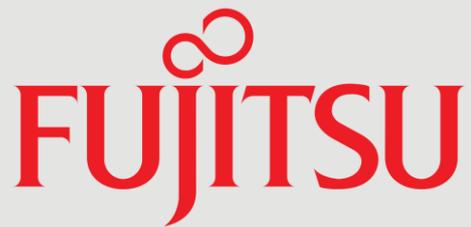




Inside Track Research Note

In association with



## NVMe: a fresh start for storage

It's time for a real change, not  
just iterative AFA evolution

Freeform Dynamics, 2018

---

#### About this Inside Track

The insights presented in this document are derived from independent research conducted by Freeform Dynamics. Inputs into this include in-depth discussions with IT vendors and service providers on the latest technology developments, along with intelligence gathered from mainstream enterprises during broader market studies.

---

*As business demands ramp up the expectations you have of storage are changing.*

---

*NVMe is an open collection of standards and information to fully expose the benefits of non-volatile memory in all types of computing environments.*

---

## In a nutshell

Until a few years ago acquiring storage was a relatively simple matter. Storage platforms were based on hard disk drives, and the only major questions were what speed drives to use and how to connect them together. But then enterprise flash arrived on the scene and suddenly it became possible to greatly increase storage performance.

However, to avoid having to modify other elements of the platform, the majority of flash storage simply emulated established hard disk protocols, taking the form of solid-state drives (SSDs). This encouraged adoption but limited the functionality and performance of all-flash arrays.

This may sound just like standard technological evolution and simply a background issue, but as business demands ramp up, the expectations that we as users have of storage are changing. For example, many new workloads springing up from the digital transformation of business processes require parallel access to data. Some of these are pushing existing storage, even all-flash array (AFA) technology, to its limits.

To tackle this, the industry has been working on new technologies to really open up the power of flash storage. This paper will take a high-level look at the development of a key one of these: non-volatile memory express storage, usually abbreviated to NVMe. We will consider how it differs from traditional storage and give some idea of how it is being adopted.

If your storage pros or IT team are forward-thinking, they will soon be telling you they want to acquire NVMe storage – if they haven't already done so – and we will give you the information you need to think about when looking at their plans. If, on the other hand, they are more conservative, you may want to ask them when they plan to consider NVMe.

## What is NVMe?

In the words of the industry association steering the development of NVMe storage, “NVM Express® is an open collection of standards and information to fully expose the benefits of non-volatile memory in all types of computing environments from mobile to data centre.”

From an architectural point of view, the key element here is that NVMe is not attempting to fix the problems of storage designs that were laid down decades ago. Instead it is designed from the ground up as a new technology, able to deliver high bandwidth and low latency storage access for current and future non-volatile memory.

This last point is extremely important. The improvement in the performance of hard disk drive systems has been tailing off and much of the storage industry is now switching attention, and budgets, to the development of flash-based alternatives. In the short term, the primary route for flash adoption will continue to be SSDs with existing connectivity technologies, with NVMe solutions only employed with SSDs and new storage media for extreme performance workloads. That said, it is only a matter of time before NVMe starts to be adopted more widely, making this a good time to start thinking about when and where you want to consider exploiting the advantages it brings.

## NVMe and traditional storage

### The problem today

*The storage schemes commonly used today were designed to handle spinning hard disks.*

The storage schemes commonly used today were designed to handle spinning hard disks, and use SAS or SATA interfaces. These are essentially serial technologies where data is sent or read one bit at a time, and they can only talk to the CPU through many layers of software.

In more recent times a physical interface called PCIe, which links directly to the CPU and was originally for graphics cards, has also been used to help improve the performance of storage. But PCIe is not a native storage interface, so while its connection into the CPU gives much improved storage bandwidth and latency compared to the existing SAS or SATA interfaces each PCIe storage device needed its own drivers, as there was no standardisation.

### Why build NVMe from the ground up?

To optimise performance, improve resilience and ultimately reduce costs and risk, a whole new approach was needed, hence the development of NVMe. To minimise disruption and speed up potential adoption, NVMe retains the established PCIe hardware layer but discards the old software in favour of a new native storage interface. Rather than attempt to speed up SAS and SATA by using faster media and connections (which is more or less what early SSD / AFA systems did), NVMe therefore replaces them entirely with a standard driver for all storage devices, removing the need for legacy software and protocols.

*NVMe has been developed specifically with solid-state/flash media in mind.*

To get technical for a moment, NVMe has been developed specifically with solid-state/flash media in mind. It has a much smaller and more efficient command set and supports far deeper and more numerous queues, thereby making much better use of modern multi-core processors. Modern CPUs can process commands in parallel, and of course solid-state media is not limited to serial access, unlike hard disks with their mechanical read/write arms. The result is that, compared to current technology, NVMe makes it cheaper to build very, very fast storage systems for a broader range of workloads.

### Does NVMe really make a difference?

As mentioned earlier, as a senior manager or executive you may think that these technical matters are not something with which you need to be concerned. Five or ten years ago, you would have been right. Today that is certainly unlikely to be the case if you think about just how important, and how visible, access to data and information has become to almost every business process.

*As a senior executive you may think technical matters are not something to concern you. Today that is unlikely to be the case.*

### Digital transformation, NVMe and parallel data access

The number of workloads deemed to be performance sensitive is rising. Good, or even very good performance, is now essential to many more applications, in all areas of the organisation, where in the past only a few, mission and business critical services could afford to have low latency and fast response times.

---

*Many new workloads associated with digital transformation create a need for simultaneous, parallel access to data.*

---

While the development of faster, cheaper servers helped address some application performance issues, they often then pushed the bottleneck to the storage systems, especially those where parallel access is key. Performance here was improved with the creation of devices combining flash and disk, and then of all-flash systems, but, as we discussed, even all-flash systems were limited by their legacy heritage. As a very rough guide, where a traditional PCIe SSD might be three to four times faster than a SATA SSD, an NVMe SSD is expected to be at least that much faster again.

However, challenges still exist, especially with new workloads associated with the digital transformation of business, many of which can create a need for simultaneous, parallel access to data. Another factor at play is the growing requirement for multiple applications to be able to access storage simultaneously without degrading performance.

To tackle these technology constraints, almost the entire storage industry has been working together on the NVMe standards, to improve storage performance, security, availability and costs. These standards are now rapidly maturing, and NVMe will increasingly be deployed in mainstream use. In particular, NVMe has been designed with parallelism in mind as it supports far deeper queues (where storage commands wait to be executed) and more of them – up to 64,000 queues, with up to 64,000 commands per queue. The comparison with traditional storage stacks, where AHCI allowed only a single queue of up to 32 commands is clear, especially on modern multi-core/multi-threaded processors.

---

*Almost the entire storage industry has been working together on the NVMe standards.*

---

## How is adoption shaping up?

Every standard takes time to evolve. And when talking about a package of standards designed for a major section of the IT industry, as we are with NVMe, it is no surprise when the first solutions to become economically effective for vendors to build and customers to acquire are niche ones to meet very tactical needs.

While some advanced NVMe solutions have already been adopted by leading edge organisations looking to meet very specific objectives, the main shape of adoption has roughly taken the following path.

### First NVMe came to servers

This involved fitting flash cards or flash drives directly to the memory bus of a server via PCIe, usually to act as a fast cache layer to improve performance of latency sensitive applications.

Before NVMe, this relied on server manufacturers creating cards specific to their servers and using bespoke driver software. Such systems were, therefore, complex and costly to implement and operate, and were used only where the business benefits of the faster service speed were very obvious. This approach also encompassed clustered server systems to be able to handle larger workloads.

As the standards have matured, NVMe has become widely accepted in server systems, displacing hard disks for a growing range of workloads. NVMe storage is also now becoming available in high end workstations and laptops. However, it is noteworthy that in all these cases, NVMe is used as local storage, for example to replace a directly-attached hard disk or SSD.

---

*First use of NVMe was in servers as a local cache, then in storage arrays.*

---

---

*NVMe drives are now widely deployed in all-flash arrays, but mostly relying on traditional protocols to handle the data flows.*

---

## Next came “NVMe / traditional AFA hybrid solutions ”

SSD drives using NVMe are also now widely deployed in all-flash arrays, but again, mostly relying on traditional protocols to handle the data flows between servers (or virtual machines) and the storage array. This approach delivers considerably lower storage latency than hard disk-based arrays, and provides excellent backwards compatibility with existing IT systems.

However, it relies on server software stacks that are not optimised for flash storage. The storage protocols consume more CPU cycles than necessary, and there are too many intermediate communication layers that were needed in the past but are now redundant.

## Bespoke all-NVMe architecture

As discussed earlier, the greatest performance benefits come with systems that are built from the ground up to use the latest NVMe protocols from end to end – from server to storage. While fully standards-based solutions are only just becoming commercially available, it is expected that their take up will increase quite rapidly.

It is clear that true enterprise-class NVMe storage systems, built from the ground up, will soon be ready for commercial use. As well as exceptional performance, such systems will have all of the capabilities and storage services that applications require.

These include high resilience and the ability to withstand component failure without restricting access to data or seriously degrading performance. They will also provide the data integrity, data encryption and system level security capabilities that business executives know they must have for some applications.

---

*The greatest performance benefits come with systems that are built from the ground up to use the latest NVMe protocols from end to end.*

---

## NVMe over Fabric

There is another extension of the NVMe standards that is getting ready for commercial deployment. This is NVMe over Fabric, usually abbreviated to either NVMe-oF or NVMe-f.

NVMe-oF makes it possible to extend the NVMe storage protocol over a range of networking connections, called fabrics. These include both specialist storage networking protocols such as Fibre Channel and InfiniBand or standard networking such as Ethernet, but crucially still with low latency.

What this means is that NVMe-oF allows storage systems to scale out to large numbers of NVMe devices, making it possible to have very fast performance even for systems that need to access very large amounts of data. Crucially NVMe-oF allows the storage to be sited not just locally within the data centre but also to be held in other data centres some distance away.

---

*Crucially, NVMe-oF allows the storage to be sited not just locally within the data centre but also to be held in other data centres some distance away.*

---

## The potential benefits of NVMe

This paper has mentioned many of the potential benefits available to business systems. The table below gives a quick way of understanding how technology capabilities can be understood in terms of business benefit (Table 1).

**Table 1**  
Business benefits associated with NVMe technology

NVMe Technology	Business Benefit
Greater capability to handle requests in parallel	Make full use of processor technology advances – minimise hardware footprint and latency
Massive parallelism and multiple threads	Faster performance, better hardware resource utilisation
Maximise performance and throughput	Simpler command set
Multiple namespace support	Simpler management, ability to mix multiple workloads easily
QoS Management / Data location	Service performance consistency, legal compliance
NVMe over Fabric	Support for large data sets and distance

*The NVMe ecosystem is growing quickly.*

*Storage and system vendors are committed to NVMe. It is likely to become the storage standard for the next decade.*

*NVMe is designed to optimise performance as storage technology develops, without having to rewrite applications.*

## The bottom line

The NVMe ecosystem is growing quickly as vendors see it as offering many potential benefits going forwards. There is also a growing recognition of those benefits amongst IT professionals.

NVMe has already become part of servers running workloads that demand the best performance. It has also begun to move into mainstream storage systems, although mostly using a mix of new NVMe technology and existing protocols. It is also likely that, as the surrounding ecosystem is established in the coming months, we will start to see the first pure-NVMe systems, designed from the ground up to exploit all its potential performance capabilities.

Looking slightly further ahead NVMe-oF will be ready for enterprise use where large data sets are required. From a business point of view, better performance, resilience, scale and security will soon be available to the applications you run.

## Five things to think about NVMe

If you only take away a few things from this paper, please remember these

- Storage and system vendors are heavily committed to NVMe, and it is likely to become the storage standard for the next decade.
- With access to data being fundamental to almost every business process, you want to invest in storage standards that will be around for a long time, and to which the storage community is committed.
- Business performance is essential, and NVMe was designed to optimise performance using the best technology available today, and as it develops, without having to rewrite applications.
- Many applications, especially those associated with digital transformation, require parallel but fast access to data with minimal latency as do scenarios where multiple applications need to access storage simultaneously without

---

*As NVMe usage increases, costs will fall.*

---

degrading performance. NVMe was constructed to be able to handle these workloads.

- NVMe is designed to work with all device types, from mission critical data centre and cloud storage platforms to the storage in your tablet and smartphone. Volume will ensure acquisition costs fall.

This paper is not intended to give you all the technical details of NVMe, but to position how it will fit into the expanding storage landscape in the coming years. We hope this paper has helped you assess where and how NVMe might fit into your plans and activities.

## 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, and access to our library of free research, please visit [www.freeformdynamics.com](http://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.

Fujitsu's storage portfolio balances storage capacity, performance and costs for the complete lifecycle of data - from production, business analytics and big data to backup and long-term archiving.

For more information, please visit <http://www.fujitsu.com/global/products/computing/storage/all-flash-arrays/index.html>

## 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 document 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.