

IDC TECHNOLOGY SPOTLIGHT

Sponsored by: NetApp

Increasingly, enterprise workloads are requiring more storage performance than SCSI-based solutions can provide. Emerging persistent memory technologies will address that requirement.

True Enterprise-Class Persistent Memory Solutions Will Open Up a New, Higher-Performance Storage Tier

October 2018

Written by: Eric Burgener, Research Vice President, Server and Storage Infrastructure

Introduction

Evolving enterprise workloads are continuing to push the need for higher storage performance than SCSI-based technologies can provide. Emerging memory technologies such as Persistent Memory (PMEM) and the Non-Volatile Memory Express (NVMe) storage protocol will resolve these issues by providing higher-performance options. With NAND flash in broad usage throughout enterprise datacenters today, the next step up in performance will have to come from memory-related technologies. These emerging memory technologies will drive order of magnitude improvements in what SCSI-based systems can offer in terms of latencies, throughput, bandwidth, and performance consistency in the face of widely unpredictable workload demands.

Within IT infrastructure, there have always been some number of "tier 0" workloads where additional performance (however that is measured)

AT A GLANCE

WHAT'S IMPORTANT

Emerging memory technologies will become necessary to meet performance requirements in modernized infrastructure and will increasingly require enterprise-class capabilities.

KEY TAKEAWAYS

Emerging memory technologies include Persistent Memory (PMEM) and the Non-Volatile Memory Express (NVMe) storage protocol.

These technologies will be used to turbocharge both existing and next-generation applications starting this year.

drives additional bottom-line business benefits. High-performance relational databases, often transactional in nature, are a classic example of a "tier 0" workload, and as new technologies become available that enable higher performance, these types of applications are typically the first applications to move to them. There are many "tier 0" relational database workloads, typically running on Oracle, Oracle RAC, or SQL databases today, that will benefit from more performance than SCSI-based all-flash arrays (AFAs) can provide. Other newer workloads that are also demanding "tier 0" performance include both real-time and batch-oriented big data analytics and many technical computing applications in financial services, healthcare, retail, oil and gas, computational fluid dynamics and other design applications, security, and forecasting (such as for weather events). To get the storage latencies needed, many of these newer workloads have been hosted using local (DAS) instead of networked (SAN/NAS) storage.

While local storage has always provided very low latency storage options, it lacks many features that storage managers rely on to effectively handle large and growing data sets. First, it has limited capacity, particularly if workload data protection expectations require some form of local data redundancy (e.g., mirroring or RAID). Second, server-resident

data cannot be easily shared without undue performance impacts (i.e., the network hop to another server), resulting in poor capacity utilization and overprovisioned (and hence wasted) local capacity. High-performance storage capacity cannot be cost effectively shared because each application server requiring that performance must have its own high-performance storage. Third, storage infrastructures that leverage local storage also limit the ability to scale compute and storage resources independently to achieve an optimal and cost-effective balance between the two. Moreover, local storage lacks the kind of enterprise-class data services that shared storage solutions make available, such as inline data reduction, thin provisioning, snapshots, encryption, replication, and many others that provide key performance, availability, and recovery options necessary to meet stringent enterprise service-level agreements (SLAs).

The next evolution of storage, in the minds of many customers, is a persistent memory solution that offers the low latencies of local storage with all the availability, capacity, and manageability advantages of enterprise-class shared storage.

Definitions

PMEM uses dual inline memory modules (DIMMs), backed by NAND flash, to deliver DRAM performance in a persistent storage device. PMEM devices can be either server-side or networked (shared) storage devices, and when they appear in shared storage solutions, they can offer all the advantages of enterprise-class arrays. There is clearly a latency difference between server-side and networked storage, although with emerging NVMe over Fabric (NVMe-oF) options that support remote direct memory access (RDMA), the additional latency imposed by the network hop can be as low as 5–10 microseconds (making it negligible even for many workloads that require DRAM-like latencies).

NVMe is a storage protocol optimized for use with solid state media. Unlike SCSI, NVMe was developed specifically for memory-based media and uses solid state performance and capacity much more efficiently than SCSI (which was originally developed for hard disk drives [HDDs] but can be used with flash). NVMe is a more lightweight I/O interface that delivers lower latencies, much higher bandwidth and throughput, and up to two orders of magnitude higher parallelism than SCSI (making it much better suited for dense mixed workload consolidation as well as the data access patterns of big data analytics and other next-generation applications). Emerging memory technologies such as PMEM require the NVMe protocol, which supports the data access semantics (load/store) of byte addressable (i.e., memory) instead of block-addressable (i.e., HDDs) storage. Because NVMe supports direct memory access, systems that use it dispense with all the complexity of block-addressable storage (host bus adapters, storage controllers, SCSI I/O stack, etc.). NVMe over Fabric extends the direct memory access capability to switched fabrics (hence the term remote direct memory access), allowing high-performance storage based on NVMe to be effectively shared across hundreds of servers.

Software-defined storage (SDS) refers to a standalone storage platform with all of the storage functionality implemented in software. SDS platforms typically depend on the underlying hardware to provide just two things: raw performance and persistence. SDS solutions run on commodity off-the-shelf hardware, which can provide flexibility to customers to run the software on the hardware of their choice. SDS products include block-, file-, and object-based offerings and are sold by vendors as appliances (with hardware bundled in by the vendor), software-only (providing customers the option to use the hardware of their choice), converged infrastructure (rackmounted solutions that include servers, storage, and networking all under a single SKU supported by a single vendor), and cloud-based services (which are based on SDS software). The SDS market will grow from \$10.12 billion in 2017 to \$16.14 billion in 2021 at a 13.5% compound annual growth rate.



Benefits

The key benefits of emerging memory technology products revolve around performance. Server-side PMEM products deliver near-DRAM access latencies from persistent storage while supporting somewhat higher capacity than actual main memory and, depending on how they're implemented, can offer the additional benefits of enterprise-class shared storage. As data sets requiring extreme performance continue to grow and big data analytics workloads come into broader use, IT organizations clearly need the performance, persistence, and access to enterprise-class features that emerging memory technologies provide.

The benefits of enterprise-class shared storage include access to higher capacities, more efficient sharing of storage, access to key data services such as inline data reduction, thin provisioning, RAID, snapshots, encryption, and replication, and the resiliency and recoverability of arrays that can be configured to meet "five nines" (99.999%) availability. Arrays that leverage these emerging memory technologies exhibit much higher infrastructure densities in terms of performance and capacity than SCSI-based arrays and when deployed at scale will reduce energy and floor space consumption as well as drive higher CPU utilization (leading to a more streamlined and cost-effective server infrastructure).

Emerging memory technologies are not just a technical play. For the most performance-sensitive workloads, these technologies will help IT organizations overcome a number of critical business challenges. Their capabilities will better meet market requirements around predictable performance and availability, and their efficiencies will help IT organizations undergoing digital transformation cost effectively modernize while streamlining infrastructure. Their characteristics will also improve IT agility for common tasks such as provisioning, helping make businesses more responsive to market demands as well as helping alleviate a key concern many businesses have around continuing to meet SLAs as they scale.

SDS brings its own set of benefits. It allows customers to select the hardware platform of their choice, whether driven by certification, vendor lock-in, or economic concerns. It also allows customers to take advantage of newer hardware technologies more quickly as they become available. Most SDS products have been designed within the past five years, leveraging newer design tenets around ease of use, self-managed storage, and autonomous operations. The different consumption models supported by SDS products also provide customers with more financial options when deploying new IT infrastructure.

Trends

As business models come to depend more on the internet, IT infrastructure demands are becoming increasingly real time in nature. Customers expect real-time response in many cases and will penalize businesses that do not provide it. Data-centric organizations that derive value from the data they collect and analyze are demanding rapid insight that can be used to drive the business forward, even as the amount of data they must regularly handle grows. These trends are driving the need for a new, higher-performance storage tier that exhibits local latencies but supports enterprise-class shared storage features and efficiencies.

SDS is clearly a trend; the SDS market is expected to grow to over \$16 billion in revenue by 2021. Customers are slowly moving away from hardware-defined infrastructures to more software-defined infrastructures primarily because of flexibility, ease-of-use, and economic reasons. A software-only consumption model also makes it easier for existing customers to adopt major new features on their existing storage platforms through just a simple software upgrade.



Considering NetApp

NetApp is a \$6 billion enterprise storage vendor that has been serving the needs of enterprise customers and cloud providers for almost three decades. It is a leading vendor of unified (SAN/NAS) storage solutions and has what IDC believes to be the industry's most comprehensive vision and product set for enabling hybrid cloud environments. NetApp storage solutions include HDD-only, Hybrid Flash Array (HFA), and AFA platforms, and the company also offers software-defined storage, converged infrastructure, and cloud-based solutions. NetApp's all-flash storage platform portfolio revenue, driven primarily by the NetApp All-Flash FAS (AFF) arrays, is growing faster than any other vendor's AFA revenue, and NetApp currently holds the number 1 market share position by revenue in the AFA space. NetApp's systems are known for their easy scalability (enabled by ONTAP, NetApp's clustered storage operating environment), comprehensive enterprise-class storage management capabilities, and proven availability. NetApp ONTAP-based solutions have met a "six nines" (99.9999%) availability requirement across the entire NetApp installed base for the past three years (according to installed base statistics kept by NetApp's support organization).

In May 2018, NetApp introduced NVMe technology into its AFF product line. The A300, A700, and A700s can support either SAS or NVMe-based solid state disks (SSDs) (based on customer choice), while the high end A800 is an end-to-end NVMe-based AFA (NAFA) that also supports NVMe-oF host connections. The NVMe-oF support on the A800 allows the low latency of NVMe SSDs to be delivered all the way to the application side for workloads that require the lowest latencies. This means that the significantly increased performance density of NVMe will allow customers to more densely consolidate workloads, either legacy or new, on these systems, leading to a better overall TCO (i.e., more workloads can be effectively supported on one platform). Given the extremely high performance of NVMe, even densely consolidated workloads will not experience a "noisy neighbor" problem. This same high performance effectively eliminates the need to perform any storage performance tuning, resulting in reduced administrative costs. These systems all support NetApp's ONTAP storage operating system and offer the multitenant and hybrid cloud management capabilities necessary for dense mixed workload consolidation in enterprise environments.

In October 2018, NetApp introduced a new software product, MAX Data, that creates a persistent memory solution that can offer all the scalability, availability, and efficiency benefits of shared enterprise-class storage. The MAX Data architecture supports several different configurations, including server-side and array-side options that offer different performance, capacity, and functionality profiles. The initial configuration, generally available, uses PMEM as a near-DRAM speed storage tier on the server, backed by network-based ONTAP persistent storage. For customers that want near-DRAM performance but with all the advantages of shared storage, any NetApp ONTAP-based array can be connected to servers running the MAX Data software. They can be either AFF or FAS systems, offering several different potential performance tiers that can include NVMe- or SCSI-based SSDs, SCSI-attached HDDs, or cloud storage.

With this type of configuration, MAX Data manages data placement between the "tier 0" and the "tier 1" storage, putting the "hot" data in the server-resident storage and the "warm" data in the appropriate back-end tier. In this configuration, customers have access to the full complement of ONTAP-based data services, which can be used to manage the data in both tiers. A server-side only (i.e., DAS) configuration will be available in the future as well and can use DRAM, NV-DIMM, or Optane DIMM devices, making that capacity persistent through the use of mirroring over RDMA links to similar capacity in another locally networked server.

It is important to note that this approach does *not* use caching technology. In server-side only configurations, the data is mirrored in real time to another server. In shared storage configurations, NetApp's optimized snapshot technology is used to maintain current copies of the data on each attached server in the shared array. In the event of any kind of



failure, data can be quickly recovered from one of these secondary locations. Because of this, MAX Data offers the full capacity of server-resident, solid state storage to be used to drive application performance (rather than provide data redundancy for recovery).

It is extremely important that both of these types of configurations will deliver near-DRAM performance without requiring application changes of any kind. In the past, certain relational databases were specifically altered to support in-memory execution (Oracle Database In-Memory, Microsoft SQL Server In-Memory Optimization) for customers that needed DRAM-level performance. These types of changes were made by vendors only for a very limited set of applications, resulting in products that included much new code that was less proven than the core database platforms from which they were developed. Because MAX Data delivers DRAM-like performance without requiring any applications changes at all, it can be used with a much broader set of workloads without compromising maturity.

The fact that MAX Data is a software product is strategically important. With all the storage management functions in software, customers can rapidly take advantage of advances in underlying hardware technologies as they become available. With many new storage technologies on the near-term horizon, including denser NAND flash media based on QLC, PMEM products such as Intel Optane, and broader use of NVMe, it will be critical to be able to integrate them into storage systems as needed to meet performance requirements. Solutions based around custom hardware, while they may meet requirements within a single generation, make it much more difficult to incorporate new capabilities as the underlying hardware evolves. Solutions based around commodity hardware are available from multiple sources, tend to be more reliable because of higher shipment volumes, and tend to cost less because of standards-based competition. When solutions based on commodity hardware deliver all the performance of custom products, they present a better solution.

Target workloads for MAX Data include performance-sensitive databases (relational and NoSQL), real-time applications such as big data analytics, and in-memory databases, or MAX Data can be used to ensure predictably consistent low latency and high throughput for mixed workloads that may or may not include any of these "tier 0" applications. Use cases include financial trading and real-time market analytics, fraud, and cyberthreat detection; healthcare diagnostics, therapeutics, and genomics workflows; weather forecasting; social media environments leveraging personalization algorithms; Internet of Things (IoT) workloads; artificial intelligence/machine learning—driven inference models; or any environments where customers need ultra-low latency or the ability to support huge data sets or want to enable enterprise-class data services for these demanding "tier 0"applications.

Challenges

Although there has been much industry discussion around PMEM technologies over the past two years, NetApp is the first vendor to announce and ship a generally available, enterprise-class storage solution based around them. There is always an awareness challenge with these types of new technologies. Customers that think of PMEM as only a local storage option with limited capacity and no data services will likely see expanded use opportunities once they understand how NetApp has created a shared storage solution that delivers near-DRAM latencies. NetApp's reputation as a proven reliable vendor of highly available enterprise storage solutions with an installed base of tens of thousands of customers will help overcome concerns some customers may have about moving to new, extreme performance storage solutions based around PMEM and NVMe. Where these technologies are in their life cycle also means that they are available at a price premium relative to conventional AFAs, so customers will need to consider using solutions such as MAX Data where they can easily justify their deployment for performance, efficiency, or TCO reasons.



Conclusion

Market demands for higher storage performance will drive the penetration of NVMe technology, and many vendors are already selling storage platforms based on it. NVMe is being used with NAND flash-based storage, but it also opens up a new, higher-performance storage tier option based around persistent memory technologies that can deliver DRAM-like performance at scale. As more workloads come to demand this type of "tier 0" performance, products that combine it with all the advantages of enterprise-class shared storage will become popular. NetApp is the first vendor to introduce a storage solution set that leverages PMEM to deliver this performance in a true enterprise-class product that can be deployed in either DAS or shared storage configurations. It is significant for end-user customers looking to deploy this solution that it requires no application changes and, when deployed in configurations with NetApp ONTAP-based storage platforms, delivers all the efficiencies of highly available, massively

For now, NetApp is the only game in town for customers that want near-DRAM performance in a solution that is truly enterprise class.

scalable shared storage that supports the full set of ONTAP enterprise-class data services. This software-based solution, called NetApp MAX Data, leverages commodity persistent memory hardware as server-resident storage and, in shared storage configurations, NetApp AFF and FAS arrays. Ultimately, NetApp plans to extend MAX Data to support shared storage configurations that support additional vendor arrays.

IDC expects solutions such as MAX Data to be deployed with both existing, latency-sensitive relational database applications and newer workloads that include real-time and batch-oriented big data analytics and many technical computing applications in financial services, healthcare, retail, oil and gas, engineering, security, and forecasting (such as for weather events) where their performance drives compelling, bottom-line business benefits. IDC believes this "extreme performance" market will grow over time as more next-generation application workloads evolve into requiring ever lower latencies at scale. NetApp's "first mover" status will help the company build experience with these types of solutions faster than its competitors, and its ability to support both DAS and SAN/NAS-style configurations that deliver DRAM performance in a persistent storage platform provides deployment flexibility to customers. In the next year, NetApp is likely to be challenged by other enterprise storage vendors in this space, but for now, it is the only game in town for customers that want near-DRAM performance in a solution that is truly enterprise class.



About the analyst:

Eric Burgener, Research Vice President, Storage

Eric Burgener is Research Vice President for Storage in the Enterprise Platforms, Storage & Network Infrastructure Group. Mr. Burgener's core research coverage includes Storage Systems, Software and Solutions, quarterly trackers, and end-user research as well as advisory services and consulting programs.



O IDC Custom Solutions

IDC Corporate USA

5 Speen Street Framingham, MA 01701, USA T 508.872.8200 F 508.935.4015 Twitter @IDC idc-insights-community.com www.idc.com

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

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.



