

Storage in the Era of NVMe

An Evaluation Guide for IT Teams

It's time to ride the storage express

How and where you use solid-state memory is the new differentiator for enterprise storage



As an IT professional dealing with storage, you will already have a good idea of how the demands of an application and its associated data will shape the type of storage required. For instance, cold data might fit well on cloud storage, an archive platform or tape. Similarly, disk arrays remain a strong option for secondary storage or even for primary applications where capacity and cost matter more than performance.

The game-changer though is solid-state storage such as Flash, thanks to its high performance and its non-mechanical nature. The latter also brings a range of other benefits including lower power consumption, smaller size and greater density, and of course it produces no noise or vibration. Expensive at first, it was initially added to disk arrays as a performance-enhancing layer to create hybrid arrays, an approach that is still cost-effective today in many scenarios. Technology evolution and cost reductions then brought us the enterprise-grade All-Flash Arrays (AFAs) that have become the “new normal” for high-performance primary storage.


Now the rapid maturation of the NVMe (Non Volatile Memory Express) storage standard means that everything is changing again. Fortunately though, if you're familiar with the concept of tiering from the disk world then it's similar – except that the type of disk involved no longer matters. Instead, the key performance differentiator is what type of memory is used and how it is implemented:

Hybrid Disk	This combination of disk storage with performance-boosting Flash (or even NVMe) still has its uses, for example where low cost and high capacity are paramount, where physical space is not at a premium, and where the slower read speed of disk can be tolerated or worked around.
All-Flash Arrays	The new mainstream option for general-purpose use, All-Flash Arrays are not just fast, but consistently fast, serving multiple applications simultaneously with relative ease. They have all the advanced storage management features enterprises have come to expect, plus capabilities that disk arrays cannot easily support such as in-line data reduction.
NVMe end-to-end	Its leaner software protocol and use of solid-state storage on the fast PCIe bus allows NVMe to deliver super-fast parallel data access direct to the processor (or shared via a PCIe switch). Although it is more expensive than general-purpose Flash, this makes it ideal for workloads needing the highest performance and simultaneous access to data.
NVMe over Fabric	Can extend the NVMe storage protocol over a range of low-latency links (a minimum of 25Gbps is preferable) such as InfiniBand, Fibre Channel and Ethernet. It both offloads storage processing from the server CPU for efficiency, and allows NVMe storage to be accessed remotely, giving systems fast access to very large volumes of data.

We will see more combinations and versions appear, especially as new classes of memory are developed and brought to market. As a very rough guide to performance, where a Flash SSD might be at least five times faster than a disk drive, an NVMe SSD could be five times faster than Flash. So where the original hybrid arrays added Flash SSDs to disk for improved performance, the next-generation adds NVMe SSDs to All-Flash or disk arrays. NVMe can also incorporate different media types. Most NVMe modules today use Flash media, but there are already examples based on faster and even more expensive Storage Class Memory (SCM), such as 3D XPoint from Intel and Micron, and Samsung's Z-NAND.

NVMe: 6 key things to look or plan for

Most storage in use today, even SSDs and the general purpose All-Flash Array, relies on at least some technology that was developed with spinning disks in mind. Under the covers, NVMe is very different and it represents a break with the past. From a high-level perspective it simply extends the Flash continuum upwards. After all, it is still based on solid-state storage - that is, Flash and its successors such as storage-class memory.



Look closer though, and it has the potential to change almost everything. For example, its support for massive parallelism and multiple threads provides better resource utilisation and potentially moves the bottleneck back from the storage to the server layer. Similarly, NVMe over Fabric allows larger data sets and long-distance clustering, while still supporting key features such as QoS management and defined data locations for compliance.



Performance

NVMe connects via the processor's fast PCIe bus, not a legacy serial connection designed for disk drives. This removes layers of latency-inducing software left over from the disk days and eliminates the serial bottleneck of interfaces such as SAS/SATA. This allows NVMe to provide very fast parallel data access.



Capacity and scalability

NVMe can be installed as a PCIe card or module for direct-attached storage, within an All-Flash or hybrid array to enhance its performance, or as scalable end-to-end all-NVMe storage arrays. These arrays may in turn be connected together via the emerging NVMe over Fabric (NVMe-oF) standard for networked storage.



Availability and resilience

Flash may well be non-volatile in normal use, but any storage device can fail so NVMe still needs backup and protection. This may include employing flexible RAID, or asynchronous replication to a second site via NVMe-oF. Look too for systems that provide fast hardware recovery via hot-swappable NVMe modules.



Intelligent data management

Data reduction with minimal impact on performance is essential to minimize NVMe capacity cost, as is a high degree of RAID level flexibility to support rapid data growth and use expensive capacity efficiently. Similarly, multi-node NVMe-oF systems must intelligently balance I/O operations across the nodes.



Skills and operation

We need to break with the past at some point, but preferably via a smooth migration, not by adding a new storage silo. Ultimately you will want to optimize the whole storage stack to fully exploit NVMe, but in the meantime, add NVMe as part of a wider storage infrastructure with common management skills and software.



Future-proofing

Obtaining visibility of your future business needs is essential. The chances are that there will be more digital workstreams, shared with more partners, or greater demand for parallel data access. Planning for NVMe is part of planning for uncertainty, because you don't want to be constrained by your storage infrastructure.

Understanding the array landscape

Nine key criteria that matter

Turning specifically to storage arrays, making effective investment decisions today means matching the essential business imperatives of growth, risk mitigation and cost reduction to all the available options, even as you plan to add the latest NVMe technology. The strengths and opportunity matrix below offers some guidance on what matters here (key: more stars is better).

		Disk	Hybrid Disk	All-Flash Arrays	NVMe Arrays
Characteristics	Practical scalability	★★	★★★	★★★★	★★★★★
	Minimise space/heat/power	★	★★	★★★★	★★★★★
	Low cost of capacity	★★★★★	★★★★	★★	★
Performance	Parallel access to data	★	★	★★	★★★★★
	High IOPs & low latency	★	★★	★★★★	★★★★★
	Low cost of performance	★	★★	★★★★★	★★★★★
Operational	Workload versatility	★	★★★★	★★	★
	Current adoption trend	★	★★★★	★★★★	★
	Low cost of operation	★	★★	★★★★★	★★★★★

For guidance purposes only

Locate your application’s practical and budgetary requirements on the matrix, then use our rankings to help guide your thoughts and plans. For instance, if high capacity at a low cost is key, then disk may still be the best option – unless it must also be versatile enough to handle for example workload consolidation. In the latter case, hybrid disk may be a better option. Similarly, if you require massive scalability and very good performance, with cost as no barrier, then NVMe-oF could be appropriate, while All-Flash Arrays are a strong contender for the best overall cost balance, covering capacity, performance and operation.

Note that we also expect to see intermediate ‘next generation’ hybrid arrays, where a relatively small proportion of NVMe is added to older disk or Flash-based array technology.

In conclusion

We hope that the above helps you crystalize your understanding of where All-Flash and NVMe fit in the ever-evolving continuum that is enterprise storage.

For more information on NVMe specifically, read our introductory guide ‘NVMe: a fresh start for storage’.

About Freeform Dynamics

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

For more information, please visit www.freeformdynamics.com.

About Fujitsu

Fujitsu is the leading Japanese information and communication technology (ICT) company offering a full range of technology products, solutions and services. Approximately 155,000 Fujitsu people support customers in more than 100 countries. We use our experience and the power of ICT to shape the future of society with our customers.

For more information, please see: www.fujitsu.com.

Terms of Use

This document is Copyright 2018 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process. Hosting of the entire report for download and/or mass distribution by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd or Fujitsu. The contents contained herein are provided for your general information and use only, and neither Freeform Dynamics Ltd nor any third party provide any warranty or guarantee as to its suitability for any particular purpose.