

IDC Innovators

IDC Innovators: NVMe over TCP, 2019

Eric Burgener

IDC INNOVATORS IN NVME OVER TCP

There are multiple NVMe over Fabric (NVMe-oF) host connection options available using different transports (Fibre Channel, Ethernet, InfiniBand), but to date, all these have required custom content on the server side (proprietary drivers, special host bus adapters). NVMe over Terminal Control Protocol (TCP) (NVMe/TCP) is a new NVMe-oF implementation that does *not* require custom content on the server side but still provides the performance benefits of NVMe-oF. It is the NVMe-oF version that IDC believes will ultimately enjoy the broadest adoption in the industry. For a broader discussion of the importance of NVMe-oF, see *IDC TechBrief: NVMe Over Fabric* (IDC #US43854018, September 2018). NVMe-oF host connections are used on NVMe-based all-flash arrays (NAFAs).

IDC has identified three start-ups that are taking different approaches with their NVMe/TCP (and related NVMe-based systems) implementations, each of which offer customers different options. This IDC Innovators document reviews each of these three vendors, calling out the differentiation that their approach (and their products) provide to potential customers (see Figures 1-4).

FIGURE 1

IDC Innovators in NVMe over TCP, 2019

		IDC Innovators are emerging vendors with annual revenue <\$100 million that have an innovative new technology or a groundbreaking business model.	
NVMe over TCP			
Company Name:		Founded:	Headquarters:
Lightbits Labs		2016	Kfar Saba, Israel
Excelero		2014	San Jose, California
Pavilion Data		2014	San Jose, California

Source: IDC, 2019

FIGURE 2

Lightbits Labs

Why Lightbits Labs Was Chosen as an IDC Innovator

Lightbits Labs introduced a software-defined storage solution that is built entirely around NVMe and is the primary (although not the sole) author of the NVMe over TCP (NVMe/TCP) NVM Express standard. It is a target-only solution with no host touch points and offers a variety of consumption models that allows the company to target both enterprise and cloud customers.

Company Name				
 Founded 2016	 Number of Employees 70	 Headquarters Kfar Saba, Israel	 Number of Customers Unavailable	
 Product Names LightOS, LightField, SuperSSD		 Founders Avigdor Willenz, Eran Kirzner, Muli Ben-Yehuda, Ofir Efrati, Sagi Grimberg		
 Profiled Products LightOS, LightField, SuperSSD		 Funding Lightbits Labs has raised \$50 million in funding across 2 rounds (with the most recent one being a Series B round in December 2018)		 Typical Deal Size \$100,000 to \$5 million (varies by customer type)

IDC Innovator Assessment

- Lightbits personnel led the development of the NVMe/TCP standard, which first released from the NVM Express organization in November 2018.
- Lightbits Labs' offers a software-defined storage solution that can run on industry-standard x86-based storage servers over standard Ethernet networks and requires no host-side footprint. It's products include the LightOS storage operating system (which includes an NVMe-based global flash translation layer), an optional LightField hardware acceleration card, and the SuperSSD storage appliance. Features implemented in the Global Flash Translation Layer (GFTL), including thin provisioning, erasure coding, RAID, compression, quality of service, and replication, do not require the LightField card, but they can run on it if customers want to offload the storage server CPUs for better performance.
- Lightbits Labs has announced an OEM agreement with Dell EMC, a partnership that gives the company significant credibility in the enterprise storage industry. With the March 2019 introduction of an appliance (the SuperSSD), the company now has offerings to target OEMs, cloud providers, and enterprises.

Key Differentiator

While the Lightbits team's involvement in leading NVMe/TCP specification development (and its extensive prior experience with switched fabrics associated with high-performance shared storage) is important, its focus on providing a 100% target-side solution (with no host software footprint) using commodity off-the-shelf server-based storage hardware is key.

Challenges

Lightbits Labs does have a small direct sales force that covers the United States, Europe, and China but has not yet made the full transition to more scalable go-to-market models that will help it increase awareness among target customers and drive revenue through indirect (i.e., reseller) channels.

Source: IDC, 2019

FIGURE 3

Excelero

Why Excelero Was Chosen as an IDC Innovator

Excelero offers a software-defined, distributed, block storage product (NVMesh) that runs on commodity server-based storage hardware and was the first NAFA vendor to base its offerings around a software-only product. Excelero was an early NAFA market entrant. It has garnered significant investment and customer support and remains the premiere example of a distributed NAFA design.

Company Name			
 Founded 2014	 Number of Employees 65	 Headquarters San Jose, California	 Number of Customers Unavailable
 Product Name NVMesh	 Founders Lior Gal, Yaniv Romem, Ofer Oshri, Omri Mann		 Geographic Mix (% of Revenue by Major Region) 80% United States and 20% in the rest of the world
 Profiled Product/Service NVMesh	 Funding Excelero has raised \$35 million across two funding rounds (with the most recent one being in August 2018)		

IDC Innovator Assessment

- Excelero's distributed design, which has a host-side software footprint, avoids the possibility that controllers in a shared storage array may become a performance bottleneck. The NVMesh software offers several modules, which provide significant configuration flexibility – MeshProtect supports RAID and erasure coding data protection options, MeshConnect supports four different NVMe-oF options (FC, RoCE, TCP, InfiniBand), and MeshInspect provides storage analytics.
- Excelero's software-only approach gave the firm significant flexibility in working with channel partners to create different storage appliance offerings as well (which could target the enterprise market). Today, customers can create storage platform solutions from some combination of the software-only product, a SmartNIC hardware acceleration card to offload host CPUs, and any of the NVMe-oF options.
- Of the vendors profiled here, Excelero has gotten the most market traction to date (in terms of installed base). Under Excelero's 100% indirect model, Excelero has sold through Lenovo, Supermicro, Boston, and Penguin Computing and vertical market partners like Pixit Media and CMA. Excelero is one of the more mature NAFA start-ups with venture capital investments from both Micron and Western Digital.

Key Differentiator

Excelero stands out for its software-only deployment model on any of thousands of configurations (which is less expensive and eliminates hardware vendor lock-in), distributed design, and support for a SmartNIC hardware acceleration card to offload server CPUs in the hosts (a particular advantage for NVMe/TCP).

Challenges

While Excelero's distributed design does support better scalability, particularly for the types of workloads in large webscale infrastructures, some customers prefer shared storage solutions with no host-side footprint.

Source: IDC, 2019

FIGURE 4

Pavilion Data

Why Pavilion Data Was Chosen as an IDC Innovator

Using commodity off-the-shelf components, Pavilion Data has created a unique hardware architecture that enables direct, any-to-any connectivity between up to 20 controllers and 72 NVMe solid state disks (SSDs) in a compact 4U system. It's approach is very attractive for both enterprises and large cloud providers that are looking for a simple, highly scalable shared NVMe solution that requires no host touch points.

Pavilion Data			
 Founded 2014	 Number of Employees 70	 Headquarters San Jose, California	 Number of Customers Unavailable
 Product Name NVMe-oF Storage Platform	 Founders V.R. Satish and Sundar Kanthadai	 Typical Deal Size Close to \$1 million	
 Profiled Product/Service NVMe-oF Storage Platform (using NVMe/TCP)	 Funding Pavilion Data has raised \$33 million across four funding rounds (with the most recent one being in May 2018)	 Geographic Mix (% of Revenue by Major Region) 65% North America and 35% Europe	

IDC Innovator Assessment

- Most of the NVMe system start-ups are building storage appliances using server-based storage, but Pavilion Data leverages a unique hardware architecture built primarily from commodity off-the-shelf hardware. It offers what is arguably the highest-performance NVMe-based array targeted for enterprise use today, supporting up to 20 controllers, which service I/O in parallel.
- Pavilion Data shipped its initial systems with an RDMA-based host connection (RoCE) that required custom drivers on the host side. While this approach provided a very low-latency shared storage solution, many customers do not want to have to buy and maintain custom configurations. The NVMe/TCP option that Pavilion Data now offers its customers works with widely available Ethernet drivers and adapters, is easier to maintain, and enables access to high-performance storage for any server.
- Those customers requiring the absolute lowest-latency storage access will consider RoCE, while those that care more about ease of purchase, installation, and maintenance will likely go with the NVMe/TCP option. Pavilion Data offers both options in a turnkey storage appliance deployment model.

Key Differentiator

Three features of Pavilion Data's NVMe-oF Storage Platform are called out by customers: it is standard based and requires no specialized host components (for NVMe /TCP), it is looking to the future (Ethernet) rather than the past (FC) for host connections, and it is the unique and highly parallel hardware architecture that delivers industry-leading performance in a simplified, compact, low-cost package.

Challenges

The lack of awareness of Pavilion Data's unique approach, which is generating significant interest among customers with large webscale infrastructures, is holding the vendor back more than anything else.

Source: IDC, 2019

TECHNOLOGY DEFINITION

NVMe is a new storage protocol, developed and optimized specifically for persistent solid state storage, that delivers lower latencies, significantly greater throughput and bandwidth, and hugely greater parallelism than SCSI (a storage protocol that was developed specifically for hard disk drives [HDDs] nearly 40 years ago). Persistent solid state storage has come into wide use in the enterprise for general-purpose workloads, and the SCSI storage protocol was leaving a lot of the performance, capacity, and endurance potential of this new media type on the table. NVMe unlocks the full potential of persistent solid state storage (and in fact doesn't even support HDDs).

NVMe-oF enables the use of NAFAs over switched fabrics, a feature that allows high-performance storage systems based on NVMe to be efficiently shared across multiple hosts while equaling the latencies of server-based storage. This effectively creates an NVMe-based storage area network (SAN). Today, NVMe-oF connections can be based on Fibre Channel (FC), Ethernet (RDMA over Converged Ethernet [RoCE] or Internet Wide Area RDMA Protocol [iWARP]), or InfiniBand. When those customers with preexisting FC SANs move to NVMe-oF, they will often want to use the FC transport for NVMe-oF. Most newer applications being built are being connected to servers using Ethernet, so those customers will often be interested in one of the Ethernet options, and those few enterprise customers that already have InfiniBand networks in place may want to use the InfiniBand option (although InfiniBand is generally more of a massively parallel computing market play).

NVMe/TCP also uses the Ethernet transport, but it offers some significant differences. RoCE and iWARP both require custom HBAs and custom host-side drivers, but they support remote direct memory access (RDMA) for the absolute lowest latencies. NVMe/TCP supports NVMe over a switched fabric without the RDMA protocol support, but it supports latencies that are almost as low. Most importantly, it does not require custom HBAs or custom host-side drivers. This makes NVMe/TCP much cheaper to buy and easier to deploy and maintain than the RDMA-based options, and it is the NVMe-oF option that IDC predicts will quickly come to dominate the NVMe-oF market. NVMe/TCP also supports much larger switched networks than the RDMA-based versions, a feature of particular interest to customers with large webscale infrastructures. It is the newest NVMe-oF protocol, however, and has just started to become available from vendors.

Persistent memory (PM) and storage-class memory (SCM) are two types of persistent storage devices built out of a new type of solid-state media that is faster than NAND flash but not as fast as dynamic random-access memory (DRAM). PM is accessed using a memory interface (DDR4/DDR5), whereas SCM is accessed over a block interface (like NVMe). The actual solid-state media in each of these device types may be the same (e.g., Intel Optane), but PM is faster because the memory interface is faster than the block interface. The NVMe block interface offers much better performance than the SCSI block interface though. For a more complete explanation of PM and SCM, see *IDC TechBrief: Persistent Memory and Storage-Class Memory* (IDC #US44891819, March 2019).

IDC INNOVATORS INCLUSION CRITERIA

An "IDC Innovators" document recognizes emerging vendors chosen by an IDC analyst because they offer an innovative new technology or a groundbreaking business model, or both, and were approved by the IDC Innovators Review Panel. It is not an exhaustive evaluation of all companies in a segment or a comparative ranking of the companies.

An IDC Innovators document highlights vendors that meet the following criteria:

- In IDC's opinion, the company exhibits innovative technology or a new business model.
- The company has annual revenue <\$100 million at the time of selection.
- Customers are currently using the company's products and services (i.e., the products and services are not conceptual or in the process of being released).
- The product, service, or business model must solve or help alleviate an IT buyer challenge.

In addition, vendors in the process of being acquired by a larger company may be included provided the acquisition is not finalized at the time of publication of the document. Vendors funded by venture capital firms may also be included even if the venture capital firm has a financial stake in the vendor's company.

LEARN MORE

Related Research

- *IDC TechBrief: Persistent Memory and Storage-Class Memory* (IDC #US44891819, March 2019)
- *Enterprise NVMe Growth and Use Cases* (IDC #DR2019_BS2_EB, March 2019)
- *Enterprise NVMe Market Developing Quickly with Offerings from Start-Ups and Established Players Alike* (IDC #US44852719, February 2019)
- *IDC TechBrief: NVMe Over Fabric* (IDC #US43854018, September 2018)

Synopsis

IDC Innovators are emerging vendors with revenue <\$100 million that have demonstrated either a groundbreaking business model or an innovative new technology – or both. This IDC Innovators study reviews three vendors, Lightbits Labs, Excelero, and Pavilion Data, calling out the differentiation that their approach (and their products) provides to potential customers. As digital transformation drives the need for higher storage performance, more customers are looking at NVMe and related technologies. When customers can deploy a shared storage solution that delivers the same latencies as local storage, there is a strong preference for it. This is because a shared storage environment offers easier scaling, higher capacity, better capacity utilization, and access to the types of enterprise-class data services typically only available in SAN-based solutions. For the most latency-sensitive customers, these shared storage solutions will be connected over an NVMe over Fabric host connection. NVMe/TCP, a variant of this that is just starting to become available, will ultimately dominate the NVMe over Fabric market because of several distinct advantages that are discussed in this document.

"Ease of management, simplified maintenance, and lower cost are the key advantages that will drive NVMe over TCP deployments in both enterprises and cloud service providers," said Eric Burgener, research vice president, Infrastructure Systems, Platforms, and Technologies. "This technology is just starting to become available now, but IDC expects to see very rapid adoption of it over the next two to three years."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright and Trademark Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights. IDC Innovator and IDC Innovators are trademarks of International Data Group, Inc.

Copyright 2019 IDC. Reproduction is forbidden unless authorized. All rights reserved.

