

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

EDSFF April 2023 Panel Update

Moderated by Jonmichael Hands,
SNIA SSD Special Interest Group




Today's Speakers



Jonmichael Hands

Chia Network
Co-Chair
SNIA Solid State Drive
Special Interest Group



Jonathan Hinkle

Micron
Distinguished Systems
Architect



Paul Kaler

HPE
Future Storage Architect



Kevin Mundt

Dell
Server Systems Architect



Lee Prewitt

Microsoft Azure
Director, Cloud Hardware
Storage

Agenda

- Jonmichael Hands – Welcome and what we'll cover today
- Jonathan Hinkle – EDSFF architectural considerations and CXL update
- Paul Kaler – HPE E3.S systems overview and benefits
- Kevin Mundt – Dell E3.S systems overview and benefits
- Lee Prewitt – E1 update for Microsoft and OCP platforms
- Jonmichael – SSD vendors, Supermicro EDSFF platforms, and industry updates

SNIA COMPUTE + MEMORY
+ STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

EDSFF Explosion

Presented by

Jonathan Hinkle

Distinguished Systems Architect -
Micron SBU



Proven Promise of EDSFF (Enterprise and Datacenter Standard Form Factor)

Optimized for scaling solid-state modules 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 PCIe Gen6 and CXL for high performance storage and new applications like main memory expansion




Industry standard EDSFF connector



Improved airflow through connector



E1.S: 1U height-optimized



EDSFF improved cooling



E3: 2U height-optimized

Now what can we do with these EDSFF slots?

- Scale out drive count density with smaller EDSFF like E1.S and E3.S 1T



1U Rack System

2U Rack System



(A couple extreme examples ;)


- Scale up to higher capacity and power with larger form factor like E1.L and E3.S 2T



E3.S 2T more easily enables
power support for >25W




E1.L enables new capacity points higher
than most drives – 32TB and beyond!



Key tradeoffs in system designs



- Cost is king
 - Performance increase per device ultimately increases value per \$ for a given capacity at high volume
 - Backplanes, cables and connectors are a key cost factors: size appropriately for performance and capacity
 - Faster signaling requires more expensive materials for signal integrity, but SI is much improved with EDSFF.
- System storage performance scales out and scales up with EDSFF:
 - Smaller form factors allow for drive density, so can multiply saturated x4 PCIe link performance (8-32X)
 - Larger form factors provide highest performance per drive without hitting power limitations
- Space is still a precious commodity
 - Front system space is still critical for hot-swap devices, so need to make the most of it.
 - EDSFF provides choices for device density and capability
- Thermal performance
 - EDSFF improves cooling with smaller and new connector options
 - Must be very careful about system pre-heat, especially in servers
(25W x 24 = 600W in front of CPU and memory: already very difficult to cool!)

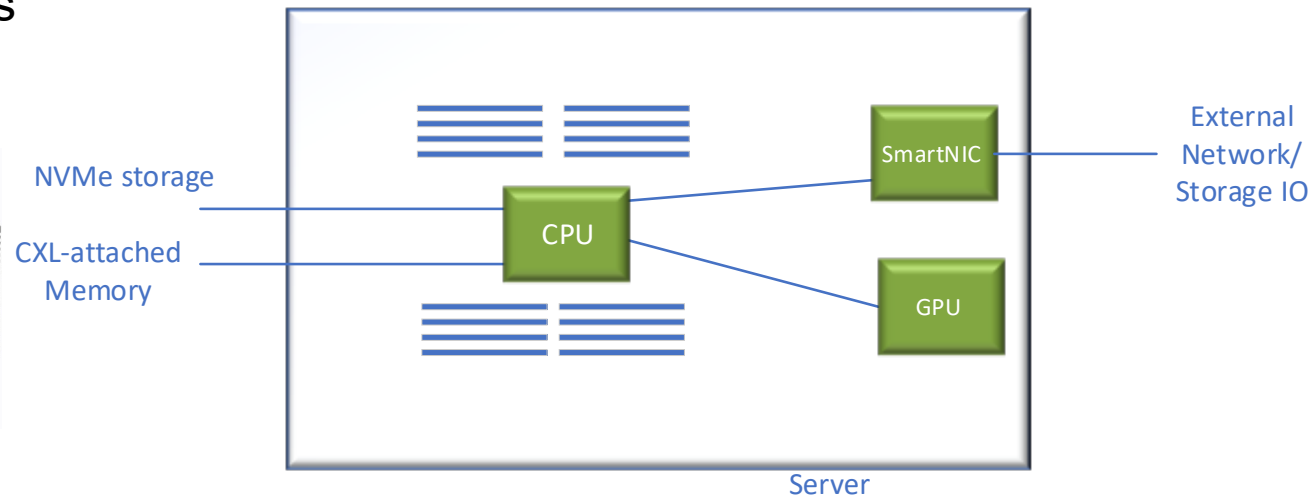



New uses for EDSFF

- EDSFF form factors are now being adopted for all sorts of other applications in addition to storage
- It helps they are so flexible with good thermals and SI, high power and x4, x8 and x16 support
- Similar benefits to scale up and scale out modular devices with standard connector and FF.
- A few examples:
 - AI Acceleration – can leverage for scale-out or scale-up devices with ASIC or FPGA compute
 - CXL Memory Expansion – complement to DIMMs




Example: Blaize AI accelerator



CXL Memory Modules (CMMs) – new standard

- EDSFF form factors are also now leveraged for main memory expansion with CXL interface
- In fact, JEDEC has recently published the first CXL Memory Module Specification (CMM) with significant industry support: JESD317
- Standard CMM targets specified: E1.S x8, E3.S 2T x8, E3.S 1T



JEDEC STANDARD

Compute Express Link (CXL™)
Memory Module Reference Base
Standard

JESD317

March 2023


JEDEC SOLID STATE TECHNOLOGY ASSOCIATION



Specification available now for download: <https://www.jedec.org/system/files/docs/JESD317.pdf>

CXL Memory Modules

- Enable memory expansion physically removed from CPU or even outside of host system
- Can leverage the same EDSFF slots for CXL in system as used for PCIe-attached NVMe drives



High memory bandwidth

Up to 790 GB/s total¹

Low latency²

Ultrahigh memory capacity

Up to 3.2 TB total¹

Emerging data center application workloads

Micron CXL memory expansion is ideal for data-intensive applications



High-performance computing



In-memory database



Business analytics



Virtualized platforms

Benefits of Micron CXL memory

Increase performance

Achieve up to 40% faster execution time³

Boost utilization

Balance memory bandwidth to CPU cores

Optimize allocation

Scale up memory without overprovisioning


¹ Memory capacity and bandwidth values based on AMD's next-generation SP5 high-end platform that offers 1P and 2P support, 12-channel DDR5 memory and 64 CXL V1.1+ lanes

² Equivalent to single NUMA hop

³ Micron internal testing results based on CXL evaluation platform and emulation


EDSFF form factors and beyond

- The ecosystem for EDSFF has grown so significantly with E1 and E3 devices, further leverage is also being explored in industry.
- Most recent is work in SNIA SFF on new SFF-TA-1034 module: Pluggable Multi-purpose Module
- Could be leveraged for many different applications



- EDSFF-compatible connector
- Supports up to 32 lanes of PCIe/CXL
- High power support (400W+)

Targeting ~38-39mm height max
and 16.8mm, fits in 1U rack space



SNIA
COMPUTE + MEMORY
+ STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Thank you



SNIA COMPUTE + MEMORY
+ STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

HPE EDSFF E3.S Solutions

Presented by Paul Kaler
Future Storage Architect, HPE



HPE 1U ProLiant Server EDSFF E3.S Adoption

HPE ProLiant DL320 Gen11
1U, 1P



Up to 8 EDSFF E3.S or 4 E3.S 2T NVMe PCIe Gen5 SSDs

Software Defined Compute, Data Collection and Management, Cold Storage workloads requiring a cost optimized, compact form factor

HPE ProLiant DL360 Gen11
1U, 2P



Up to 20 EDSFF E3.S or 10 E3.S 2T NVMe PCIe Gen5 SSDs

Compute-dense solution for high-performance workloads such as VDI, EDA, CAD, or virtualization workloads that require increased compute density with built-in security and flexibility

HPE 1U ProLiant Server EDSFF E3.S Adoption

HPE ProLiant DL325 Gen11
1U, 1P



Up to 20 EDSFF E3.S or 10 E3.S 2T NVMe PCIe Gen5 SSDs

Software Defined Compute, Data Collection and Management, CDN, VDI, and Cold Storage workloads requiring a cost optimized, high density, compact form factor

HPE ProLiant DL365 Gen11
1U, 2P



Up to 20 EDSFF E3.S or 10 E3.S 2T NVMe PCIe Gen5 SSDs

Compute-dense solution for high-performance workloads such as VDI, EDA, CAD, or virtualization workloads that require increased compute density with built-in security and flexibility

HPE 2U ProLiant Server EDSFF E3.S Adoption

HPE ProLiant DL380 Gen11
2U, 2P



HPE ProLiant DL380a Gen11
2U, 2P



Up to 36 EDSFF E3.S or 18 E3.S 2T NVMe PCIe Gen5 SSDs

Performance, expandability, and scalability for diverse workloads and environments, e.g., Collaborative, CRM, SCM, ERM, Data & Analytics, AI, VDI, and Content Mgmt

Up to 8 EDSFF E3.S or 4 E3.S 2T NVMe PCIe Gen5 SSDs

4DW or 8SW GPUs in a 2U2P ProLiant DL design for emerging AI workloads
AI Training & Inference, MCAD, HPC, Engineering Apps, & Network SW

HPE 2U ProLiant Server EDSFF E3.S Adoption

HPE ProLiant DL345 Gen11

2U, 1P



HPE ProLiant DL385 Gen11

2U, 2P



Up to 36 EDSFF E3.S or 18 E3.S 2T NVMe PCIe Gen5 SSDs

Data-intensive workloads such as software-defined storage, video transcoding, and virtualized apps

Up to 36 EDSFF E3.S or 18 E3.S 2T NVMe PCIe Gen5 SSDs

Compute and data storage demanding workloads requiring increased core count, and storage and I/O scalability, e.g., AI, ML, Big Data analytics

HPE Alletra 4000 Storage Server EDSFF E3.S Adoption

HPE Alletra 4110

1U, 2P all-NVMe data storage server



20 EDSFF E3.S NVMe PCIe Gen5 SSDs

For the most **performance-demanding** data storage-centric workloads:

- Data stores for machine learning
- Distributed and NoSQL databases
- High-performance Software-Defined Storage (SDS)
- Data-heavy hyperconverged infrastructure (HCI)

HPE Alletra 4120

2U, 2P hybrid-NVMe data storage server



12 EDSFF E3.S NVMe PCIe Gen5 SSDs (rear mount)

For the **broadest range of** data storage-centric workloads including:

- analytics data lakes
- general purpose Software-Defined Storage (SDS)
- converged data protection
- active archives
- and many more!

SNIA COMPUTE + MEMORY
+ STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

EDSFF E3 Intercepting the PowerEdge Portfolio

Presented by Kevin Mundt,
Senior Distinguished Engineer,
Dell Technologies



Solving System Integration Challenges

- Continual increases in CPU and GPU power challenge server and storage cooling.
 - The smaller size of E3 devices enables more airflow through a chassis while maintaining storage capacity and device count.
- Smaller Size enables increase device density in a chassis
 - Any chassis form factor can double its device count or storage capacity compared to U.2 devices.
 - A 10 drive U.2 chassis can now support up to 20 E3 devices.
- A family of E3 form factors enables server portfolio simplification.
 - A configurable chassis can now support storage devices and emerging CXL memory devices.
 - Fewer unique chassis helps with solution cost reduction and greater flexibility for customers.
- Covering future needs
 - Supporting PCIe/CXL x2, x4, x8, and x16 connections
 - U.2 supported 25W. The E3 family supports up to 70W per device enabling higher power devices.
 - Higher performing interconnect enables more robust signal integrity to the channel.
 - Better channel performance for PCIe Gen5 & Gen6

EDSFF E3 Across the Dell PowerEdge 16G Portfolio

- 1U Servers – 14 E3 & 16 E3 configurations
 - R660 – 4th Generation Intel Xeon
 - R6615, R6625 – 4th Generation AMD EPYC
- 2U Servers
 - R760 – 4th Generation Intel Xeon
 - R760xa – 4th Generation Intel Xeon – 6 E3
 - R7615, R7625 – 4th Generation AMD EPYC – 32 E3
- Four Socket – 4th Generation Intel Xeon
 - R860 - 8E3, R960 – 16E3
- C6620 – Multi-Node - 4 nodes in 2U
- Acceleration – 4th Generation Intel Xeon
 - XE8640 – 8 E3
 - XE9640, XE9680 – 16E3
- MX760c Modular – 4th Generation Intel Xeon – 8E3



SNIA COMPUTE + MEMORY
+ STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023


EDSFF at Hyperscale

Presented by Lee Prewitt
Director, Cloud Hardware Storage
Microsoft Azure



E1.S Use Cases - Density Matters

- Multiple swim lanes
 - Data
 - Cache
 - Boot
- Multiple standardized widths
 - 5.9mm
 - 8mm
 - 9.5mm
 - 15mm
 - 25mm
- IOPS matters for high performance compute and storage
- Front panel efficiency matters for dense 1U systems



Concept Project Olympus Design

HDDs Versus E1.L SSDs – A Quick Comparison

HDDs

- Pros
 - Cost
- Cons
 - Everything Else

SSDs

- Pros
 - Everything Else
- Cons
 - Cost



SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

EDSFF – Drive Vendors and Adoption


Jonmichael Hands,
SNIA SSD Special Interest Group




SSD EDSFF Ecosystem – Drive Vendors




KIOXIA



SAMSUNG



SOLIDIGM™



FADU



SMART
Modular Technologies




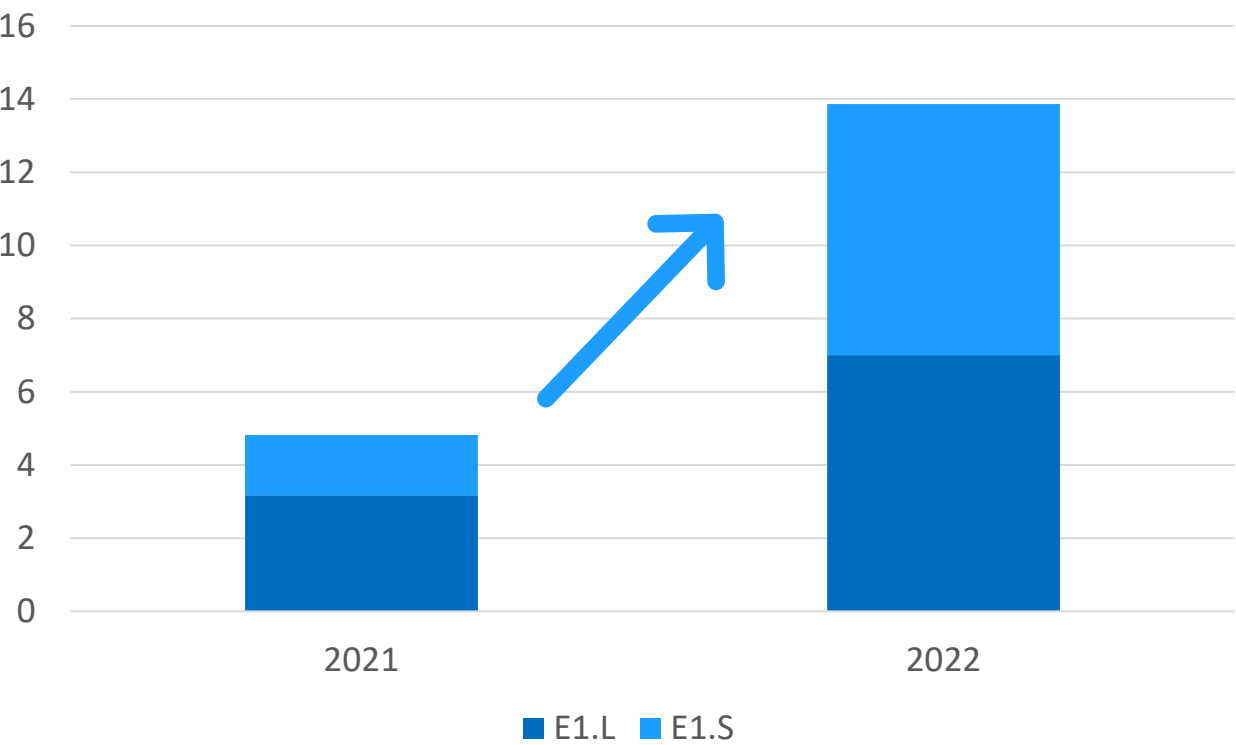
Micron



SK hynix

EDSFF Market Adoption

EDSFF Exabytes



Source: IDC Worldwide Quarterly SSD Shipment Results, CY 4Q22 Doc #US49401423, February 2023

EDSFF Server and Storage



SOG-221E-NE324R

ASG-1115S-NE316R



16x Hot-swap EDSFF E3.S (1T) NVMe slots, support E3.S (2T)




SOG-121E-NES24R – Available Now



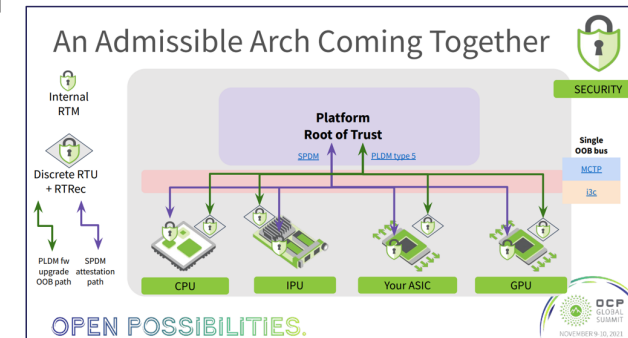

24x Hot-swap EDSFF short (E1.S) 9.5 or 15mm NVMe slots

SOG-121E-NE316R – Support E3.S 2T x8 (CXL) Devices




EDSFF Specification Update

- E3, SFF-TA-1008, 2.03
 - Support for 4C+ (OCP NIC), and 2x1C
 - added new use cases (CXL)
 - label positioning
 - LED definition (power/activity & fault/identify)
- SFF-TA-1023: Thermal Characterization
- SFF-TA-1009: Pinout
 - I3C support



"Composable Security Architectures" at OCP Global Summit, November 9-10, 2021
Co-presented by Google, Intel, Microsoft, Meta, AMD
Replay at <https://www.youtube.com/watch?v=iZw0xhNgyYQ>





COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023



Please take a moment to rate this session.

Your feedback is important to us.