PCI-SIG® Educational Webinar Series 2020

Emerging Form Factors: EDSFF Overview

Presented by Tom Friend, Jonmichael Hands and Jonathan Hinkle

Copyright © 2020 PCI-SIG® - All Rights Reserved

Meet the Presenters

Tom Friend, Illuminosi

Jonmichael Hands, Intel

Jonathan Hinkle, Lenovo

Copyright © 2020 PCI-SIG® - All Rights Reserved

PCI-SIG® Snapshot

Organization that defines the PCI Express® (PCIe®) I/O bus specifications and related form factors.

800+ member companies located worldwide. Creating specifications and mechanisms to support compliance and interoperability.

PCI-SIG member companies support the following industries

- Virtual reality
- Automotive
- Artificial intelligence
- Enterprise servers

Board of Directors
2019 – 2020

AMD
arm
DELL EMC
IBM
intel
KEYSIGHT TECHNOLOGIES
NVIDIA
Qualcomm
SYNOPSYS
Silicon to Software

Copyright © 2020 PCI-SIG® - All Rights Reserved

PCIe® Technology:
One Interconnect, Infinite Applications


Storage


Enterprise Servers


Aerospace


Automotive


AI/ML

IoT


Mobile


Cloud


Industrial


PC

Copyright © 2020 PCI-SIG® - All Rights Reserved

PCI-SIG® Roadmap

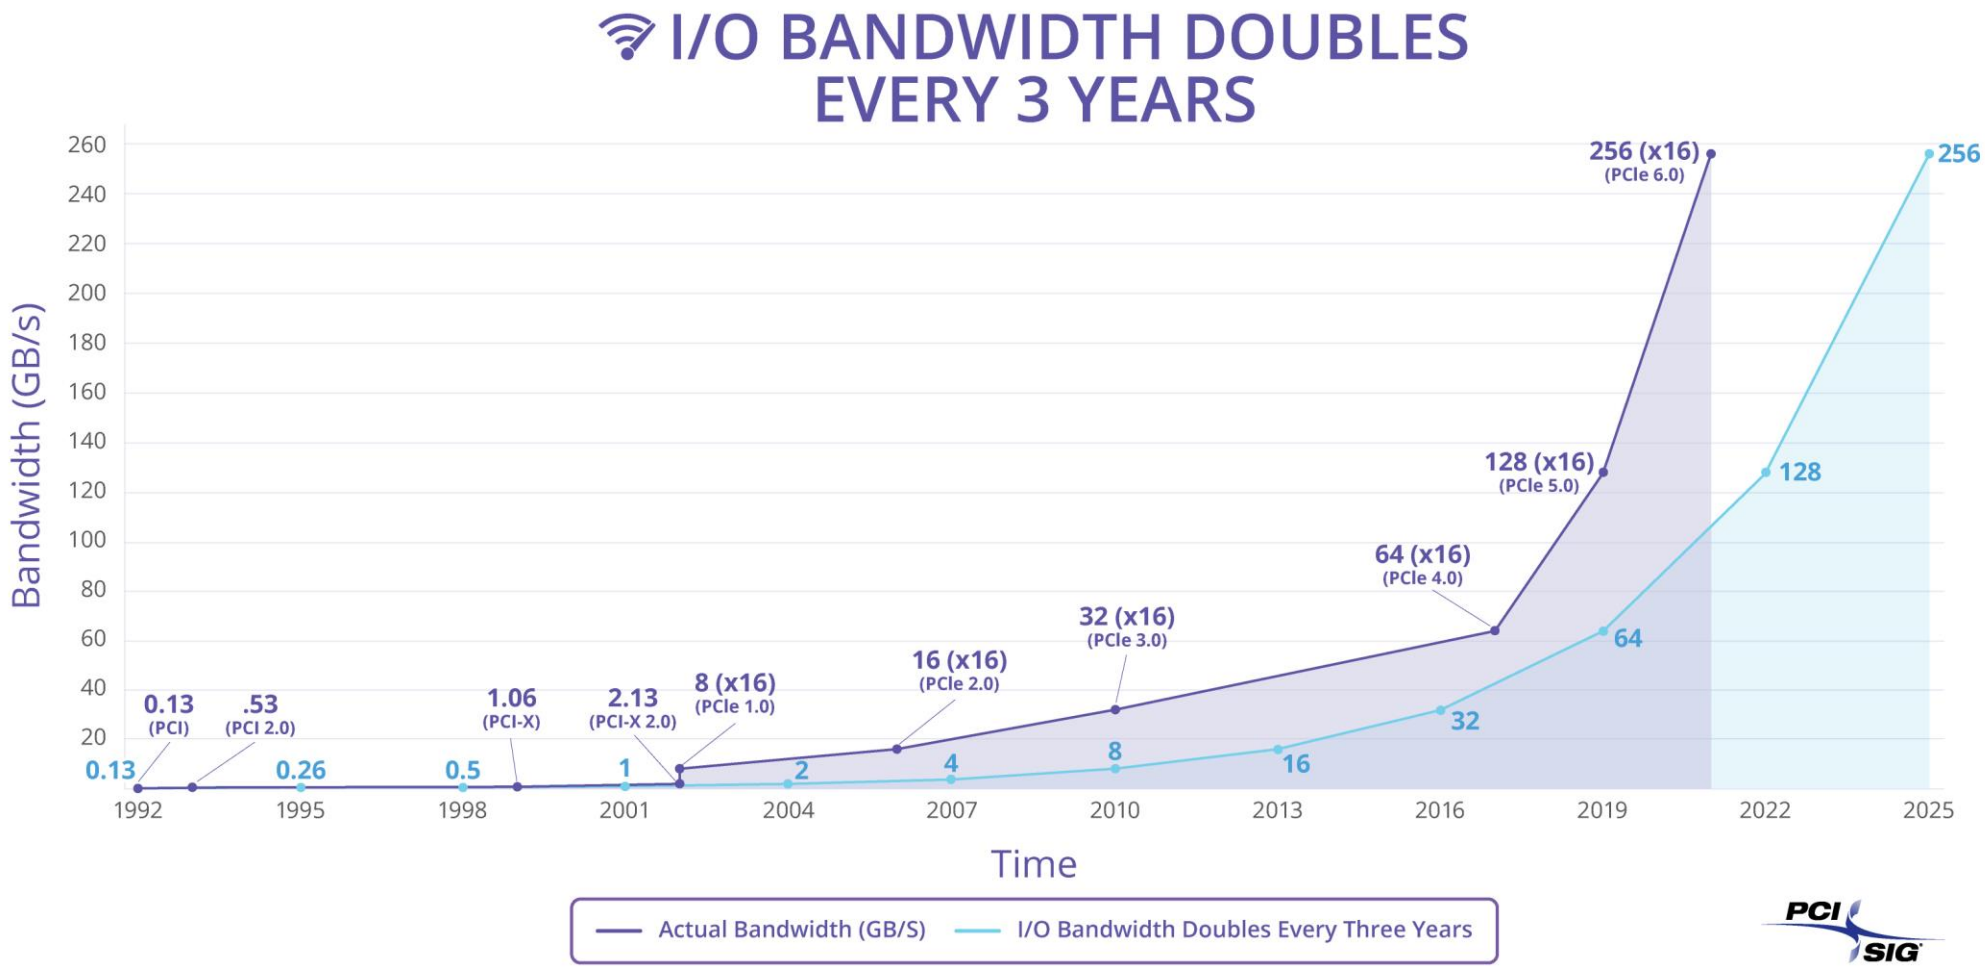

Copyright © 2020 PCI-SIG® - All Rights Reserved

Data Center / Enterprise SSD Units Trend

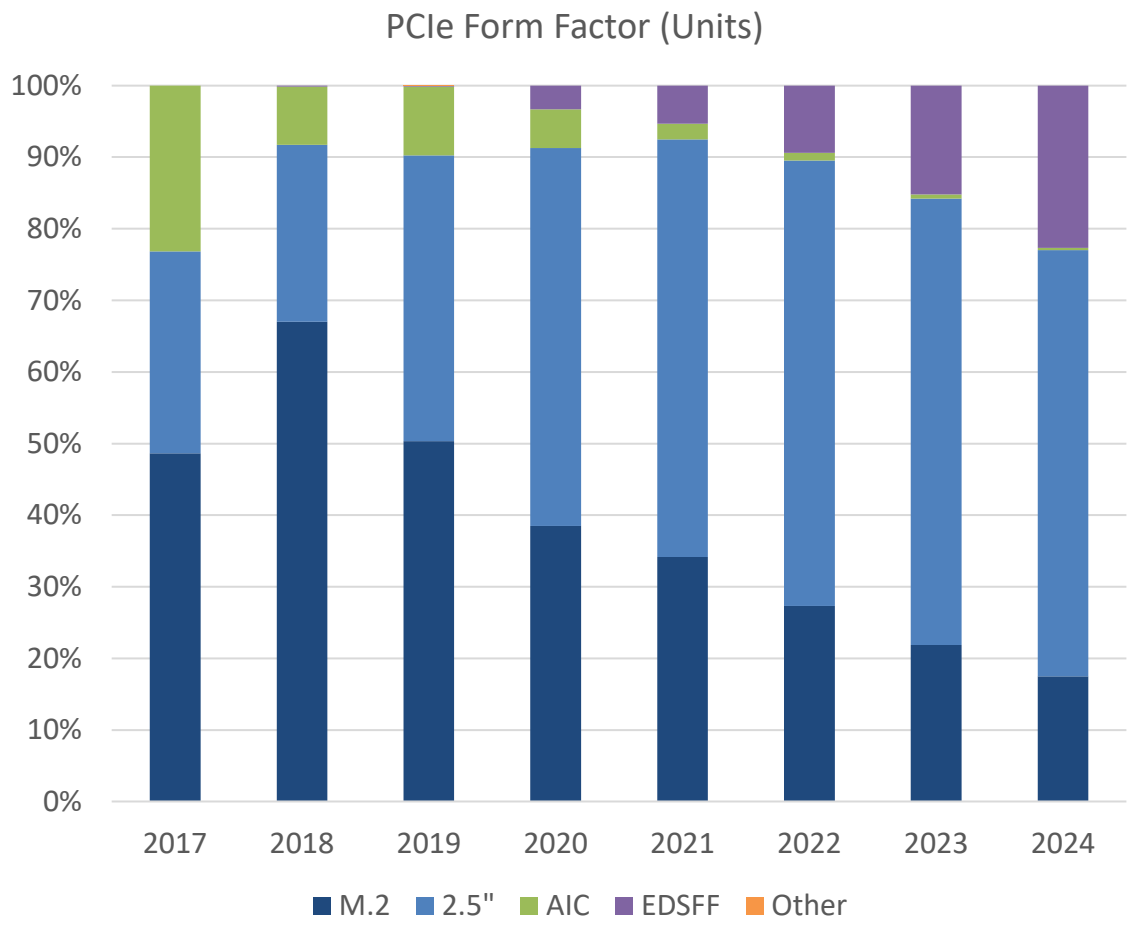

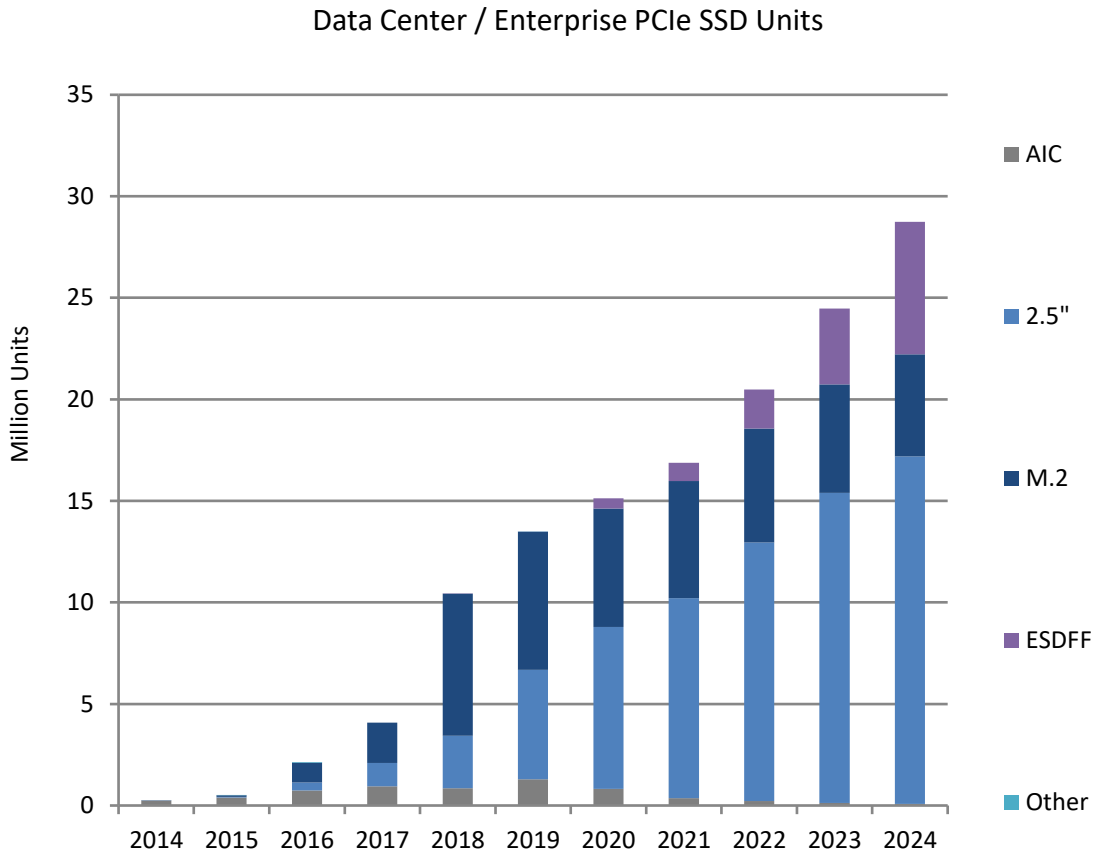

Source: Forward Insights, Q1'20

Copyright © 2020 PCI-SIG® - All Rights Reserved

PCI Express® Form Factors

M.2

42, 80, and 110mm lengths, smallest footprint of PCIe connector form factors, use for boot, for max storage density, for PXI/AXIe ecosystem

U.2

2.5in makes up the majority of SSDs sold today because of ease of deployment, hotplug, serviceability, and small form factor Single-Port x4 or Dual-Port x2

CEM Add-in-card

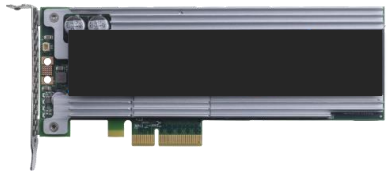

Add-in-card (AIC) has maximum system compatibility with existing servers and most reliable compliance program. Higher power envelope, and options for height and length

Copyright © 2020 PCI-SIG® - All Rights Reserved

Data Center SSDs: Previous and Current Options

AIC / CEM - Generic	M.2 – Consumer	2.5in Form Factor
		
Good: High-performance, general compatibility	Good: Small and Modular	Good: Hot-plug, Storage features
Bad: need PCIe® AIC slots for other devices, limited hot-plug	Bad: Low capacity, no hot-plug	Bad: Mechanical design descended from HDD
Ugly: consumes lots of space	Ugly: limited power and thermal scaling for data center use	Ugly: Blocks airflow to the hottest components in server

Copyright © 2020 PCI-SIG® - All Rights Reserved

Today's Challenges

Capacity Scaling	→	Density: Drive ease of expansion and optimize costs
Performance Scaling	→	Scale low latency bandwidth to the media
Thermally Efficient	→	Optimized for airflow to maximize cooling
Future Proofing	→	Support current and future interfaces, memory and devices
Solution Range	→	Flexibility: Drive consistency across family of devices

Copyright © 2020 PCI-SIG® - All Rights Reserved

What is EDSFF?

- Enterprise and Data Center SSD 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 toolless serviceability

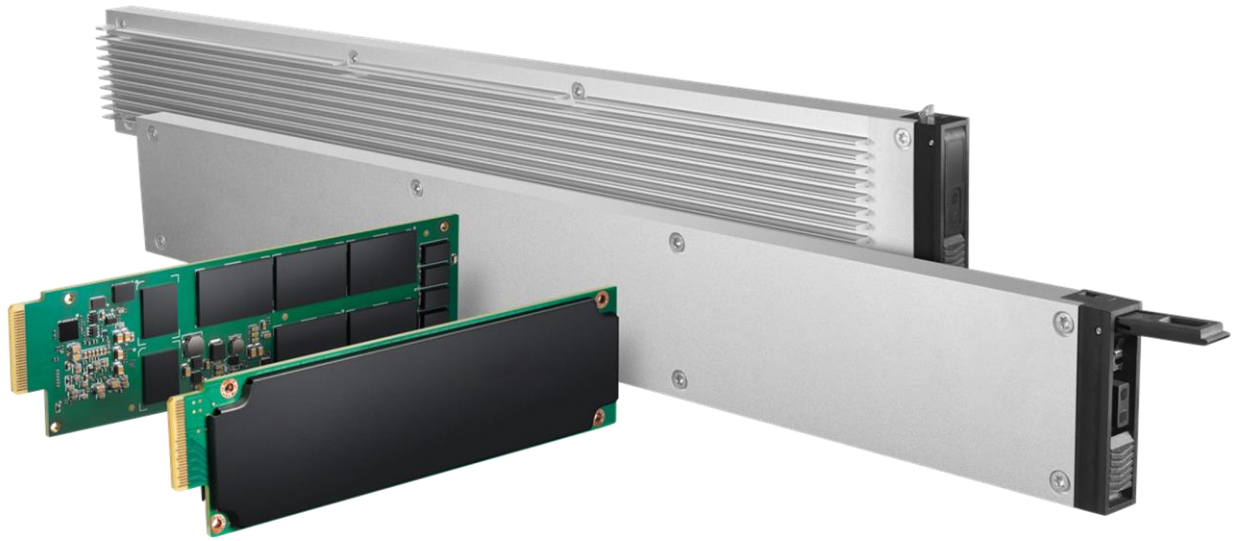

Copyright © 2019 PCI-SIG® - All Rights Reserved

The ESDFF WG/SNIA Form Factors

- Family of form factors and standards for data center devices
- E1.S for scalable & flexible performance
- E1.L for high capacity storage
- E3 for 2U higher power/performance or 1U horizontal

E1.S

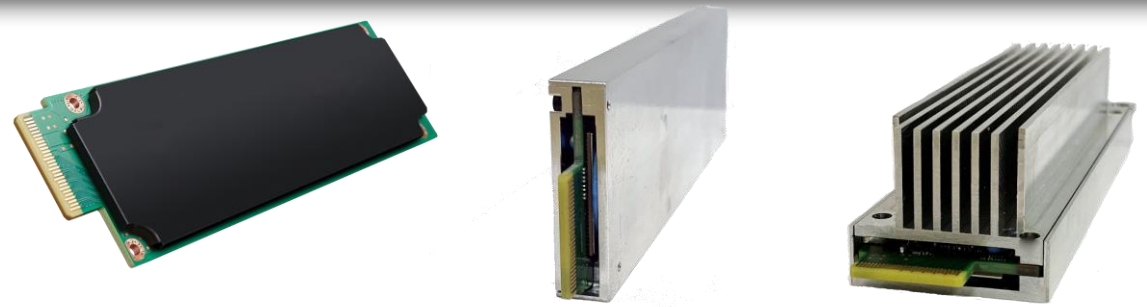

5.9mm 9.5mm 25mm

E1.L 9.5mm

E1.L 18mm

E3.L, E3.S

Thin/thick, short/long

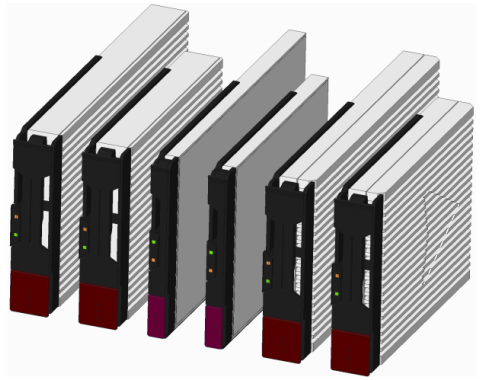

Copyright © 2020 PCI-SIG® - All Rights Reserved

How To Make an EDSFF Device

Features	Related Org. or Standard	Specification
Electrical, PHY, Channel, Link, Retimers, Transaction, Config	PCI-SIG®	**PCIe® Base Specification**
Command Set	NVMe, CXL	**NVMe Specification** **CXL Specification**
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

Copyright © 2020 PCI-SIG® - All Rights Reserved

Connector Ecosystem

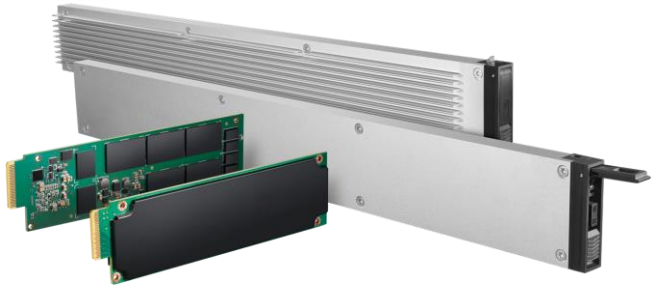

EDSFF

Native NVMe-oF™ Drives

High Speed Cables

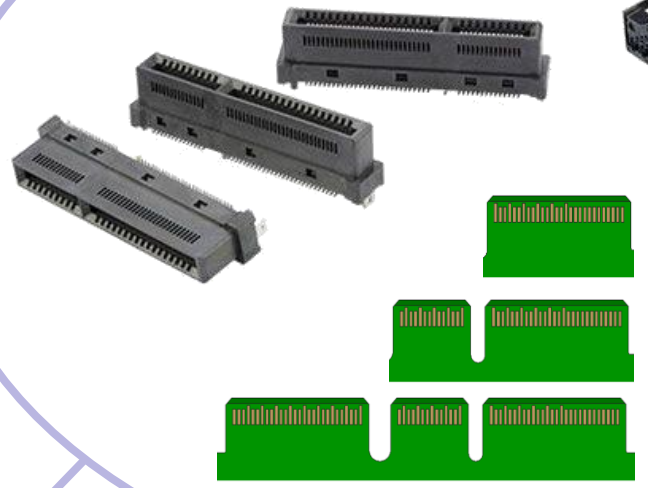

X4
X8
X16

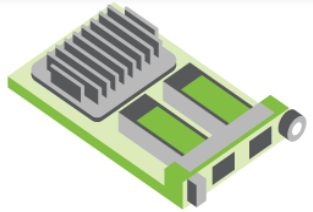

OCP NIC 3.0

PECFF / SFF-TA-1021

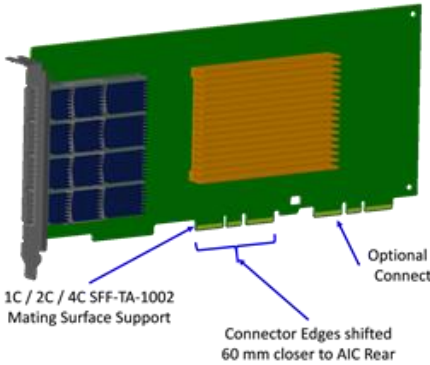

SFF-TA-1002 Connector versions: Vertical, Right-Angle, Orthogonal, Straddle, Cable. High speed up to 112 GT/s PAM4

SFF-TA-1012 shows pinout differences between EDSFF, OCP, etc.

Copyright © 2020 PCI-SIG® - All Rights Reserved

Ecosystem Work

Storage is driving the new form factor

- Every major SSD vendor supporting EDSFF (all have E1.S designs)
- Volume being generated from multiple system manufacturers
- SSD & system and thermal level collaboration happening at OCP
- EDSFF WG disbanded, all working being done in SNIA SFF

Copyright © 2020 PCI-SIG® - All Rights Reserved

E1.L

E1.L: Storage Reimagined

Optimized PCIe® technology routing allows for better signal integrity and lower cost attach per drive

More board space for higher capacity SSDs

Thermally efficient design with enclosure as heatsink

Designed to fit in 1U Server

Scalable power and performance

Carrier-less design with integrated latch and LEDs

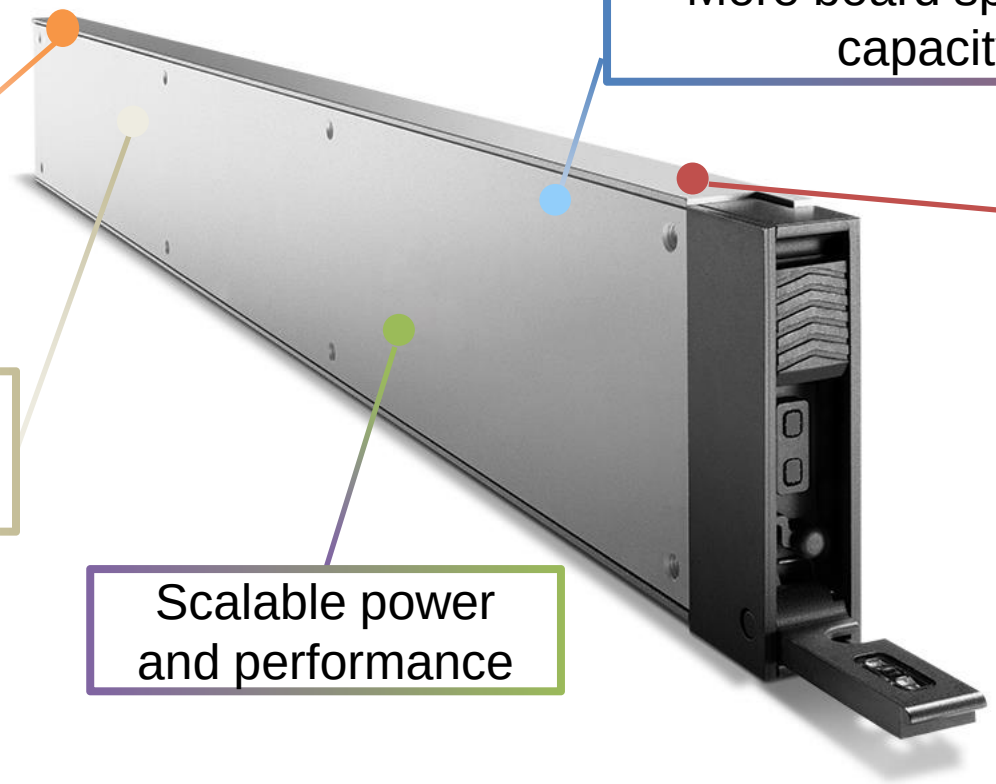

Power	Length	Width	Thickness
25W max	318.75mm	38.4mm	9.5mm
40W max	318.75	38.4mm	18mm

Scalable, thermally efficient, and dense, E1.L is designed for high capacity storage with superior serviceability

Copyright © 2020 PCI-SIG® - All Rights Reserved

System Designs for E1.L

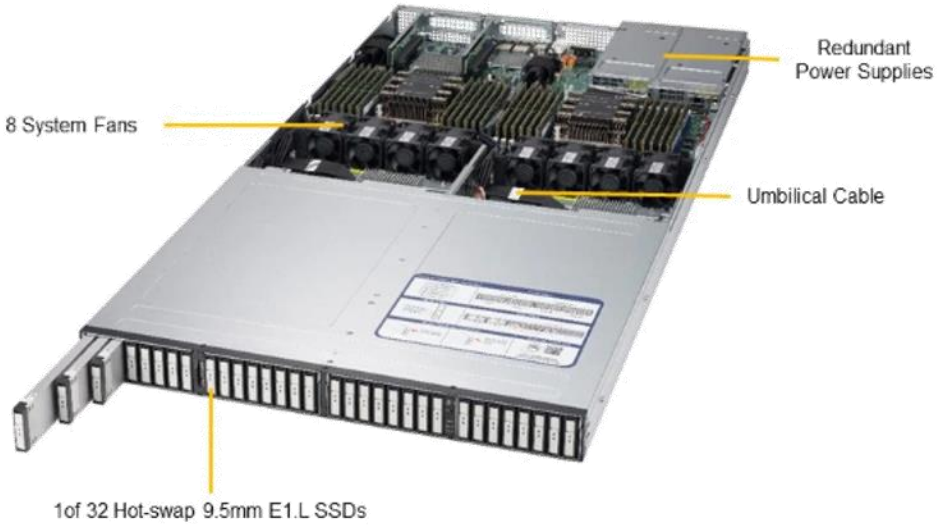

E1.L optimized for TB/rack unit & performance.
High capacity per drive and system

E1.L optimized for lowest TCO on QLC NVMe technology

https://www.supermicro.com/en/products/system/1U/1029/SSG-1029P-NEL32R.cfm

Copyright © 2020 PCI-SIG® - All Rights Reserved

https://www.wiwynn.com/schinese/product/type/details/2079?ptype=29

E1.S

E1.S: SFF-TA-1006 Industry Standard Form Factor

Vision:
Create a smaller, high density solid state drive standard that is optimized for the data center

E1.S (EDSFF 1U Short):

- Mainstream NVMe™ drive
- Compact, modular form factor
  - Vertical fit in 1U height (44.45mm)
  - 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

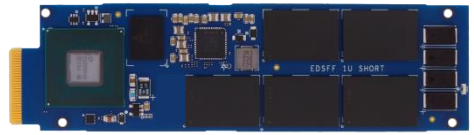

Optimized for the NVMe drive design and use across **all data center** and **edge** systems to scale as mainstream storage.

Copyright © 2020 PCI-SIG® - All Rights Reserved

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

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

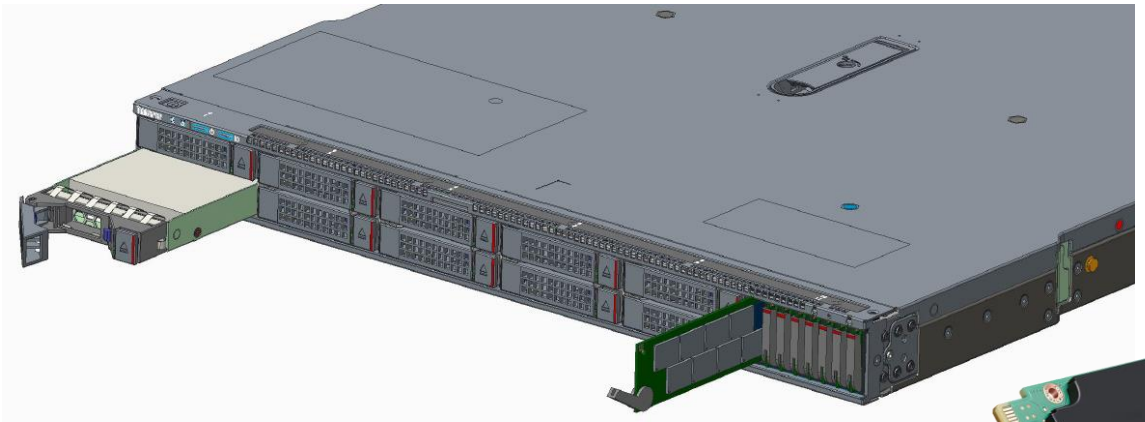

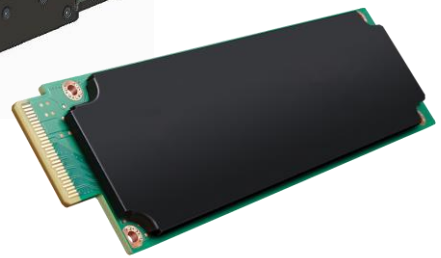

- Key benefits:
  - Much **smaller** enabling high **density** storage
  - Significantly improved **system airflow** and thermal solution
  - Most **efficient** modular **scaling** of NVMe capacity and performance
  - Enhanced feature set in **space-constrained** servers
  - **Lower** base system infrastructure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

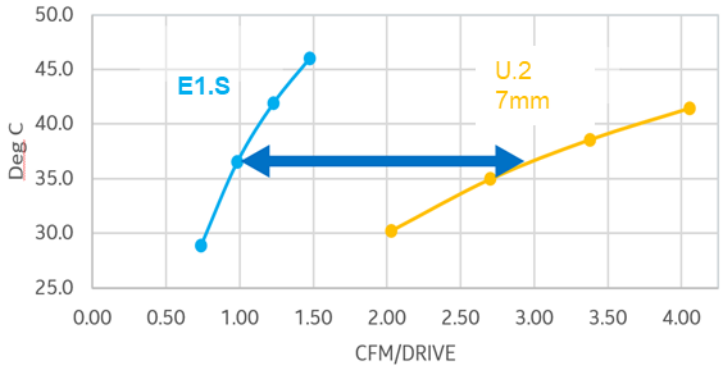

Copyright © 2020 PCI-SIG® - All Rights Reserved

E1.S Options: Dimensions and Power

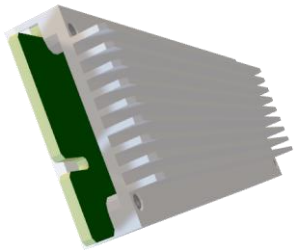

Enclosure Parameter	5.9mm Device	Device with Heat Spreader (8.01mm)	Device with Symmetric Enclosure (9.5mm)	Device with Asymmetric Enclosure (15mm)	Device with Asymmetric Enclosure (25mm)
Recommended sustained power (W)	12	16	20	20	25
Enclosure Max Inlet air temperature, 950 m to 3050 m (° C)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)	35 - (1° C for 175 m of elevation gain)
Add in card to add in card pitch (mm)	9	11	13	17	26
Recommended Fan Pressure loss across device (Pascal)	83	52	64	40	21
Airflow, average min per device (CFM). 1 CFM = 1.7 m3/h)	1.41 – (0.01 CFM for every 1° C below 35° C inlet temp)	1.71 – (0.06 CFM for every 1° C below 35° C inlet temp)	2.02 - (0.02 CFM for every 1° C below 35° C inlet temp)	1.5 - (0.02 CFM for every 1° C below 35° C inlet temp)	4.10 - (0.04 CFM for every 1° C below 35° C inlet

Copyright © 2020 PCI-SIG® - All Rights Reserved

E1.S Future Use Cases

- E1.S is ideal with its versatility as a common, high volume building block FF across systems and use cases
- In addition to scaling resources in datacenter use, **edge** and **IoT** use cases are also ideal to leverage its small, modular form factor
- Optimal future use focuses on system scaling of devices with low to moderate power (6-25W)
- Future uses beyond NVMe storage include DRAM memory, Persistent Memory, computational storage, PCIe® accelerators and Ethernet NICs

Copyright © 2020 PCI-SIG® - All Rights Reserved

E3

E3 Options: Dimensions and Power

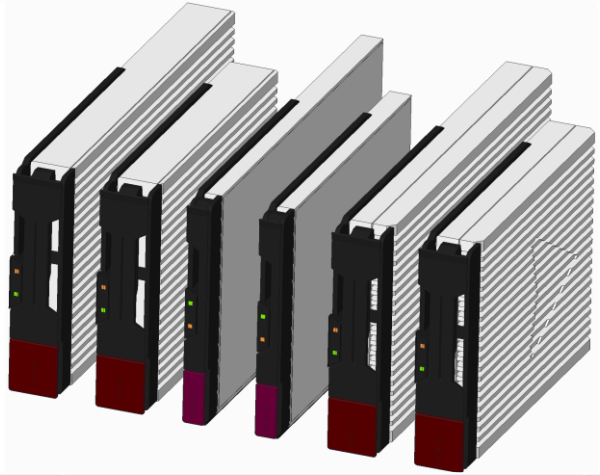

	Short/Thin	Short/Thick	Long/Thin	Long/Thick
Dimensions*	76mm x 104.9mm x7.5mm	76mm x 104.9mm x 16.8mm	76 x 142.2mm x 7.5mm	76 x 142.2mm x 16.8mm
Power	20W	40W	35W	70W
Recommended drive Pitch	9.3mm	18.6mm	9.3mm	18.6mm
Airflow (min per device) CFM – x for every 1° C below 35° C inlet temp	1.8 CFM	2.8 CFM	2.7 CFM	4.2 CFM

*These dimensions are expected to change with pending spec updates

Copyright © 2020 PCI-SIG® - All Rights Reserved

E3 Use Cases

- NVMe™ SSD devices with very high capacity and performance are possible due to larger package size, 16 lanes and increased power

-also-

- Persistent memory devices using E3 will appear with the coming PCIe® 5.0 platforms
- GPUs up to 70W are possible
- 400GB NICs

Copyright © 2020 PCI-SIG® - All Rights Reserved

E3 Opportunity in Servers

Potential Reasons to Move Away From U.2		
Signal Integrity for PCIe® 5.0 spec @ 32GT/s 	Multi PCIe device (beyond SSD) support 	Decreased server infrastructure cost 
Improved thermals & cooling 	Higher power & perf modes 	Improved capacity 

Copyright © 2020 PCI-SIG® - All Rights Reserved

Real Performance

- Once NVMe controllers with 16 lanes of PCIe® technology are available, they will have superior burst transfer rates
- PM (Persistent Memory) cards in E3 would replace PM in DDR slots
- PCIe 5.0 architecture is essential for 400G NICs

Copyright © 2020 PCI-SIG® - All Rights Reserved

EDSFF Device Solutions

Capacity Scaling	E1.L: 3x more capacity than U.2, 4x the drives in system of U.2 E1.S: 2x more capacity than M.2, 3-4x the drives in system of U.2
Performance Scaling	EDSFF Connector Support for x4, x8, x16. Up to 112Gbps signaling E3: 2x+ higher power support and performance than U.2
Thermally Efficient	E1.L: Up to 2x less airflow required per drive vs. U.2 15mm E1.S: Up to 3x less that U.2 7mm
Future Proofed	PCIe 4.0® and 5.0 specification ready - enabling scalability & Interoperability to be the innovation form factor for the next 20 years
Solution Range	E1 and E3; low, med, and high power; case and caseless designs

Copyright © 2020 PCI-SIG® - All Rights Reserved

Join PCI-SIG®

PCI-SIG members have access to the PCIe® specification library. If you would like to learn more about joining, please visit the PCI-SIG website:
https://pcisig.com/membership/become-member

@PCI_SIG

https://www.linkedin.com/company/pcisig/

PCI-SIG

Copyright © 2020 PCI-SIG® - All Rights Reserved

Questions?

Copyright © 2020 PCI-SIG® - All Rights Reserved

Thank You For Attending