Flash Memory Summit

SSDS-201-2:
Form Factors and Interfaces Part 2

The Challenges of EDSFF in Enclosure Design

Brent Yardley
IBM STSM, Storage Hardware Architect, Master Inventor
yardleyb@us.ibm.com

35 YEARS OF NAND FLASH MEMORY INNOVATION

Brent Yardley is a Senior Technical Staff Member and Master Inventor with IBM System Storage, where he focuses on developing All Flash Arrays (AFAs). Brent is currently the overall Chief Hardware Architect responsible for the hardware architecture, design, and integration of IBM's storage hardware products and planning future generation storage platforms. He specializes in system designs that integrate multiple I/O protocols, FPGAs, and ASICs, and is an expert in both hardware and software design and system integration.

A 22-year veteran of IBM, Brent holds multiple patents focused on storage architectures and solutions. He has an extensive background and understanding of the architecture, design, and implementation of highly available storage-based systems and solutions. He has earned BS degrees in both Software and Hardware Engineering from the Oregon Institute of Technology.

The Layout of IBM FlashCore™ Module

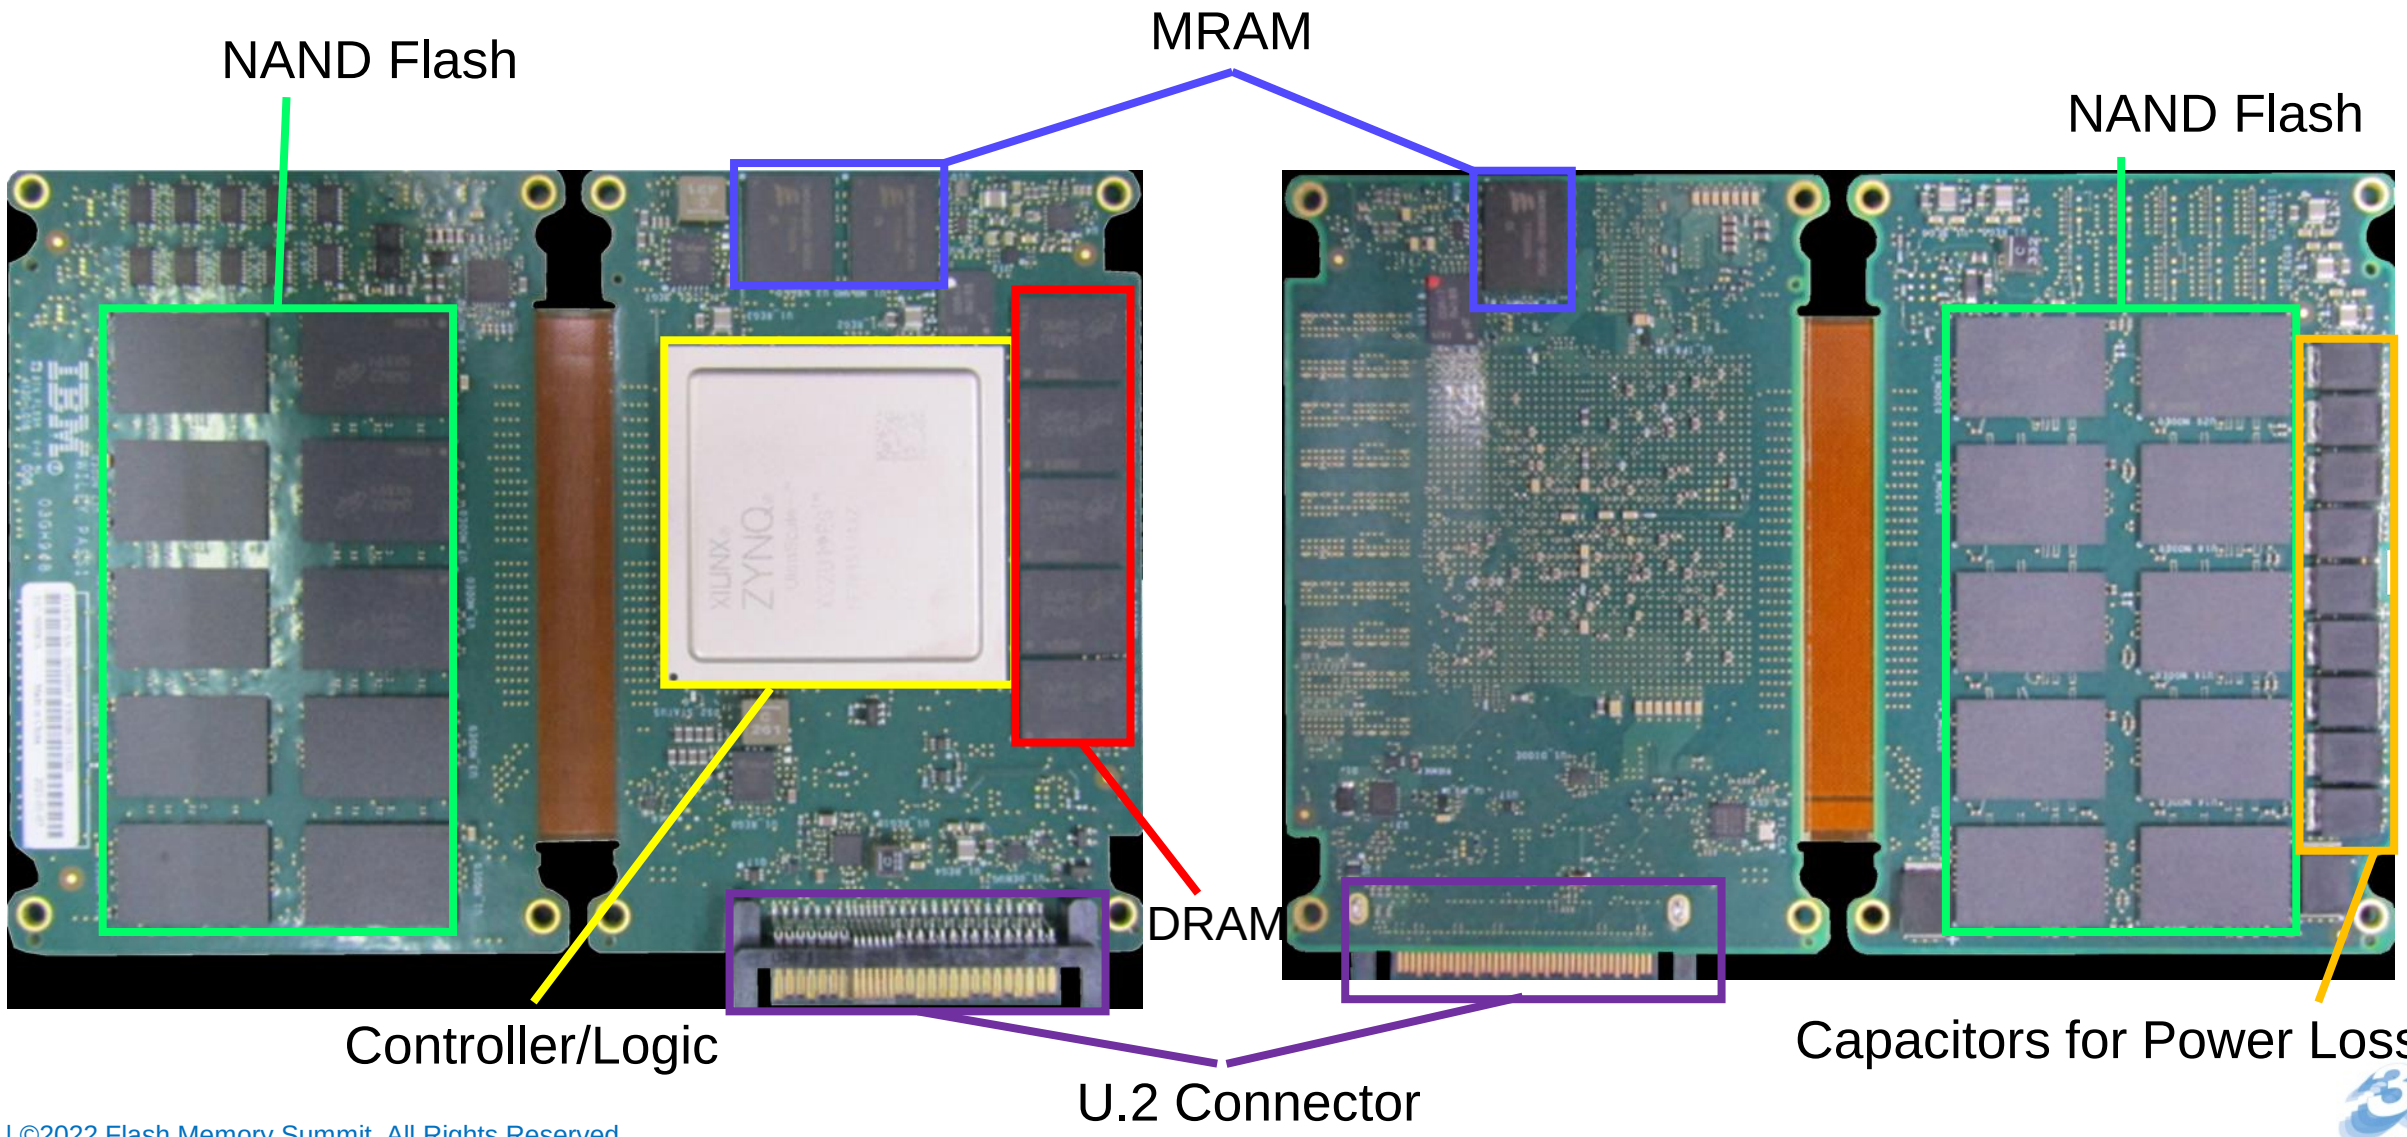

E3.S/L Floor Planning Concept

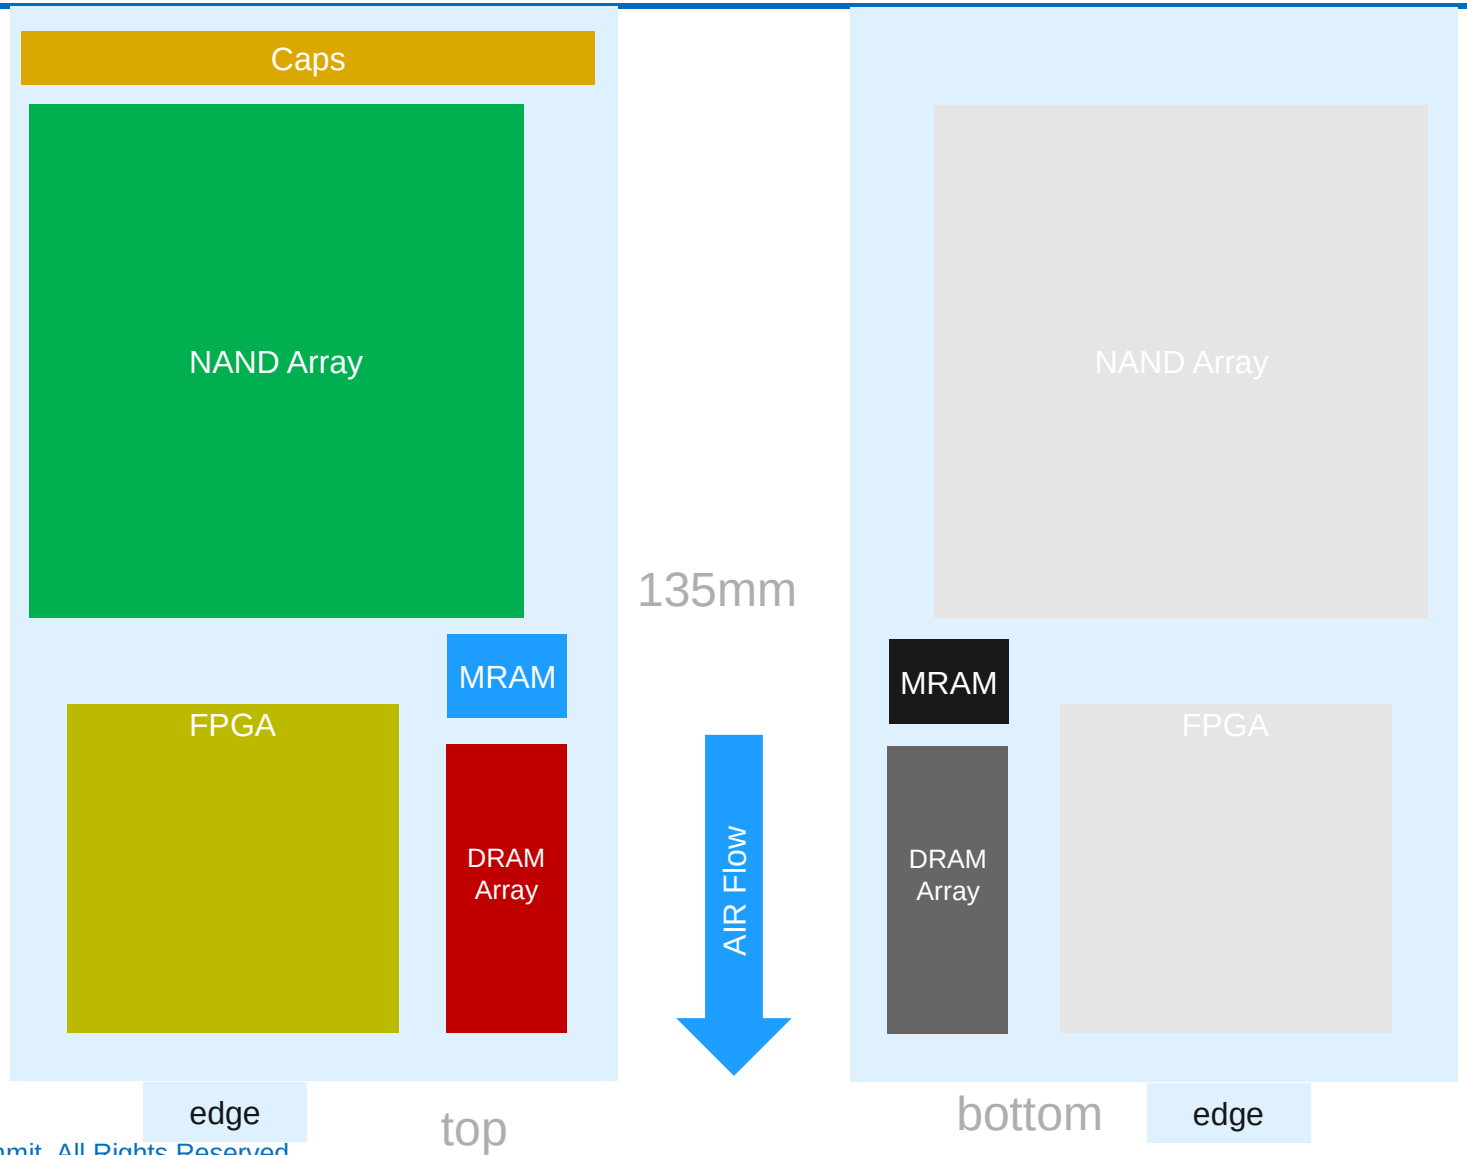

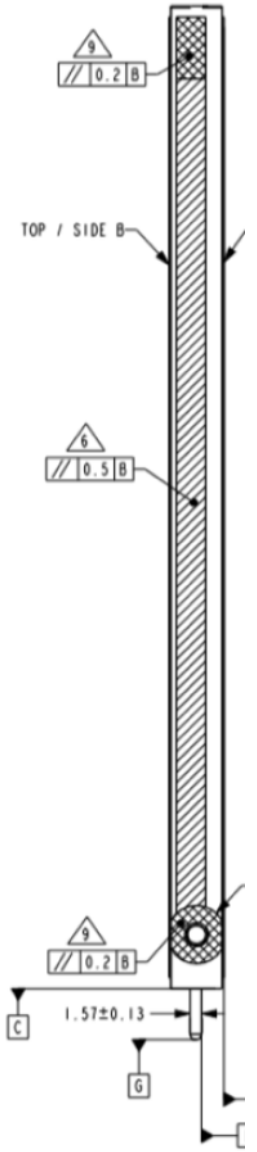

Stack Up Challenges Moving form U.2 2.5” 15mm to E3.S/L

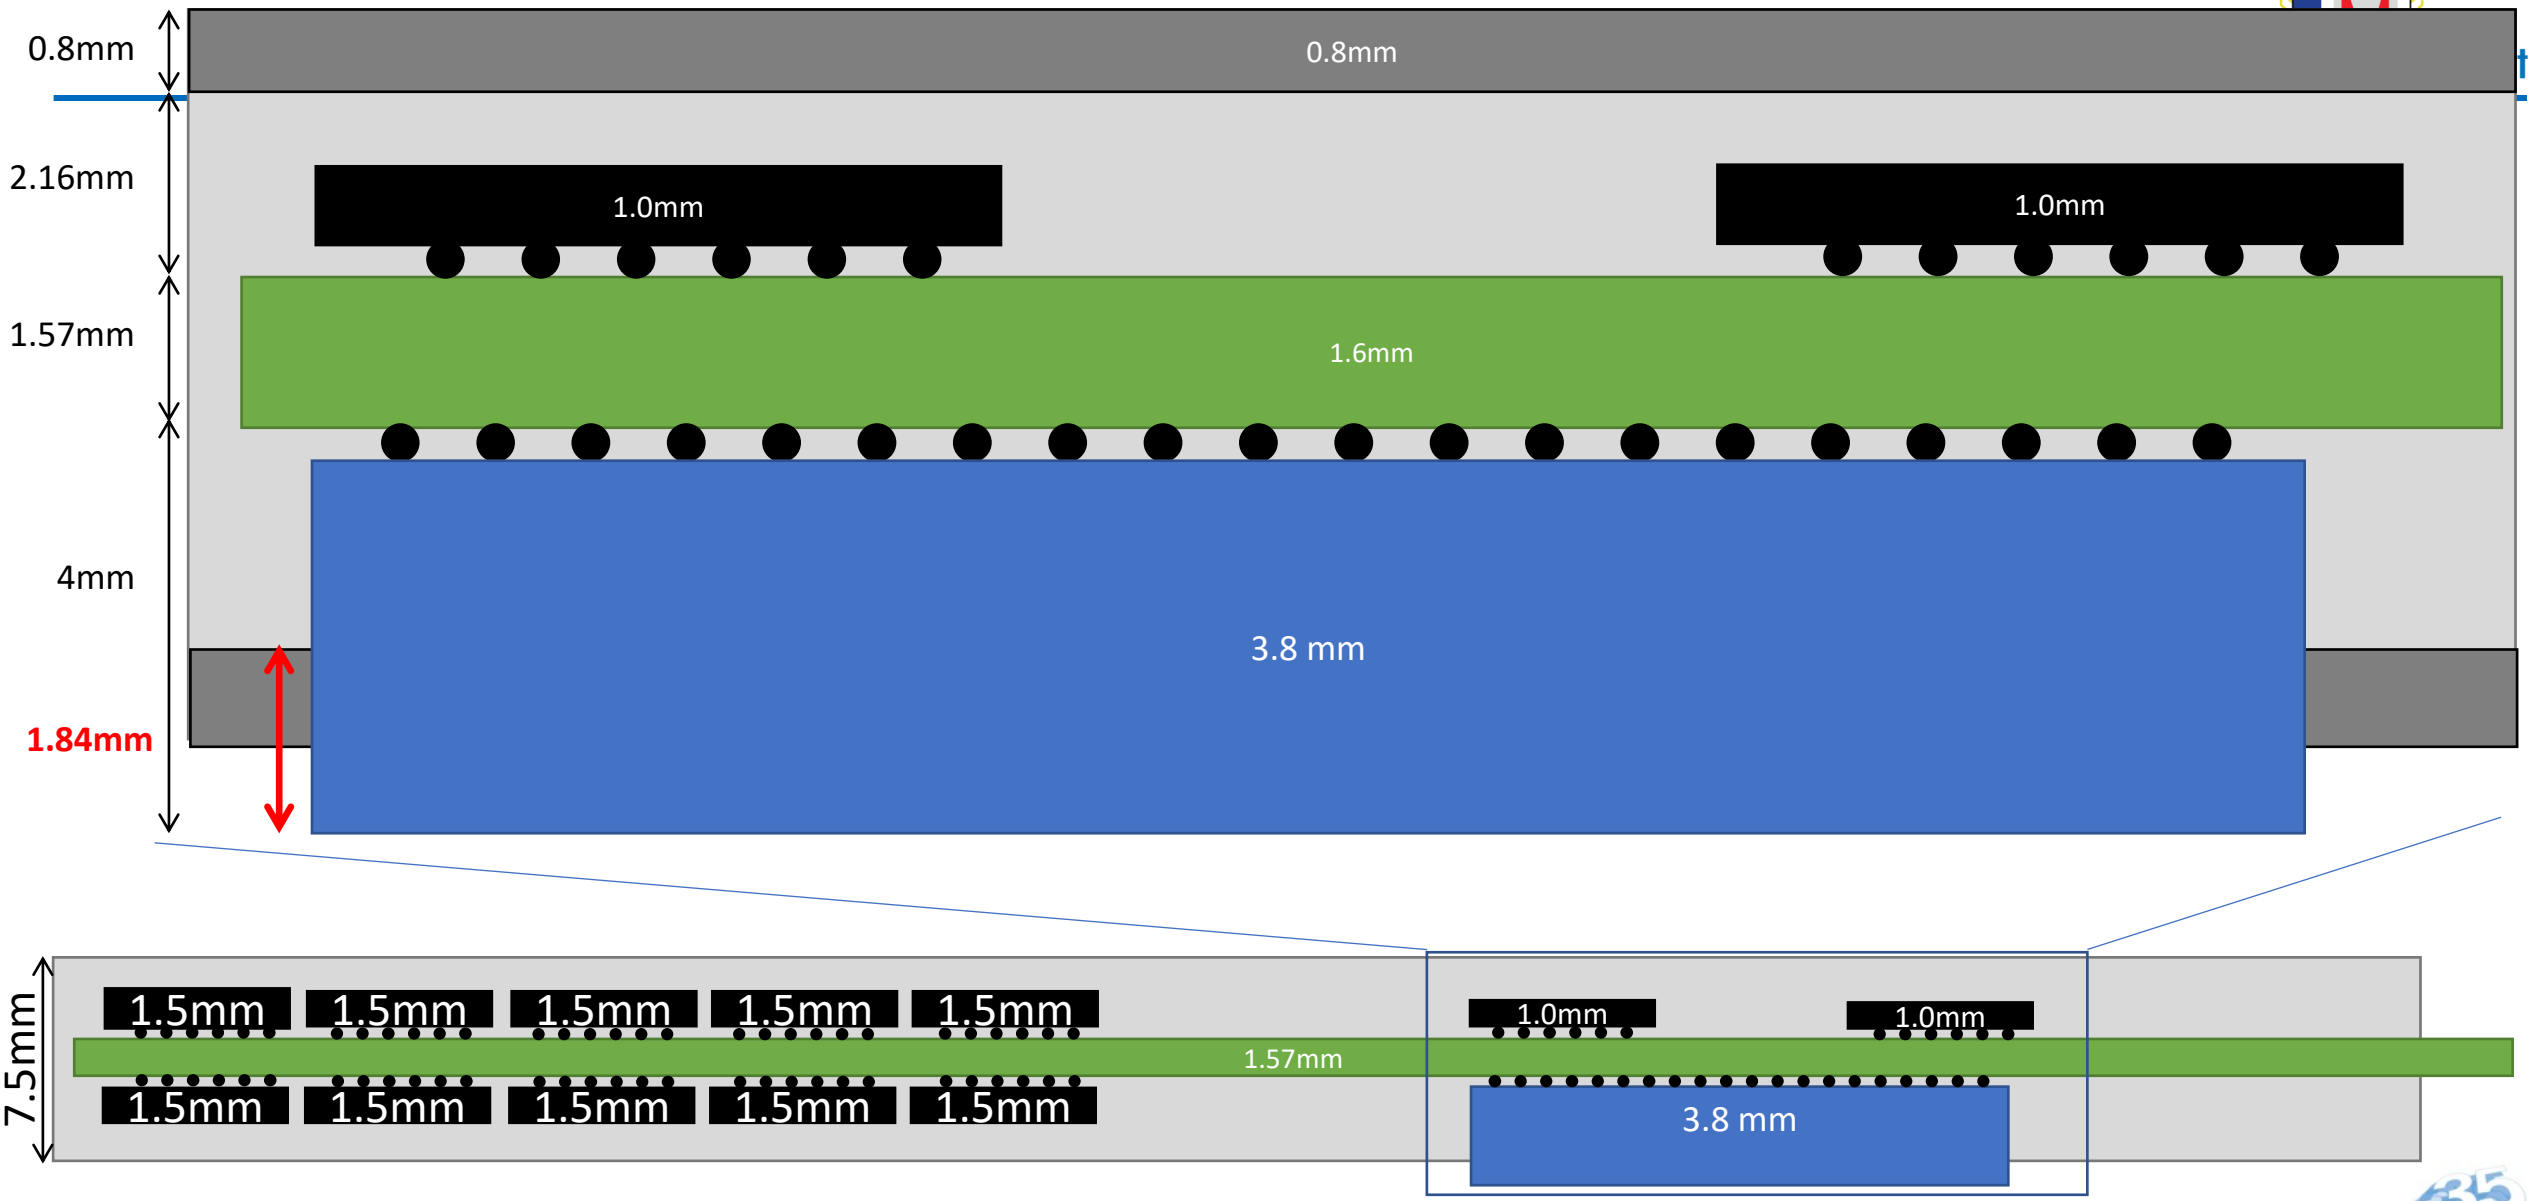

Power Design

- U.2
  - 15mm drives had been designed to up to 25W, however it was found that 25W was a limiting factor for the design
  - Some solved this issue by designing enclosures to support 30W per slot, to provide extra power density.
- E3
  - The design point for E3.L is up to 40W, and up to 70W for 2T. However, with the number of slots we are working to support, the enclosure design will likely be designed for a target of ~35W per slot for E3.S/L, and ~60W for 2T.
  - The biggest challenge is trading of the power and thermal requirements for the additional slots, while supporting next generation CPUs TDP.

EDSFF Thermal Design Considerations

- SFF-TA-1023 recommends E3 device to operate within the shaded design space envelope (Level3 in the left picture)
  - Spec defines the SSD-SSD pitch For E3.S/L to be 1.8mm
- Ta 35C 2.5CFM can be translated to 1700 LFM (8.6 m/s)
  - 2.5CFM results 1697LFM as follows using Equation 5-1.
  - Compared to 2.5" 7mmT (pitch 12.5mm, gap 5.5mm), LFM is much higher from smaller gap (5.5mm vs. 1.8mm). But regarding 2.5" 15mmT (pitch 16.5mm, gap 1.5mm), we will see little differences
- Channel velocity will vary based on chassis design SSD to SSD pitch but can be calculated easily.

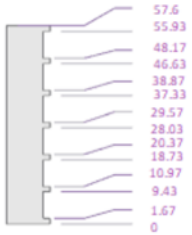

[SFF-TA-1023 R1.0a] Figure 4-3: MaxTherm and DTherm Levels

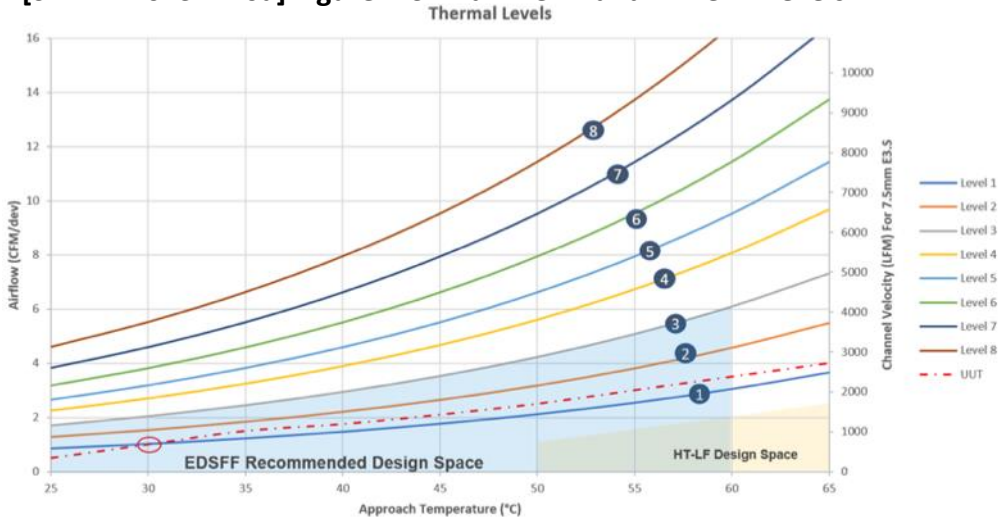

Equation 5-1: $Channel\ Velocity\ (LFM) = \frac{\frac{CFM}{Dev} * 92\,903}{(Drive\ Pitch\ (mm) - Drive\ Thickness\ (mm)) * Drive\ Width\ (mm)}$

Example: $Channel\ Velocity\ (LFM) = \frac{2.5\ CFM/dev * 92\,903}{(9.3mm - 7.5mm) * 76mm} = 1\,697\ LFM$

[SFF-TA-1023 R1.0a] Figure 5-5: Required Dimensions of E3 1T Test Fixture

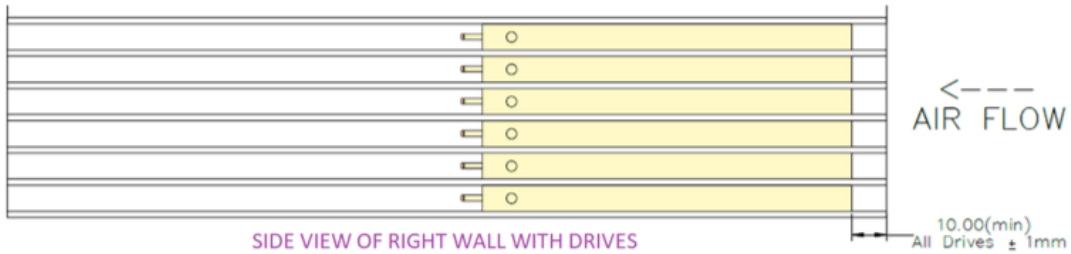

Mechanical Enclosure Slot Design

- To meet the thermal envelope requirements, a design is proposed using an approx. 3mm gap between E3.S/L devices.
- Using this same spacing, for 2T device in the same slot as two E3.S/L devices this provides for a slightly larger gap of 4.2mm.

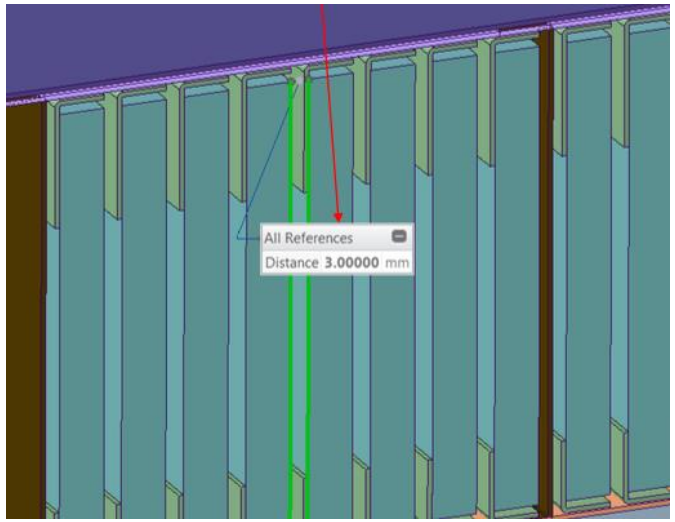

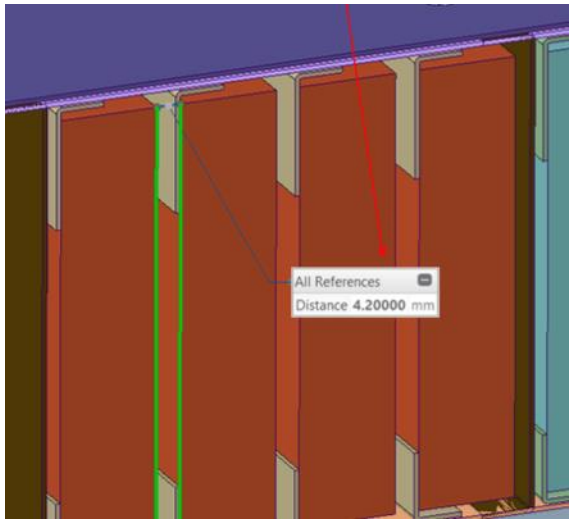

E3 1T Carrier Designs

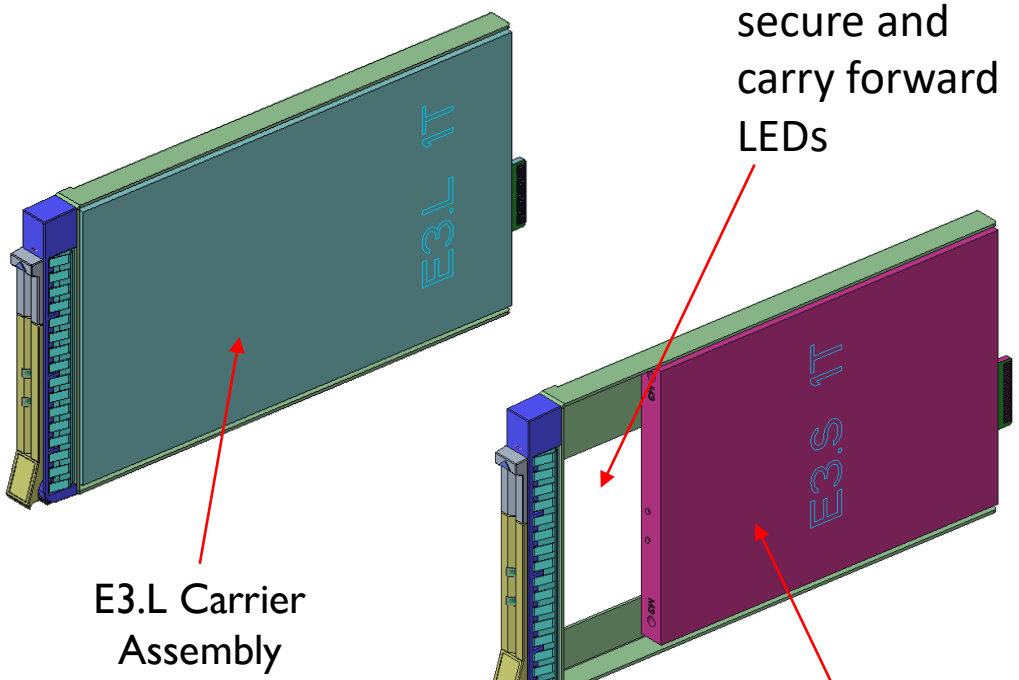

Spacer, to secure and carry forward LEDs

E3.L Carrier Assembly

E3.S Carrier Assembly

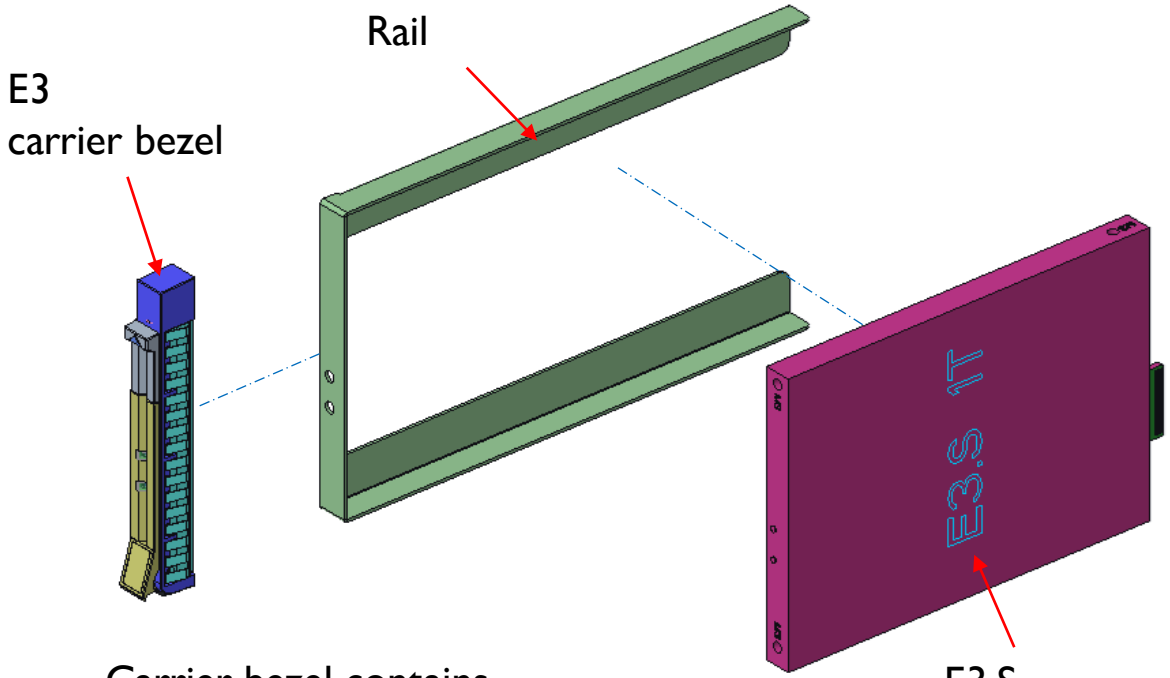

E3 carrier Rail

E3 carrier bezel

Carrier bezel contains lever and EMI gasket

E3.S SSD

E3 2T Carrier Design

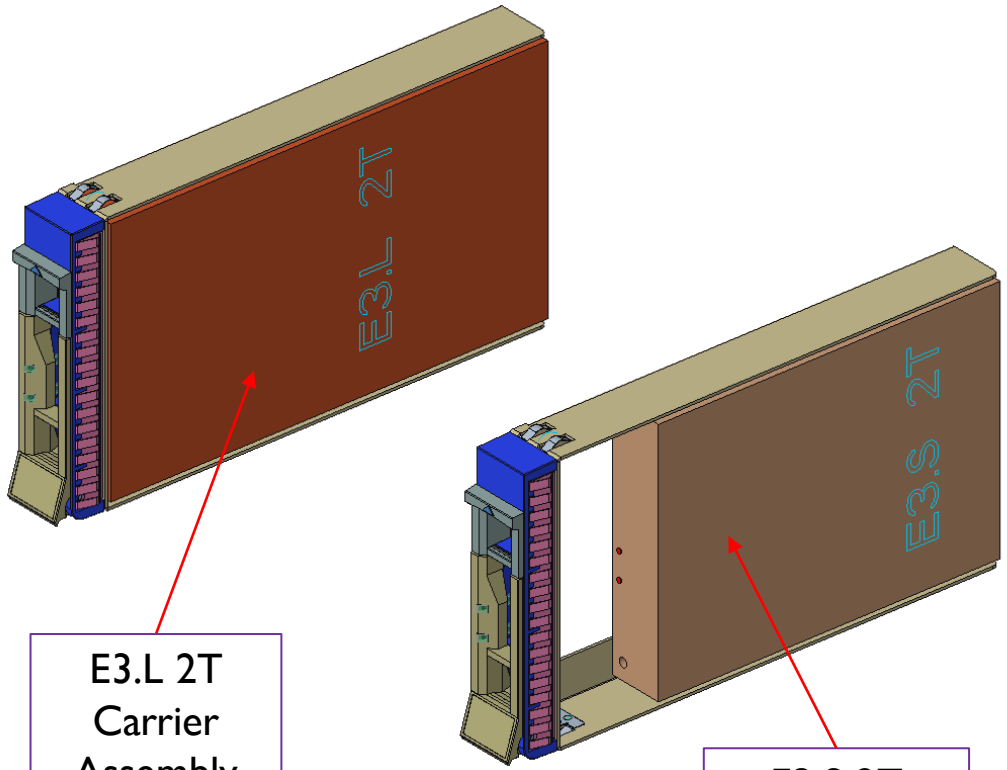

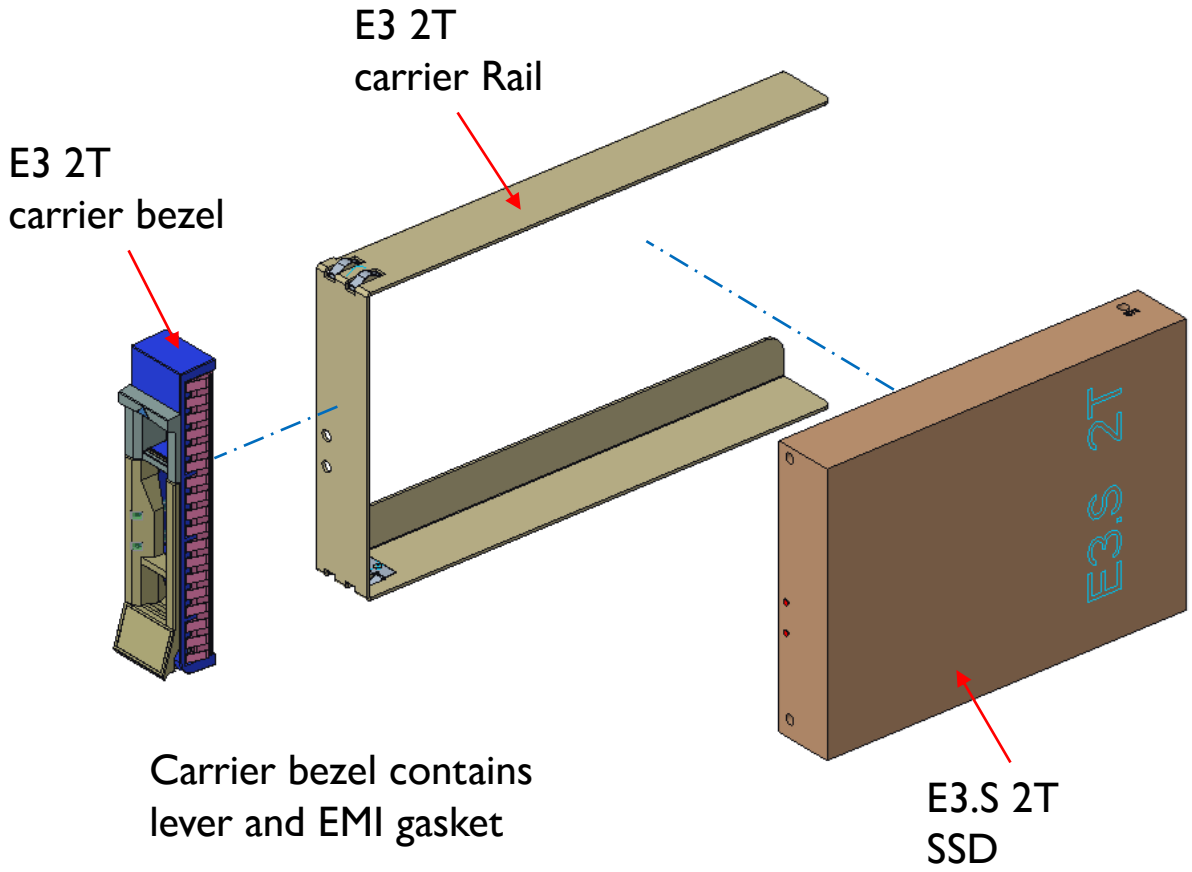

Enclosure Configurations

- In a white paper, publish in Sept, 2021, B. Lynn, P. Kaler, and J. Geldman, several enclosure configuration options were explored.

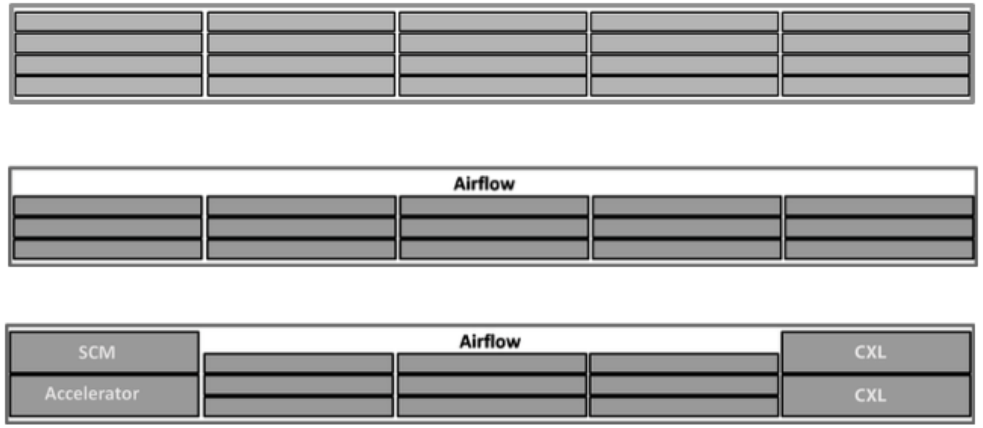

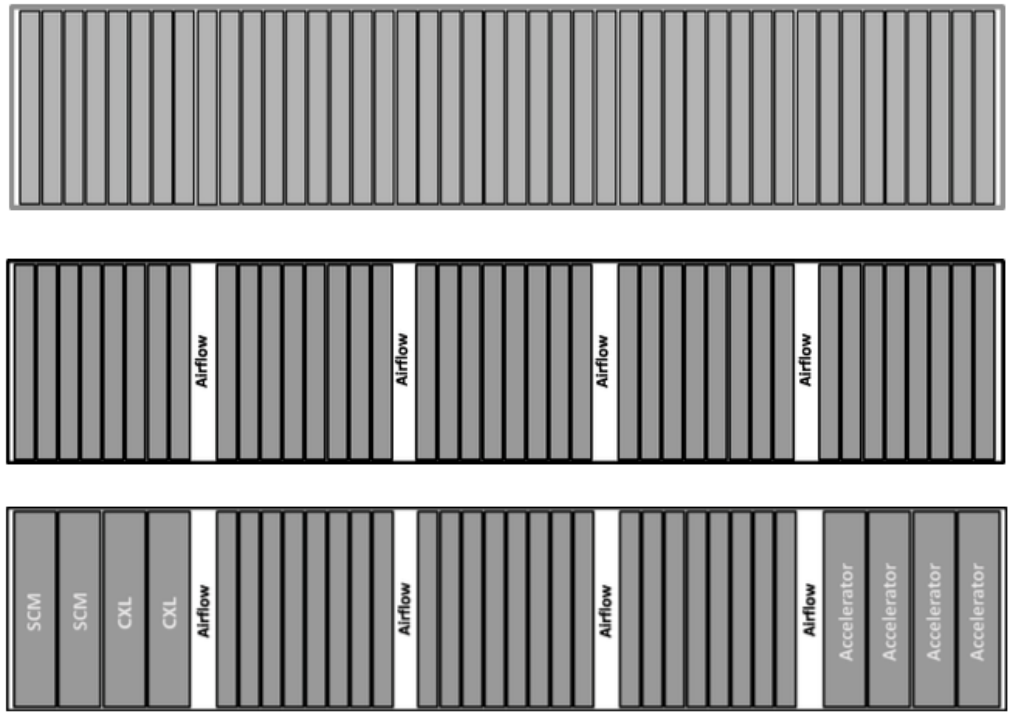

"The flexibility of the E3 form factor gives platform architects a wide range of options when it comes to supporting multiple system use cases. The ability to optimize around either density, host bandwidth, system power or device type makes the E3 form factor an ideal choice for platform architects and system designers. "

https://business.kioxia.com/content/dam/kioxia/ncsa/en-us/business/asset/KIOXIA_EDSFF_Intro_White_Paper.pdf

Enclosure Configuration Considerations

- We are looking at specifically 3 different enclosure types.
  - 1U – Entry and Expansion
  - 2U – Midrange
  - 4U – High End
- Want to maintain a common design for storage
  - E3.S/L is primary design point
- Flexibility to support E3.S/L and 2T in the same enclosure
- Meet the thermal design requirements based on anticipated slot to slot pitch

1U Mechanical Concept

- Support up to 20 is possible, however, the ability to maintain a 3mm is not possible for all slots. To support a true even gap, between all devices, a gap spacing of 2.4mm would need to be used.
- 2T would be possible, however, uniform spacing gaps is not maintained. The gap spacing would likely lower: 2.4mm, middle: 3mm, and top: 3mm.

2U Physical Mockup Concept

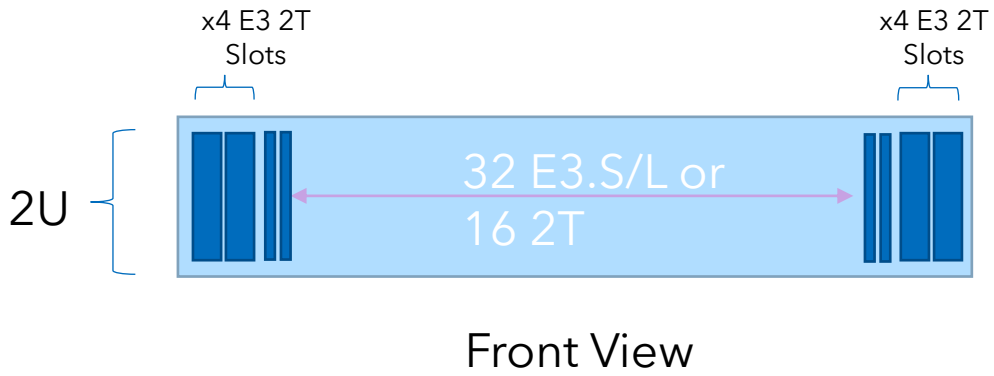

Front View

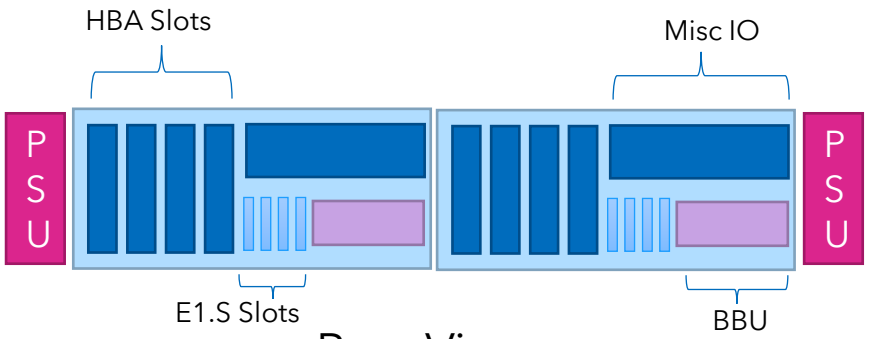

Rear View

- 2U enclosure with 2U half wide canisters configured in a side-by-side configuration
- Each canister supports single socket w/ up to 12 DIMM slots per socket
- Up to 4 PCIe® slots per canister
  - All mechanically x16, supporting a mix of x8 and x16 electrical
- PCIe® 5.0 throughout
- Hot plug/replicable, E1.S boot drives
- Hot plug/replicable power loss protection
- 1+1 redundant CRPS power supplies
- Configuration using up to 32 E3.S/L 2x2 slots and 4 E3 2T 2x4 or 1x8 slots

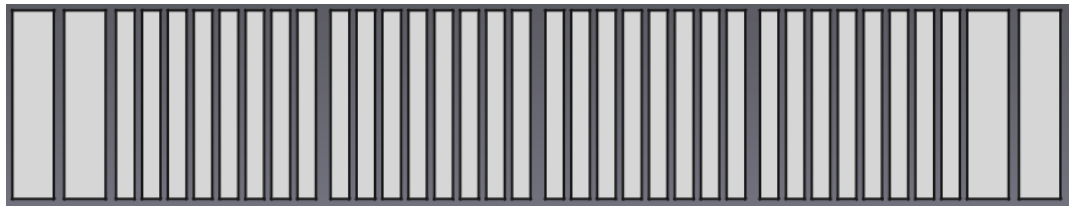

4U Physical Mockup Concept

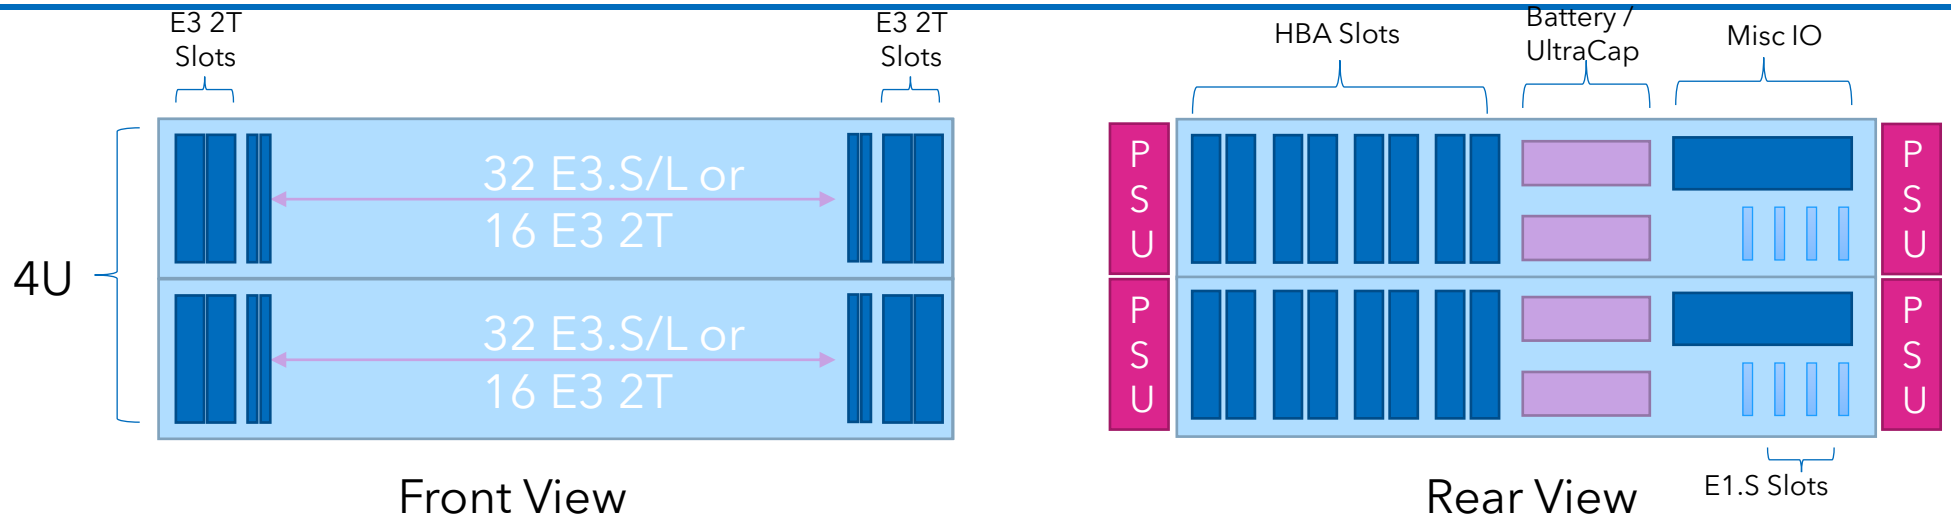

- 4U enclosure with 2U canisters configured in a top/bottom configuration
- Each canister supports dual socket w/ up to 12 DIMM slots per socket
- Up to 8 PCIe® slots per canister
  - All mechanically x16, supporting a mix of x8 and x16 electrical
- PCIe® 5.0 throughout
- Hot plug/replicable, E1.S boot drives
- Hot plug/replicable power loss protection
- 2+2 redundant CRPS power supplies
- Configuration using up to 64 E3.S/L 2x2 slots and 8 E3 2T 2x4 or 1x8 slots

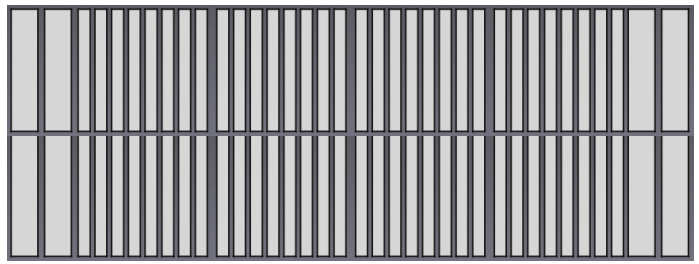

4U Enclosure 2U Full Wide Canister PCIe® Topology 2S/CPU 48 Drive U.2 Support

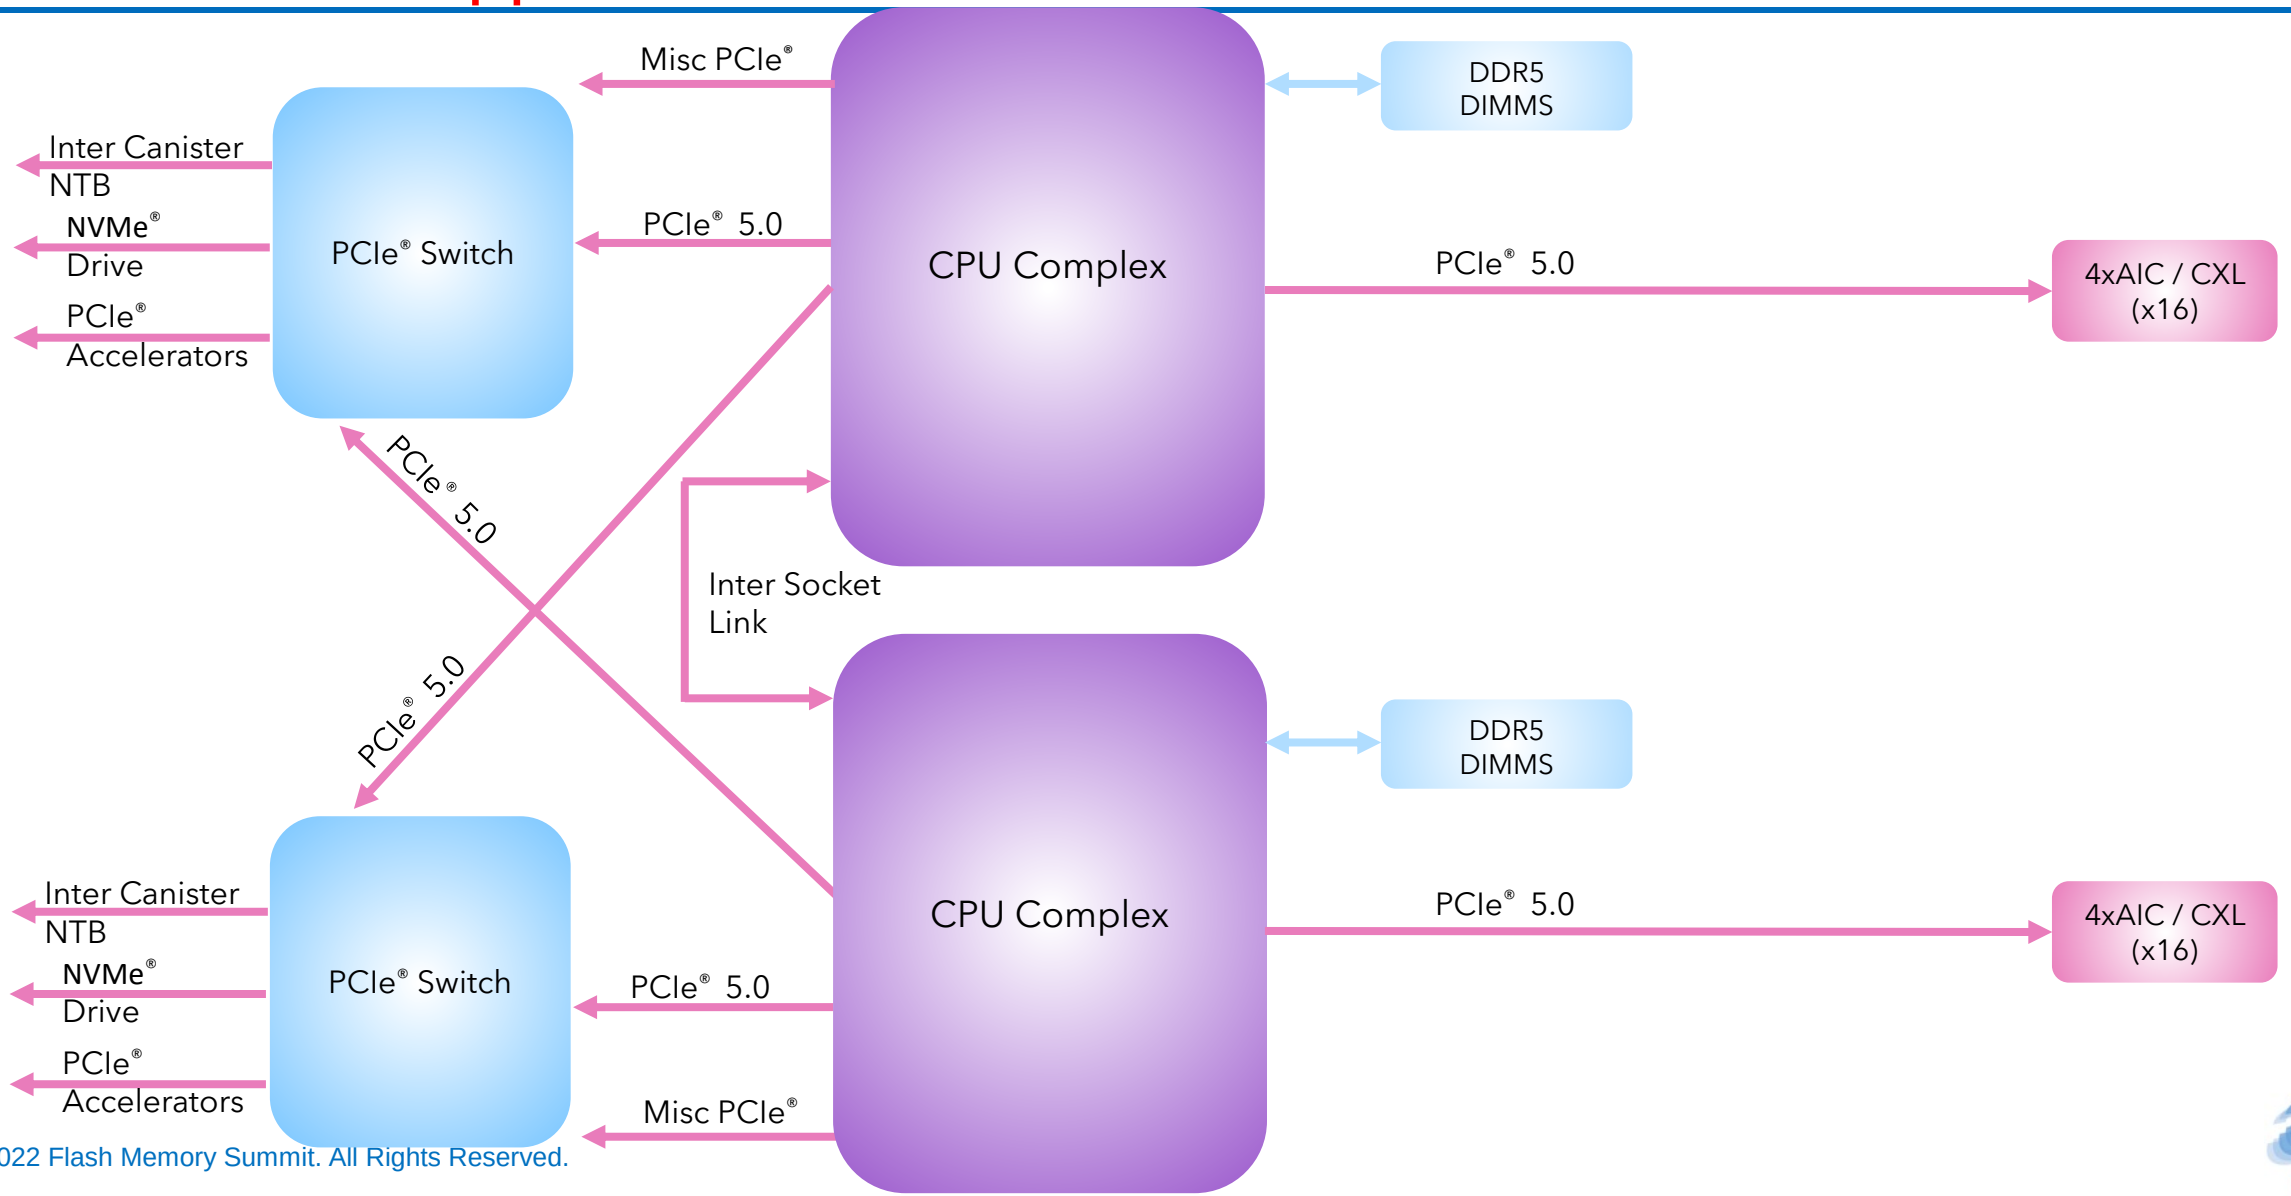

4U Enclosure 2U Full Wide Canister PCIe® Topology 2S/CPU 64 2x2 E3.S/L Support + 4x4 2T

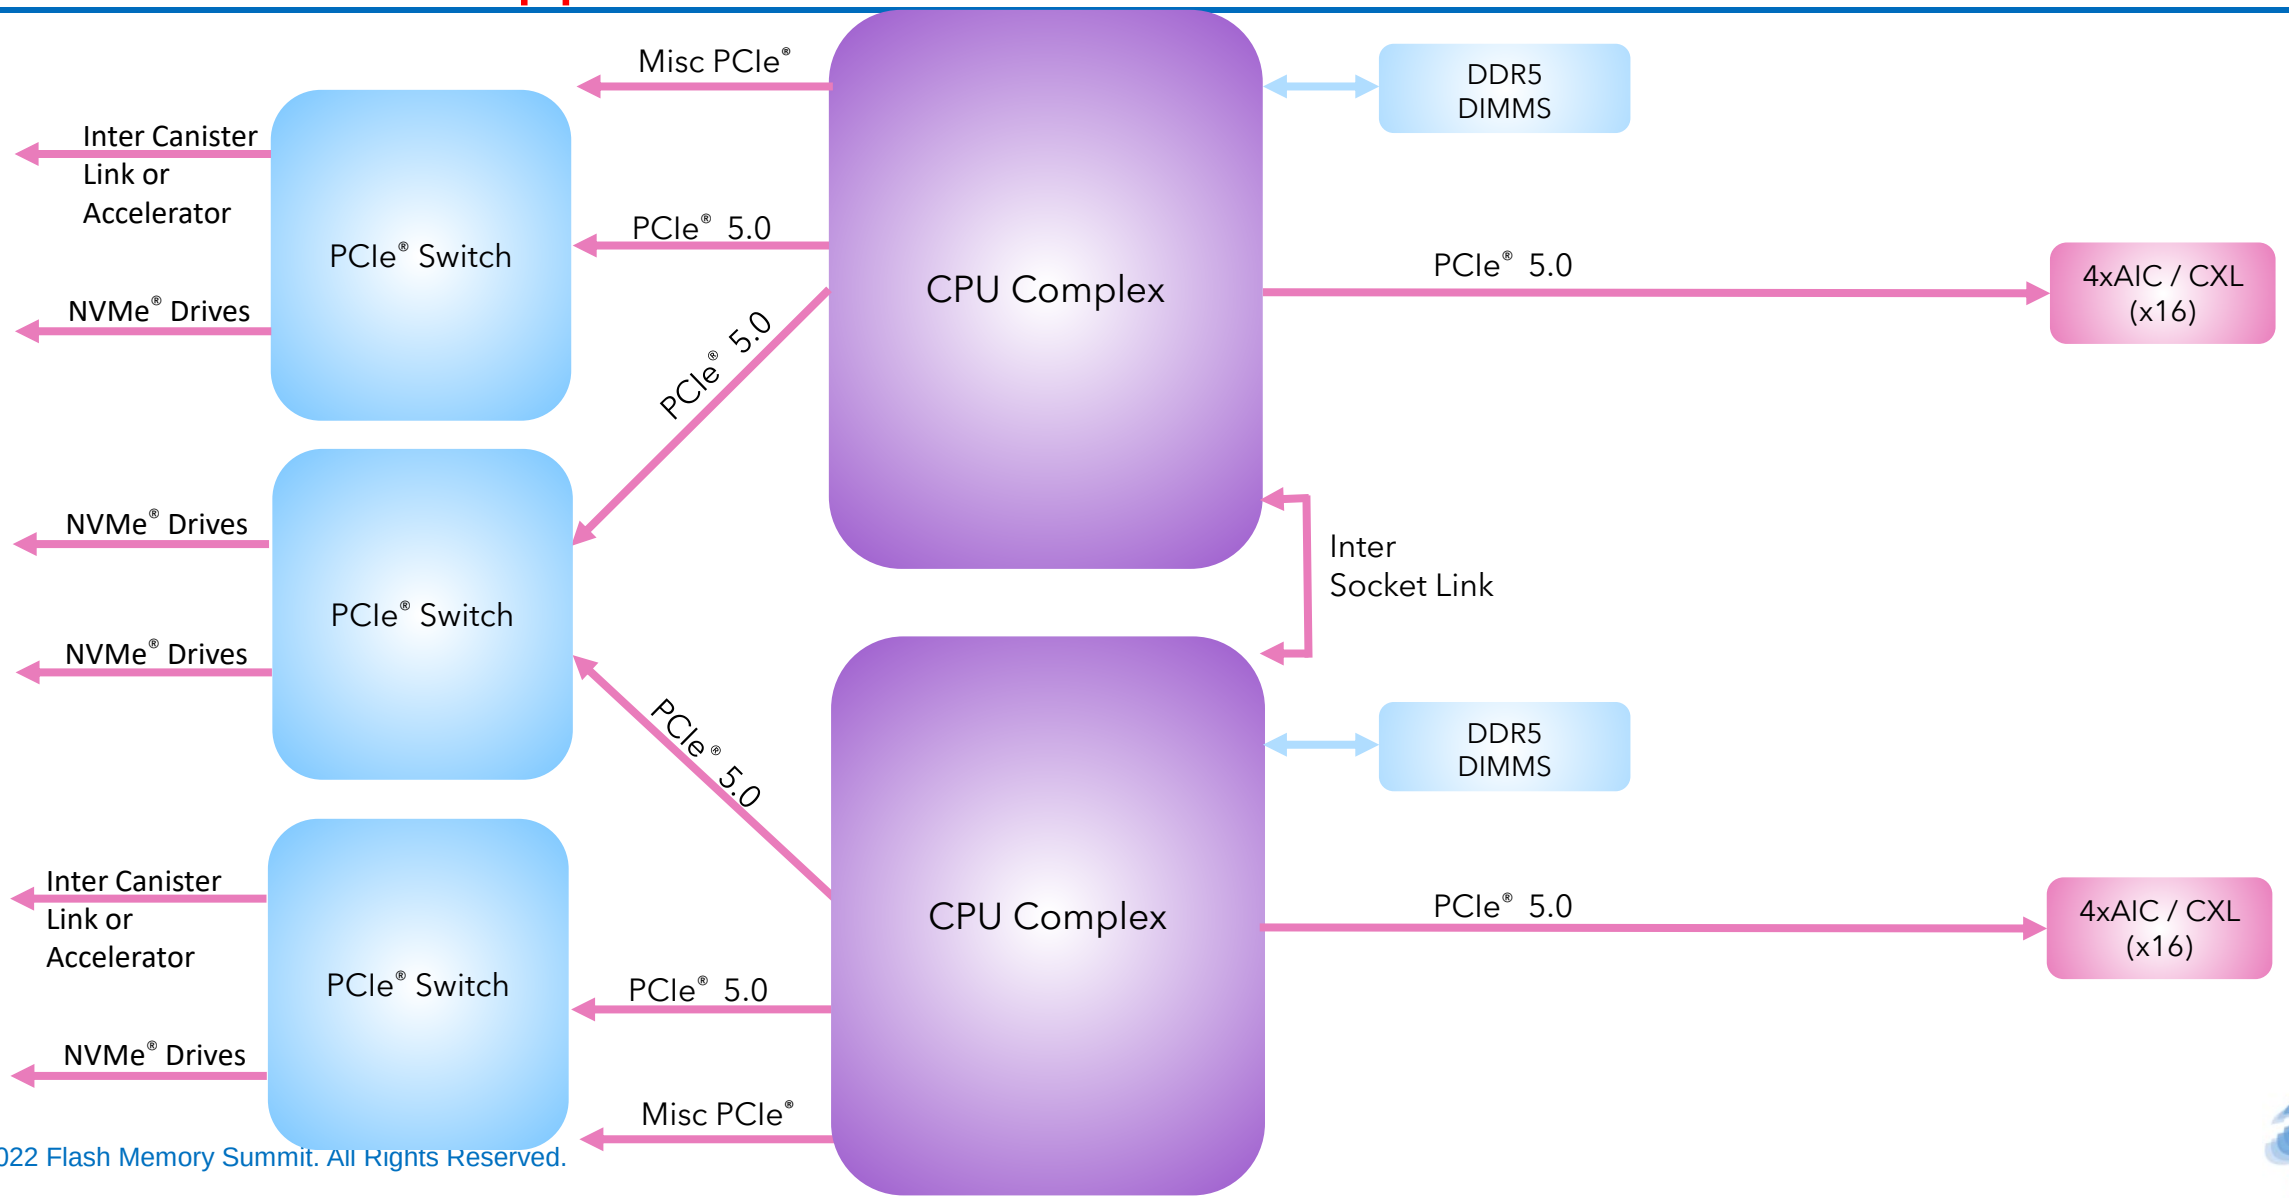

When you interact with IBM, this serves as your authorization to Flash Memory Summit or its vendor to provide your contact information to IBM in order for IBM to follow up on your interaction.

IBM's use of your contact information is governed by the IBM Privacy Policy.