

## White Paper

# Enterprise Storage Evolves Beyond All Flash with Storage-Class Memory, Unleashing a New Generation of Innovation

Sponsored by: HPE

Eric Burgener  
November 2018

### IDC OPINION

---

Many enterprises today are undergoing a digital transformation (DX) as they move to business models that collect and use data as a critical strategic asset to drive better insights and improved agility. New analytics-oriented workloads that leverage artificial intelligence (AI) and machine learning (ML) are helping fuel these changes as businesses move to a real-time orientation for data collection, analysis, and response. The need for higher server, networking and, in particular, storage performance to service these workloads is driving a requirement for information technology (IT) infrastructure modernization. Older storage technologies just aren't able to meet the evolving storage performance requirements of these newer workloads. This in turn is driving significant interest in new types of very high-performance persistent memory-based media (commonly referred to as storage-class memory [SCM]) and a new memory-optimized protocol designed for use with solid state storage called nonvolatile memory express (NVMe).

While a number of enterprises are already experimenting with SCM, this has primarily been through the use of products designed for deployment as local storage. Most of this testing uses a product from Intel called Intel Optane SSD, which is the most widely available SCM technology in the market today. Local storage solutions, regardless of how fast they are, pose challenges to broader use across the enterprise such as limited capacity, a lack of enterprise-class data services (which typically require a shared storage solution like an array), and an inability to efficiently share high-performance storage capacity. What many enterprises would prefer are true enterprise-class solutions that deliver not only SCM performance but also all the other benefits of shared storage like scalable capacity, a broad and proven set of storage management tools and services, and efficient capacity utilization.

Hewlett Packard Enterprise (HPE) has just introduced a solution that leverages SCM (Intel Optane technology) and NVMe for HPE's 3PAR and Nimble Storage enterprise storage platforms and will be starting to ship these solutions in calendar 4Q18. This white paper takes a look at the evolving enterprise storage market, and which workloads tend to require low latency and very high throughput and bandwidth, and provides a brief look at one of the first enterprise storage offerings based around these two technologies today, HPE Memory-Driven Flash.

## IN THIS WHITE PAPER

---

As enterprises undergo their digital transformation journeys, more and more of them are deploying real-time analytics workloads. It is proving difficult for enterprise arrays to provide the type of low-latency performance these workloads need, and customers are looking to NVMe and next-generation media like SCM to provide this performance. HPE is an early entrant in this space, providing an enterprise-class storage solution optimized for Intel Optane technology. This white paper discusses the evolving low-latency market requirements in the enterprise space and then turns to a review of HPE Memory-Driven Flash, a new enterprise storage architecture accelerated with SCM and NVMe for HPE's 3PAR and Nimble external storage platforms.

## SITUATION OVERVIEW

---

DX is under way in most enterprises today. Decisions about IT infrastructure modernization are fueled by the fact that, along with legacy workloads like relational databases, CIOs are also deploying more next-generation applications (NGAs) that challenge traditional SCSI-based storage systems. Many of these NGAs are looking to leverage artificial intelligence and machine learning algorithms, along with big data and analytics, to drive business insights that better inform operational decisions. While today the lion's share of these types of systems is based around batch-oriented workflows, leveraging real-time analytics offers opportunities to move faster than competitors, and IDC is seeing more of this method being deployed across all industries. In many industries, this ability to be more responsive to market developments is driving clear differentiation. As enterprises come to depend on IT infrastructure to drive an increased level of business agility, these workloads will become more critical to ongoing business success. IDC expects that by 2020, 60-70% of Fortune 2000 companies will have at least one real-time big data and analytics workload that they view as *mission critical* to their business (driving additional requirements for high availability).

In this information-driven era, the scale of data that must be handled has significantly increased relative to what most enterprises had to deal with just a few short years ago. While all-flash arrays (AFAs) significantly increased storage platform performance relative to hard disk drive (HDD)-based arrays, the fact that they are based around SCSI is becoming an issue not only with these new, much more latency-sensitive NGAs but also for enterprises looking to lower costs by more densely consolidating workloads onto fewer storage platforms. While throughput, bandwidth, and latency are all identified as issues in these types of more performance-sensitive environments, the single biggest concern enterprises identify is performance predictability. Enterprises not only require predictably low latencies but also need to ensure that they can continue to meet defined service-level agreements (SLAs) as their environments scale. These concerns around performance predictability at scale have encouraged many enterprises to look at NVMe, a storage protocol alternative to SCSI, and the much higher-performance emerging memory technologies it supports (like SCM).

A number of storage vendors have started to ship enterprise-class storage platforms built around NVMe, but their continued use of NAND-flash based storage devices leaves a significant amount of potential NVMe performance on the table. SCM products like the Intel Optane technology for the datacenter (which leverages a new memory media called Intel 3D XPoint [CrossPoint]) can deliver latencies that are an order of magnitude better (at the device level) than NAND flash-based solid state disks (SSDs) with similar storage capacities – and deliver this lower latency with much more predictability under load than NAND flash-based devices. The combination of SCM and NVMe provides a much more capable technology combination than NVMe and NAND flash for delivering extremely low latencies at scale on mixed read/write, enterprise-style workloads.

Customers that have specific requirements to either deliver consistently lower latencies at scale or enable more densely consolidated mixed workloads without fear of "noisy neighbor" problems should look at those vendors offering tiered storage platforms that leverage SCM technology. As the industry moved away from HDD-based arrays, we first saw the rise of "hybrid flash arrays" that included a higher-performance flash-based tier and a lower-performance, higher-capacity HDD-based tier, but over the past three years, AFAs with a single tier of storage have come to dominate primary storage shipments. With the advent of NVMe and SCM, we will again move toward tiered storage platforms, but this time, they will be leveraging a higher-performance SCM-based tier with NAND flash making up the lower-performance, higher-capacity tier. Using intelligent caching and data placement algorithms (many of which will be driven by AI/ML), the best new "hybrid" designs will be able to deliver SCM performance most of the time at a blended cost per gigabyte that will be much closer to the lower cost of NAND flash.

While SCM products have become available for use on the server side in the past year, it will be important for vendors to offer true enterprise-class shared storage solutions that leverage the technology as well. This not only gives customers much better capacity utilization of the high-performance storage but also provides higher availability and potentially a full suite of enterprise-class data services (RAID, inline data reduction, snapshots, quality of service [QoS], encryption, replication, etc.). As the types of workloads that require SCM and NVMe performance become more critical to business, they will also require storage platforms that can deliver very high levels of availability. With respect to shared storage solutions, customers can expect to see SCM technologies first appear as cache tiers in enterprise storage arrays, and then (as prices drop and volumes increase) as persistent storage tiers supporting much higher capacities.

## HPE Memory-Driven Flash: A New Class of Enterprise Storage

HPE is a leading enterprise IT vendor based in Palo Alto, California. With overall company revenue of roughly \$29 billion in 2017, HPE is one of the top enterprise storage vendors and was an early entrant into the enterprise AFA market segment in 2014. What was significant about its entry into that market at the time was that it was the first of the leading enterprise storage providers to flash optimize its flagship storage offering, based on the 3PAR product line. This was effectively the first AFA product line that offered a full complement of enterprise-class data services, based on its mature, proven 3PAR operating environment, and really began an evolution in the use of AFAs from dedicated extreme performance platforms (at the time) to more general-purpose use with mixed enterprise-class workloads. Today, of course, AFAs dominate primary storage revenue, and HPE's systems (as well as those of the other AFA market leaders) are all used extensively for general-purpose mixed-workload consolidation.

In November 2018, HPE announced the integration of SCM and NVMe technology into its two main enterprise storage platform product lines – the 3PAR line, which provides massive parallelization and is targeted for use with consolidated, mission-critical workloads at scale, and the Nimble Storage line, which is a highly efficient architecture targeted for general-purpose mixed workloads in the midrange. The new solution, dubbed HPE Memory-Driven Flash, uses Intel Optane technology as a caching tier in the shared storage array controllers, delivering acceleration for potentially all workloads running on that platform. The Intel Optane Data Center line consists of two products today: Intel Optane DC Persistent Memory and Intel Optane DC SSDs. Both use the Intel 3D XPoint persistent memory media technology, but the Persistent Memory products are DDR4 pin-compatible memory devices that plug directly into the memory bus of an x86-based server while the SSDs are block devices that are connected over an NVMe bus and offer significantly expanded capacities (relative to the Intel Optane DC Persistent Memory products).

The 3PAR solution is available this quarter (calendar 4Q18), while the Nimble Storage solution will ship early in 2019. Several aspects of this announcement are important:

- **Deterministic low latency at high IOPS intensity.** This solution has already been tested at scale by HPE on the 3PAR platform and can deliver an average of sub-200 microsecond latencies (with nearly 100% of all I/O coming in under 300 microseconds). Similar latencies have been obtained on Nimble Storage platforms. The latency is measured from the application perspective using a Fibre Channel (FC) transport to the array. This level of performance clearly indicates that HPE can deliver extremely low latencies at scale, and customers can use the solution to handle real-time analytics and other NGAs that absolutely demand the highest performance or to enable denser mixed enterprise workload consolidation – or some combination of the two.
- **True enterprise-class shared storage solution.** HPE's two product lines are proven mature offerings, with over 100,000 systems deployed in production across both 3PAR and Nimble Storage. They support very large shared storage capacities as well as very efficient capacity utilization across networked servers. They boast a comprehensive suite of enterprise-class data services, including various RAID options, inline data reduction, thin provisioning, snapshots, encryption, quality of service, and a suite of replication options. Both 3PAR and Nimble Storage platforms deliver guaranteed "six-nines" (99.9999%) availability.
- **Mature storage tiering technology.** Both the 3PAR and Nimble platforms have a long history of efficiently supporting tiered storage environments. While in the past the autotiering features leveraged a combination of NAND flash and spinning disk, the storage operating environments of these platforms are very adept at getting the most out of a high-performance caching tier. The earlier systems were able to provide NAND flash performance nearly 100% of the time at a blended cost per gigabyte that was very close to that of spinning disk (for the entire system), and testing by HPE has shown that, on its newer systems, it will be achieving SCM + NVMe performance most of the time at a blended cost per gigabyte that is very close to that of NAND flash. AFA vendors that have not had prior experience with storage tiering will not be able to bring the same tiering expertise to their systems once they also offer a combination of the two technologies.
- **Nondisruptive upgrade path.** Customers can nondisruptively field upgrade both 3PAR and Nimble Storage systems to consistently deliver low latency at scale. The Intel Optane SSDs are packaged as add-in cards that can be inserted into each controller (up to eight cards for a 3PAR system and up to two cards for a dual-controller Nimble Storage array). Customers should ensure that they are running the right array software release, which includes a set of Intel Optane technology-optimized read caching algorithms to maximize read cache hits with the technology, and should note that the solution can transparently accelerate read performance for any and all applications. This approach makes it very fast and easy for customers to integrate this solution as needed.

Both Nimble and 3PAR arrays still redundantly write to controller-based main memory (NV-DIMM), send an acknowledgement back to the requesting application after the write is recorded in two separate controllers, and then periodically coalesce writes before they are written through to persistent storage. HPE Memory-Driven Flash is deployed as a separate SCM read cache, and the new caching algorithms determine when and how to move data from persistent storage into this tier. It is connected to main memory across NVMe (PCIe), which offers near-DRAM latencies for data already residing in the cache. SCM can be enabled at the volume level if desired, allowing customers to focus the performance acceleration on specific applications, a very nice feature when using the solution in mixed-workload configurations. Given the 3PAR and Nimble Storage systems' proven multitenant management capabilities, this is likely how most customers will use it – HPE Memory-Driven Flash will allow very performance-sensitive workloads to be hosted on systems running other workloads without the fear of running into the "noisy neighbor" problem.

It is significant that this announcement brings low-latency capabilities to HPE customers along with all the other trappings of true enterprise-class storage. Existing HPE customers have access to a fast, easy and, most importantly, nondisruptive upgrade with their current systems that preserves their existing investments and in fact improves the efficiency of their HPE storage solutions. New customers now have a very low-latency option from an established enterprise storage leader that can also help drive denser, more cost-effective workload consolidation to improve management efficiencies and reduce administrative costs. And it is important that all of these systems are covered by InfoSight, the industry-leading (in IDC's opinion) cloud-based predictive analytics platform – a feature clearly not available on AFAs from other vendors. IDC has tracked and written about InfoSight over the past five years, commenting on the ability of this cloud-based predictive analytics platform to drive compelling value for customers through improved performance, higher availability, faster problem resolution, greater efficiencies, and lowered cost.

## CHALLENGES/OPPORTUNITIES

---

SCM is still a relatively new technology, and products based on it are not yet widely available. In July 2018, IDC performed an extensive worldwide survey around the use of NVMe and found that there was high awareness of SCM technology (70.4% of respondents were familiar with it) and high interest in exploring what SCM might offer enterprises' highest performance workloads (54.4% of respondents were very interested in testing it, and 40.6% were somewhat interested). Performance was clearly the driving factor for SCM interest, with respondents most interested in predictably low latencies at scale, higher throughput and/or bandwidth, and an ability to enable real-time workloads that just could not be supported today (in that order). IDC's July 2018 survey also indicated that the two biggest concerns around SCM technology were cost and maturity. HPE's strong reputation as a reliable provider of enterprise storage solutions can help put customers at ease, and as Intel moves into more enterprise shipments with the Optane technologies in 2019, this will help allay any concerns customers may have about deploying SCM in production. Although the technology is relatively new, HPE has chosen to go with the supplier most associated with SCM in the market today – Intel – and will work jointly with the company to get the HPE Memory-Driven Flash and Intel Optane technology performance messages out there.

The early availability of these low-latency technologies (SCM, NVMe) on proven, true enterprise-class arrays that are known as storage consolidation workhorses presents significant opportunities for HPE to bring new customers into the fold. The fast, easy, and nondisruptive upgrade path to this new level of storage performance for a very high percentage of HPE's installed base on both the 3PAR and the Nimble Storage platforms presents significant opportunities for existing customers to stay on the performance leading edge.

## CONCLUSION

---

Enterprise workloads are evolving to more real-time orientations and are demanding more storage performance than ever before. For workloads like in-memory and other high-performance transactional databases, real-time analytics, and high-performance computing, the combination of SCSI and NAND flash is just not able to deliver the latencies, throughput, and performance consistency at scale that many enterprise SLAs require. Enterprises are also looking to streamline operations and exploring storage workload consolidation where possible, raising concerns about performance predictability at scale and "noisy neighbor" problems. All of these issues are driving interest in new, higher-performance storage technologies like SCM and NVMe. With HPE Memory-Driven Flash, HPE is providing a "very low latency at scale" option that addresses these concerns with true enterprise-class storage platforms that have

been proven across well over 100,000 production deployments. Customers interested in this type of low latency will want to take a look at either the 3PAR or the Nimble Storage platforms, depending on their performance, capacity, and resiliency requirements.

Given that HPE is among the first movers in getting an SCM-based solution to market, it will have to effectively communicate the value the solution provides and the types of workloads that can most benefit from it in the near term. IDC's survey data indicated that the types of workloads customers are most interested in using SCM for include in-memory databases, real-time analytics, and high-performance computing – findings that map to the targets HPE has identified for its HPE Memory-Driven Flash solution. This early-to-market offering does provide opportunities that HPE's traditional competition may not yet have to deploy a solution that delivers both extreme performance and cost-effective enterprise workload consolidation. By being very prescriptive about where this technology can best be used and by providing a nondisruptive upgrade path that makes it easy for its existing customers to adopt, HPE provides incentives for customers that need SCM today to look at its offerings.

## 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 Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2018 IDC. Reproduction without written permission is completely forbidden.

