Inside Track
Executive Brief

FREEFORM DYNAMICS

IN ASSOCIATION WITH

FUJITSU

EDSFF: a new standard for enterprise SSDs

Bringing NMVe up to date with the needs of the data center

Freeform Dynamics, 2021

Introduction

Enterprises have rapidly adopted solid-state storage, such as Flash and now storage-class memory (SCM), initially to augment hard disks and often then to replace them entirely with SSDs. This has brought significantly greater storage performance, and in some cases has enabled new ways of designing and building applications.

As part of this, storage designers have developed new technologies to allow fast modern servers to take advantage of the performance characteristics of SSDs, characteristics that differ significantly from those of spinning disks. Key among these new technologies is NVMe, a lightweight protocol designed to be highly future-tolerant, so it will work with forthcoming storage technologies as well as today's.

However, there are issues in scaling-out NVMe storage within an enterprise server. NVMe uses the well-established PCIe standard for physical connectivity, but servers have a limited number of PCIe slots and connectors. Worse, most of these are internal to the system, which limits physical access, and NMVe devices must compete for them with the likes of graphics cards and network interfaces.

Something new is needed, therefore – a way to bring NVMe, and by extension PCIe, to the front of the data center rack. Enter EDSFF, a new small scale-out format for NVMe SSDs that allows them to plug right into the front of a server or storage system.

The need for new storage designs

The story of computing is, in many ways, one of evolution and adaptation. We are constantly looking for ways to make new technology look like old technology, so we don't have to change too much of what's around it. There always come a time though when big changes must be made – when we have to stop compromising and start afresh.

So it is with enterprise storage. We have moved from spinning disks to SSDs, which use Flash memory to emulate a disk, and we have moved from parallel SCSI connections to SAS, which sends SCSI commands over a faster serial link. Now, we are moving from disk-optimized SAS to Flash-optimized NVMe, which uses the PCIe interface originally meant for graphics cards, disk controllers and the like.

At every step though, there are compromises when it comes to server and data center functionality. For example, the U.2 SSD interface is PCIe-based but it uses 2.5" disk drive enclosures which are not density-efficient packaging for Flash, plus they are relatively power-hungry and hard to cool. Meanwhile, the M.2 mini-card SSD format is a better way to connect Flash, as it can support NVMe, but is still not optimal for data center use: it has a limited power budget and, like the U.2, it is not hot-swappable.

There is an opportunity, therefore, for a new SSD format. One designed to package Flash efficiently, optimized for cooling and power, enterprise-oriented, and purpose-built for very high density in the racks of today's space-constrained data centers. The result is EDSFF: the Enterprise and Data Center SSD Form Factor specification.

What is EDSFF?

The EDSFF Working Group is part of SNIA, the trade association which oversees and co-ordinates the storage industry's work on standards, interoperability and education. Its eponymous specification covers a set of new high-capacity storage formats for plug-in modules, all purpose-designed to meet data center needs such as high density, thermal efficiency, fast connectivity, and easy hot-swappability.

EDSFF currently defines two widths (E1 and E3) and lengths (L and S) for modules, plus three different connectors for PCIe x4, x8 and x16. The connectors are forward and backward compatible, so x4 devices will work in x16 slots and x16 devices should work in x4 slots, though in both cases it will be at x4 speed of course. There are also multiple thicknesses to allow for modules that need bigger heatsinks, for example compact SSDs that use PCIe x8 instead of x4.

The narrower E1 modules are 32mm wide which allows them to fit vertically in a 1U rack, hence the name, while the wider 76mm E3 versions (the 3 in this case is the width in inches) are a similar size to a U.2 2.5" SSD – but can contain more Flash and a faster PCIe-based NVMe interface than a U.2 device.

EDSFF use cases and benefits

Let's look at the E1 and E3 formats, and their potential, in more depth. To begin with, EDSFF is designed for serviceability – each module has LED status indicators and as mentioned above, modules are hot-swappable (hot-plug) without the need for tools.

E1: the new measure for server SSDs

Based on Intel's earlier Ruler format, it is a flat stick that resembles a measuring ruler and is wide enough to fit vertically in 1U. It has a tongue-like PCIe x4 or x8 connector at the back end, and status LEDs on the front. The SSD slides lengthways into a slot in a server or storage system. Typically, there will also be a lever on the front of the SSD to latch it in place, and when unlatched this of course acts as a handle to pull it out.

Figure 1: E1 SSDs in a dense 2U rack unit

As well as supporting NVMe and packaging Flash more efficiently than legacy SSD types, the two E1 formats are designed to permit good airflow and cooling. All this makes them well suited to the requirements of scale-out and hyperscale data centers, and to the broader need for dense SSD capacity both in servers and in disaggregated storage infrastructure.

Storage infrastructure

The E1.L 'Long' version is intended to prioritize high capacity, for use in storage infrastructure. At 318mm in length, it is almost three times as long as the 112mm 'Short' E1.S and can hold much more Flash – for example, Intel has announced 30TB SSDs using the E1.L format.

Depending on their thickness, you can slot as many as 36 E1 SSDs side by side across a standard rack. This means that a single 1U rack unit could contain almost one petabyte of raw front-accessible SSD storage, or 3PB to 4PB once you add data reduction.

Dense and space-constrained servers

E1.S, on the other hand, is designed to balance high performance with high capacity for space-constrained devices such as 1U servers. In effect it is a more capable and enterprise-grade replacement for the M.2 format, as well as for 2.5" U.2 SSDs: not only can each E1.S SSD hold (and power) more Flash than those earlier formats, but the E1.S drives can pack more densely too. For example, a cage containing six hot-swap E1.S SSDs can fit in the chassis space formerly occupied by two 2.5" drives.

E3: enterprise server expansion

The upcoming E3.S and E3.L formats are larger than their E1 equivalents, and will offer a broader range of options where enterprise servers are concerned. Here, 2U is a more common system height – E3 can fit vertically into 2U, or horizontally into 1U.

The primary use-case for E3 is still for NVMe-connected storage devices such as SSDs and storage-class memory (SCM). However, its size and its ability to support PCIe x16 connectors means that other devices, such as network adapters or accelerators, could also be built as EDSFF modules. Indeed, when revision 2.0 of the E3 spec was released in 2020, it added compatibility with the Open Compute Project's NIC 3.0 spec. Future servers could therefore have slots capable of accepting a range of storage and non-storage devices – E3 rev 2.0 hardware is likely to reach the market by 2022.

Compatibility with existing chassis designs

As with E1.S, the E3 module sizes have been defined to make it easier for system builders to adapt existing designs. An enclosure for E3.L modules can fit within a space that formerly held 3.5" drives (note here that despite both being 'long', E3.L is less than half the length of E1.L), while a cage holding E3.S short SSDs is as deep as a cage holding 2.5" drives. Again, their optimized design means these SSDs allow considerably greater storage density than was previously possible – for example, a 2U rack unit could hold up to 48 front-facing E3.L SSDs.

What to watch out for

Systems: As with any new format, there will be an adjustment period as the industry absorbs the new technology. On the server side, this initially involves system builders replacing existing SSD mounting points with the cages to hold E1.S SSDs, and of course adding any additional PCIe connectivity that might be required.

On the storage system side, new rack units will be needed to hold E1.L SSDs, and they may require high-capacity, resilient power supplies to meet the needs of such dense storage. But of course, other design aspects of both servers and storage systems will remain the same: it is still NVMe and PCIe, both of which are now well understood.

As E3 devices come into commercial production, there will be more scope for change and innovation. System designers could have new degrees of freedom, for example to build servers able to accept GPU-based compute accelerators for AI/ML or HPC work, or DPU cards for network and security offload.

Storage: Several companies have developed or are developing EDSFF SSDs. The differentiators between their products will of course include the type and amount of memory they contain, and also their power consumption. Higher-powered devices will need to use the thicker EDSFF profiles in order to include larger heat sinks. In memory-intensive applications, using fewer but thicker SSDs may yield better TCO.

Suppliers: It is likely that generic EDSFF-compatible systems will be widely available, especially those using E1.S as the new standard for hot-pluggable enterprise-grade SSD. However, the range of suppliers able to meet more complex data center needs – properly assessing the balance between storage density and heat dissipation, for example, and ensuring effective utilization of the upcoming PCIe gen 5 standard – may be rather smaller.

So the need for solid technical skills and local support capabilities will not change. Indeed, as EDSFF evolves and expands, there will be scope for additional support, as enterprises learn to make use of the new capabilities it brings. Users will also want to check that their supplier can test and qualify their systems with new EDSFF modules as they reach the market, too.

Summary

NVMe was, and is, a huge step forward for the software side of storage connectivity. It enabled system designers to discard much of the remaining hard disk 'baggage' in favor of a storage protocol designed from the ground up to work with today's – and tomorrow's – solid-state storage technologies.

However, on the hardware interface side it had limitations. Few systems have enough spare PCIe slots to support more than one or two NVMe drives. In addition, while drive formats such as M.2 and U.2 could yield good performance, they still included legacy design elements as they were built to fit in with existing system architectures.

EDSFF is in many ways the data center-oriented hardware upgrade that NVMe was missing. Its physical design owes little or nothing to the storage devices that went before it. Instead, like NVMe on the software side, it is optimized for Flash, SCM and whatever might come next.

It also allows designers to exploit the falling price of Flash and target high capacities, rather than the “just enough to be cost-effective” approach taken in previous SSD generations. And of course it is not intended to be the only drive in a system – as M.2 is the boot drive in a laptop, say. In contrast, EDSFF defines hot-swappable storage modules with data center applications in mind.

In short then, EDSFF is the hardware specification that solid-state storage was waiting for. As NVMe did on the software side, it jettisons all the leftovers from the days of the hard disk, and instead it focuses on taking what’s special about the SSD and making that data center-ready. It does mean change, and that will take time, but the promise is a strong one.

About Freeform Dynamics

Freeform Dynamics is an IT industry analyst firm. Through our research and insights, we help busy IT and business professionals get up to speed on the latest technology developments and make better-informed investment decisions.

For more information and access to our library of free research, please visit www.freeformdynamics.com or follow us on Twitter @FreeformCentral.

About Fujitsu

At Fujitsu, we're passionate about using technology to create a more inclusive, sustainable and trusted future. It drives everything we do. Throughout our history, we've supported businesses and society through delivering robust and reliable IT systems.

Find out more about our business, our history, our philosophy and the countries we operate in. Please visit www.fujitsu.com