



Technical White Paper

New Frontiers in Solid-State Storage

By Rodney DeKoning, VP and Chief Architect,
and Ravi Kavuri, VP and Chief Architect
April 2017 | WP-7248

New innovations in storage technology are changing the game for data centers, and NetApp is working to help our customers capitalize on them. The introduction of faster and faster media types and more efficient mechanisms to access those media across well-defined SAN and Ethernet infrastructures will unlock unprecedented speeds, lower latencies, and dramatic improvements in system and application efficiency.

Table of Contents

What You Will Learn	3
Delivering on the Promise of NVMe, NVMe over Fabrics, and Storage-Class Memory	3
Implications of Emerging Storage Technologies.....	4
Inside NVMe	5
NVMe over Fabrics	5
NetApp Is Leading the Way in Storage Innovation.....	6
The NetApp Advantage	7
NetApp: The Right Partner for Storage Innovation.....	8
Learn More.....	8

What You Will Learn

New innovations in storage technology are changing the game for data centers, and NetApp is working to help our customers capitalize on them. The introduction of faster and faster media types and more efficient mechanisms to access those media across well-defined SAN and Ethernet infrastructures will unlock unprecedented speeds, lower latencies, and dramatic improvements in system and application efficiency. These benefits are based on three advances: Nonvolatile Memory Express (NVMe), NVMe over Fabrics (NVMe-oF), and new storage-class memory (SCM), also known as PMEM, such as Intel 3D Xpoint and Samsung Z-NAND media.

NetApp is helping our customers navigate these fast-coming changes by introducing new storage media and interconnecting technologies in a way that maximizes existing investments. In the near future, look for new NetApp® All Flash FAS and E-Series products that use NVMe-oF to drive down CPU cycles per I/O in NetApp all-flash arrays. We will also be enhancing our flagship All Flash FAS (AFF) platform with NVMe connectivity to leverage the highest-performance solid-state drive (SSD) devices, including those using innovative SCM media. NetApp customers will gain immediate value from SCM SSDs as a read/write cache to achieve the lowest possible latencies. Over time, as new SCM media become more affordable, NetApp will make it possible for customers to create pools of high-performance SSDs using SCM that are tightly integrated with application requirements and deliver order-of-magnitude latency improvements.

With these and other NetApp innovations, our customers will be able to:

- Use NVMe-oF to drive lower CPU overhead in the I/O stack and more efficient protocol interconnectivity to external storage arrays, which translates to significant improvements in the cost and performance of data center infrastructures
- Use NVMe-oF for the device connect as well, allowing for scalable connectivity of storage OS compute to SSDs and ultimately enabling more flexible and powerful cluster architectures
- Use NVMe-attached SCM initially as Flash Cache™ intelligent caching, and ultimately as pooled storage, to accelerate I/O and provide the lowest possible latencies for demanding real-time applications

NetApp is delivering all of these innovations in a way that lets our customers continue to use the full suite of NetApp data management capabilities on which they rely in enterprise-grade storage arrays and protects and extends their existing investments. Read on to learn more.

Delivering on the Promise of NVMe, NVMe over Fabrics, and Storage-Class Memory

Much has changed in the world of storage in the past decade. Traditional spinning hard disk drives (HDDs) have given way to incredibly fast (and increasingly affordable) SSDs. Now, we're gearing up for the next big evolution in memory with NVMe, and especially NVMe-oF. NVMe is a new, faster memory interconnect protocol that's being manufactured in a number of interfaces and form factors. NVMe is already hitting the consumer market, providing faster direct CPU access to SSDs in laptops. The recently formalized NVMe-oF specification will extend these efficiencies to large-scale storage arrays. Now, organizations are wondering what NVMe and NVMe-oF will mean for data centers. Many organizations saw a huge boost in performance for postprocessing large files when they made the move from HDDs to SSDs. Now that flash memory is more common and nonvolatile memory and SCM technologies are being integrated into enterprise storage systems, should organizations expect similar benefits? Yes, and much more.

The combination of NVMe-oF and SCM will soon create a tipping point in storage. Emerging data center fabrics that support exponentially more I/O operations per second (IOPS) and device latencies measured in microseconds (μ s) will allow organizations to do many of the things they're doing today (business intelligence, data warehousing, and data analytics) much more quickly, with much less CPU overhead. And taking full advantage of the massive IOPS that new architectures are capable of delivering makes many new applications—advanced real-time analytics and decision support, more efficient machine learning, and faster processing of very large datasets—become possible.

NetApp has been participating at the highest levels of the NVM Express Organization and has been actively involved in the development of the NVMe standard and the NVMe-oF specification. We were early pioneers in nonvolatile memory interconnect technologies and in fact have already shipped over 3 petabytes (PB) of NVMe flash. We're currently building NVMe-oF capabilities into our All Flash FAS and E-Series platforms that will allow organizations to unlock the full capabilities of their networking infrastructure and emerging SCM media in their data centers. We are also adding NVMe connectivity for attached SSD devices to our flagship AFF platform.

Much work remains to be done across the industry to tap the full potential of NVMe, NVMe-oF, and SCM, and organizations should expect these innovations to emerge incrementally in the coming years. NetApp is working closely with industry partners to make sure that when they do, our customers will be able to take maximum advantage of them. We are prioritizing NVMe-oF, because it will provide immediate performance and

efficiency benefits in real-world storage environments, even as the industry at large continues the broader work needed to unlock these technologies' full transformative potential. At the same time, we are working to implement NVMe and SCM technologies across our portfolio in a way that allows customers to gain maximum advantage from them in the short term, while making sure of long-term flexibility and investment protection. Most importantly, we're working to make sure our customers can realize the benefits of these innovations in their existing storage environments nondisruptively, without having to sacrifice all of the features and rich storage and data management capabilities that caused them to choose NetApp in the first place.

Implications of Emerging Storage Technologies


The new class of superfast memory technologies is making it possible to envision system latencies lower than anything possible in the past. Flash and NAND flash were the first iterations, driving down device latencies from the approximately 8 milliseconds associated with HDDs to 120µs to 150µs. Now, new SCM technologies such as 3D XPoint and Z-NAND SSDs are upending the status quo again. These memory innovations will offer sub-10µs device latencies along with tremendous improvement in IOPS when connected over new NVMe and NVMe-oF technologies: 10 times faster than today's fastest NAND-based SSDs.

What will these order-of-magnitude improvements in bandwidth and latency mean for organizations? Traditional data center applications (high-volume transactional

databases, CRM, and analytics) will run much more quickly and more efficiently, as organizations extend speeds and latencies previously available only with local server based memory technologies to shared network storage at scale. Organizations will be able to support the same or better application performance with fewer CPU cycles, translating to immediate, concrete capital and operational expense savings in large-scale data centers. Emerging applications will benefit even more. Organizations will be able to apply machine learning with real-time training and real-time inference to much larger datasets. They'll be able to analyze massive amounts of information coming in from Internet of Things (IoT) sensors. A wide range of problems that could previously be tackled only by high-performance computing environments will now be able to live inside the data center.

All of these advances will be powered by new nonvolatile memory and fabric technologies. Vendors across the industry are now working to implement them in many different places in system and fabric infrastructures (Figure 1). However, to unlock their true potential, vendors can't focus on just one part of the Non-Volatile Memory (NVM) story. They will need to be able to execute across the entire vision: SCM, NVMe, and NVMe-oF. They will need to deliver these innovations in a way that customers can take advantage of them incrementally so that each new element drives immediate, tangible gains. And they will need to empower customers to capitalize on NVM and SCM innovations cost-effectively, nondisruptively, and with the full suite of data management features required for enterprise-grade systems.

Figure 1) Implementing new NVMe and NVMe-oF architectures.



NVMe replaces the Small Computer System Interface (SCSI) standard on which systems have long relied to facilitate communication between CPUs and local or remotely attached disk drives. SCSI emerged in the early days of computing and has proven extremely durable.

Inside NVMe

NVMe replaces the Small Computer System Interface (SCSI) standard on which systems have long relied to facilitate communication between CPUs and local or remotely attached disk drives. SCSI emerged in the early days of computing and has proven extremely durable.

However, the SCSI stack was designed for HDDs, with many elements still focused on the requirements of rotational media accessed using drive heads that can service only one request at a time. As the industry has moved to flash SSDs, which are faster and have fewer management requirements, many elements of the SCSI stack are no longer necessary. Indeed, the ability to realize the full speed and capacity of solid-state storage is now bumping up against the limitations of the SCSI protocol and SCSI software stacks. Modern SAS interfaces are limited to 1.1Gbps to 1.2Gbps, while SSDs can deliver much more.

NVMe eliminates this bottleneck by introducing a more streamlined stack, with a new set of commands and a massive number of I/O queues for direct-connected storage. As a result, NVMe devices can achieve significant improvements over SAS in latency, IOPS, and bandwidth, enabling up to 4000MBps through the PCI bus. NVMe can provide better throughput than SATA, lower compute overhead, and substantially reduced latency. When implemented with a carefully crafted software stack, NVMe can enable 30% to 40% lower I/O latencies at reasonable concurrencies.

It's important to recognize that NVMe is just a starting point. Reducing access latency to solid-state media on a PCIe-attached fabric does yield benefits, but they don't translate to huge gains for data centers using existing flash media. In fact, highly tuned existing SAS solutions, such as the NetApp E-Series, can already deliver efficiency comparable to that provided by NVMe. The moment organizations begin considering new SCM technologies, however, NVMe becomes mandatory. In a conventional system, SAS interconnects add 20 μ s to 40 μ s. With access latencies for current NAND storage around 70 μ s, that additional overhead is largely unnoticeable. However, when moving to SCM storage, with access latencies of just a few microseconds (or less), adding 40 μ s of overhead for the interconnect protocol becomes a huge bottleneck. So as organizations look to implement new applications that require the ultralow latencies afforded by SCM, faster device-to-memory interfaces such as NVMe become the only option for achieving them.

NVMe over Fabrics

NVMe can offer real benefits for direct-attached storage, and vendors are already exploiting it. Laptops with NVMe SSDs are now available, and even in the data center market, NetApp is already shipping FAS controllers that use NVMe for Flash Cache intelligent caching. The major implications for data centers, however, will arise only when NVMe bandwidth and latency improvements can be extended to larger-scale storage arrays "outside the box." The recently formalized NVMe-oF specification aims to make this approach a reality.

NVMe-oF defines a common architecture to extend the NVMe block storage protocol over a range of storage networking fabrics (Ethernet, Fibre Channel, and Infiniband) using low-latency RDMA interfaces. Whereas NVMe back-end system improvements will exist largely "under the hood," NVMe-oF enables a host-side interface into storage systems and the ability to scale out to large numbers of NVMe devices within a data center, even over distances.

NVMe-oF can allow systems to access a remote storage device with less than 10 μ s additional latency compared to locally attached storage. NVMe-oF topologies will allow systems to share SSDs with more than two storage controllers. As a result, organizations will be able to apply more CPUs to pools of SSDs and take much better advantage of the very high performance capabilities that flash and emerging NVM storage arrays can enable. When NVMe-oF is implemented with 100 Gigabit Ethernet (GbE), for example, the entire storage fabric will add only a few microseconds to the system. Whereas previously such latencies were achievable only on limited "islands" within a server, they will soon be applied to shared system storage at scale. Ultimately, organizations will be able to achieve near DRAM speeds not just in a single node, but across a fabric, and solve problems that couldn't be approached before.

A number of applications can potentially benefit from these ultralow latencies. Emerging applications in big data, real-time analytics, and machine learning/deep learning will need to handle much larger datasets with higher throughput and lower latencies. To do it, they'll need multiple compute servers and fast access to persistent data. NVMe-oF and SCM storage make this approach possible. With faster transport, organizations can bring large amounts of persistent data "closer" to their applications, while retaining the rich data management capabilities on which they rely from enterprise-grade storage fabrics.

NVMe-oF will also enable new use cases. With NVMe-oF, it will be possible to access storage in remote storage servers or direct-attached NVMe storage without requiring remote CPU cycles or remote CPU memory when accessing NVMe SSDs. This accessing will be done with a latency increase of just 3µs to 7µs compared to local PCIe-attached NVMe devices. Often referred to as “NVMe direct,” this model will allow organizations to efficiently utilize software-defined storage and hyperconverged infrastructure (HCI) architectures utilizing pooled storage in servers or in top-of-rack direct-attached storage systems.

Of course, these new NVMe-oF architectures are very different than current storage systems and can present challenges if they are not implemented effectively. Latencies of just a few microseconds are indeed possible, but if they are based on the I/O controller communicating directly with the NVMe device—bypassing the CPU—customers will stand to lose many of the essential data management capabilities on which they rely. That’s one of the reasons that NetApp is implementing NVMe-oF in phases. Customers using our A700, A700s, and E-Series platforms will gain the immediate performance efficiencies and cost savings of reduced CPU cycles per I/O. In the coming years, when customers are ready to begin taking advantage of more advanced NVMe-oF host connectivity and remote access use cases, NetApp will make sure that they can do so while retaining the full features and rich data management capabilities that their applications require.

NetApp Is Leading the Way in Storage Innovation

As work progresses in emerging persistent memory, NVMe, and NVMe-oF technologies, we can begin to see how these trends will open up new possibilities for data centers. The NVMe storage protocol enables more bandwidth and much lower latencies over physical PCIe fabric. NVMe-oF, along with faster and less expensive Ethernet fabrics such as 100GbE, extends these benefits to large-scale storage arrays. Organizations will have the ability to analyze and process much larger datasets more quickly and more efficiently than ever before.

Of course, organizations won’t overhaul their entire storage architectures overnight. In the near future, they will want to add SCM for targeted applications: for example, putting the hottest metadata on SCM storage, while the rest goes to flash SSDs. The only practical way to use this approach is to use a data fabric that can support tiers of SCM storage alongside other kinds of storage, with the data management features to move data back and forth as needed. That’s exactly what the NetApp Data Fabric delivers.

NetApp is investing significant resources into developing technologies and architectures that capitalize on NVMe, NVMe-oF, and SCM. NetApp serves on the board of the NVM Express Organization and played a key role in the development

of the NVMe standard and the NVMe-oF specification. In fact, working with Samsung, NetApp has added new multistream write features to the NVMe standard that enable higher-performance communication with SSDs and reduced write amplification, which extends the life of NAND storage. (Previously, NetApp worked with Samsung to add these features, developed by NetApp, to the T10 and T13 standards for SCSI, and implemented them into our products. Through this joint work, multistream write capabilities have now been adopted for NVMe as well.)

Today, NetApp is already shipping NVMe technology as Flash Cache intelligent caching in existing products. NetApp ONTAP® is also ready to capitalize on NVMe and NVMe-oF, with significant work done to optimize our software for latency. That includes addressing “tail latencies” as well as average latencies, so that our systems can meet latency requirements measured in percentiles, with the reduced variability needed to support more rigorous service-level objectives (SLOs). These software changes not only will support NVMe-based SSDs, but also will allow our customers to take advantage of low-latency SCM media technologies, along with support for 100GbE and NVMe-oF.

Work is under way at NetApp right now in:

- **NVMe-oF.** New NetApp All Flash FAS and E-Series solutions are ready now for NVMe-oF. In the future, all NetApp AFA platforms will support NVMe-oF host connectivity.
- **NVMe.** NVMe might be a new protocol, but NetApp has been shipping solutions with direct CPU-to-flash storage access over a PCIe bus since 2009. NVMe merely replaces different vendors’ proprietary approaches with an industry-wide standard. The latest NetApp FAS systems replaced our proprietary Flash Cache protocol with standard NVMe. Today, we’ve already shipped over 3PB of NVMe-enabled flash media to customers. We are now working to extend NVMe storage across our portfolio. Our flagship AFF platform will provide NVMe connectivity to the highest-performance SSD devices. We are also incorporating dual-ported NVMe drives in our All Flash FAS and E-Series solutions and single-ported NVMe drives in SolidFire® platforms.
- **SCM.** New technologies with sub-10µs latencies, such as Optane SSD and Z-SSD, will begin reaching the market in 2017. NetApp is working to develop faster transport stacks and drive down system-wide latencies to allow our customers to take full advantage of them. Initially, SCM will be used for caching purposes, much like NAND flash was in the past for HDD-based arrays. In the coming years, our customers will be able to use the scale-out capabilities of ONTAP and SolidFire architectures to integrate tiers of ultralow-latency storage into their existing environments and data fabrics.

Ultimately, NetApp customers will be able to capitalize on NVMe, NVMe-oF, and SCM as part of a cloud-connected Data Fabric. They'll retain total control over their data and workloads, with the flexibility to move data wherever they need it across hybrid cloud resources.

The NetApp Advantage

Different NetApp solutions will benefit from NVMe, NVMe-oF, and SCM in different ways, but all will share some common underlying principles:

- **You won't have to sacrifice enterprise-class features and resiliency to use new storage technologies.** There are many storage vendors from which to choose today, most of which use similar underlying drive technologies. When our customers choose NetApp, it's not because we're the only ones who can offer high-performance storage arrays, but because we implement them within innovative architectures with industry-leading feature sets and one of the industry's broadest portfolios. All storage vendors will look to support NVMe drives, SCM, and NVMe-oF technologies in the coming years. However, when customers choose NetApp, they won't have to sacrifice the core value they derive from our solutions—scale-out architectures, mature and stable software, proven resiliency and redundancy, flexible data management, simplified application integrations, and more—to take advantage of them.
- **You will be able to add new storage technologies nondisruptively to your existing environment.** NetApp takes a “software-first” approach to storage innovation, as we've shown time and again with our ONTAP products (moving from SCSI to fiber-attached disks to SAS-attached disks), as well as with SolidFire. There's a good reason for this fact: underlying hardware technologies can change quickly. If you design a storage architecture from the ground up for a specific new technology, you've effectively created custom hardware, which means it will be much more difficult to incorporate new capabilities as the underlying hardware evolves (in some cases, within a matter of months). That's why NetApp follows a software-centric design path across our portfolio. When new technologies such as NVMe-oF and SCM emerge, you can continually integrate faster, denser, more cost-effective solutions into your existing architecture and operations, without making wholesale changes.
- **You will be able to easily integrate new NVM technologies through a scale-out architecture.** NetApp has long differentiated itself in the market with scale-out architectures. It's

this unique scale-out capability that allows our customers to add heterogeneous nodes to their existing storage clusters and easily balance capacity and performance nondisruptively in a fine-grained way. We previously demonstrated this capability when we introduced our all-flash storage arrays, allowing customers to seamlessly integrate flash into their existing infrastructures and operations. As organizations adopt NVMe drives, NVMe-oF, and SCM, scale-out capabilities become even more important. Indeed, it's the only way to add these new technologies to an existing cluster in a nondisruptive way. NetApp scale-out capabilities across our All Flash FAS and SolidFire solutions support both mixed controller types and mixed drive types and sizes, allowing our customers to add different hardware architectures to the same cluster. We also make it easy to cost-effectively upgrade older controllers to newer architectures after three years through the NetApp FlashAdvantage and Flash Forward Program.

Ultimately, NetApp customers will be able to capitalize on NVMe, NVMe-oF, and SCM as part of a cloud-connected Data Fabric. They'll retain total control over their data and workloads, with the flexibility to move data wherever they need it across hybrid cloud resources. In practice, this fact means that customers will be able to use tiering to move data back and forth across media, including NVMe-attached SCM and NAND storage, capacity flash arrays, HDDs, and all the way out to the cloud. They can choose the right media, with the right mix of price and performance for each workload, with the ability to seamlessly move it to a faster tier—or a slower, less expensive one, or even to the cloud—whenever they choose.

NetApp solutions can provide extremely dense, high-performance, cost-effective on-premises storage environments that can now be augmented with near-cloud and in-cloud capabilities as needed. Our Fabric Pools solution, for example, will allow customers to use NVMe drives and NVMe-oF architectures for high performance on the premises, while seamlessly tiering inactive data to the cloud. As new innovations expand the range of applications that can benefit from faster, lower-latency SSDs, this flexibility will allow our customers to more easily take advantage of them.

We offer some of the fastest flash storage platforms available, including new systems that were recently ranked among the fastest in the industry in independent benchmark testing.

NetApp: The Right Partner for Storage Innovation

NetApp has a long, storied history of leadership in storage technology. More recently, we've introduced a series of groundbreaking innovations in both software and hardware architectures, exemplified by our SolidFire, E-Series, and All Flash FAS solutions. In many ways, the capabilities of flash storage have spurred NetApp to demonstrate to the market and our customers what innovation is all about.

We offer some of the fastest flash storage platforms available, including new systems that were recently ranked among the fastest in the industry in independent benchmark testing. We are releasing new features and capabilities for our products multiple times per year. And, because of our focus on software innovation and scale-out architectures, we're helping our customers continually add new capabilities into their existing environments, extending the value of their NetApp investments.

NetApp is bringing this same approach to new NVMe, NVMe-oF, and SCM innovations. As new advances allow organizations to get more performance, density, and cost efficiency from solid-state storage, NetApp customers will be able to take advantage of them easily, nondisruptively, and with the full set of enterprise-class features they expect.

Learn More

Come talk to NetApp about the latest flash innovations and where we're taking the market for enterprise-class flash storage. Be data ready with the world's fastest enterprise all-flash storage. [Find out more.](#)

About NetApp

Leading organizations worldwide count on NetApp for software, systems and services to manage and store their data.

Customers value our teamwork, expertise and passion for helping them succeed now and into the future.

The information is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. NetApp makes no warranties, expressed or implied, on future functionality and timeline. The development, release, and timing of any features or functionality described for NetApp's products remains at the sole discretion of NetApp. NetApp's strategy and possible future developments, products and or platforms directions and functionality are all subject to change without notice. NetApp has no obligation to pursue any course of business outlined in this document or any related presentation, or to develop or release any functionality mentioned therein.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.