Flash Memory Summit

EDSFF Update: Industry Leaders Explain E1and E3 Innovations Specification Updates

Moderated by: Cameron Brett (SNIA SSD SIG and KIOXIA)

Panel members:

John Geldman (KIOXIA), Paul Kaler (HPE), Ross Stenfort (Meta), Lee Prewitt (Microsoft), Anthony Constantine (Intel)

35 YEARS OF NAND FLASH MEMORY INNOVATION

EDSFF Update: Industry Leaders Explain E1&E3 Innovations & Specification Updates

Flash Memory Summit

Panelist speakers and topics:

- John Geldman – E3 form factor overview
- Paul Kaler (for Bill Lynn) – E3 use cases, device types and applications
- Paul Kaler – E3 specifications list and recent updates
- Anthony Constantine – E1 specifications list and recent updates
- Ross Stenfort – E1 form factor overview
- Lee Prewitt – E1 use cases, device types and applications

35 YEARS OF NAND FLASH MEMORY INNOVATION

What is EDSFF?

- Enterprise and Datacenter Standard 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 tool-less serviceability

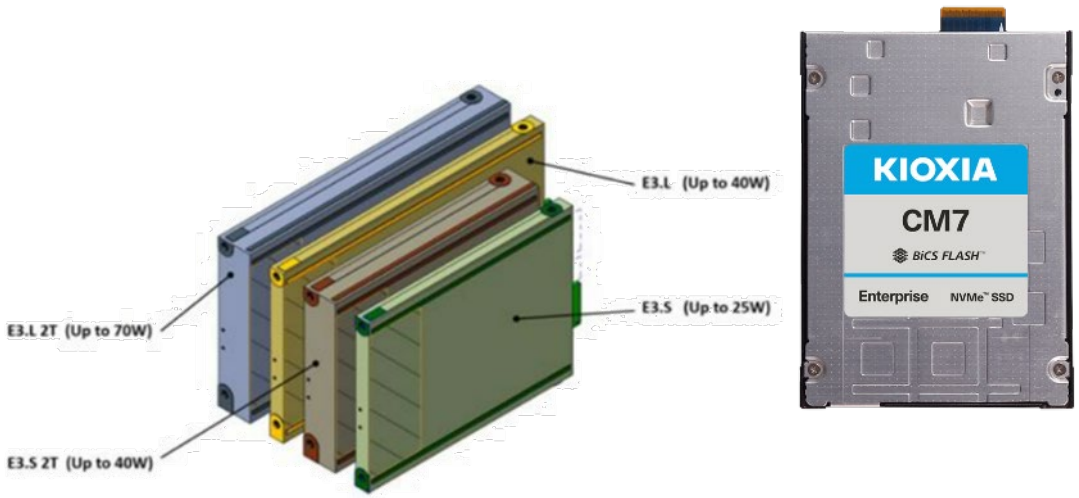

Visit the SNIA SSD Form Factors Web Page for EDSFF Updates

Flash Memory Summit

E3 form factor overview

John Geldman, KIOXIA

35 YEARS OF NAND FLASH MEMORY INNOVATION

E3 – A design scheme from the past, restarting

- In the beginning, 2.5 inch devices evolved into PCI Express U.2®
  - From SATA, from SAS, from HDDs
  - That venerable and useful form factor came with legacies
- Core principles for re-imagining of 2.5 inch devices
  - SDDs (and other functions) don't need rigid spinny things, but do need area flexibility
  - Signal integrity of one or two 6 Gb/s lanes is one environment to design for, but four 128 Gb/s lanes?
  - And the power and cooling for 4 lanes of that performance?
- The E3 family was developed to rebuild from the concept of PCI Express U.2 from the ground up
  - Designed for chassis flexibility
  - Designed for a family of device sizes to meet the needs of multiple functions and capacities
  - Designed for Signal Integrity
  - Designed for a specific cooling flow
- EDSFF started as Enterprise and Data Center SSD Form Factor
  - It's not just for NVMe SSDs any more
  - Now it is: Enterprise and Data Center Standard Form Factor

E3 is a Family of Four Form Factors

- The Family:
  - E3 short single thickness device (E3.S)
  - E3 short double thickness device (E3.S 2T)
  - E3 long single thickness device (E3.L)
  - E3 long double thickness device (E3.L 2T)

Device Variation	Height	Length	Thickness
E3.S	76mm	112.75mm	7.5mm
E3.S 2T	76mm	112.75mm	16.8mm
E3.L	76mm	142.2mm	7.5mm
E3.L 2T	76mm	142.2mm	16.8mm

- Fun Facts
  - Note the SFF-TA-1002 PCB card edge
  - Note the latch friendly housing (like U.2)
  - Note the flexibility for volumetric heat sinking or for multiple boards
  - Note what you don't see: this Form Factor was designed for a planned cooling air flow

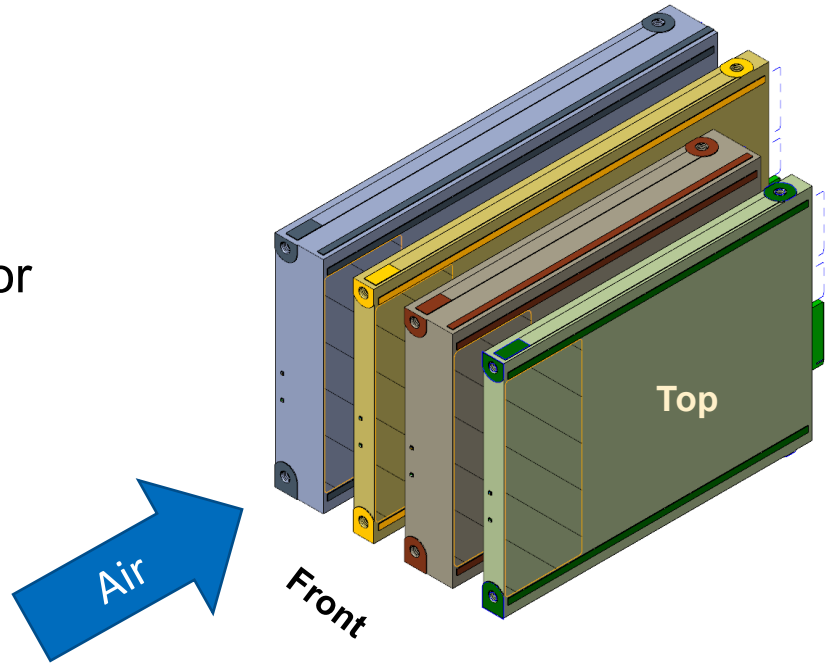

The E3 family is a member of the EDSFF family

- All of EDSFF support the SFF-TA-1009 mapping
  - and the REF-TA-1012 options
- All of EDSFF share the SFF-TA-1002 connectors
  - A system designed for signal integrity
- E3.S and E1.S were designed for the same depth, making it easy for a chassis to support both
  - An E1.S PCB can even be mounted in an E3.S housing
- All of EDSFF types were designed with a specific air flow in mind
  - SFF-TA-1023 provides a cooling strategy, simplifying both device and chassis design

E3 is well positioned with a lot of options

- A rainbow of storage
  - From Low Power Storage Nodes – e.g., E3.S for a 48 device rack
  - To Performance Nodes – e.g., E3.S 2T w/16 lanes & 9 mm of heat sink fins

- Other functions
  - REF-TA-1012 collects pinouts
    - Including: SFF-TA-1009, OCP-NIC, Gen-Z, PECEFF, and SNIA Native NVMe-oF
  - PCIe xPUs (accelerator engines with or without built-in storage/memory)
  - CXL memory pools
  - 25G or 50G Native NVMe-oF eBOFs

E3 – Connector Options

- SFF-TA-1002 Connector with SFF-TA-1009 PCIe® mapping
  - 1C (4 lanes), 2C (8 lanes), 4C (16 lanes), or 4C+ (16 lanes + power tab)
    - ‘Vanilla’ E3 devices typically have 4 lanes
  - Vertical AIC, Right Angle, Orthogonal connector options
    - ‘Vanilla’ E3 has a midplane with Vertical connectors

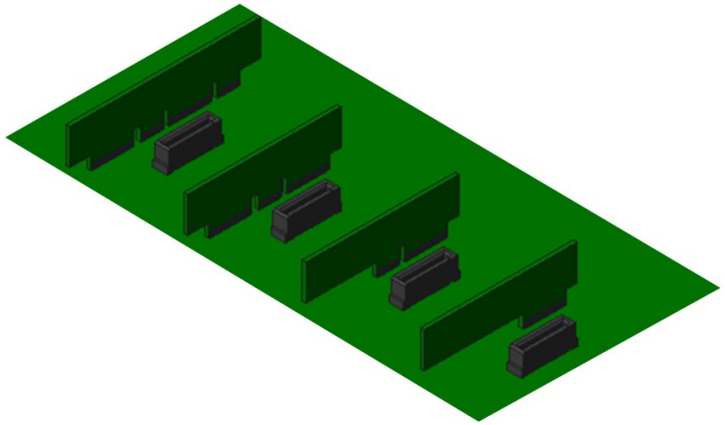

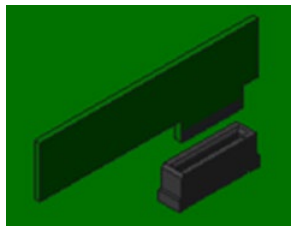

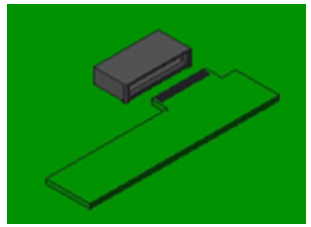

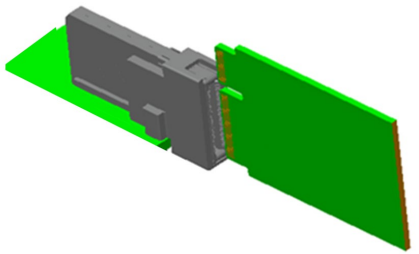

Flash Memory Summit

E3 use cases, device types and applications
E3 specifications list and recent updates

Paul Kaler, HPE

(also presenting for Bill Lynn, Dell EMC)

35 YEARS OF NAND FLASH MEMORY INNOVATION

Cohesive Family of Devices

Flash Memory Summit

High Power Devices

E3.L 2T (Up to 70W)
- Higher end accelerator (DPU)
- Computational storage

E3.S 2T (Up to 40W)
- High-capacity NAND
- High-performance NAND and SCM
- Mid-range accelerator
- Computational storage

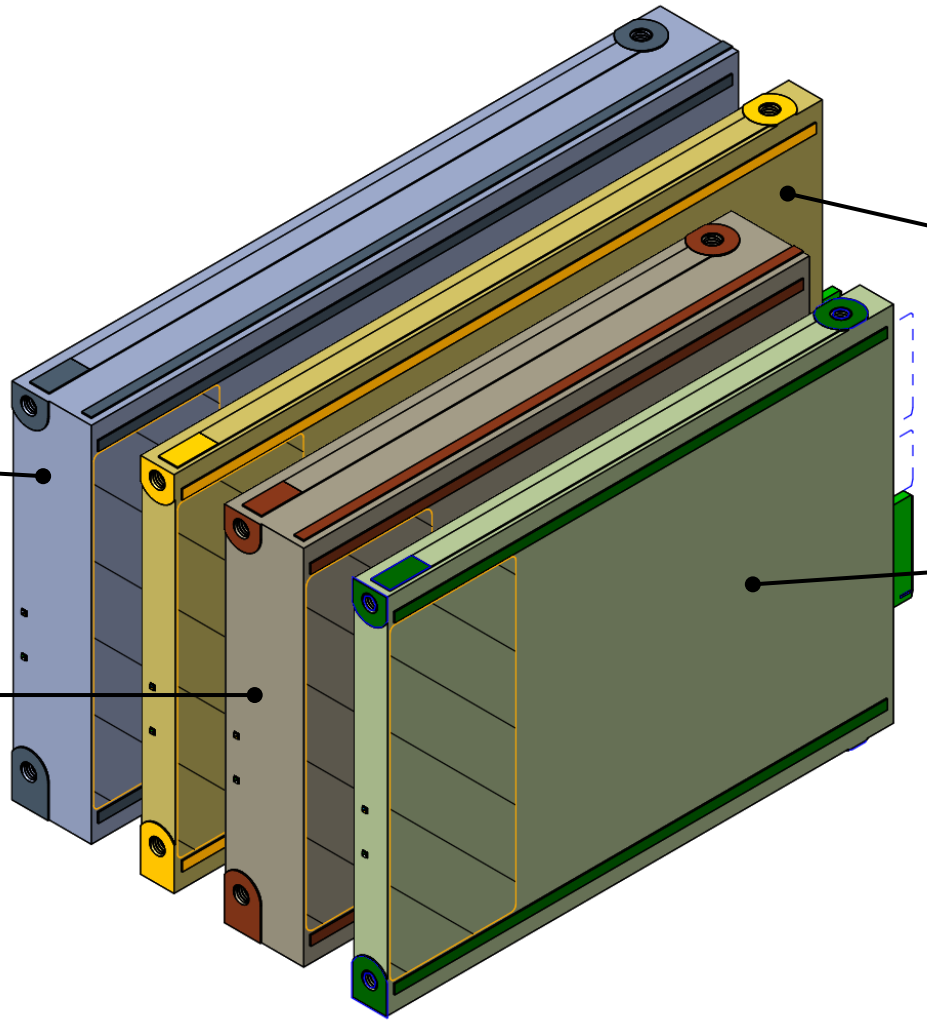

Storage Devices

E3.L (Target 35 to 40W)
- Highest capacity NAND / SCM (Block Storage Device)
- Dense and storage centric chassis

E3.S (Up to 25W)
- 1U & 2U Mainstream NAND-based NVMe
- Similar size to 2.5” 7mm SATA enables transition to NVMe

Potential E3 Chassis Configurations (1U)

Flash Memory Summit

Storage Config

20 E3.L Storage Devices

High Power Server Config

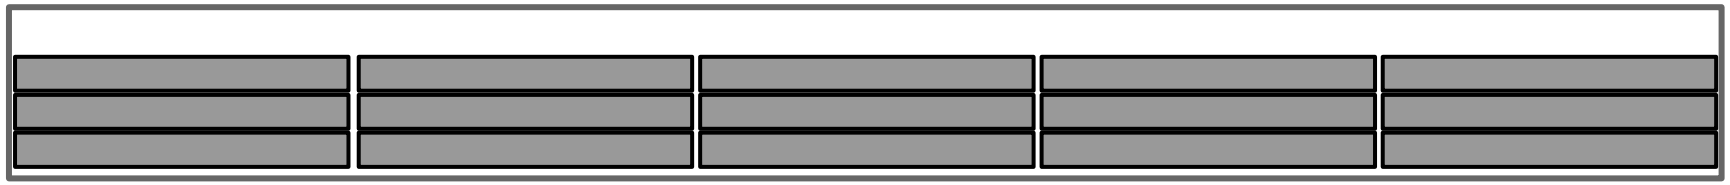

15 E3.S Storage Devices with an air channel above

Alternate Device Config

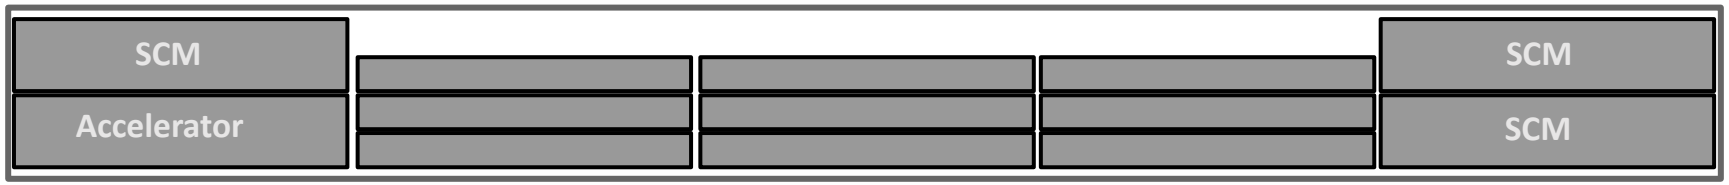

9 E3.S Storage Devices and 4 E3.S 2T SCMs or accelerators

35 YEARS OF NAND FLASH MEMORY INNOVATION

Potential E3 Chassis Configurations (2U)

Flash Memory Summit

Storage Config

44x E3.L Devices

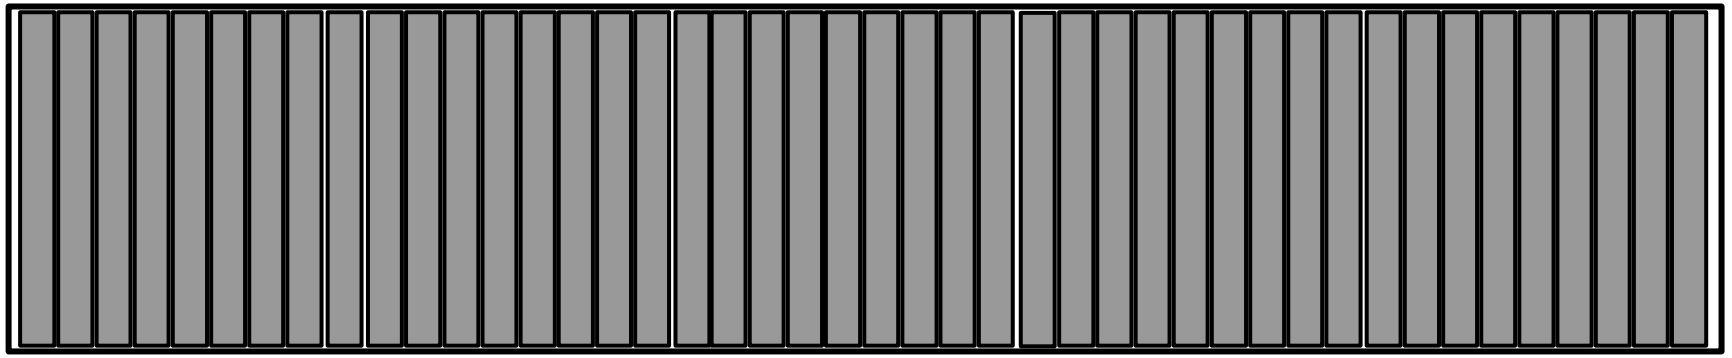

Alternate Device Config

24x E3.S Devices and 8x SCMs or accelerators

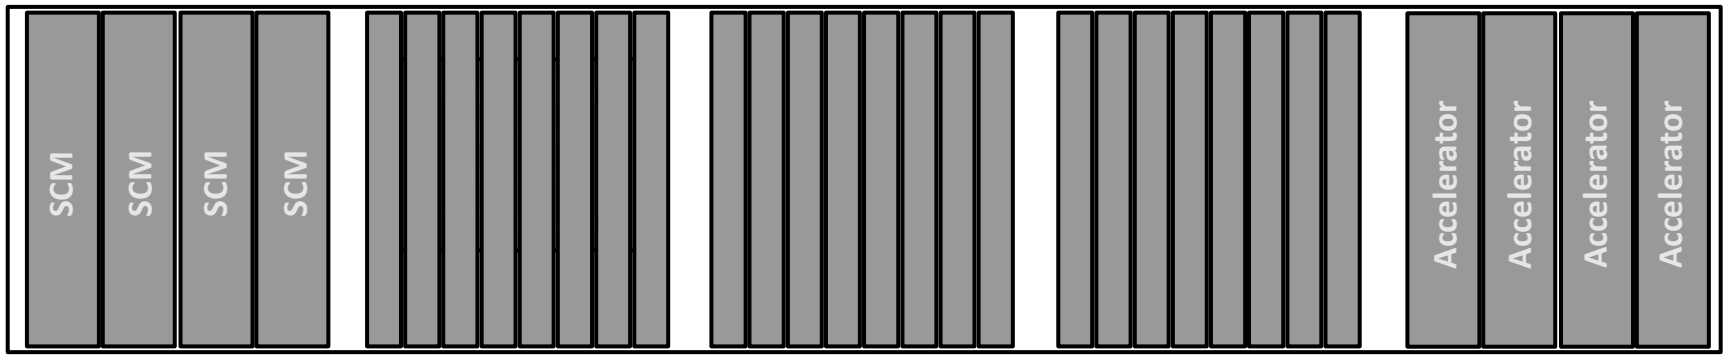

35 YEARS OF NAND FLASH MEMORY INNOVATION

E3 Electrical and Mechanical Specs

- Electrical related specs are in SFF-TA-1009
  - Signals and pinout
  - Power Requirements
  - Signal Integrity
  - LED requirements
- Mechanical related specs are in SFF-TA-1008
  - Covers complete E3 family

Table 5-1 Nominal Device Form Factor Dimensions

Device Variation	Height	Length	Thickness
E3.S	76mm	112.75mm	7.5mm
E3.S 2T	76mm	112.75mm	16.8mm
E3.L	76mm	142.2mm	7.5mm
E3.L 2T	76mm	142.2mm	16.8mm

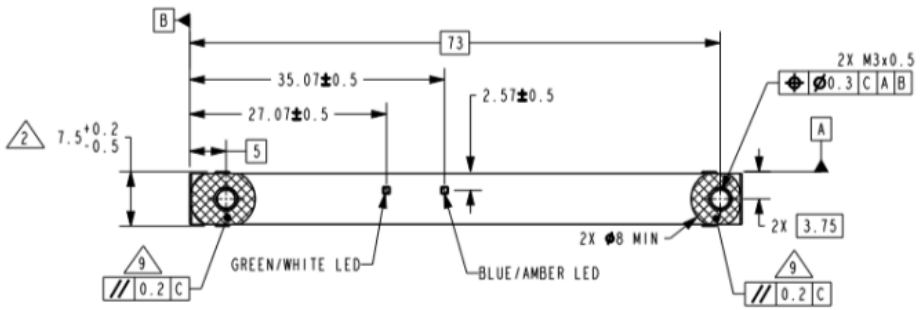

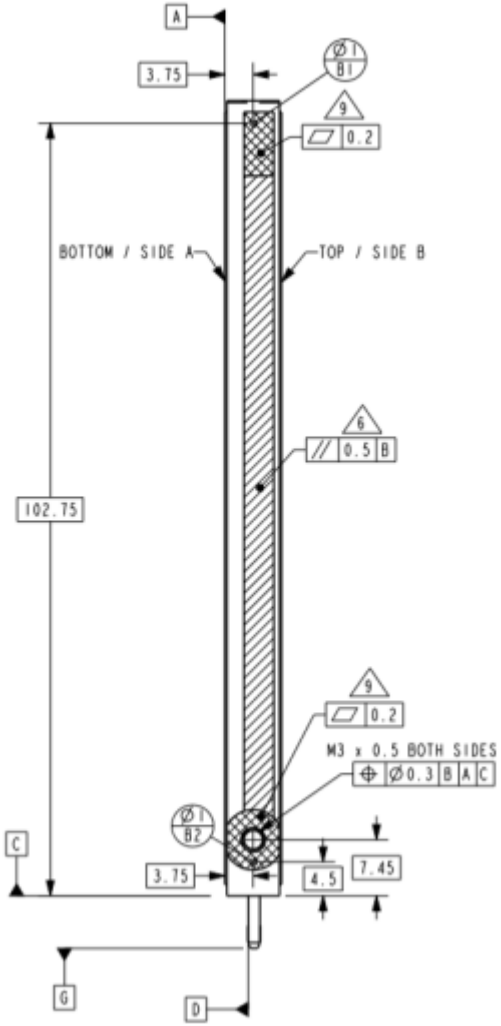

E3.S Form Factor

E3 Thermal Specs

- To Characterize Thermals for E3 use SFF-TA-1023
  - Defined test fixtures to determine airflow impedance level (AFI), MaxTherm, & DTherm

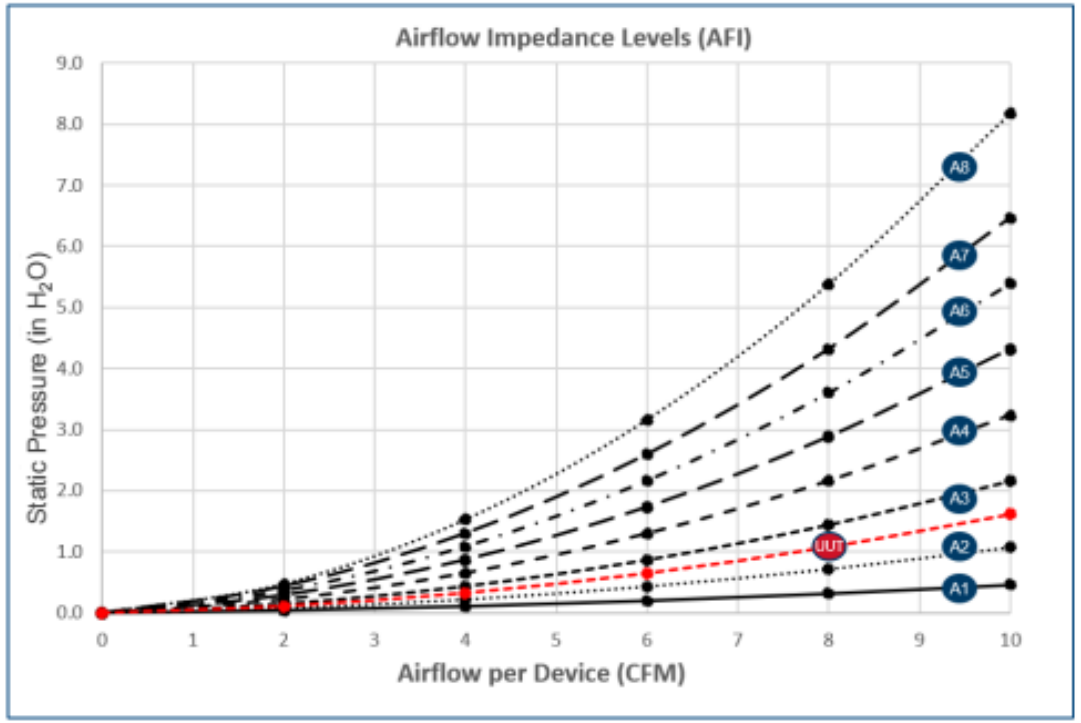

Figure 4-2: Example Device AFI Level

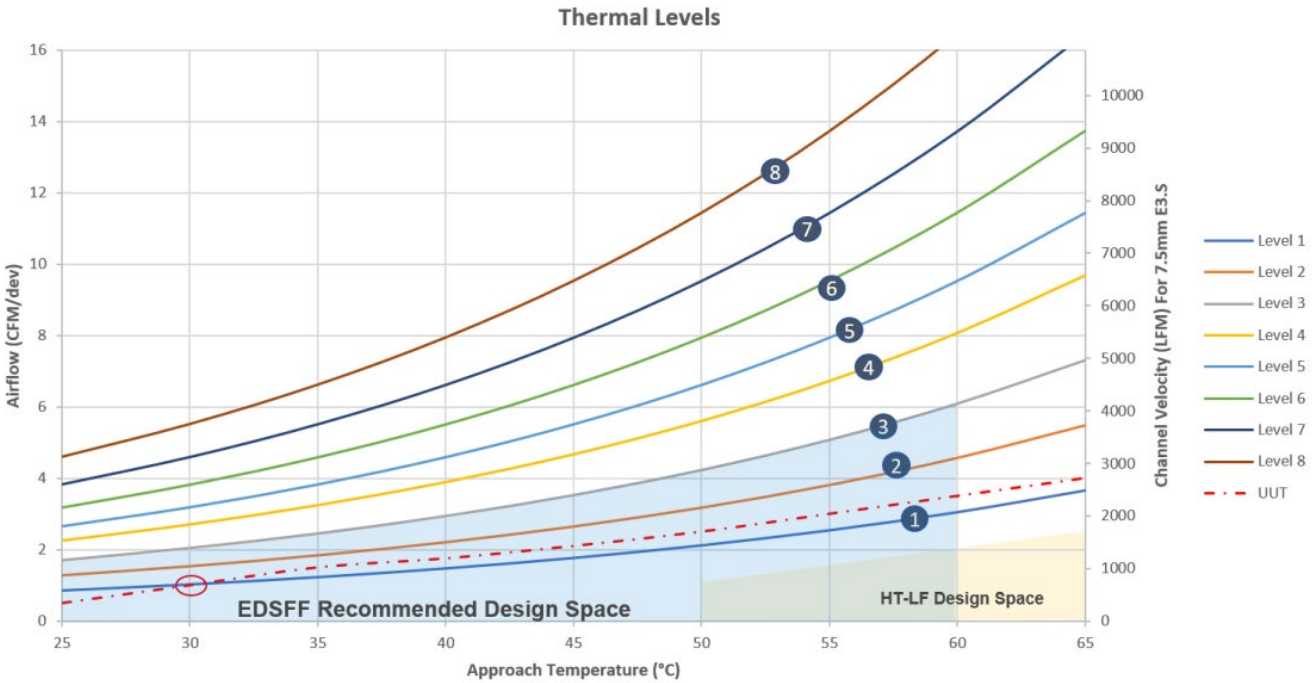

Figure 4-4: Example Device Thermal Profile

E3 SFF-TA-1008 Form Factor Updates in Progress

- Defining LED apertures to eliminate light bleed—proposing 1.8mm +/-0.20mm diameter hole
- Proposing new initial slot power requirements of 10W
- New 4C+ connector defined to enable OCP NIC 3.0 implementations
- Side-by-side 2x1C connector configuration for x8 lane width CXL applications
  - Card edge connector pitch of 9.3mm with +/-0.20mm tolerance to ensure blind-mate plugability

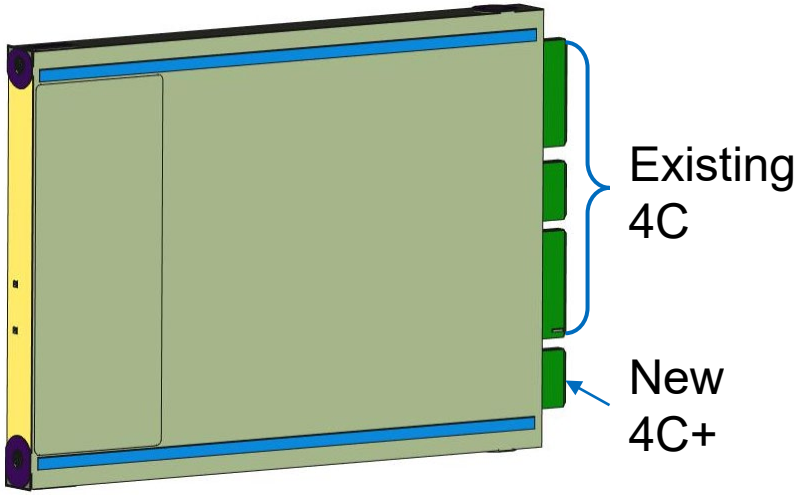

OCP NIC 3.0 Use Case

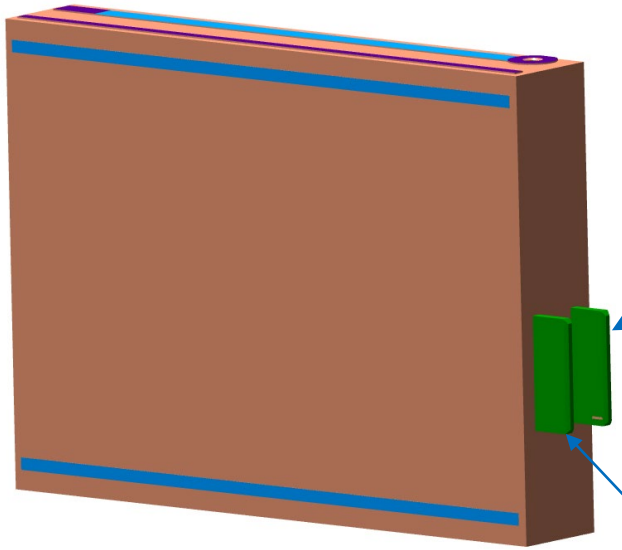

Flash Memory Summit

E1 specifications list and recent updates

Anthony Constantine, Intel

35 YEARS OF NAND FLASH MEMORY INNOVATION

E1: Optimized for 1RU

- E1.S design focuses on space, front panel efficiency to provide higher performance in a given space.
  - 3 enclosure designs for different airflow targets and density needs
    - 9.5mm, 15mm, and 25mm
    - 1 enclosure-less design for space constrained or custom applications
- E1.L design focuses on TCO by maximizing for capacity.
  - Front panel efficiency to provide more capacity for a storage focused system.
  - 2 enclosure designs (9.5mm and 18mm)

What Are The Benefits Of E1?

- E1.S with ***same PCB and firmware supports a diverse family of thermal options for the market***:
  - High density
  - High performance
- Hot plug support
- Works for both storage and compute in 1OU
- Support for PCIe® 5.0 and 6.0
- Fully standardized in SNIA/SFF:
  - SFF-TA-1006 Revision 1.5

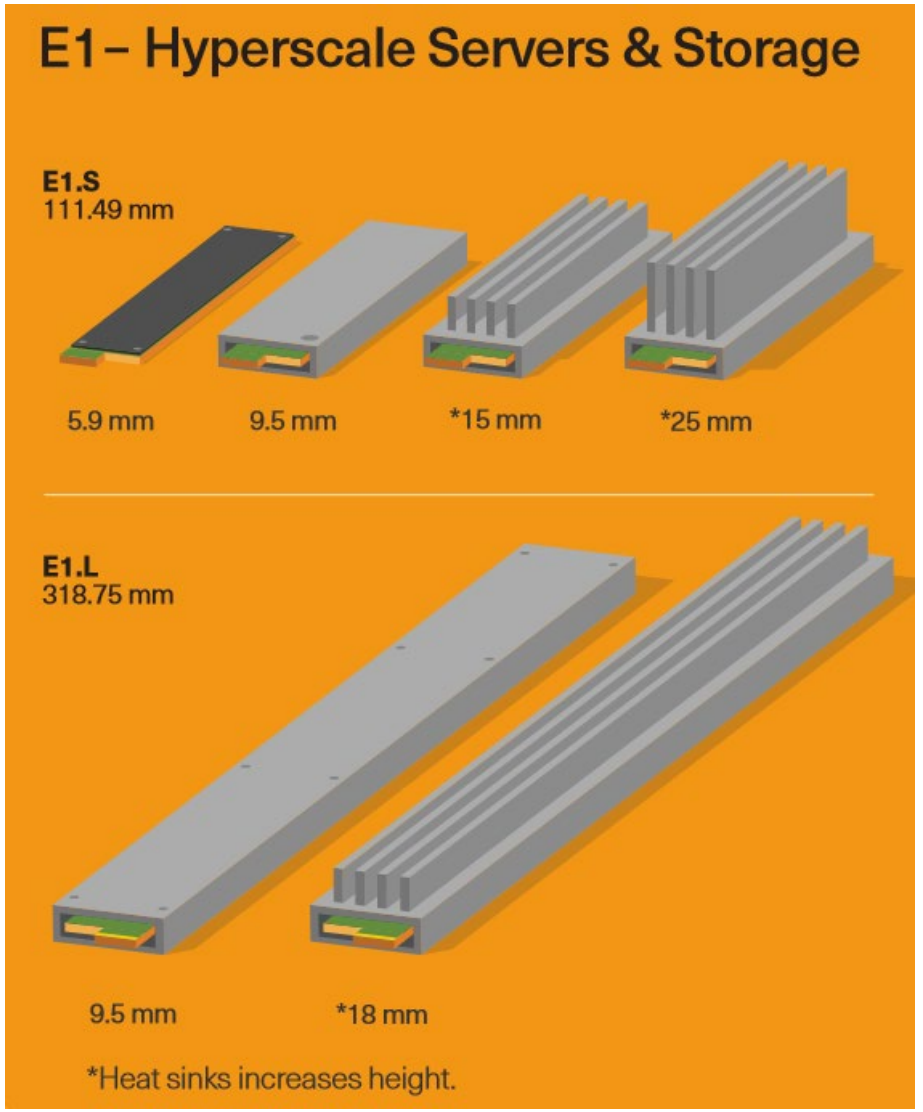

How To Make an EDSFF Device

Flash Memory Summit

Feature	Related Org. or Standard	Specification
Electricals, PHY	PCI-SIG	**PCIe Base Spec**
Command Set	NVMe or CXL	**NVMe Spec**; **CXL Spec**
Pinout, Power	SNIA SFF-TA	**SFF-TA-1009** Pin/Signal Spec
Connector	SNIA SFF-TA	**SFF-TA-1002** Connector Spec
Form Factor	SNIA SFF-TA	**SFF-TA-1006** E1.S Mechanical; **SFF-TA-1007** E1.L Mechanical; **SFF-TA-1008** E3 Mechanical
Other	SNIA SFF-TA	**SFF-TA-1023** Thermal spec

E1.S History

Flash Memory Summit

SNIA

Q1 2018
SFF-TA-1006, rev 1.0

Q2 2018:
SFF-TA-1006, rev 1.1
-Errata and minor cleanup

Q2 2019
SFF-TA-1006, Rev 1.2
-"E1.S"
-Added 9.5mm and 25mm versions
-Errata and cleanup

Q3 2019
SFF-TA-1006, Rev 1.3
-Added x8 connector option
-Errata and cleanup

Q1 2020
SFF-TA-1006, Rev 1.4
-Added 15mm version

Q3 2021
SFF-TA-1006, Rev 1.5
-Reference new thermal spec (SFF-TA-1023)
-Errata and cleanup

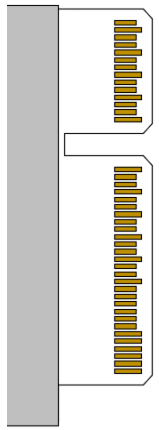

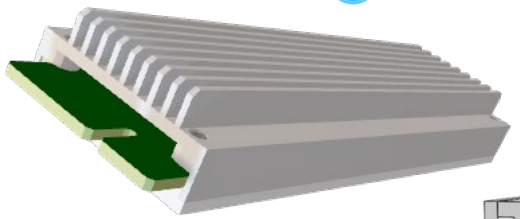

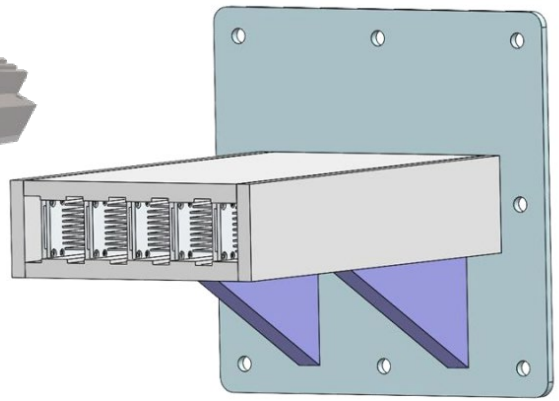

35 YEARS OF NAND FLASH MEMORY INNOVATION

Coming soon to EDSFF: Optional I3C support

- Needed for sideband usages requiring better performance than SMBus
- Uses same signals as SMBus
- SMBus is still required to keep backwards compatibility.
- Draft specification is posted
  - https://www.snia.org/sff/specifications

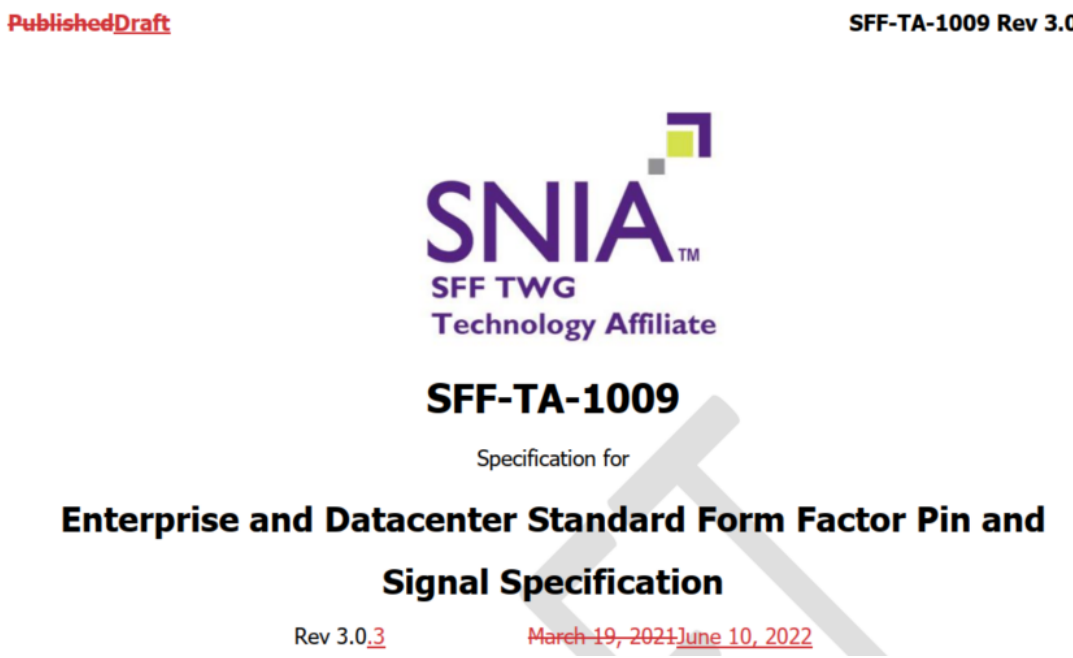

Flash Memory Summit

E1 form factor overview

Hyperscale Form Factor Update

Ross Stenfort, Meta

PCIe® Enterprise/Datacenter Form Factor Trends

TRENDFOCUS

2021

2026

Flash Memory Summit

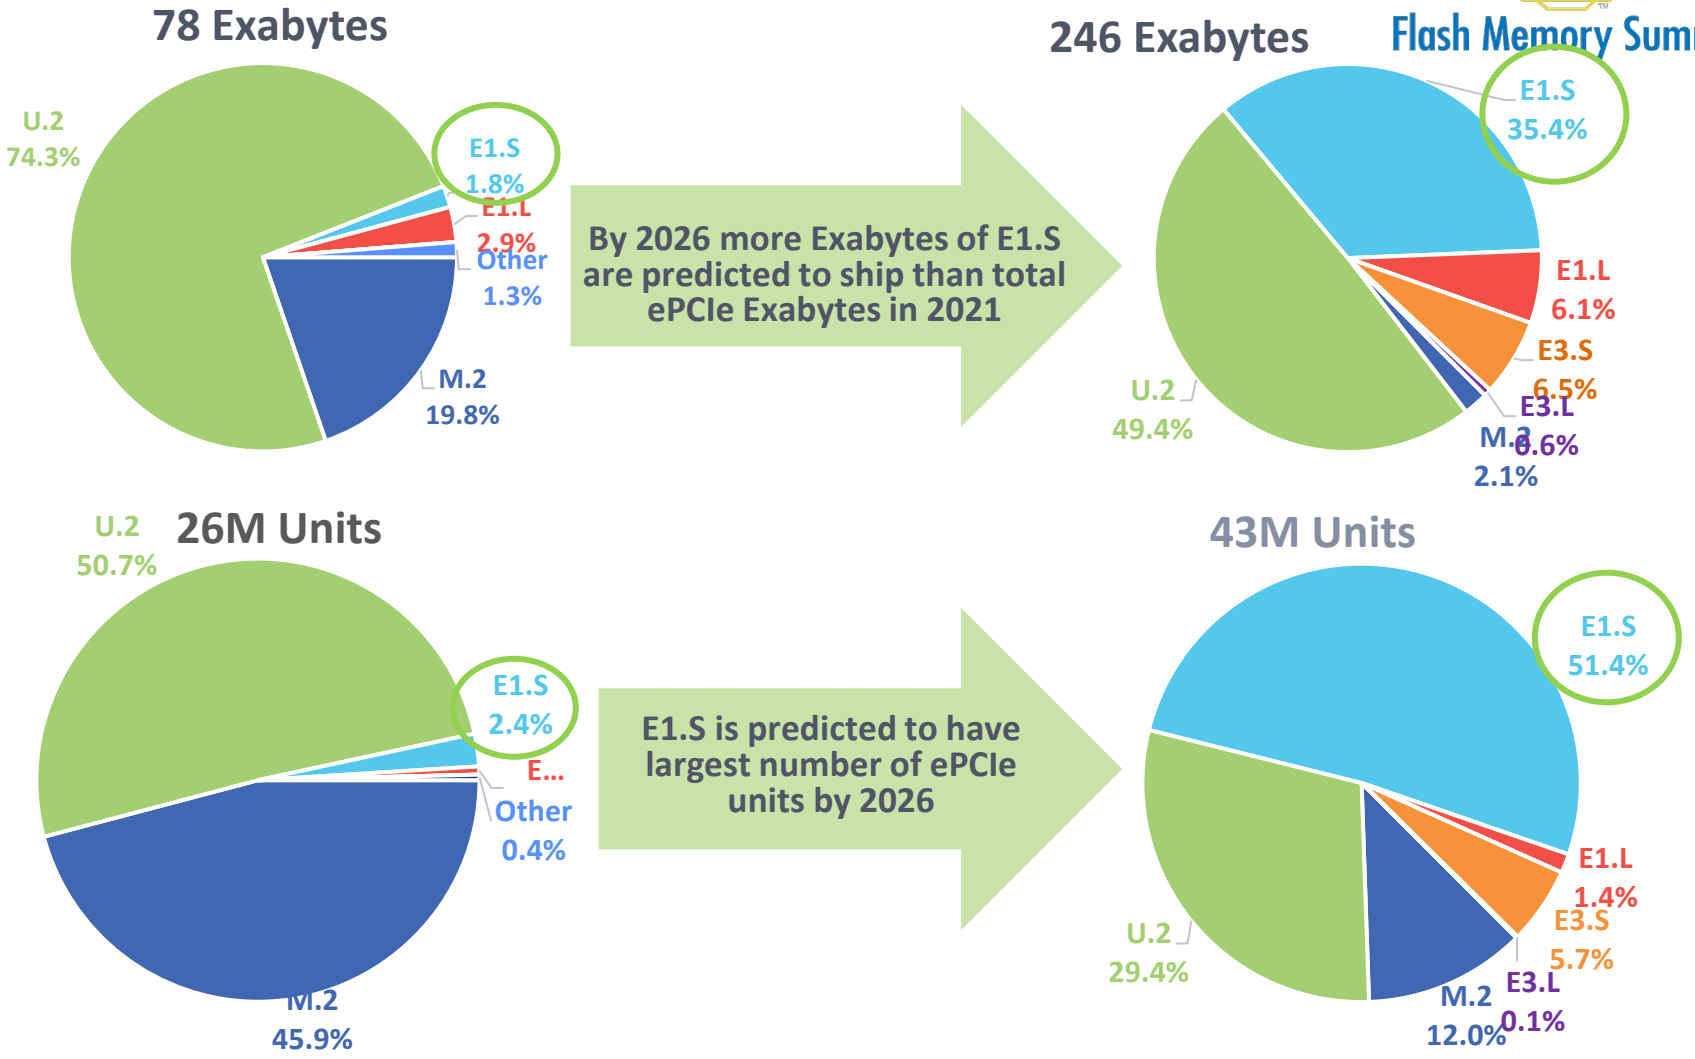

E1.S market growth is very significant.

Note: Data excludes SSD consumption where companies buy NAND and build SSDs for internal use.

35 YEARS OF NAND FLASH MEMORY INNOVATION

Real World E1.S Systems

Flash Memory Summit

Hyperscale OCP Based Systems

4OU Chassis with 48
25mm E1.S SSDs

4OU Chassis with 24
25mm E1.S SSDs

Link to OCP YV3 Contributions:

Yosemite V3: E1.S Faceplate:

https://www.opencompute.org/documents/e1s-faceplate-reference-design-specification-pdf

Yosemite V3: Vernal Falls E1.S 1OU Flash Blade and Expansion Board

https://www.opencompute.org/documents/e1s-expansion-1ou-1s-server-design-specification-pdf

Yosemite V3: Sierra Point E1.S 2OU Flash Blade and Expansion Board

https://www.opencompute.org/documents/e1s-expansion-2ou-1s-server-design-specification-pdf

Yosemite V3 Platform Design

https://www.opencompute.org/documents/ocp-yosemite-v3-platform-design-specification-1v16-pdf

Delta Lake 1S Server Design

https://www.opencompute.org/documents/delta-lake-1s-server-design-specification-1v05-pdf

Enterprise Based Systems

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

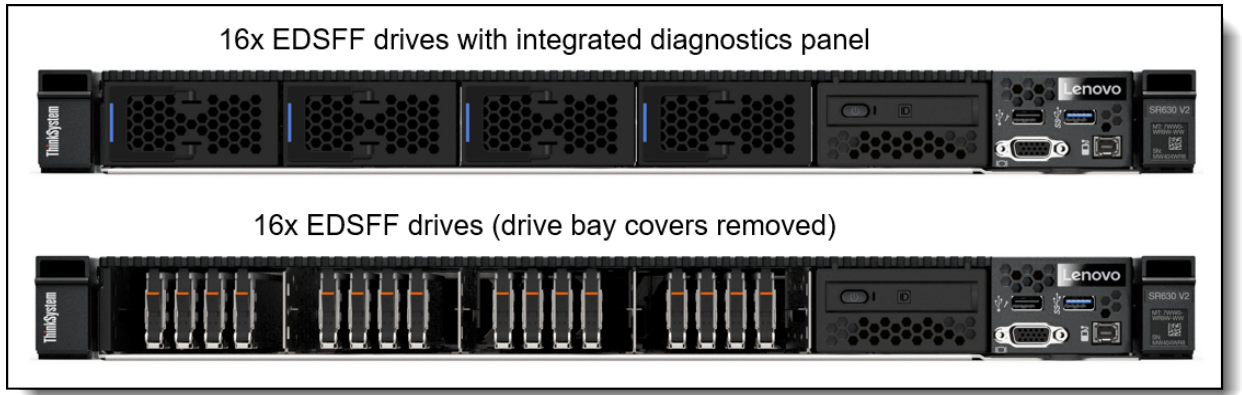

35 YEARS OF NAND FLASH MEMORY INNOVATION

Why is E1.S Growth/adoption So Rapid?

Flash Memory Summit

Hyperscale Challenge	E1.S Solution
Performance	• Designed for Gen 5 PCIe and beyond • Able to scale power/thermal for performance
Thermal	• Variety of heat sinks to select based on thermal needs • 25mm provides best thermal solution for low airflow
Density	• Small/ dense form factor fitting in 1OU
Serviceability	• Front Panel/ Hot Swap supported
Security	• Presence Detect

E1.S solves market needs

E1.S Thermal Guide

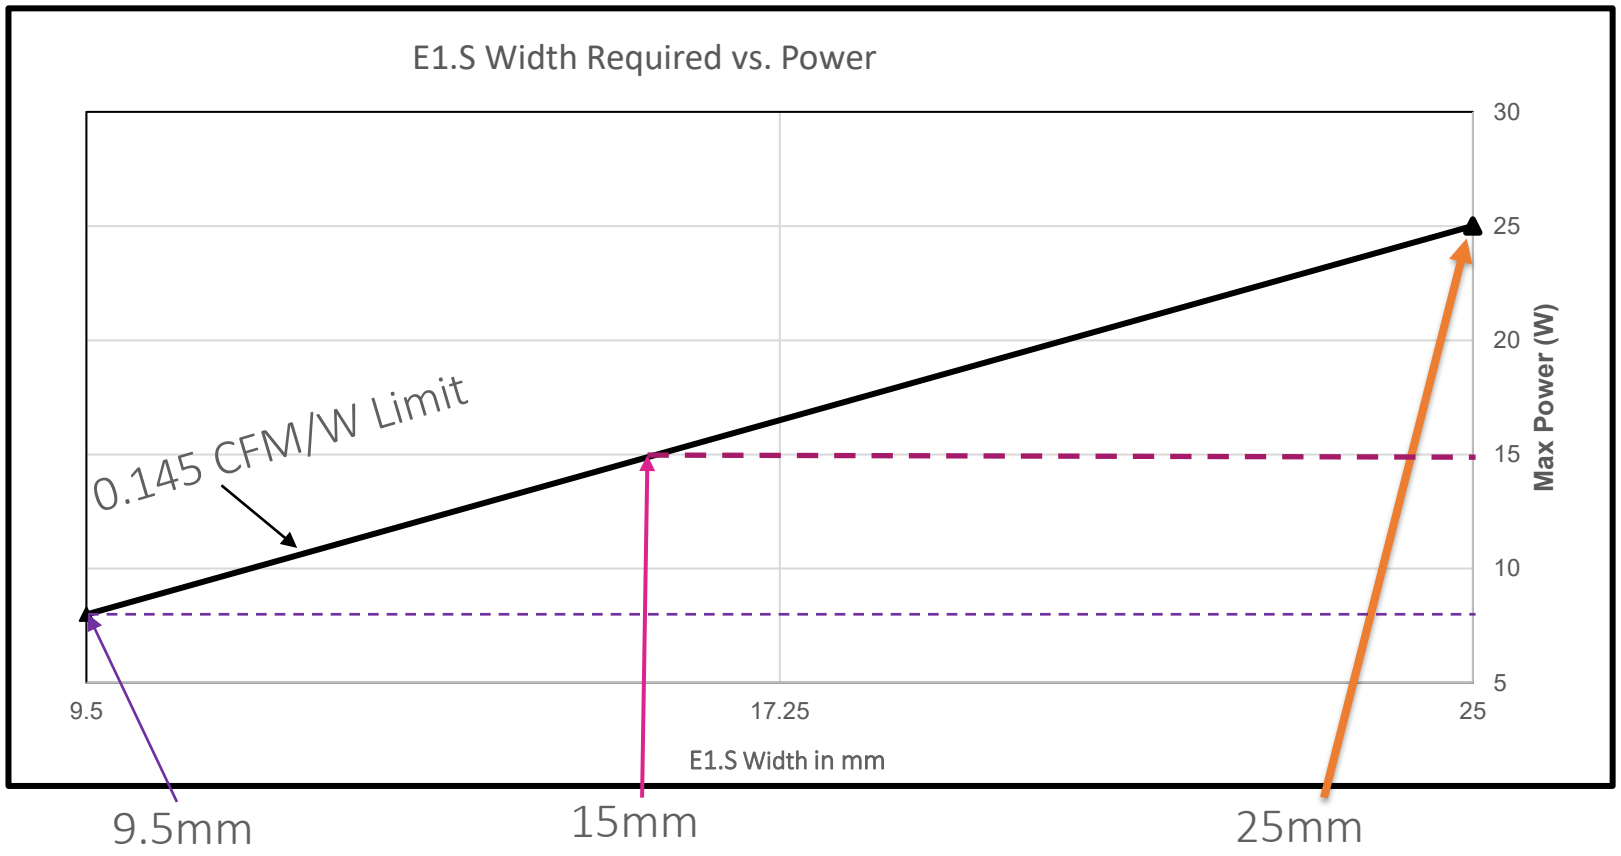

YV3 4OU Chassis

- 48 25mm SSDs (25mm)
- Up to 768 TB

YV3 4OU Chassis

- 36 25mm SSDs (25mm)
- Up to 576 TB

Graph Assumptions

- Thermal limit based on 0.145 CFM/W
- Assumes OCP Yosemite V3 (YV3) style chassis

Note: *25mm enables the best power and thermal scaling*

Based on power/thermal needs appropriate size heat sink can be used.

Flash Memory Summit

E1 use cases, device types and applications

Lee Prewitt, Microsoft

35 YEARS OF NAND FLASH MEMORY INNOVATION

E1.S Use Cases

- Multiple swim lanes
  - Data
  - Cache
  - Boot
- Multiple standardized widths
  - 5.9mm
  - 8mm
  - 9.5mm
  - 15mm
  - 25mm

Density Matters

- IOPS matters for high performance compute and storage
- Front panel efficiency matters for dense 1U systems

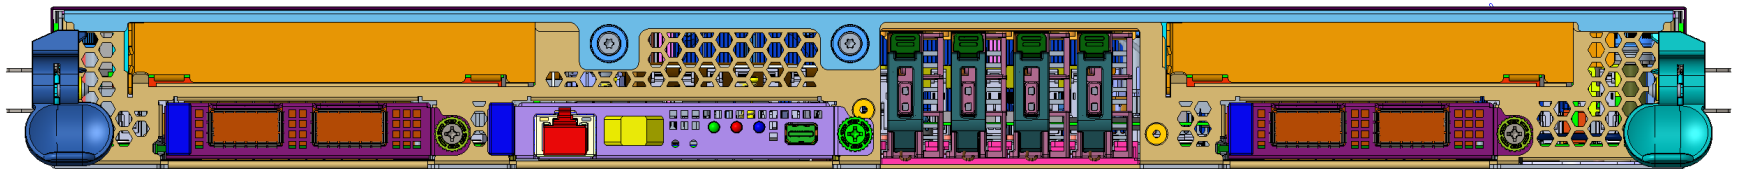

Concept Project Olympus Design

E1.S – Use Cases

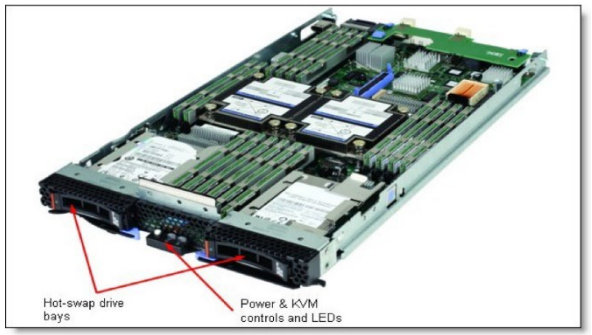

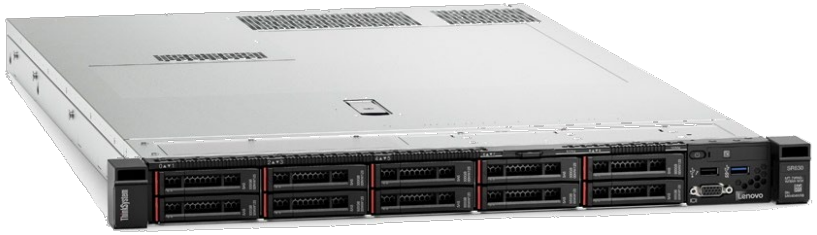

Smaller footprint systems:
- Blade servers
- Edge compute and storage
- Dense, scaled servers

NEED: Performance, excellent thermal
USE: **U.2** and **M.2** replacement
Typical: 6-12 x E1.S

1U systems:
- General purpose enterprise compute
- Rack configurations with 32-40 servers
- Hyperscale environments

NEED: scaled performance and capacity, excellent thermal
USE: mainstream **U.2** replacement
Typical: 16-32 x E1.S

2U and larger systems:
- Storage appliance
- IO and storage-rich server / database
- Performance-oriented storage systems

NEED: scalability of performance with capacity, lower cost/performance
USE: mainstream **U.2** replacement
Typical: 48-64 x E1.S

E1.S 15mm

- Best width to balance density with thermal headroom
- Works for all three swim lanes
- Broadly adopted by suppliers

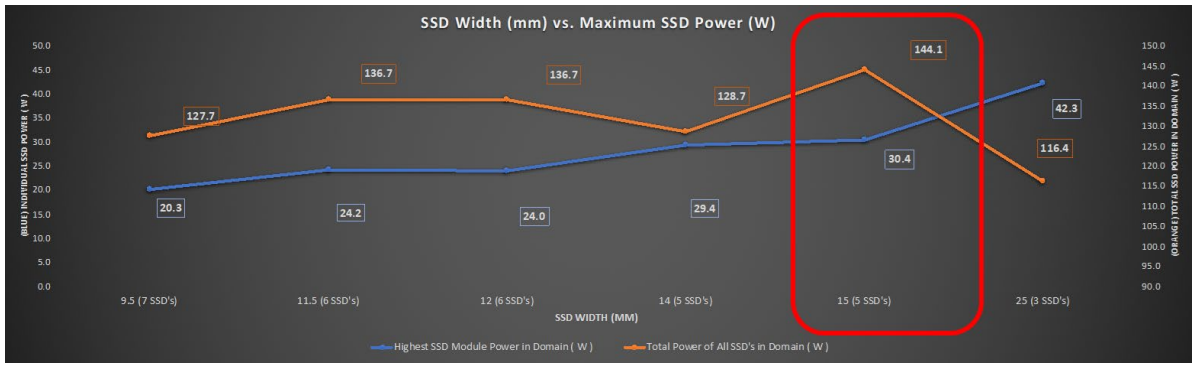

E1.S 15mm

Flash Memory Summit

- 15mm E1.S enables more flexible designs
- Ability to remove m.2's, and their limited performance due to thermals
- Can run SSDs at full performance, no throttling
- Enables hot-plug
- 15mm will support Gen 4 and Gen 5 based SSD capabilities
- In addition to traditional SSDs, other use cases can be supported
- M.2's may linger for boot devices, but for performant data storage it will fade away

Now in broad adoption across Azure

Thank you!
Please be sure to fill out the survey

SSDS-201-1: Form Factors and Interfaces Part 1

Wednesday, August 3
Sessions

EDSFF Update: Industry Leaders Explain E1&E3 Innovations & Specification Updates

Abstract:

- Enterprise and Datacenter Standard Form Factor (EDSFF) SSDs include E1 and E3 types and allow higher system density and more effective forms for hyperscale and enterprise data center servers and storage. They utilize a common connector that can also be used for GPUs, CXL and advanced network interface cards. Builders of servers, all-flash arrays (AFAs), and storage appliances can use EDSFF devices to develop systems with a new level of flexibility, allowing a single configuration to serve a variety of applications such as databases, AI, virtual reality, video/image processing, IoT, cybersecurity and more. Learn how EDSFF SSDs will replace M.2 and 2.5” form factors in hyperscale and enterprise environments, and the latest developments of the specifications.

EDSFF Update: Industry Leaders Explain E1&E3 Innovations & Specification Updates

Flash Memory Summit

- **Panelists:**
  - Cameron Brett – Moderator
  - E3 – Bill Lynn, Dell Technologies (William.Lynn@dell.com)
  - E3 – Paul Kaler, HP Enterprise (paul.kaler@hpe.com)
  - E3 – John Geldman, KIOXIA (john.geldman@kioxia.com)
  - E1 – Ross Stenfort, Meta (rstenfort@fb.com)
  - E1 – Lee Prewitt, Microsoft (leprewit@microsoft.com)
  - E1 – Anthony Constantine, Intel (anthony.m.constantine@intel.com)

- **Agenda:**
  - Welcome, format of session and panel introduction (questions at the end of the panel presentations)
  - E3 – form factor overview, use cases/applications, specifications updates
  - E1 – form factor overview, use cases/applications, specifications updates
  - Questions

35 YEARS OF NAND FLASH MEMORY INNOVATION