



EDSFF: Mainstream NVMe for the (Datacenter) Masses



Jonathan Hinkle
Sr. Principal Researcher
Lenovo Research
8/08/2018

Datacenter server systems



1U rack server

Many different types of systems, but most servers today: rack-mount and fit in multiple of the "1U" rack space.



Mission Critical



Blade and 1/2 wide Compute Nodes



Datacenter SSD – Our Previous Options

Adapter Cards



Good: High-perf Cache
Bad: Steal PCIe slots from IO
Ugly: \$\$\$, You might fit two?

Consumer SSD



Good: Small and Modular
Bad: Low capacity, no hot-plug, connector less reliable
Ugly: More expensive once adapted to enterprise needs



HDD form factors



Good: Hot-plug, Storage features
Bad: Designed for HDD, not SSD
Ugly: Blocks airflow to the hottest components in server

Need an optimized common building block across server designs.

Datacenter server systems



Server designs have evolved with Moore's law integration, except for one component: storage.

As mechanical devices, HDD have not followed the miniaturization path of semiconductors. Legacy form factors for SSD carry forward this same baggage.

Flash is now a necessity in the datacenter and quickly migrating from cache to main storage.

Legacy SSD are **not** optimized for datacenter servers and storage systems. This reduces opportunity for additional system function and density.

EDSFF 1U Short

Vision:

Create a smaller, high density solid state drive standard that is optimized for the datacenter.


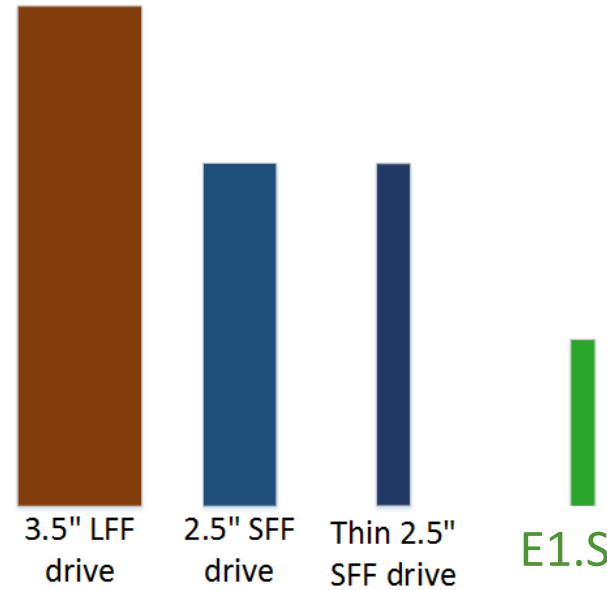

EDSFF 1U Short (E1.S):

- Mainstream NVMe drive
- Compact, modular form factor
 - Vertical fit in 1U height
 - Fits in depth of 2.5" drive
- High Performance and Capacity
- Supports hot-plug and enterprise feature set
- +12V main power for reduced system cost
- LEDs on-drive for lower cost and easier integration



Optimized for the NVMe drive design and use across all datacenter systems to scale as mainstream storage.

DRIVE VOLUME AND DATACENTER SYSTEM FRONT AREA




Drive	Dimension <small>(without cage)</small>	Volume	GB/cm ³	IOPS/cm ³	# Drives in 1U Front
3.5"	10.2 cm x 2.54 cm x 14.6 cm	378.3 cm ³	21.1	528	4
2.5" (HDD compatible bay) (thin)	6.99 cm x 1.5 cm x 10 cm	104.9 cm ³	76.3	1906	12
	6.99 cm x 0.7 cm x 10 cm	48.9 cm ³	163.6	4089	24?
EDSFF 1U Short (E1.S)	3.15 cm x 0.5 cm x 11.15 cm	17.5 cm ³	457.1	11428	36

(compares iso-capacity and performance per mainstream drive)

E1.S enables more space in servers and one use for this is density.

EDSFF 1U Short (E1.S) - Key System Benefits

Industry Standard datacenter-optimized NVMe drive that provides significant new system benefits



- Key benefits:
 - Much **smaller** enabling high **density** storage
 - Significantly improved **system airflow** and thermal solution
 - Most **efficient** modular **scaling** of NVMe capacity and performance
 - Enhanced feature set in **space-constrained** servers
 - **Lower** base system infrastructure and drive **costs**

E1.S Server System Benefits

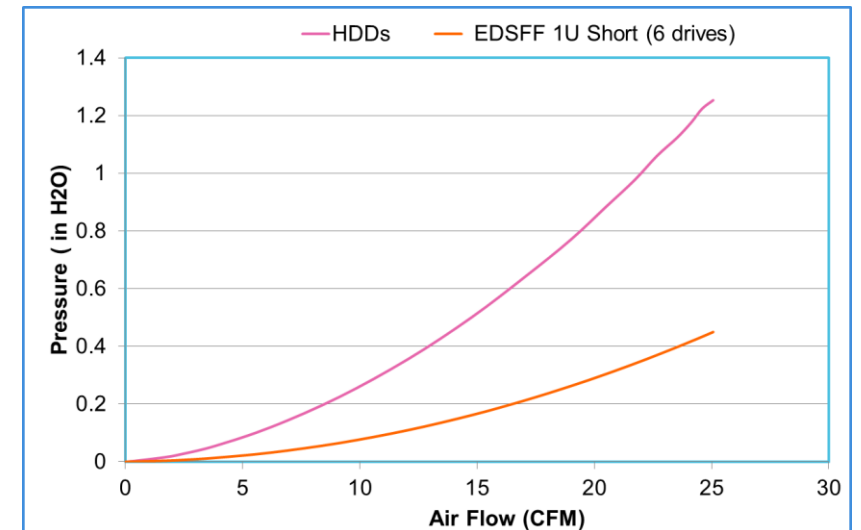
Higher drive count enables faster acceleration while leaving physical space for more slower and cheap capacity storage.

Performance scales especially well, unlocking the inherent performance of the flash media on each drive without expensive, high-end controllers.



Even in dense configurations such as six E1.S drives per 2 x 2.5" bay, additional open space enables more airflow and better system cooling.






2U Prototype system



Future EDSFF applications for system benefits



E1.S, industry standard FF: SFF-TA-1006



- EDSFF connector: SFF-TA-1002 is well suited to further extend capabilities and features of devices in the EDSFF standard form factors.
- The EDSFF connector can support over 40W in all implementations and device power is typically limited first by system thermal requirements.
- Interesting future applications to consider for EDSFF:
 - Emerging Non-Volatile Memory
 - NVMe-oF
 - Computational storage
 - Compute and accelerators
 - IO and networking

Thank you!