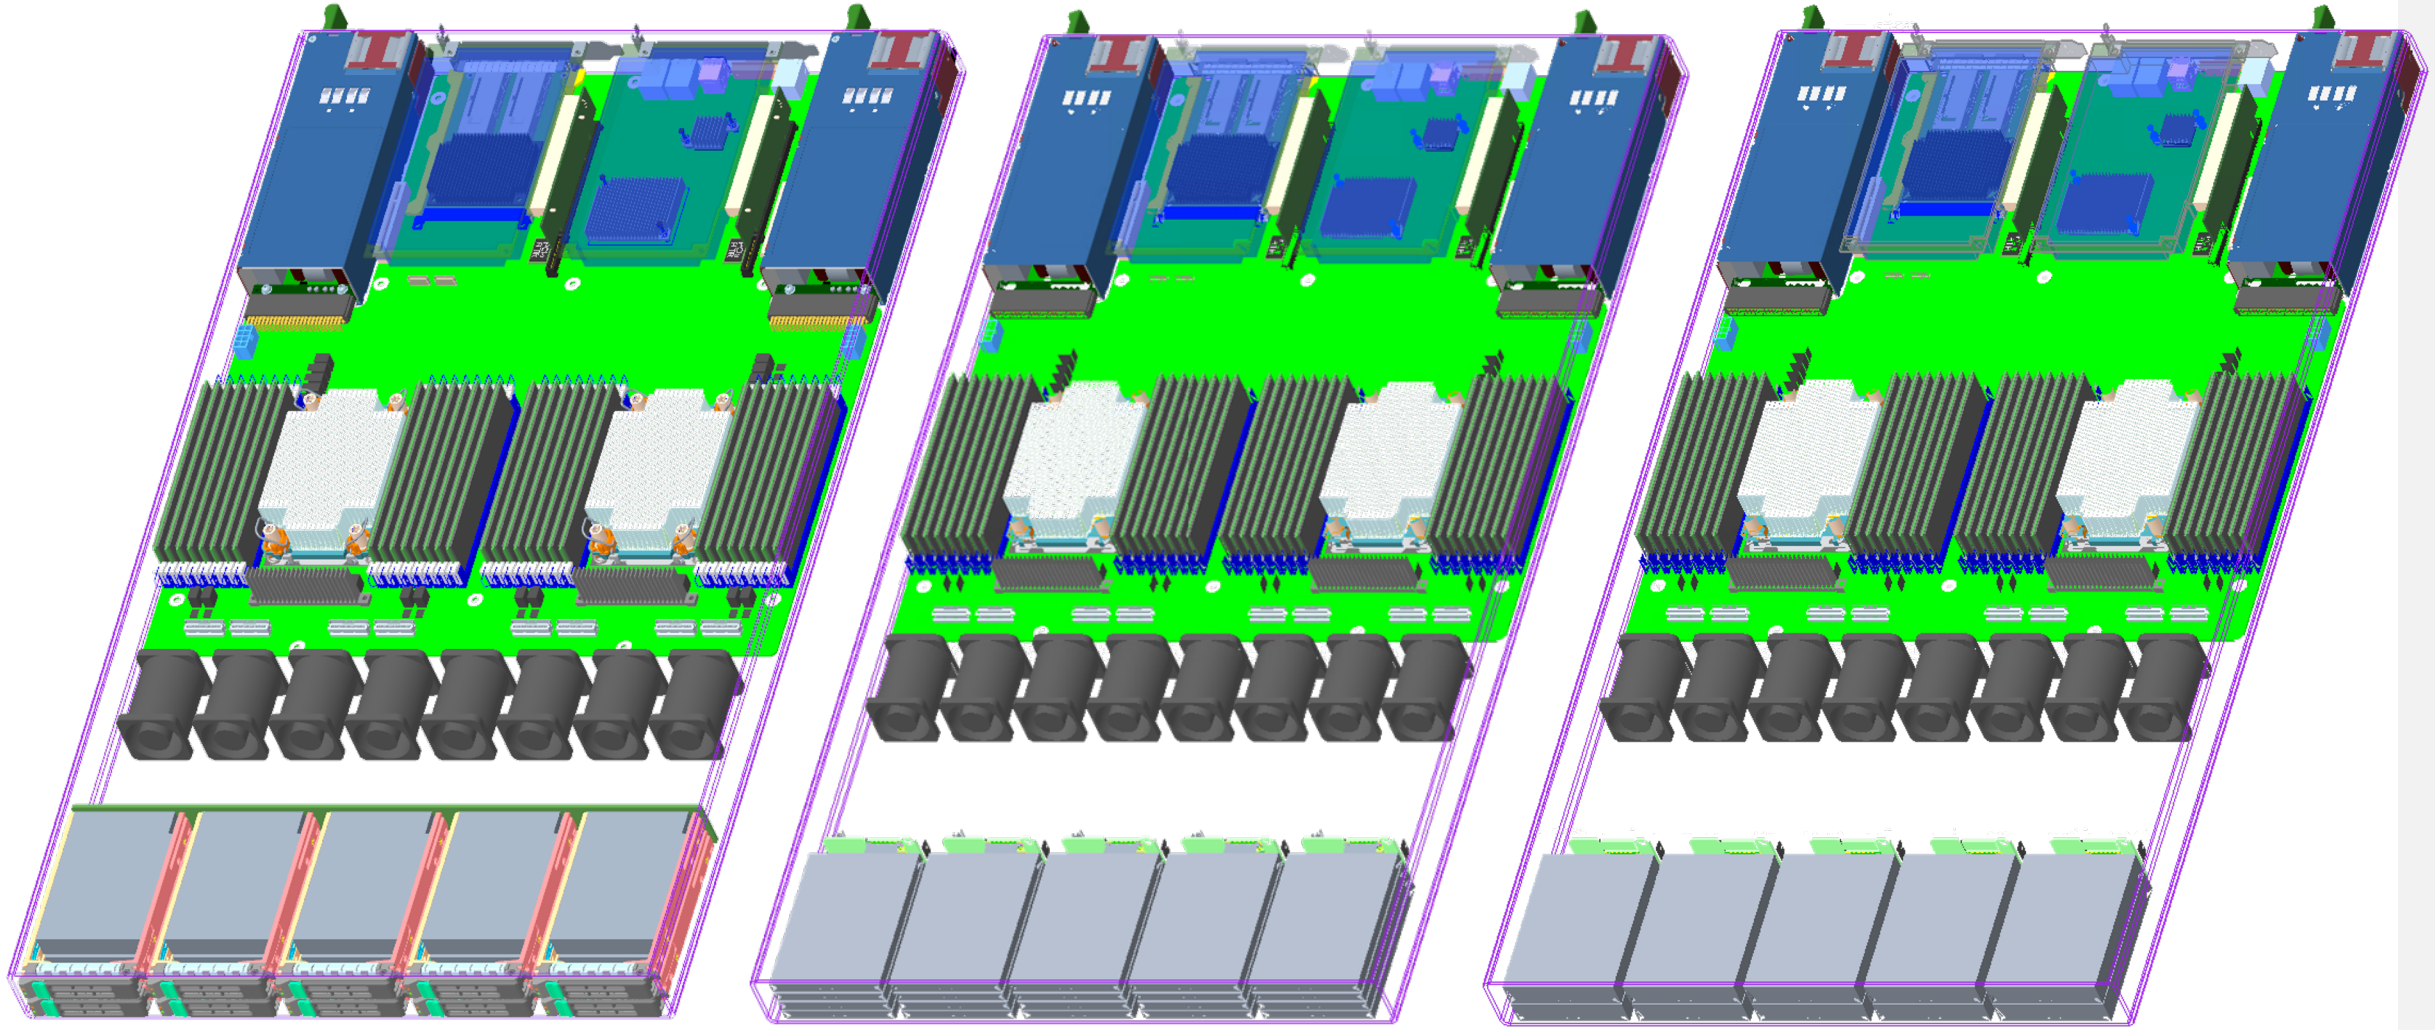

10 x U.2 SSD

20 x E3.S 7.5 mm (Vertical BP)

10 x E3.S 16.8 mm (Vertical BP)

Baseline

OCP Data Center NVMe SSD

Jonmichael Hands, Sr. Strategic Planner

Intel Non-Volatile Memory Solutions Group

OCP Cloud NVMe® SSD Spec

- NVM Express Specification Features
  - Vendor unique log pages
- PCI Express® Specification Features
- SMART Log Requirements
- Reliability, Endurance, Thermal, Power Requirements
- Management
- Security Requirements
- Labeling, Compliance
- Form Factor Requirements
- Open source tool access requirements

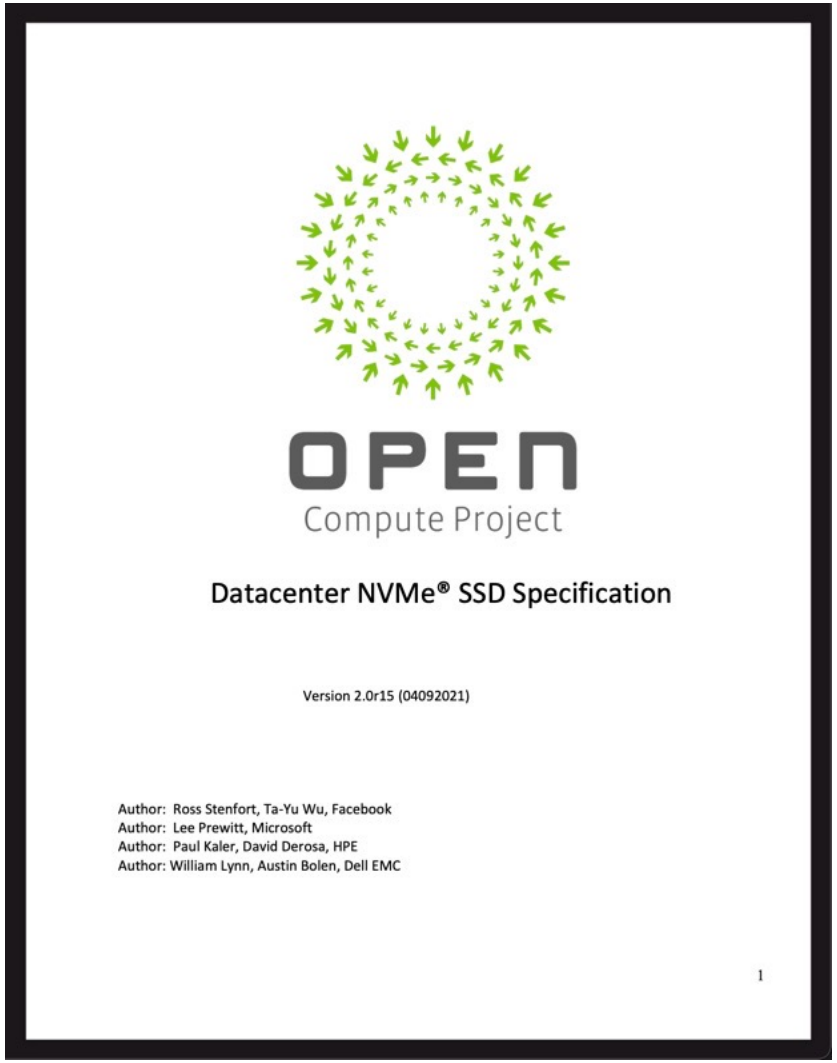

OCP Cloud NVMe® SSD Specification

- SMART Cloud Attributes Log Page, C0
  - Physical media units read/written (to calculate WAF)
  - Bad user and system NAND blocks
  - XOR recoveries
  - Uncorrectable error count
  - Soft ECC errors
  - End-to-end correction counts
  - System data % used
  - Refresh counts

- User data erase counts
- Thermal throttling status and counts
- PCIe correctable errors
- Incomplete shutdowns
- % free blocks
- Capacitor health
- Unaligned IO
- Security version
- PLP status
- Endurance estimate

C0 log page allows for deeper predictive analytics and health monitoring

Intel EDSFF SSDs

Jonmichael Hands, Sr. Strategic Planner

Intel Non-Volatile Memory Solutions Group

Intel. A History of Advancing EDSFF.

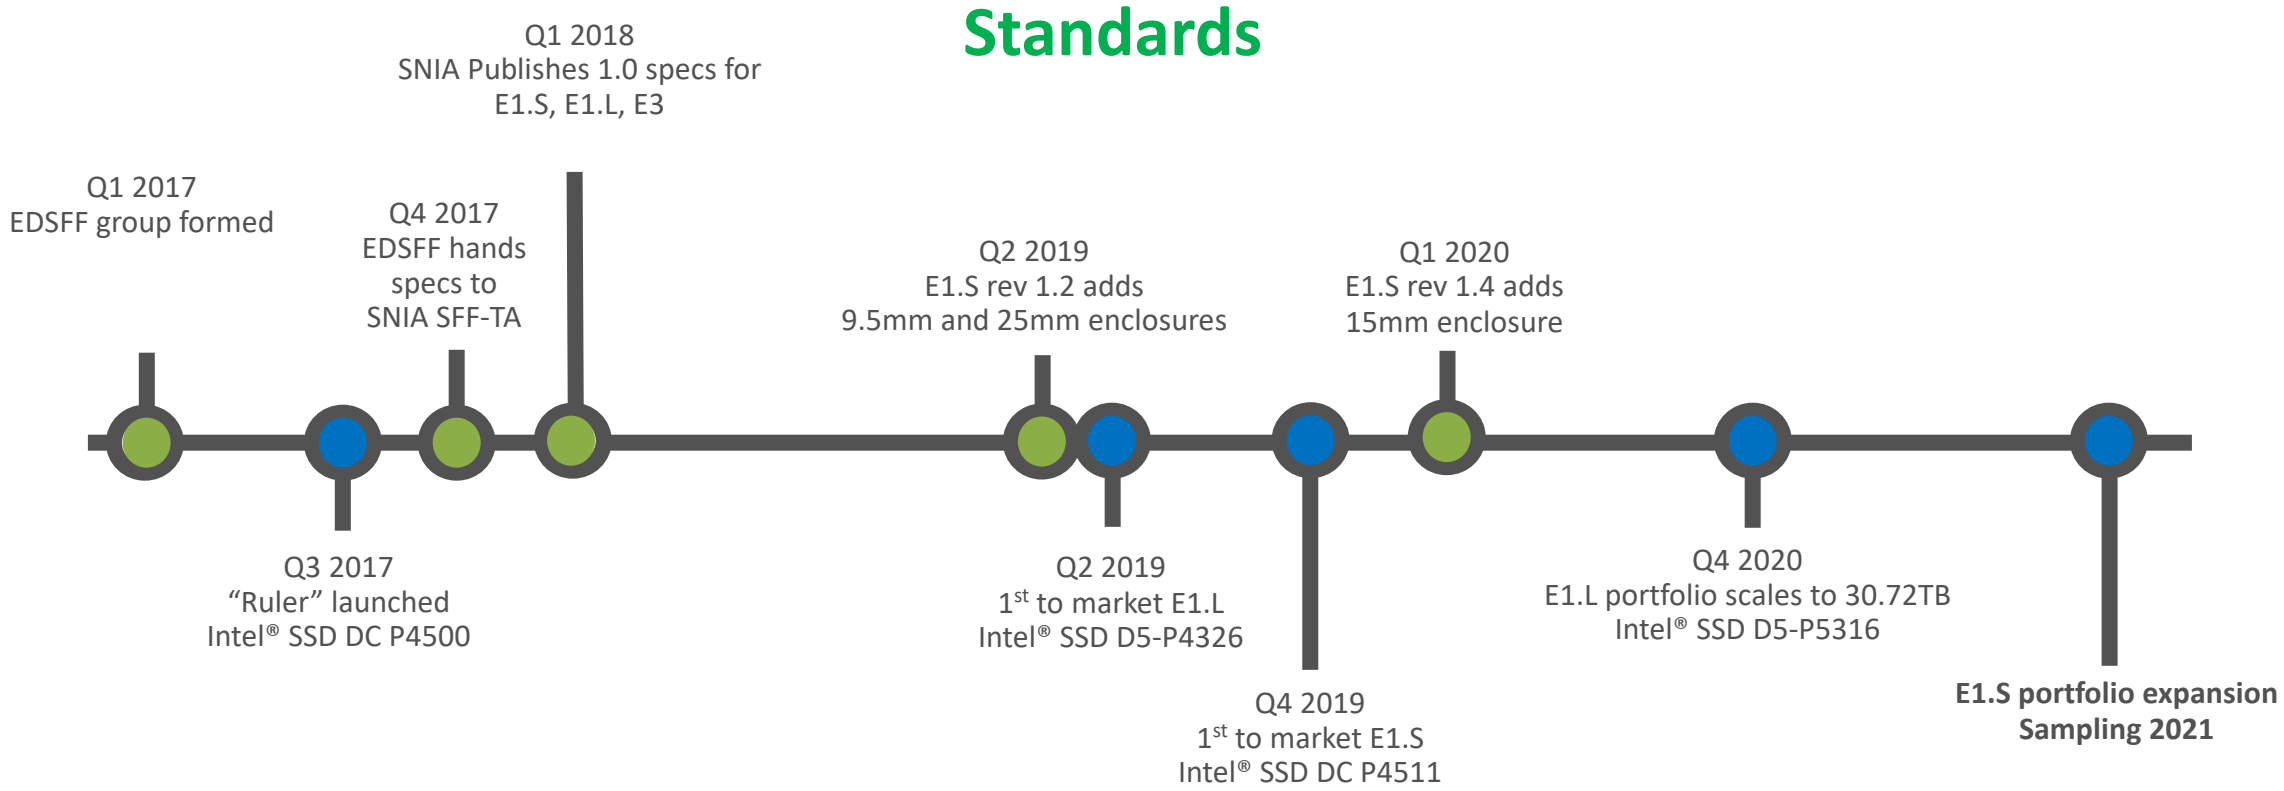

Intel EDSFF products in production today

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

Intel® SSD D5-P5316 Key Specifications – QLC

Comparison	Spec	Gen to gen
4K Rand. Read	Up to 800K IOPS	up to 38% higher[20]
128K Seq. Read	Up to 6800 MB/s	up to 2x+ higher[21]
Endurance (Total PB Written)	Up to 18PB (3K P/E Cycles)	up to 4x higher[13]

Form Factor & Capacity	
Form Factor	U.2 15mm/E1.L
Storage capacity	Industry-leading QLC storage capacity[19] up to 30.72TB

See Appendix for workloads and configurations. Results may vary.

Intel E1.S Products for cloud

Intel® Optane™ SSD P5801X

Fastest SSD in the world now in E1.S[1]

400/800GB. Sampling Q3’2021

Intel® SSD D7-P5520, 1.92,3.84,7.68TB, E1.S 9.5/15mm @ 10,14,20W

OCP Data Center NVMe 2.0 compliant

Sampling in 2021

1. Source – Intel. As compared to generally available PCIe Gen x4 Enterprise and Data Center industry SSDs.

Intel Data Center Cloud Platform for E1.S

”Vulcan City”


intel

OCP NVMe SSD and EDSFF Workshop

April 27, 2021

KIOXIA

KIOXIA: OCP NVMe SSD and EDSFF Workshop

OCP

- KIOXIA has had a long-term relationship
- Active participation since inception
- Provided significant feedback for improvement

SNIA

- KIOXIA is an active participator in SNIA SFF organization, previously on EDSFF consortium
- Planning other EDSFF compliant form factor apart from E1.S

KIOXIA

- KIOXIA announced industry's first PCIe® 4.0 SSD using OCP NVMe SSD spec*
- Keenly listening to enterprise customers as well
- Opportunistically promoting OCP NVMe SSD Cloud spec

Industry Alignment on Product Requirement is a Win-Win!

*Link: https://business.kioxia.com/en-us/news/2020/ssd-20201103-1.html

KIOXIA XD6 Series Features

E1.S 25mm

E1.S 15mm

E1.S 9.5mm

Next Generation Hyperscale Data Center Optimized PCIe SSDs

- Proprietary KIOXIA architecture, controller, hardware and firmware purpose-built for hyperscale environments

OCP NVMe Cloud SSD Design

- Meets form factor, performance, power & thermal requirements
- Enables customers to take advantage of economies of scale

Excellent QoS and Reliability

- Consistent performance, latency and reliability in demanding 24x7 data center environments

Specifications			
Form Factor		E1.S 9.5/15/25mm	
Flash Memory Type		BiCS FLASH™ 3D TLC Flash Memory	
Interface Specification		PCIe® 4.0 x 4L, NVMe® 1.3c	
Performance			
User Capacity*	GB	1,920	3,840
Sequential Read 128KiB* (QD=32)	MB/s	6,500	
Sequential Write 128KiB* (QD=32)	MB/s	1,250	2,400
Random Read 4KiB* (QD=256)	KIOPS	660	850
Random Write 4KiB* (QD=128)	KIOPS	50	90

*KIOXIA Corporation definition of capacity: 4 KiB = 4,096 bytes, 128 KiB = 131,072 bytes, GB = 1,000,000,000 (10^9) bytes and a terabyte (TB) as 1,000,000,000,000 bytes (see end of presentation for full capacity disclaimer). Subject to change

"Developing and deploying flash based products is very challenging. EDSFF E1.S is the next generation of flash form factors with its superior thermals, performance, serviceability and scalability compared to solutions today. In addition to this, the OCP based NVMe Cloud SSD Specification is a great step forward for aligning the SSD providers with hyperscale needs. XD6 support of these next generation flash needs is a great step forward for meeting hyperscale needs for now and the future." Ross Stenfort , HW Storage Engineering @Facebook, Inc.

Definition of capacity: KIOXIA defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1GB = 2^30 = 1,073,741,824 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, such as Microsoft Operating System and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

All company names, product names and service names may be trademarks of their respective companies.

Images are for illustration purposes only.

Western Digital.

OCP Workshop

Eric Pike

Sr. Director, Cloud Segment Marketing

April 27, 2021

EDSFF SSD Portfolio

Single connector and pinout for FF family – overall industry benefit

E1.L 18mm & 9.5mm

High Capacity Storage – 1/2 PB in 1U

TCO
Improved TCO

Larger Capacities

FIPS 140-3 Level 1

- 1U optimized
- Improved manageability and serviceability
- High density and improved data center TCO

E1.S

Up-to 24 SSDs in a 1U Server

Scalable performance

Hot-Pluggable

Scalable Power

- High-performance server optimized
- Hot-plug and serviceability
- Power and thermals for mainstream SSDs (2TB – 8TB)

Improved SSD density, thermals, power and scalability

Datacenter NVMe™ SSD Specification

- Common drive specs help the ecosystem

Datacenter NVMe™ SSD Specification

Encourages industry collaboration and discussion

Standardization of the common features

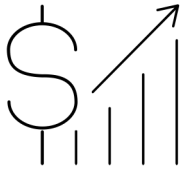

Improved Time to Market

Improved understanding of the cloud requirements

Enables open source tools to mange the SSDs

SAMSUNG

Samsung SSD
OCP NVMe SSD and EDSFF Workshop

Samsung EDSFF is ready

Model	PM9A3	PM1743
Cell Technology	V6 TLC	V6 TLC
Interface	PCIe Gen4 1x4	PCIe Gen5 1x4, 2x2
Form Factors	E1.S 9.5mmT, 15mmT, 25mmT E1.L 9.5mmT	E3.S 1T
Endurance	1 DWPD	1 DWPD
Features	 • Optimized performance and latency for hyperscale environments • Improved health monitoring and debugging features • Enhanced security features	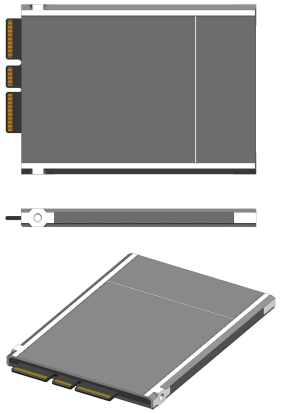 • High performance PCIe Gen5 SSD for enterprise applications • Provides enhanced data encryption and attestation
Recommendation	For hyperscale datacenter, M.2 → E1.S/E1.S	For enterprise Server/Storage, U.2 → E3.S
Schedule	Available now	Available in Q2'22

Notes: All product plans and roadmaps are subject to change without notice.

Samsung supports OCP NVMe SSD v1.0

PM9A3 Available Now!

E1.S form factor

PM9A3		Capacities	OCP v1.0
Form factors	E1.S	1TB-8TB	Support
	E1.L	16TB	Support
	M.2	1TB-4TB	Support
	U.2	1TB-16TB	Support

Health Monitoring

Telemetry for failure analysis
Capacitor health log
Firmware activation history
Performance monitoring log

Reliability

UBER < 1sector per 10^{17} bits read
End to end data protection
MTBF 2.0million hours
Global wear-leveling

Security

AES-XTS 256 bit encryption
Anti-rollback protection, TCG Opal
Secure boot, Key revocation

NVMe CLI Plug-in

- ✓ ***Less Risk during SSD dev with combined requirements***
- ✓ ***Enables to manage massive SSD at rack-scale***
- ✓ ***Enhances security and data protection to solve today's datacenter challenges***
- ✓ ***Expedites failure analysis with open-source tools***

Notes: All product plans and roadmaps are subject to change without notice.

A journey shared takes us beyond

Tremendous Growth on OCP Horizon

16.2% **Compute CAGR** (2019-2024 est.)

$ 5.7B **Storage Revenue** (2024 est.)

18.5% **Storage CAGR** (2019-2024 est.)

Source: CIO & Leader

Micron and OCP: Aligned Values

Acceleration Catalysts

Solving difficult problems with open and innovative developments. Micron continues to accelerate solutions with our public, open, validated reference architectures.

Focus on Efficiency

OCP's focus on efficiency – in energy use, cost, design and space use is tightly aligned to Micron's sustainability commitment.

Align and Contribute to Open Industry Standards

Micron demonstrated our belief in the universal benefits of open designs when we released our Heterogeneous-Memory Storage Engine to the open-source community.

Uniform, complete standards help democratize data and drive responsible business practices and continuity.

Micron

Micron and OCP: What's Next

Broad form factor range consolidates around demand-driven standards

EDSFF: Accelerating Adoption

Industry sees value in form factor optimization for flash

E1.S dominant EDSFF variant

Consolidation is coming:

Industry can't sustain offering 11 form factors (in addition to long tail of legacy form factors)

Near term: Micron sees focus around E1.S

Longer term: E1.S still dominant, E1.L support for large capacities and E3 growth aligned with PCIe Gen5

Coming Form Factor Shift

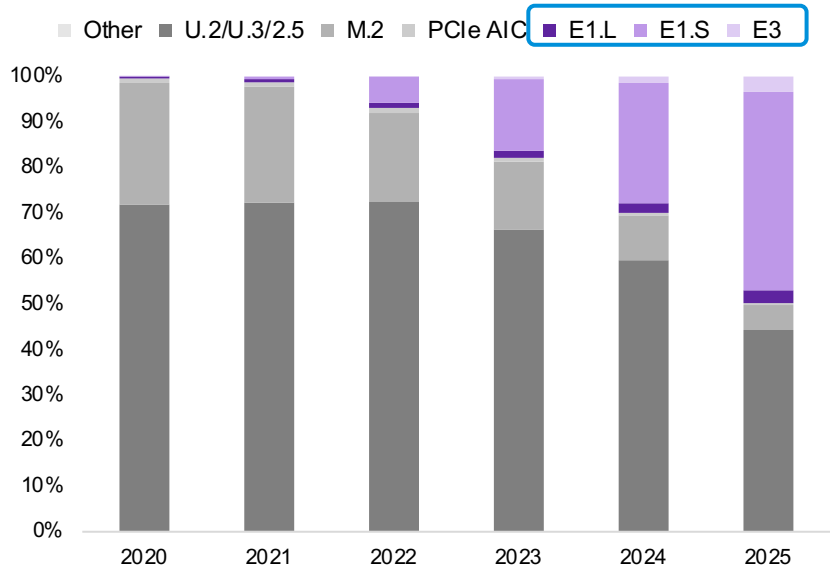

Source: Forward Insights (02/21)
Units: user PB

Flash-optimized Flexibility

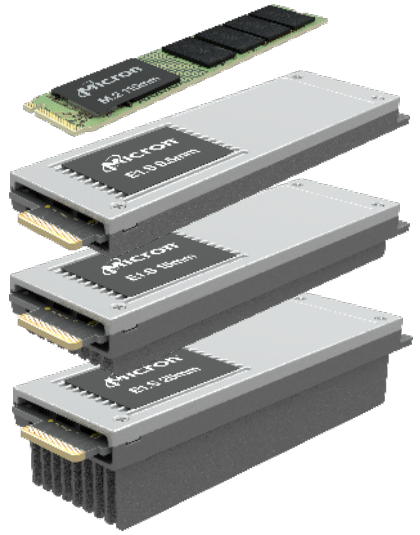

Datacenter NVMe SSD Spec Enablement

- **Datacenter NVMe SSD Spec enables a SSD supplier to build more common HW/FW for users**
  - Existing standard spec allows a high level of flexibility so users could choose many different directions from others

- **Security requirement is enhanced across the board**
  - Secure boot, Authentication and transaction path security requirements become clear and robust

- **However, there's some scope to improve going forward**
  - Depending on users, unused features could be categorized as mandatory and this may add burdens to suppliers
  - Spec update frequency may need to be managed to be predictable to help the efficient product-planning

Without Datacenter NVMe SSD spec

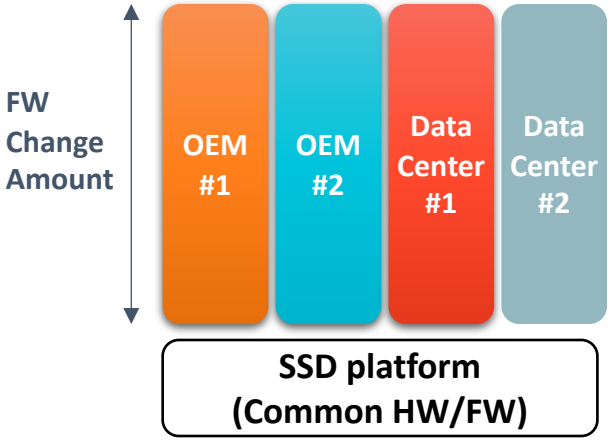

With Datacenter NVMe SSD Spec

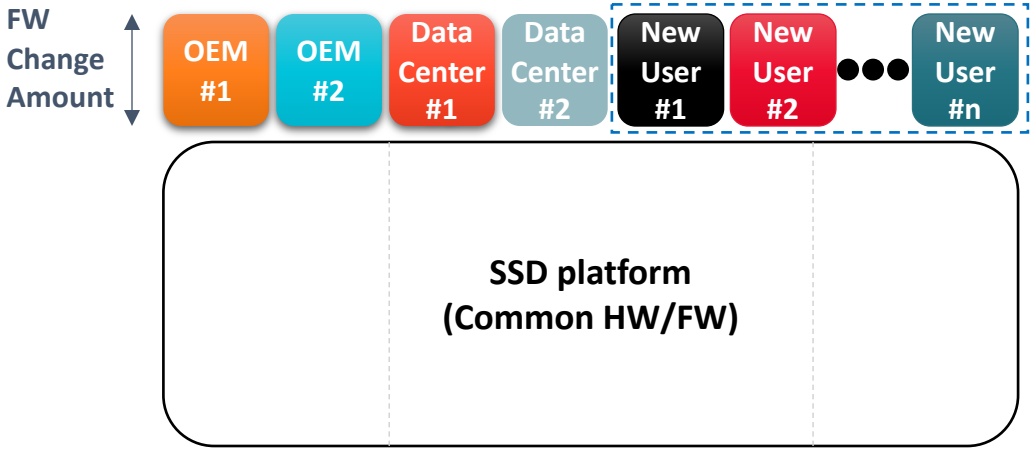

SKhynix Products Supporting DC NVMe SSD Spec

- **SKhynix has been offering products that support the Datacenter NVMe SSD spec**
  - PE8111 E1.L and PE8110 E1.S are developed based on the Datacenter NVMe SSD spec v1.0a and those are being shipped to customers
  - New products coming up next aim to meet DC NVMe SSD Spec V2.0

- **PCIe Gen5 SSDs are expected to be developed in EDSFF form factor and Datacenter NVMe SSD spec**
  - Not only E1.S/E1.L but E3 SSD is being planned for PCIe Gen5 and it will be based on version 2.0

	E1.S 15mm	E1.L 18mm / 9.5mm	E3.x (TBD)
Product	PE8110 	PE8111 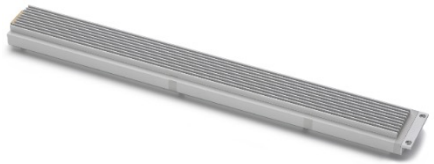	Next generation 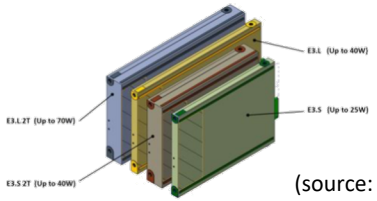 (source: SNIA)
Interface	PCIe Gen4x4	PCIe Gen3x4	PCIe Gen5
Capacity	1920GB – 7680GB	15360 – 30720GB	TBD
Read / Write Bandwidth	6500 / 4400 MB/s	3550 / 3300 MB/s	TBD
Read / Write IOPS	1100 / 160 KIOPS	750 / 105 KIOPS	TBD

FADU

Datacenter NVMe Spec and FADU SSD Offering

www.fadu.io

OCP NVMe Datacenter Specification

March 2019	Dec 2019	Nov 2020	March 2021
CP Heat sinks for EDSFF Standardized	New OCP Cloud SSD specification	HP and Dell Extend spec to E3	Boot Drive Specification
FADU Sep 2019 We start on E1.S OCP standard FFs And heat Sinks	FADU April 2020 We demo Worlds first OCP Cloud SSDs at Open compute with a Supermiro box	FADU Nov 2020 We start compliance work with UNH-IOL We announce Delta Gen4 OCP Drive	FADU Now Gen 3 OCP SSD MP Gen4 NVMe DC SSD Working Samples

FADU NVMe Datacenter SSDs

FADU OCP SSD Offering	Bravo Gen3x4 7%OP , 2 TB, E1.S	Delta Gen4x4 7%OP 4 TB, E1.S/E3?	Echo Gen5x4 7% [illegible] 8 TB E1.S/E3 /E.[illegible]
SR in MB/s	3500	7300	14,[illegible]00
SW MB/s	2700	4600	12,000
RR in KIOPS	800	1490	3500
RW in	100	180	410

Come talk to us

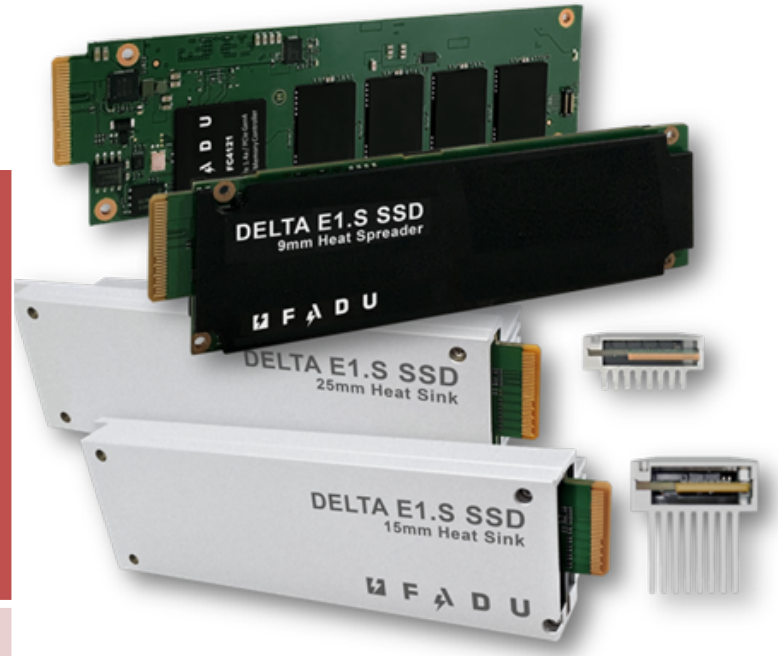

In Real workloads
We get excellent
Random Read in
Mix workloads/
Recovery after Burst/
QOS and Max latency

www.fadu.io

FADU

Thank You

www.fadu.io

Microsoft

Datacenter NVMe SSD Validation

Microsoft - Adam Guy
Teledyne LeCroy – Nick Kriczky
UNH-IOL – David Woolf

Datacenter NVMe SSD Test Strategy

- ✓ Validation against the Datacenter NVMe SSD Specification ensures the target device complies with the wide array of common requirements across multiple datacenter customers
  - ✓ Common & deterministic SSD functionality in-system
  - ✓ Numerous datacenter opportunities from one leveraged SSD design

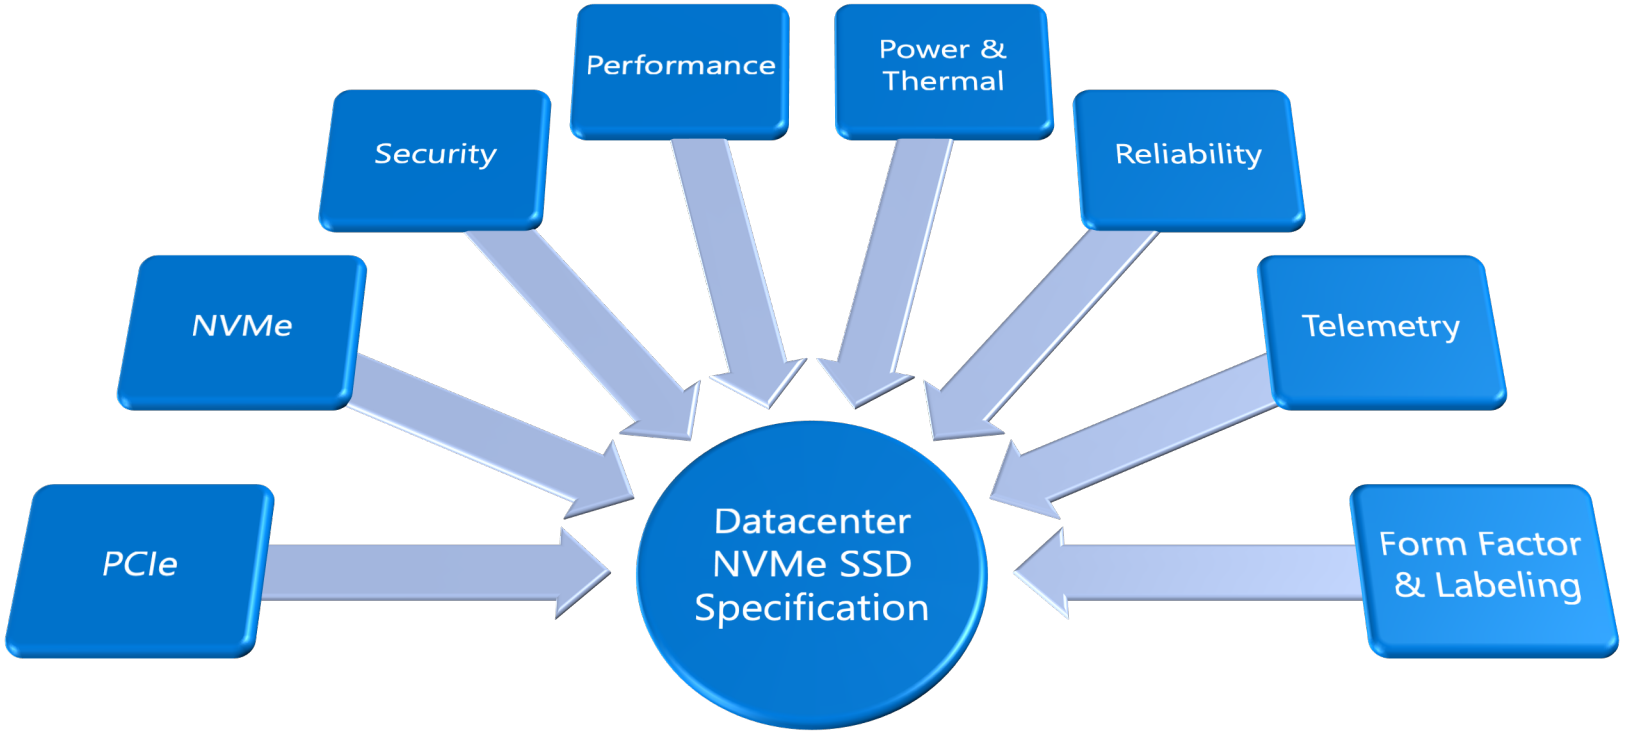

Datacenter NVMe SSD Test Strategy

- ✓ Test suites for Datacenter NVMe SSD Spec compliance are in-progress
  - ✓ First intercept is revision 1.0a
- ✓ Microsoft is generating initial test specifications mapping to requirement ID's
  - ✓ Test specifications only apply to requirements that can be validated in a black box fashion
- ✓ Test specs are subsequently reviewed & approved by key stakeholders
- ✓ Test case implementation based upon approved test specifications
  - ✓ 1.0a effort currently in-progress at Teledyne LeCroy and UNH-IOL, independently

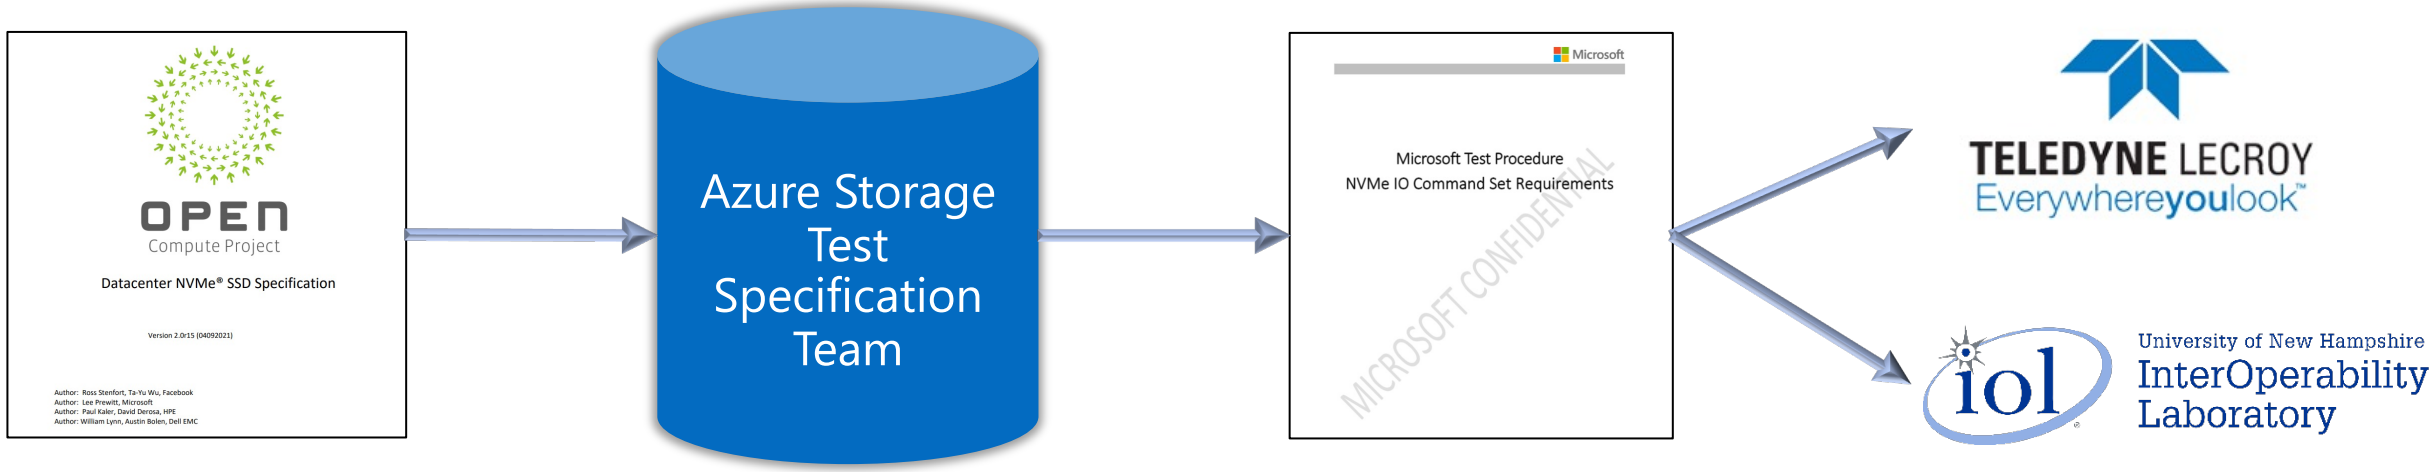

Teledyne LeCroy – Everywhere You Look

Teledyne LeCroy
Protocol Analyzers

OakGate SSD
Validation Platforms

Epic Environmental
Chamber Platforms

Austin Labs

PCI EXPRESS

CXL Compute Express Link

nvm EXPRESS

OPEN Compute Project

HI-SPEED CERTIFIED USB

FCIA FIBRE CHANNEL INDUSTRY ASSOCIATION

Serial Attached SCSI

Ethernet

Early Design & Development | Design and Development | Validation | Pre-Production | Production | Field Failure Analysis

Teledyne LeCroy Datacenter NVMe SSD Specification

OakGate TECHNOLOGY

OCP Conformance Test Equipment and Software

Austin Labs

Test Suite Development and Conformance Testing Services

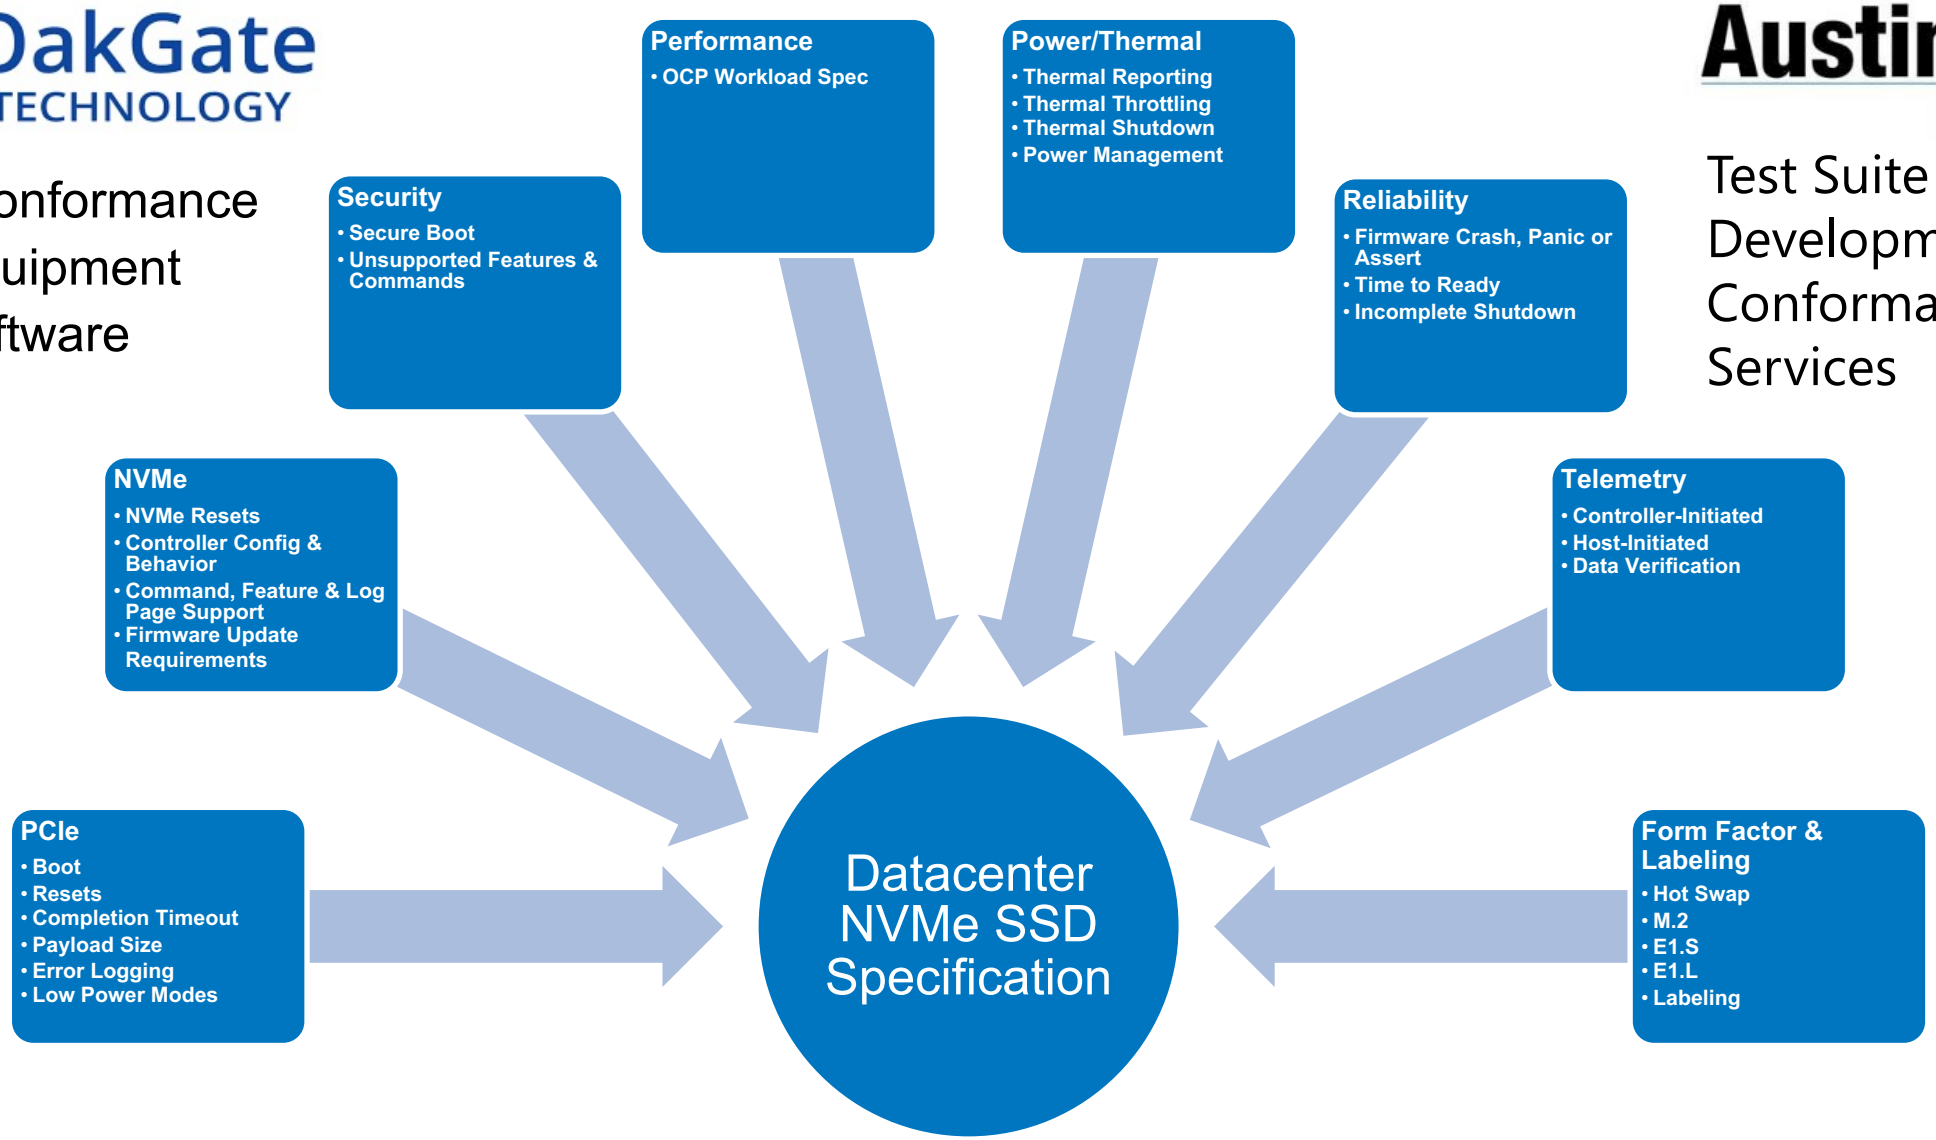

UNH-IOL Tool Support for DatacenterSSD

- UNH-IOL published test plan for DatacenterSSD in mid 2020, and is actively working with key stakeholders on adding and upgrading tests.
- DatacenterSSD spec explicitly requires NVMe v1.4 compliance, therefore, UNH-IOL INTERACT Test Tool, currently widely used for NVMe v1.4 compliance, has integrated add-on support for DatacenterSSD compliance into a single tool.

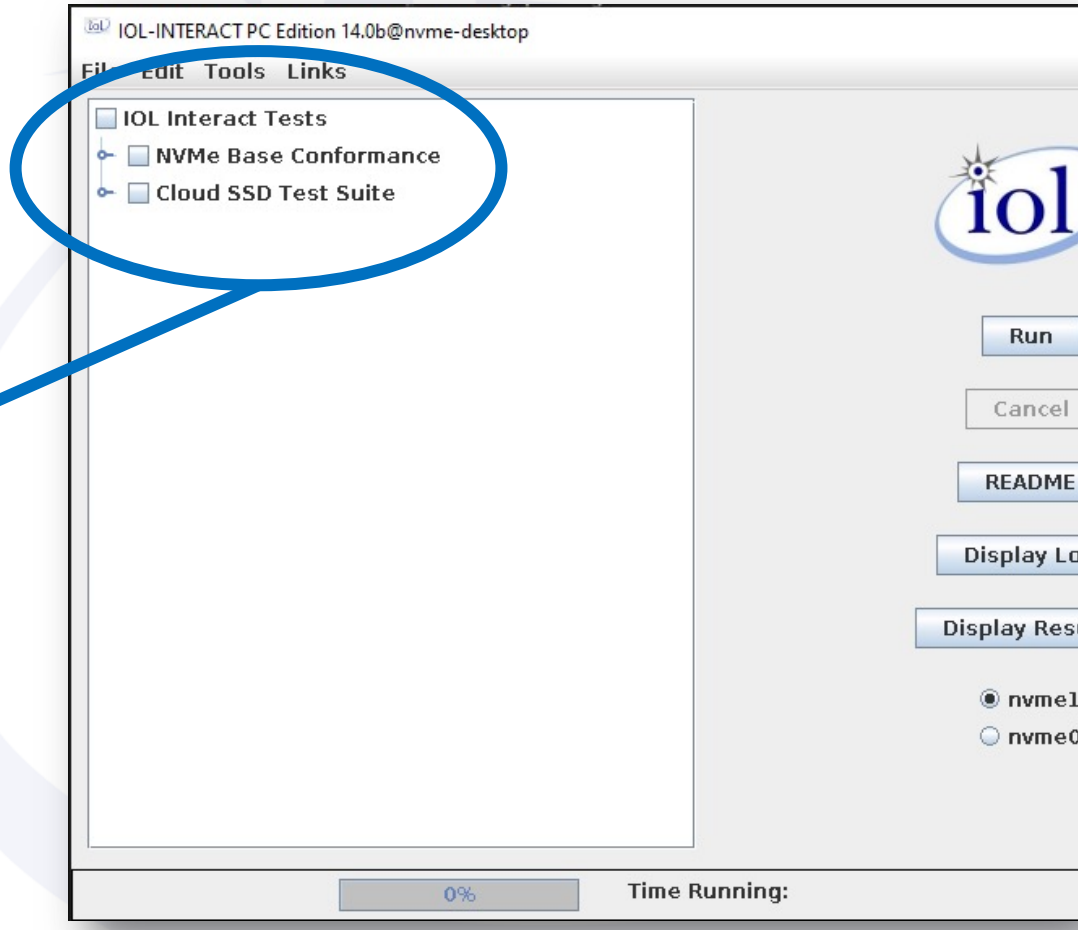

UNH-IOL Tool Support for DatacenterSSD

UNH-IOL INTERACT Test numbers correlate to Test Plan

UNH-IOL Test names correlate directly to DatacenterSSD Spec

DatacenterSSD tools currently being demonstrated at customer sites

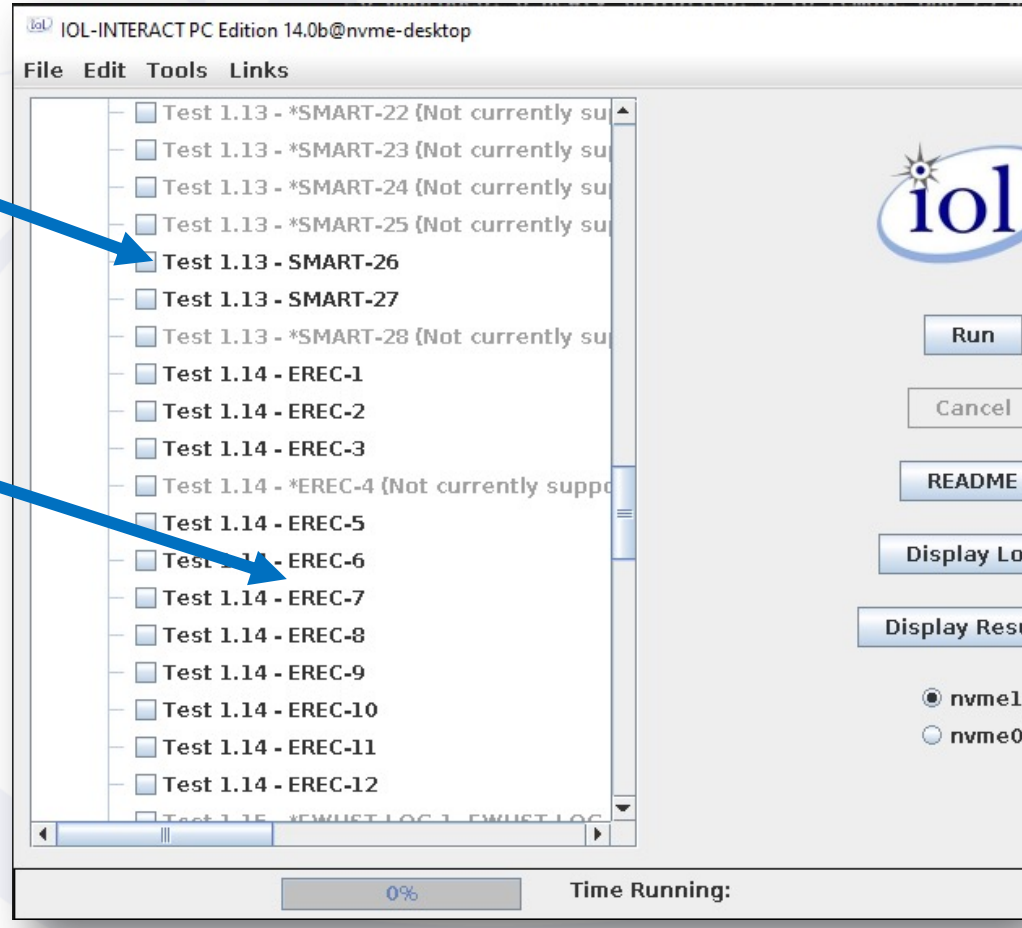

Thank You!