



SUPERMICRO X13 PETASCALE STORAGE SERVERS EMBRACE ADVANCED STORAGE TECHNOLOGY

E1.S and E3.S Drives Deliver the Performance and Packaging Required for Next Generation Compute and Storage Systems

TABLE OF CONTENTS

Executive Summary..... 1

Basics 2

Form Factors – Why E1.S Will Prevail..... 4

E1.L Description..... 6

Looking Forward..... 6

References..... 7

SUPERMICRO

Supermicro (Nasdaq: SMCI), the leading innovator in high-performance, high-efficiency server and storage technology is a premier provider of advanced server Building Block Solutions® for Enterprise Data Center, Cloud Computing, Artificial Intelligence, and Edge Computing Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green®” initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.



Supermicro X13 Petascale Storage System

Executive Summary

Storage has become a fundamental requirement of modern data centers as the amount of data generated continues to grow¹. An increasing amount of data is constantly being collected at the edge, and a percentage of that data may be moved to larger cloud or enterprise data centers. To effectively manage and store the data from the edge to the cloud, standards need to be developed that create a competitive eco-system that gives IT managers choices depending on their workloads and data access requirements. For example, the SNIA standard-based EDSFF form factors are designed so that airflow can easily reach the CPUs directly compared to earlier storage media like U.2 storage devices where a backplane is needed. The overall design of systems should not be taken lightly, as innovative mechanical designs need to consider the cooling and packaging requirements. The realization that the new storage devices can affect the CPUs' overall physical design

and cooling requirements has led to the creation of this form factor that delivers high capacities while minimizing airflow restrictions.

Basics

Storage for systems has moved from sizeable rotating disk drives (HDDs) to smaller and faster drives that fit into almost any form factor. Over the past 10-20 years, popular sizes have been the 2.5" and 3.5" HDDs, which fit into a wide array of servers and storage appliances. The Solid State Drive (SSD) has become the most popular storage device for high-powered data center servers, workstations, and laptops. With a current storage capacity in the Petabyte (PB) range per server, the SSD has become the device of choice due to data storage and retrieval speed and the lower power consumption of solid-state technology. However, a 2.5" or 3.5" device will not physically fit in many of the new types of servers used today. SSD form factors were initially developed to replace HDDs within the server. This led to inefficient designs, which were forced into a form factor that had detrimental effects on the critical airflow design within the server.

A storage device that is used within a server needs to fit within an "envelope," which will be limited by the following constraints:

1. Physical Size: Since servers are designed for different environments requiring different physical sizes, the storage device must fit within the server dimensions allocated for this function. Various form factors have been used for storage over time. However, new generations of servers, which are smaller than before that may be installed at the edge, require smaller form factor storage. The most common shape is a "stick," also referred to as a "ruler." These devices are designed to connect via a PCIe connector, utilizing the latest PCIe technology and specifications.

2. Thermal Limits: The CPUs and GPUs in a server continue to produce more heat than the rest of the components. Therefore, the mechanical design of a system is critical in allowing cooling air to reach the CPUs and GPUs with no obstructions. Supermicro systems such as the SuperStorage product line are designed to operate at the design maximum for every component. With this increase in component utilization, record-breaking performance can be achieved, but at the cost of generating more heat than a lightly used system. For this reason, innovative thermal designs are used for maximum airflow.

The EDSFF backplanes are installed in a system parallel to the airflow direction, allowing higher amounts of air to reach the CPUs than a traditional perpendicular PCBA backplane. As a result, more powerful CPUs and GPUs can be used with the EDSFF form factors, increasing performance while remaining within the thermal design envelopes.

In addition, the drive will heat up as SSDs receive data from the host CPU or deliver data back. Dissipating the heat is critical to maintaining the integrity of the drive within the specified operating performance. Likewise, the airflow over an SSD is essential to remove the heat. Thus, an SSD enclosure like E1.S includes more surface area and built-in cooling fins, sized to allow efficient airflow and heat dissipation, allowing higher TDP processors to be used in the server.

3. Capacity: Every organization will need different amounts of storage capacity, so flexibility in various capacities while retaining the form factor is essential for IT managers. As the amount of data generated and needing to be stored continues to grow, new SSD capacities must also increase. Keeping the data close to the CPU (within the same system) and in persistent memory requires increased storage density. New technologies enable higher-density drives, which are now in the double-digit TB range but increasing rapidly.

4. **Connectivity:** Standard interfaces are necessary for the broad adoption of any technology. Standards have been shown to increase markets and encourage innovation and a range of suppliers. The key to the wide adoption of this technology is using the same connectors and protocols in various form factors from multiple vendors. Therefore, connectors must be standardized, allowing different vendors to supply these SSDs, keeping competition innovative.

Current Technologies

Today's systems utilize the PCIe 5.0 bus for connecting storage devices to the CPUs. With the new PCIe 5.0 specification, the performance will increase tremendously compared to the PCIe 4.0 standard (2X throughput), requiring new storage technologies to get maximum performance from the entire system. While current storage systems may have an unbalanced architecture, as one of the main sub-systems cannot deliver or analyze data as fast as other components can, future servers that incorporate PCIe 5.0 and next-generation CPUs will be fully able to take advantage of E3.S storage technology. The challenge with the design of future high-end systems is to cool the CPUs and GPUs (if installed). While liquid cooling remains an option, most racks of servers will still be utilizing air cooling. Thus, the density and placement of the storage hardware for easy serviceability will be essential for overall system operation. Since most servers today move air from the front to the back, serviceable storage components will reside at the system's front and must not block airflow.

Supermicro X13 EDSFF Storage Systems in 1U and 2U form factors

The new X13 Supermicro EDSFF storage Platforms are designed for maximum flexibility and density in a given rack unit height (U). While E1.S and E1.L are optimized for 1U servers, E3.S is on the horizon, and this media is optimized for 2U servers. The upcoming E3.S storage system uses the same Motherboard and building block architecture, with two 4th Gen Intel Xeon Scalable processors, 32 DIMMs, and up to 24 hot-swappable EDSFF E1.S 9.5mm or 15mm) NVMe storage devices, or (future) 24 hot-swappable EDSFF E3.S 7.5 mm NVMe storage devices and 2x AIOM (OCP 3.0) networking cards as well as 2x standard FHHL expansion slots. Front IO capabilities, as well as support for CXL devices.

Supermicro Model #	Height (U)	# 4 th Gen Intel Xeon Scalable Processors	EDSFF configurations	Drive Type	Max Capacity Per NVMe Device, Total Capacity Per Server*
SSG-121E-NES24R	1	2	24	E1.S (9.5mm or 15mm)	7.68 TB, 184.32 TB
SSG-221E-NE324R	2	2	24 Up to 32 planned	E3.S (7.5mm)	>30TB has been announced >720TB




Image 1 - Supermicro Storage System SSG-121E-NES24R

Storage Networking Industry Association (SNIA)

SNIA, a non-profit organization made up of member companies spanning information technology, maintains the specifications for different form factors of storage devices. The latest one becoming widely adopted is the Enterprise and Data Center Standard Form Factor or EDSFF, described below:

Form Factors – Why E1.S Will Prevail

The current form factor gaining popularity for data center storage is designated the E1.S, which has a technical specification defined by the SNIA through the SNIA SFF Technology Affiliate (TA) Technical Work Group (TWG). The E1.S form factor is geared to replace the M.2 drive form factor in the data center due to a smaller form factor, reduced cooling requirements, and increased capacity. In addition, the E1.S form factors allow an increased amount of air to be moved through the server, resulting in cooler CPUs, less likelihood of failure, and reduced fan speed and power needed.



While a longer (and higher capacity) form factor, called the E1.L, is available, many enterprises will realize that the E1.S will satisfy capacity, longevity, and performance needs. In addition, the E1.S allows the popular 1U servers to accommodate up to 32 front-mounted drives, which are also easily serviceable. Finally, all-flash vendors produce E1.S models reflecting EDSFF acceptance utilizing different heat sink sizes. Some models are only available directly from the media vendor to the hyperscale customer.



Image 2 - E1.S – 15mm with cooling fins and 9.5 without cooling fins

Form factors for Supermicro’s E1.[S, L] are as follows:

EDSFF Formfactor	Supermicro Server Generation	Width of Enclosure	Recommended sustained power	Length
E1.S	X11 -	5.9mm	12 watts	111.49mm
E1.S	X11, X13	9.5mm	20 watts	118.75mm
E1.S	X13	15mm	20+ watts	118.75mm
E1.S	N/A	25mm	25+ watts	118.75mm
E1.L	X11	9.5mm	25 watts	318.75mm
E1.L	N/A	18mm	40 watts	318.75mm

Table 1 - Form Factor Comparison




Image 3 - From <https://phisonblog.com/nand-flash-101-enterprise-ssd-form-factors-simplified-2/>

There are various reasons why a customer would want an E1.S or an E1.L SSD. The E1.S is obviously shorter, taking up less space in a compact server. The E1.L has a larger cooling surface, allowing higher operating temperatures. Within the E1.S options, there is a tradeoff between the different widths available (which reflects the cooling capacity) and the number of drives that can be installed and easily serviced within a server.

A 1U high end server will contain up to 32 of the 5.9mm or 9.5 mm E1.S SSDs.



Image 4 - X11 Server with 32 x 5.9mm or 9.5mm E1.S

Similarly, 24 of the 15mm E1.S SSDs can be installed in a 1U server.



Image 5 – Supermicro X13 Server SSG-121E-NES24R with 24 x 9.5mm or 15mm E1.S

E1.L Description

The E1.L is a form factor optimized for data center storage systems, such as storage arrays. Longer than the E1.S at 318.75 mm length, the E1.L can store more data than an E1.S, but at a tradeoff of space in the server. In addition, the E1.L is specified at two different thicknesses, of 9.5mm and 18mm, which allow for more cooling (heat sinks) and thus increase storage capacity and activity.




Image 6 - E1.L SSDs

Related to the E1.S and E3.L is the E3 specification. The E3 form factor is designed to replace, over time, the U.2 2.5” form factor SSD.

Looking Forward

A balanced system for a given workload will not have any bottlenecks that limit performance. Since the CPUs and GPUs are the highest valued item in many servers that are designed for simulations or data analytics, keeping a system busy means getting the CPUs / GPUs the required data with few stalls. Fast storage is critical, but bus speed is vital as well. Future servers will need to incorporate the ability to communicate with other sub-systems with the latest bus technology (PCIe). The new form factors, the E1.S and E1.L, will support new generations of PCIe, which will move data center servers to the next level.

PCIe has undergone tremendous improvements in performance since 2003. As a result, PCIe 5.0 will be the ideal bus for E1.S SSDs to avoid bottlenecks.

Version of PCIe	Introduced	Transfer Rate	X16 Throughput
1.0	2003	2.5 GT/s	4.000 GB/s
2.0	2007	5.0 GT/s	8.000 GB/s

3.0	2010	8.0 GT/s	15.754 GB/s
4.0	2017	16.0 GT/s	31.508 GB/s
5.0	2019	32.0 GT/s	63.015 GB/s
6.0	2021	64.0 GT/s	126.031 GB/s

Table 2 - PCIe Generational Comparison

The general requirements for future systems will include the following:

- CPUs need to be kept busy
- Memory needs to be fast enough to feed the CPUs
- Bus speed needs to be fast enough to feed data to the memory system
- Storage needs to be fast enough so as not to be a bottleneck
- Mechanical designs that allow airflow throughout the entire system

References and Further Reading

1. <https://www.idc.com/getdoc.jsp?containerId=prUS46286020>

<https://www.servethehome.com/e1-and-e3-edsff-to-take-over-from-m-2-and-2-5-in-ssds-kioxia/>

<https://www.storagereview.com/news/ruler-is-finally-going-mainstream-why-e1-s-ssds-are-taking-over/>

<https://www.snia.org/forums/cmsi/knowledge/formfactors/>

<https://www.servethehome.com/next-gen-supermicro-ssg-121p-nes24r-1u-24-bay-nvme-fms-2022-intel-sapphire-rapids/>