



HPE 3PAR TECH PRIMER

Looking for reliable, cost-effective storage for your mission critical workloads? Look no further than the market leading solid-state storage arrays from HPE.

Introduction

Today's enterprise is increasingly dependent on storage solutions and spinning disk is no longer the "go-to" solution. While hard disk drives (HDD) price points have declined dramatically, these drives don't provide the same speed and durability as solid-state drives (SSD). Solid-state arrays are faster, smaller and more reliable. Typical customer requirements include demand for high levels of automation and guaranteed low latency.

What is a solid-state array (SSA)? It is shared storage device that uses solid state drives (SSDs) to store data. The drives use semiconductor chips to store data and have no moving parts. HPE 3PAR StoreServ has become the preferred storage choice for data centers as it delivers the following advantages:

- AI driven performance management
- Low latency
- Easy to set up
- Storage efficiency
- Easy to integrate with data center and solutions
- Increased reliability

The market for SSAs is growing with the expectation that they will displace electromechanical storage arrays and other electromechanical storage solutions such as tape. Technvavio market research analysts forecast that the global market for solid-state arrays will experience a CAGR of just over 20% from 2018 to 2022. They note the growth of hyperscale data centers as just one reason why solid-state arrays are increasing in popularity. These data centers are designed to provide a single massive and scalable computer architecture. Small individual servers, called nodes, are clustered or grouped together and operated as a single machine. The networked computer servers are used by organizations for remote distribution, processing, and storage of large amounts of data.

Differentiating SSD vs. HDD

When considering storage solutions, it's critical to align the option that best suits your anticipated workloads as well as your budget. SSDs shine for workloads with random read and write operations, while HDDs delivery great price/performance for sequential workloads like media streaming.



Here is a quick comparison of key features:

SSD	vs.	HDD
0.1ms	Access Times	5.5-8.0 ms
≥6000 IOPs	Random I/O Performance	400 IOPs
failure rate <0.5%	Reliability	Failure rate between 2-5%
SSDs consumer between 2-5 watts	Energy Savings	HDDs consumer between 6-15 watts
SSDs have average I/O wait of 1%	CPU Power	HDDs have average I/O wait of 7%
Average service time for an I/O request while running a backup is < 20ms	Input/Output Request Times	I/O request time with HDDs during backup rises up to 400-500 ms
SSD backups take ~ 6 hours	Backup Rates	HDD backups take up to 22-24 hours

In general, SSAs makes sense when workloads require any or all of the following: low latency, consistent performance, high availability, and a small environmental footprint (e.g., less power and cooling). More importantly, they are a safe bet for business or mission critical applications as they are highly secure. The architecture of the array does not allow application code or user code access to the storage software, insulating SSAs from Spectre and Meltdown security exposure. Or put more simply, the more applications that are using the same shared storage, the greater the risk of variability in latency. SSAs are a key element of a future state where latency is no longer a concern – even when storage workload demands are on the rise.

The future of SSAs is tightly aligned to the growing enterprise adoption of artificial intelligence (AI) and machine learning (ML). Gartner anticipates that these types of applications will represent 25% of SSA workloads by 2022. The information-driven enterprise is more dependent than ever real-time data and analytics workloads that drive operational insight. The scale of data has increased dramatically and all-flash arrays are under pressure to provide the throughput, bandwidth and latency required for performance sensitive environments. With this understanding, Gartner anticipates that 30% of SSAs will be based on non-volatile memory express (NVMe) technology by 2020.

NVMe - Do You Feel the Need? The Need for Speed?

Non-Volatile Memory Express is a modern storage protocol alternative to SCSI-based interfaces and has become a key differentiator for storage vendors. What makes NVMe valuable is the protocol's latency-lowering capabilities, enabling SSDs to reach their full potential. NVMe addresses challenges such as bandwidth, IOPs, and latency - making it a very good match for flash.



NVMe is a better, faster way for SSDs to communicate with their host systems. It helps alleviate bottlenecks that can occur when very fast flash is connected to systems via SAS or SATA connections. How does this work? An NVMe controller uses the PCIe bus to directly connect the storage to the system CPU, eliminating steps required by a SATA connection. In addition, NVMe's use of up to 64,000 queues eliminates the single data queue that slowed down legacy data transfer using SATA.

The bottom line on NVMe – it is designed for flash. Only SSDs run on NVMe and it makes SSDs faster. Like flash, NVMe has no moving parts. This frees up the system to read data from many different locations at the same time, providing new levels of parallelism that improve workload performance.

The Shift to Memory-Driven Storage

The HPE 3PAR StoreServ family of flash-optimized data storage helps the enterprise modernize its data center while easily handling unpredictable workloads. It's a tier-1 storage solution that starts small and scales big. With tools to accelerate and automate provisioning, it provides multi-tenant design, hardware-accelerated deduplication and compression, and sub-1ms latency.

Predictable Low Latency with HPE 3PAR

- 1** | **Start Fast** - Up to 2X improved latency and up to 50% faster than NVMe all-flash arrays
- 2** | **Stay Fast with Deterministic low latency at high IOPS intensity** -proven ability to deliver an average of sub-200 microsecond latency (with nearly 100% of all I/O coming in under 300 microseconds)
- 3** | **Nondisruptive upgrade path** – Intel Optain SSDs are packaged as add-in cards that can be inserted into each controller.

A strategic advantage for 3PAR performance is achieved through its use of Mesh-Active Architecture enabled by the 3PAR ASIC. This provides a design that allows volume to be active on every drive, every controller and every part of the host at all times. The 3PAR ASIC is a dedicated chip found in every 3PAR controller that performs functions for extremely high performance and low latency as compared to a CPU with a general-purpose processor that must be provided with instructions and pre-compiled code in order to process data. The 3PAR ASIC improves efficiency by allowing the controller CPUs to offload computationally complex workloads as the 3PAR ASIC can perform them faster than the CPU. This frees up the CPU to perform other tasks.

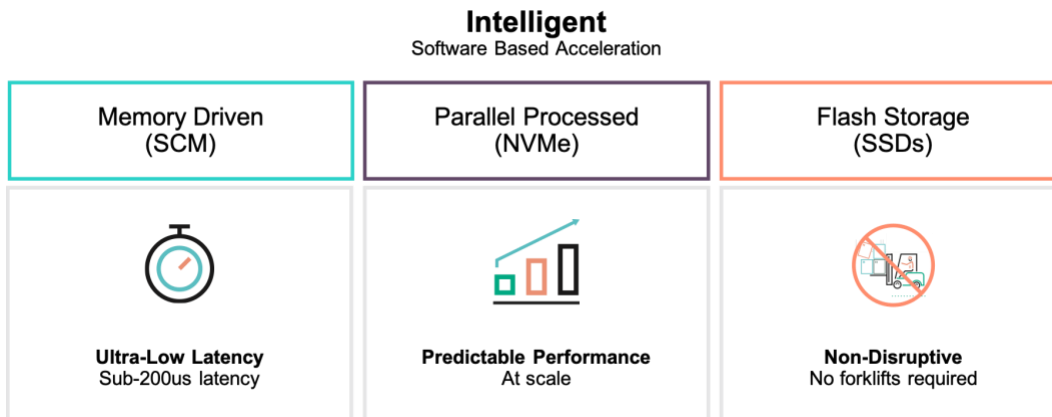
If that's not enough, the addition of HPE Memory-Driven Flash delivers a new class of enterprise storage that combines storage-class memory (SCM) and NVMe. NVMe is the accepted interface of choice for solid-state media. NVMe combined with NAND flash took applications running at high speeds to even higher heights. However, NVMe



combined with SCM delivers even more impressive performance. NAND reduced the performance gap between DRAM and mechanical storage media by huge margins, but was still slower than DRAM.

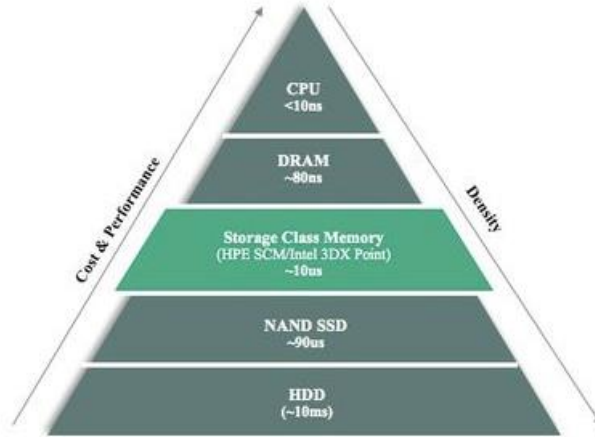
Storage Class Memory promises access times in the single digit microseconds with the endurance of DRAM and a price point closer to NAND. When SCM is combined with NVMe, enterprise IT benefits from a dramatic reduction in application latency. 3PAR SCM works as an extension of DRAM cache, and when combined with 3PAR's intelligent caching algorithm and NVMe, it delivers extreme application acceleration with HPE 3PAR StoreServ All Flash storage systems.

Breakthrough performance is derived from the HPE's combination of intelligent software-based acceleration with storage class memory, parallel processing and flash storage:





HPE Memory-Driven Flash is a new class of enterprise storage built with SCM and NVMe. It provides a memory-centric compute fabric design that allows multiple CPUs to access a central memory pool. It brings storage closer to compute, increasing the workload capability of a set of servers. HPE Memory-Driven Flash can deliver application performance up to 50% faster than all-flash arrays with NVMe SSDs. But what makes it more valuable and ideal for shared storage is that it can deliver ultra-low latency at scale, with near 100% I/O sub-300us. In short it provides predictable performance at scale, ultra-low latency and a non-disruptive path.

Over time, SCM technology may collapse the whole memory hierarchy into a single pool that will serve as both main memory and non-volatile primary storage.



3PAR's track record of incorporating the latest media technologies gives customers the benefits of these developments with minimal disruption. This is due to 3PAR's fine-grained three-layer virtualization that enables faster integration of new media types than competitive storage architectures. This makes HPE 3PAR StoreServ the optimal solution for any application that requires speed or for customers consolidating several workloads on shared storage, but don't want to worry about their mission critical applications.

Application acceleration	Multi-tenant Consolidation
 <p>Latency-sensitive applications such as transaction processing, interactive applications, and complex analytics queries</p>	 <p>In a multitenant environment, Storage Class Memory ensures your mission critical workloads are serviced with lowest possible latency</p>

AI Driven Insight

A valuable differentiator for HPE 3PAR StoreServ lies within HPE InfoSight™. HPE has incorporated analytics into the core architecture of every storage platform. 3PAR Performance Insights, powered by HPE InfoSight prevents performance bottlenecks in real time.



The HPE InfoSight data collection and analysis engine consists of three interrelated elements:

- Performance correlation analysis – quickly identifies leading factors contributing to performance or latency issues
- Detailed systems modeling – identify bottlenecks impacting workload performance
- Powerful, predictive algorithms – simplify administrator ability to visualize organic data growth and identify when arrays are approach capacity limits

By building diagnostic sensors into every module of code. InfoSight correlates thousands of sensor data points per second across the entire installed base. The learning infrastructure identifies any new problem and that problem is assigned a predictive health signature. Think of this as DNA. InfoSight uses pattern-matching algorithms to search for health signatures to either prevent an issue from occurring or to proactively repair it before it disrupts other activities

As applications increasingly drive today's enterprise, customers need self-managing IT models. InfoSight marries the power of big data with the promise of machine learning. By continuously aggregating and correlating data from across the data center, InfoSight can determine if a problem resides within the storage, or if the root cause lies elsewhere. Non-storage factors such as misconfigurations, host, network or VM problems are often the cause of bottlenecks in the data path.

HPE InfoSight is particularly valuable for virtualized environments using VMware. VMVision is a part of HPE InfoSight Predictive Analytics and provides a granular view of the resources used by every virtual machine connected to Nimble Storage. This allows you to correlate the performance of VMs in a data store with insights into hypervisor and host resource constraints such as a vCPU. The benefit is a broader view of potential latency factors, whether from the storage, the host or the network.

The value of HPE InfoSight cannot be over stated. It provides the AI driven management to predict issues and proactively resolve them before a customer is even aware of the problem.

The IIS Value Add

IIS is your trusted technology partner. When it comes to HPE 3PAR we've been aligned with HPE since it launched 3PAR. We were the first approved US based reseller. We have built up a high level of competency in designing the optimal storage architecture for your environment.

As a HPE Global partner of the Year, we have the capabilities to assess and analyze workloads and datacenter architectures to help you size the right balance of SSD and HDD resources, on-premise or in the cloud. IIS facilitates on-premise or cloud deployment leveraging its state-of-the-art Integration Center to fulfill "last mile" solution delivery through staging and testing of the end-to-end architecture.



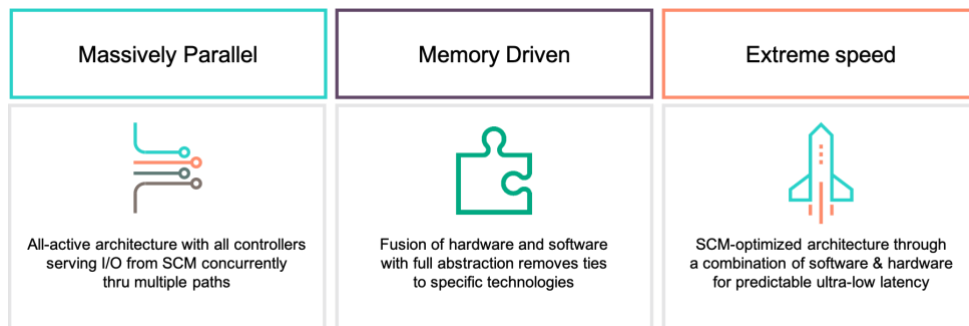
After deployment, our managed services deliver enhanced proactive, unified monitoring of your infrastructure environment. Leveraging InfoSight's predictive analytics embedded monitoring and management capabilities, IIS delivers unified monitoring of your entire HPE and non-HPE infrastructure, delivering rich dashboards, alerts, governance and custom reports. With dedicated resources to monitor, manage and maintain your cloud, hybrid or on-premise infrastructure, your team can focus on business priorities while we focus on infrastructure.

In addition, IIS can help you turn capital expenses into operating costs through the use of pay-as-you-go financing solutions. With a defined methodology, tied to your anticipated usage, a financial model for storage-on-demand is a reality.

Our goal: to share our expertise in HPE 3PAR StoreServ and complementary vendor solutions like VMware, to tailor a solution to meet your business, organizational and technology goals

Are you ready?

Choosing the right storage matters. HPE 3PAR StoreServ with HPE Memory Driven Flash and NVMe is the ideal storage platform for shared storage environments requiring extremely low latency and predictable performance.



Why is HPE 3PAR StoreServ the right solution for your enterprise workloads?

1. It's the focus on a media-independent I/O path optimized architecture – allowing applications to take advantage of both HDD and flash-based media in the same array, while already adopting new storage technologies such as NVMe.
2. It's HPE's commitment to uncompromised system resiliency – targeting very low recover point objectives (RPOs) and recovery time objectives (RTOs) to support mission critical applications within and across data centers, all backed by a 6-Nines guarantee.



3. It's about understanding that storage arrays handle a combination of workloads and HPE3PAR StoreServ offers priority optimized QoS, separating control commands from data movement to eliminate bottlenecks.
4. It's about unmatched scalability from the mid-range to the high-end and the ability to easily add capacity without disrupting existing workloads.

Are you looking for fast, reliable, predictable storage? Call us.