Flash Memory Summit

SSDS-302-1: Form Factors and Innovations Maximizing EDSFF E3 SSD Design

Trent Johnson, FlashCore Hardware Architect

trejo@us.ibm.com

Flash Memory Summit

Trent Johnson is a Hardware Architect at IBM, with a focus on the IBM FlashCore Module. He joined IBM as part of the Cleversafe Acquisition where he was the System Hardware Architect of exabyte-scale Object Storage. Prior to Cleversafe, he developed system-level manufacturing and test solutions for AMD CPUs and GPUs where he was awarded the AMD Corporate Technical Achievement Award.

He has 24 years of industry experience, holds 7 US patents and has published at the Burn-in and Test Socket Workshop as well as the Conference for Consumer Electronics. He earned BSEE and MSEE degrees from The University of Texas at Austin in Electrical Engineering with a focus on Manufacturing System Engineering.

The Layout of Today's IBM FlashCore™ Module

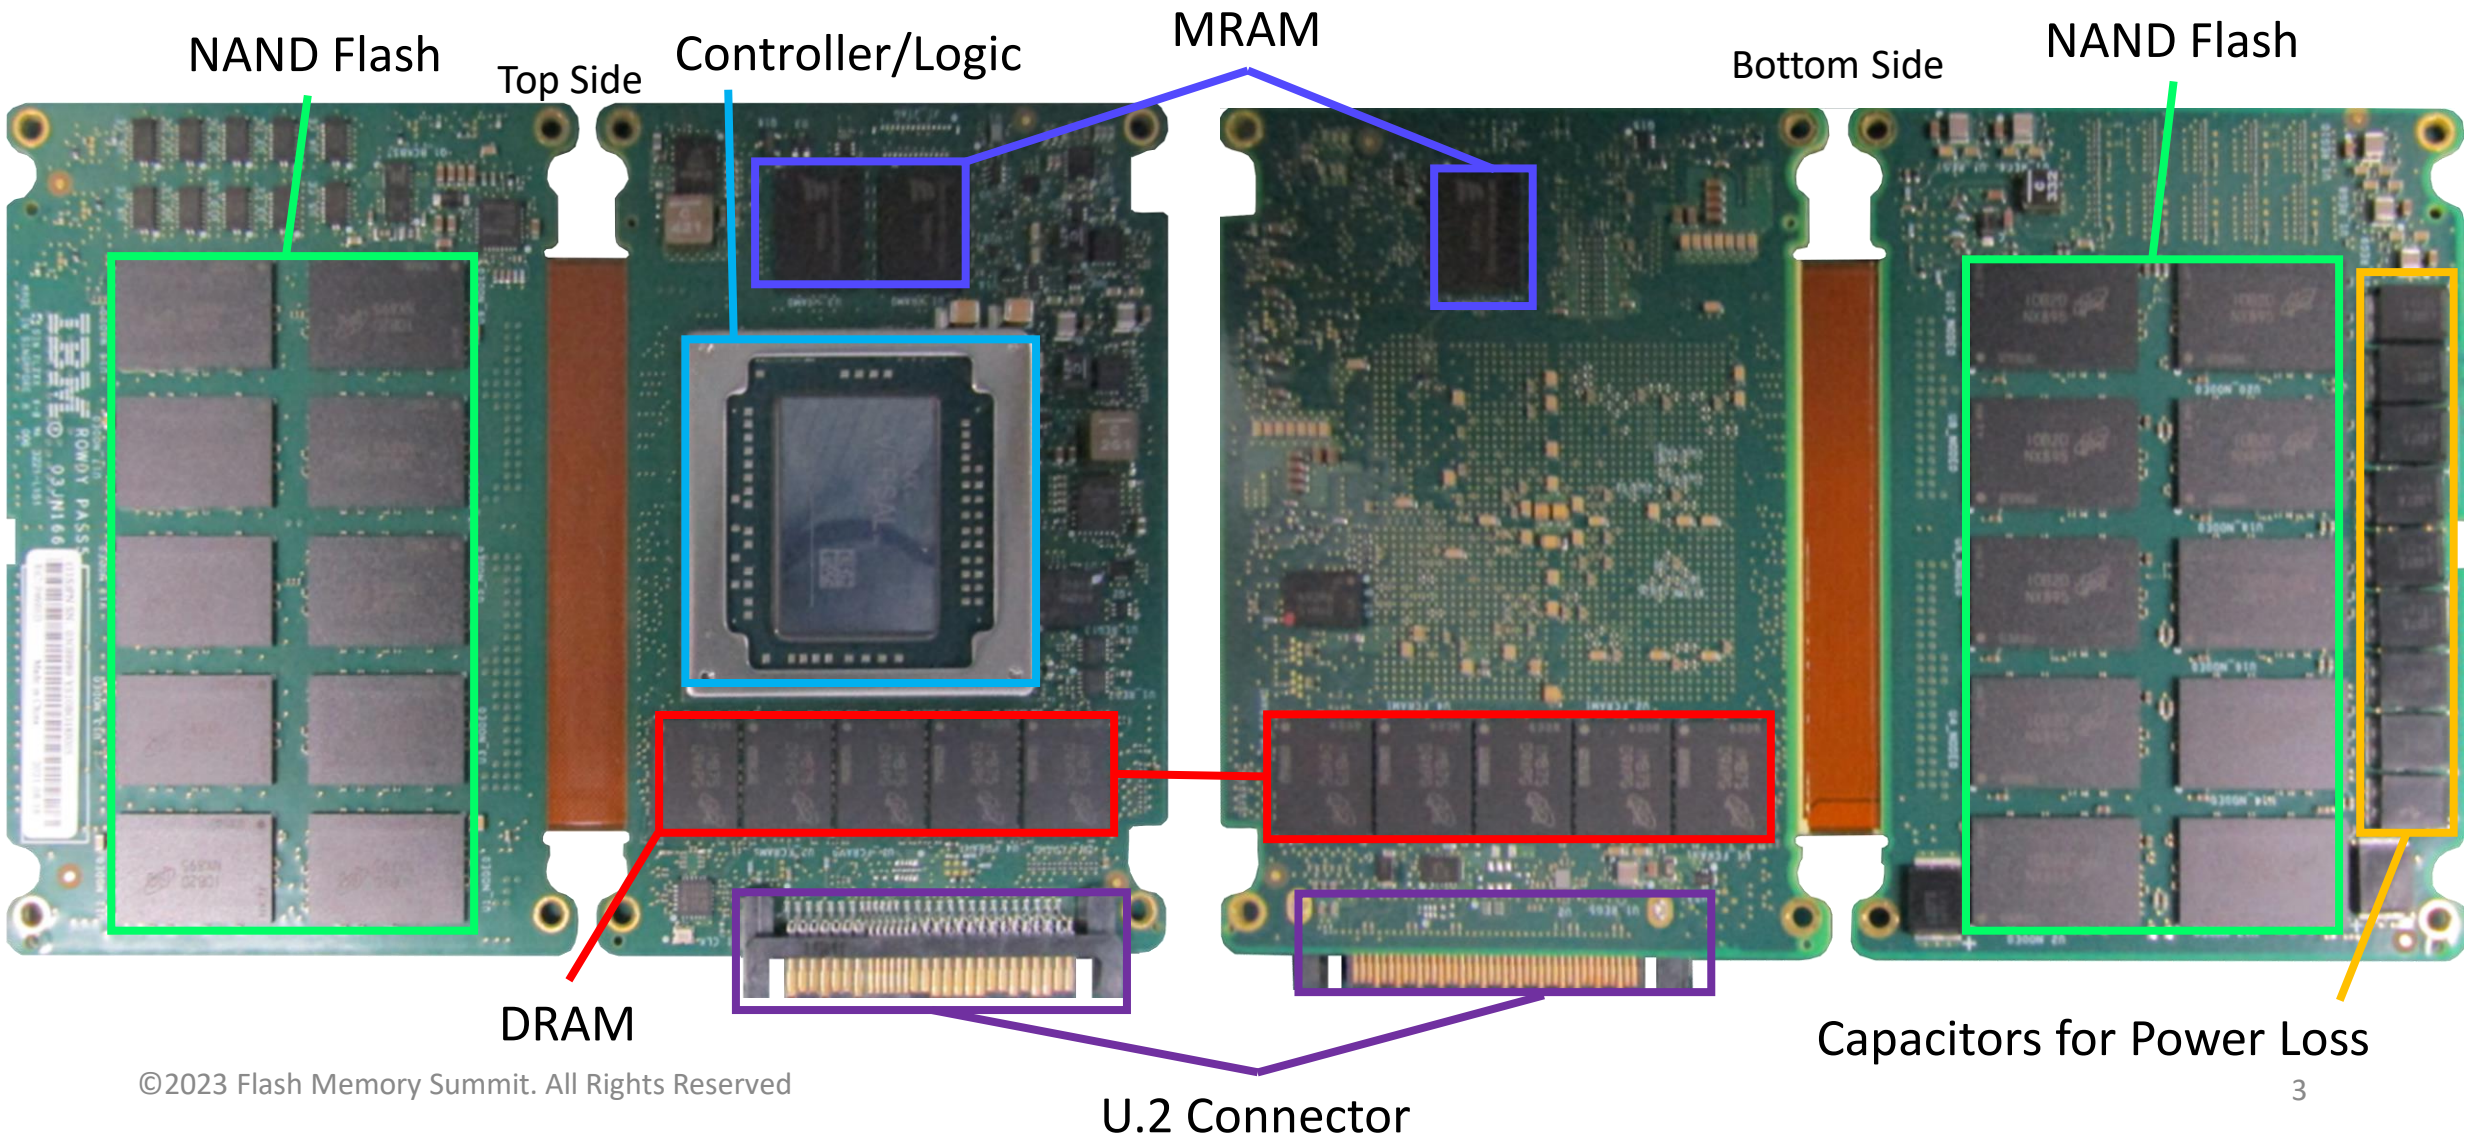

FCM Hardware Design Goals

- Migrate to a new industry standard, EDSFF E3
- Maximize the overall Terabytes per rack unit
- Minimize cost per Terabyte
- High Quality & Reliability
- Utilize an FPGA for advanced computational storage techniques like inline compression, encryption, RAID assists and future features

Today

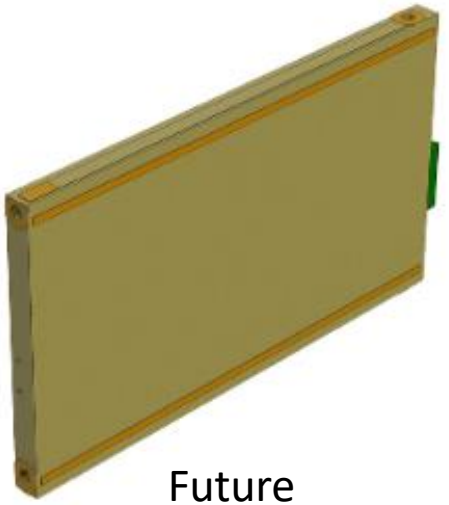

Future

Computational Storage Is Supposed To Use E3.L 2T

- E3.L **2T**: Targeted to **FPGAs** or accelerators
- Power profiles up to 70W

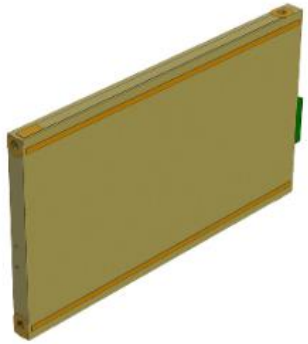

- E3.L: Targeted to be a primary form factor for storage subsystems and server platforms requiring **maximum capacity** for each ‘U’
- Power profiles up to 40W

https://americas.kioxia.com/content/dam/kioxia/en-us/business/ssd/data-center-ssd/asset/KIOXIA_EDSFF_E3_Intro_White_Paper.pdf

- What if your product is a mix of both use cases?
- If you don’t need 70W of power, you lose 50% of your density by using E3.L **2T** vs E3.L

Computational Storage On E1.L?

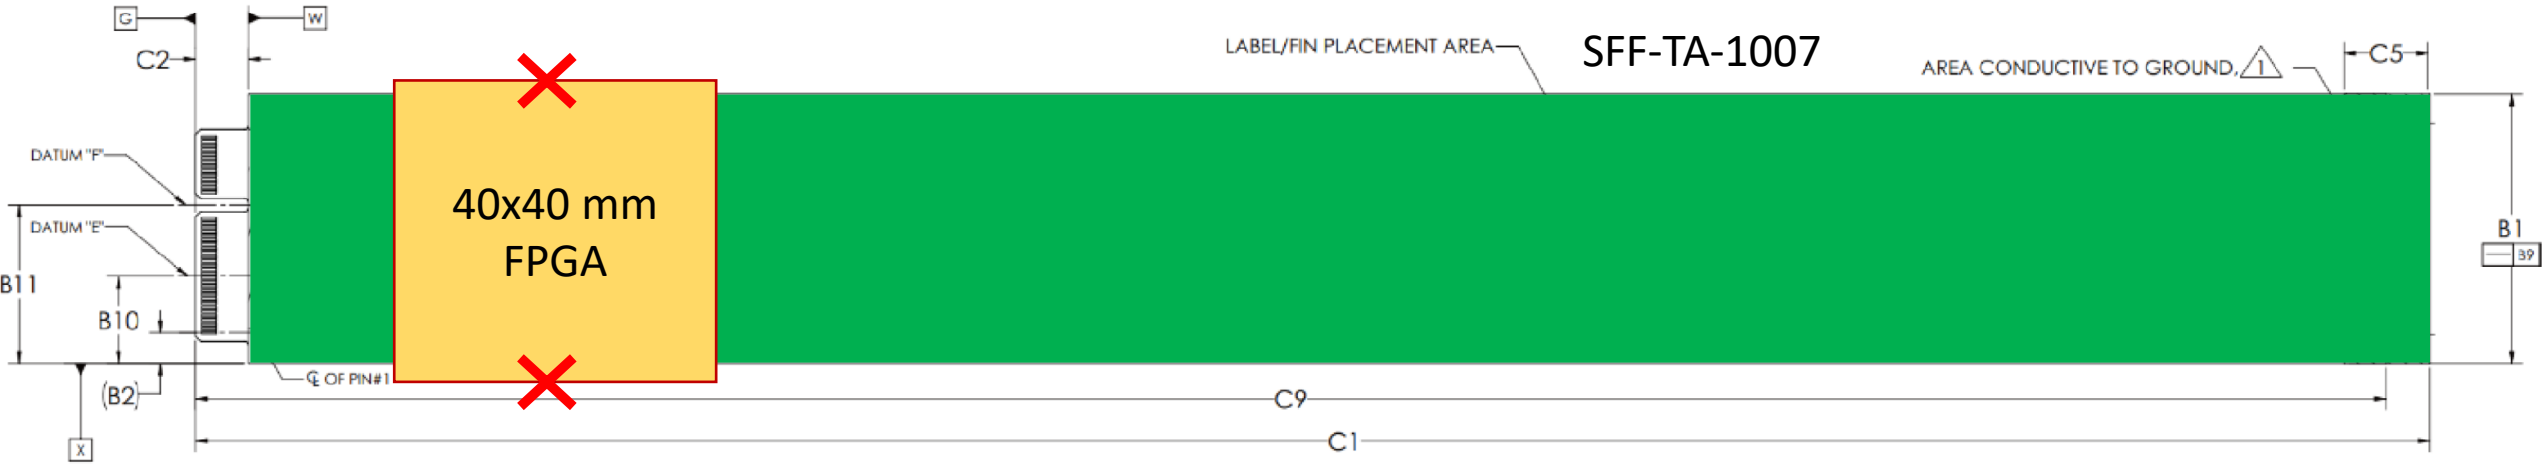

- 40mm x 40mm FPGAs clearly do not fit on a 38.4mm wide board
- 35mm x 35mm FPGAs are plausible, but signal exits can only go east/west
  - Complex routing
  - Reduced I/O for Flash

Enclosures Keeping My SSD Cool: A Leap of Faith

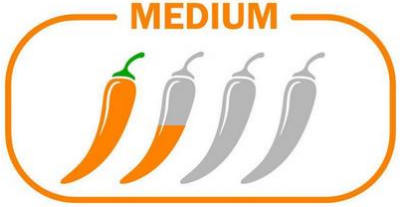

A typical 2U U.2 Storage Server:
24 SSDs at 25W max each
600W of drive power

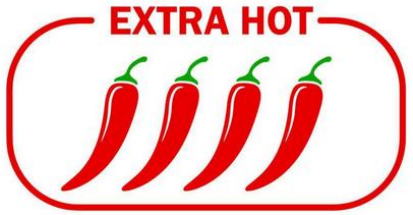

A 2U E3.L Storage Server:
44 SSDs at 40W max each
1760W of drive power

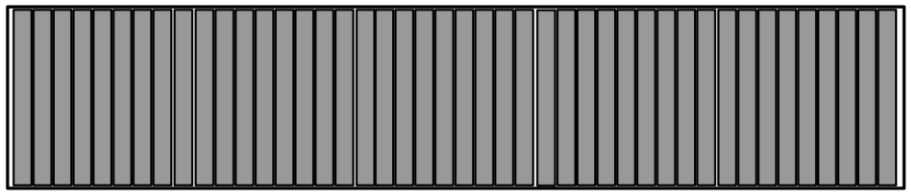

Concept from 2023 EDSFF whitepaper, B. Lynn, P. Kaler, and J. Geldman

- SFF-TA-1023 defines design criteria for both Enclosures and EDSFF devices
- Careful EDSFF design must still be followed
- Hopefully, the enclosure can keep downstream components cool, too!

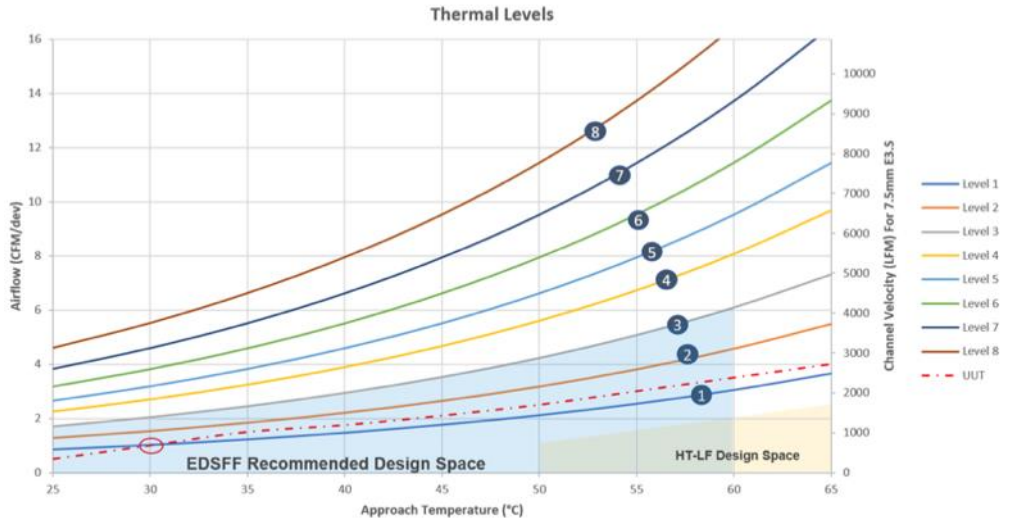

E3.L Floor Planning Concept

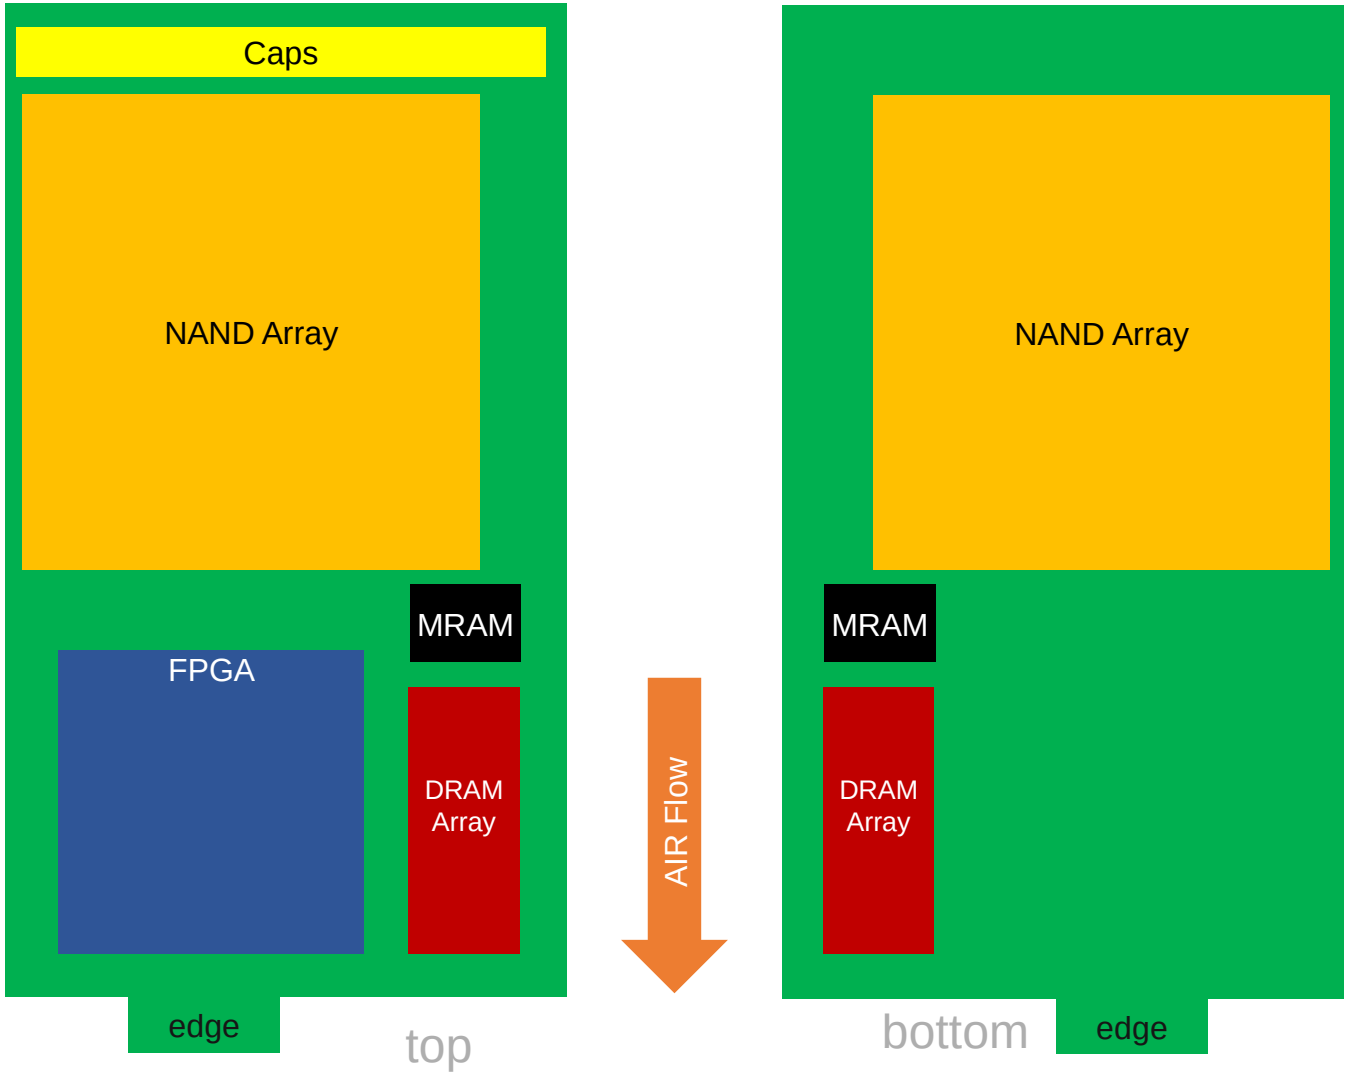

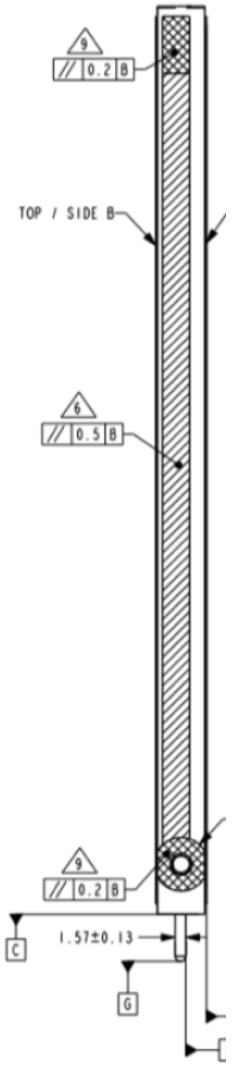

EDSFF E3.L Challenge With FPGAs

- The E3 spec is optimal for thin ASICs
- The card edge is very close to the center of the stackup

E3 Problem: FPGAs Are Big!

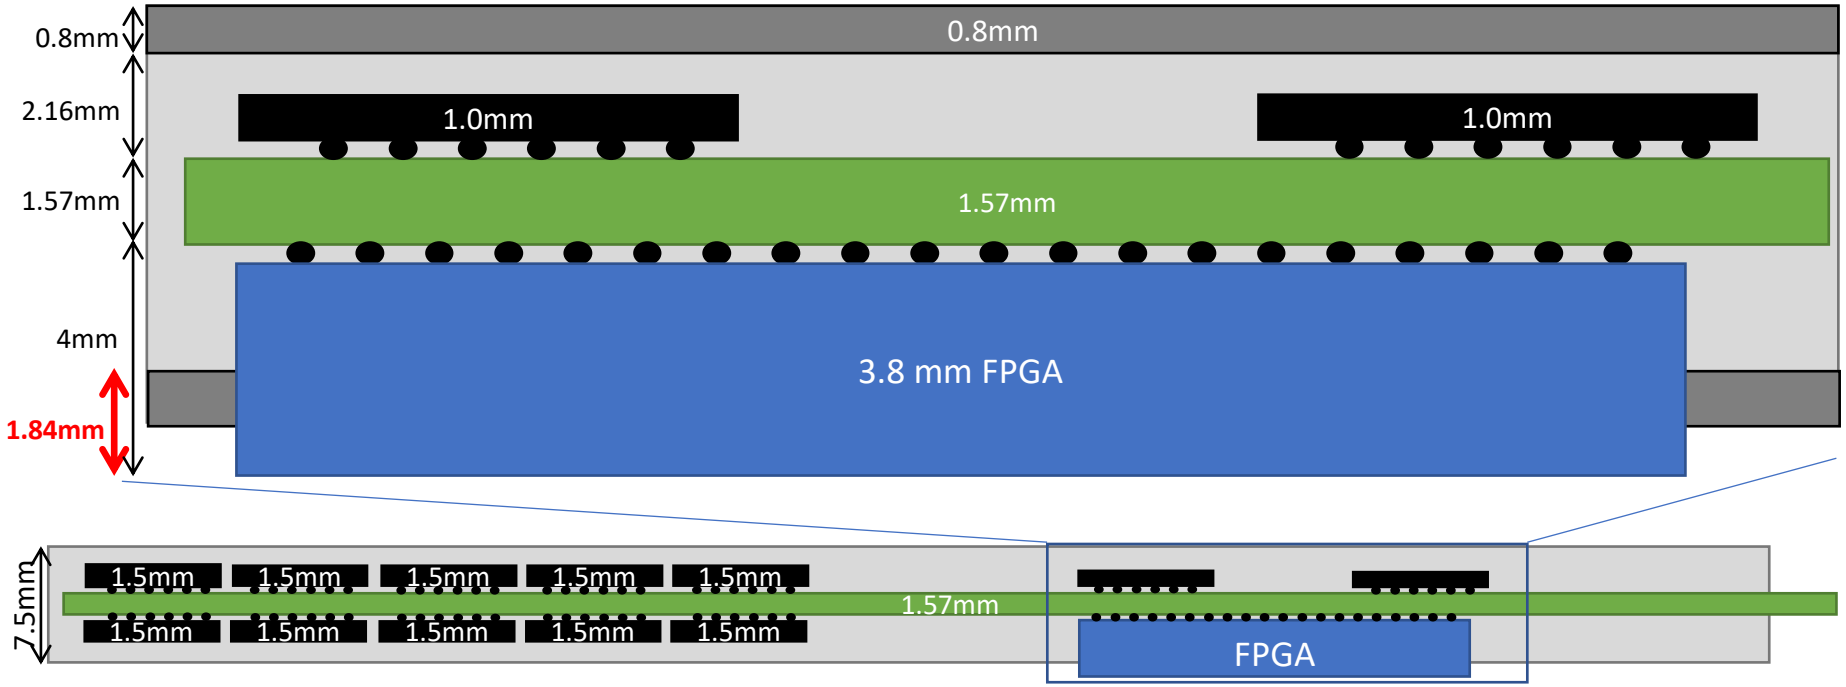

One Potential Solution: Offset The PCB

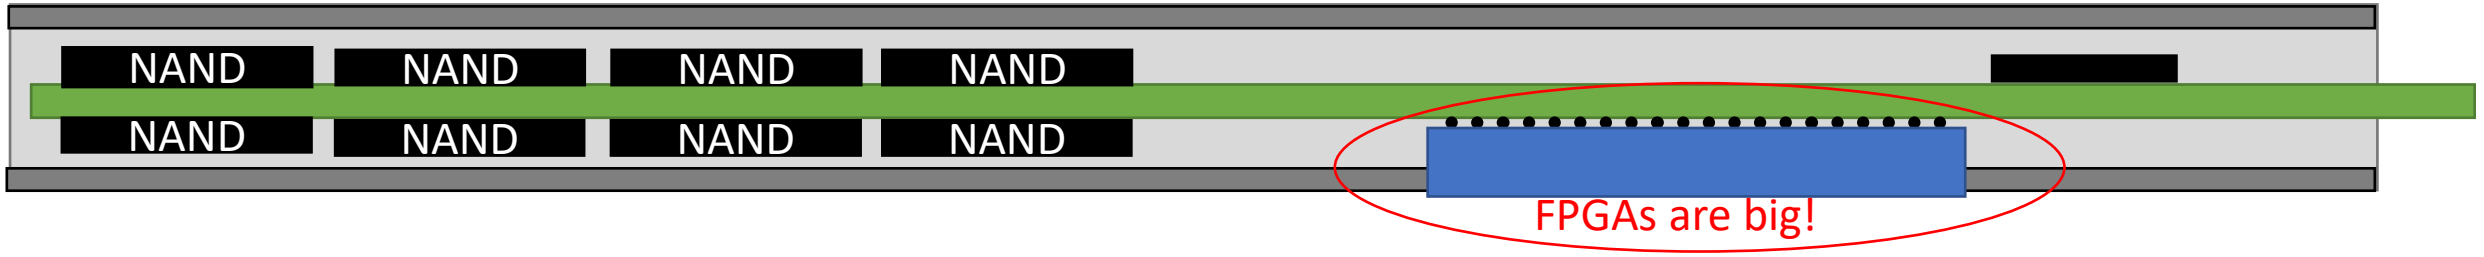

Challenges:

- Card edge alignment
- LED alignment
- Back-side components
- Thermal design

Shifting Your PCB And Maintain The Card Edge

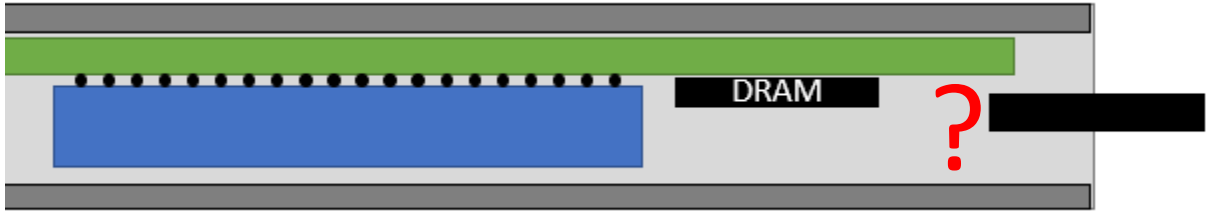

- Methods
  - Mezzanine soldering
  - Rigid flex
  - Plug/socket connectors
  - Elastomer connections
- Challenges:
  - Signal Integrity
  - Mechanical Stability
  - Tiny Dimensions
  - Reliability
  - Cost

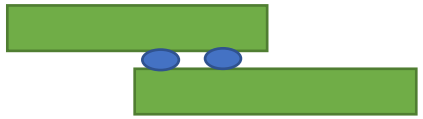

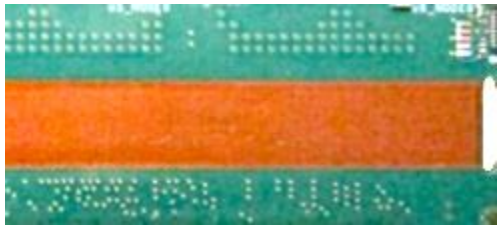

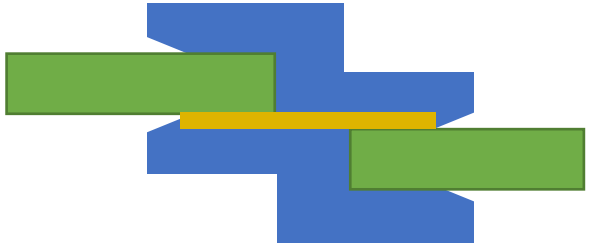

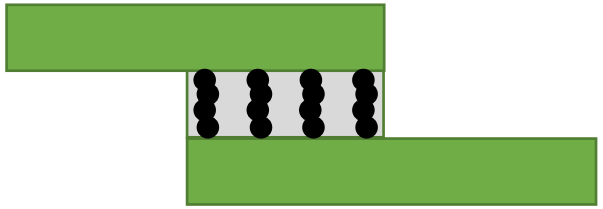

Aligning LEDs

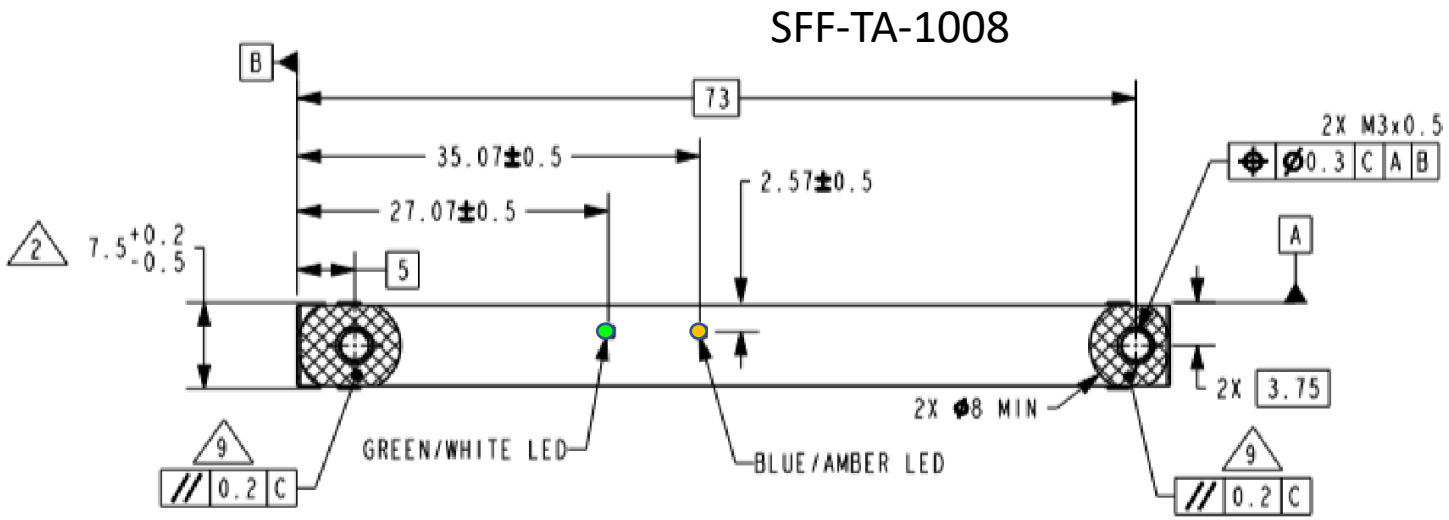

- LED hole is ideal for SMT LEDs on nominal board height
- Too high, and your LED holes are covered
- Too low, and your LED won't be visible
- Mitigations:
  - Light Pipes
  - Custom LED
  - PCB Mezzanine
  - Flare the hole

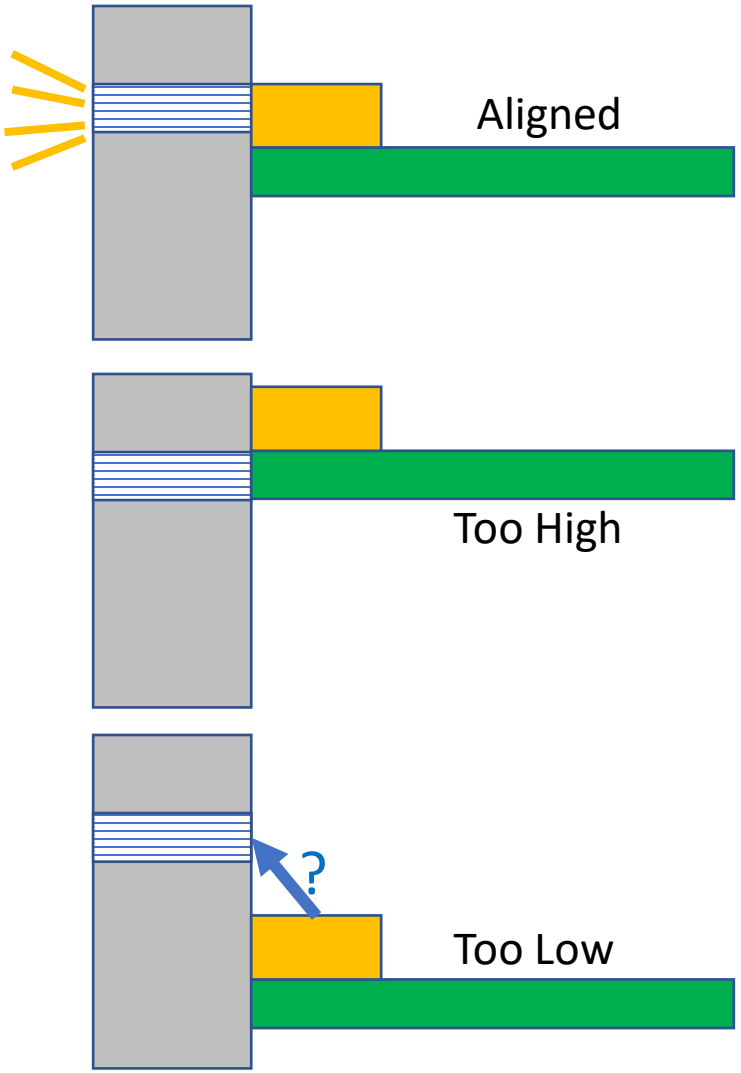

Back-Side Components

- You may find you run out of space when you adjust your Z height
- As with anything, it's a tradeoff
- Mitigations:
  - Move tall components to the taller side
  - Thin your shell
  - Use a mezzanine to put the PCB in the center again

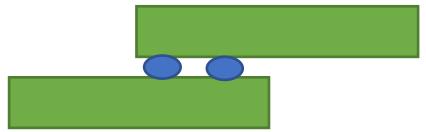

Thermal Design Considerations

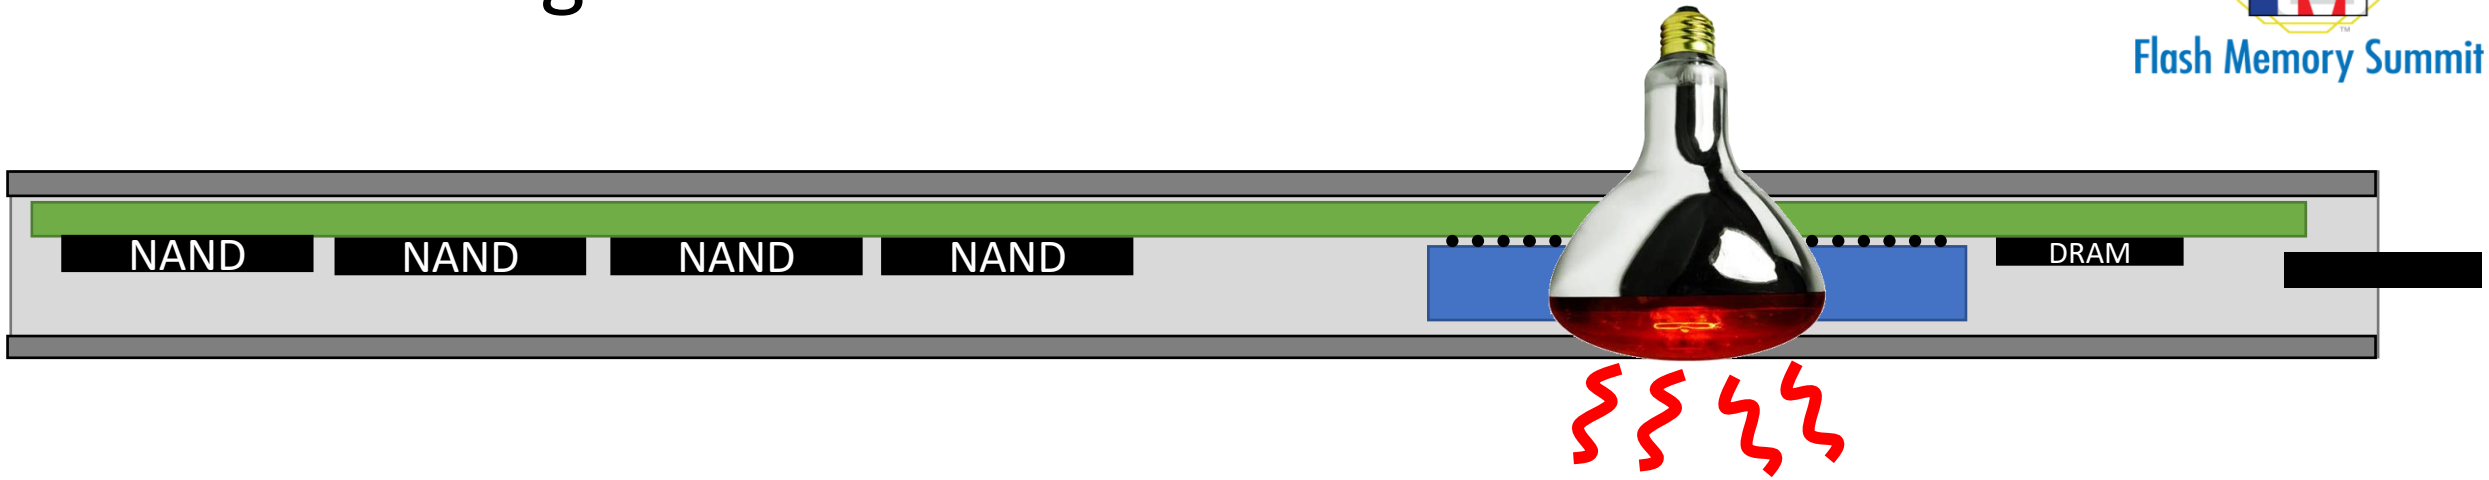

- FPGAs produce a lot of heat
  - Much of the power envelope will come from the FPGA
- Case design is very important
  - Material: Copper, Aluminum, alloys
  - Consider heat spreaders
  - Thermal Interface Material
  - Fin design and aerodynamics
- Thermal simulations are a must

Summary

- E1 is not well suited for FPGA implementation
- E3 enclosures need to be well designed to remove all the power the spec allows
- You can fit an FPGA into an E3 form factor to maximize density
  - Several exotic design decisions are needed
  - It’s not easy