

NVMe-oF Storage for Research Facilities

The most performant, dense, scalable, and dense data storage platform for research

Deliver new insights with speed and agility

Genomic sequencing, disease understanding, and cure, particle recognition, geospatial analysis, all of these big data analytics require high volume, velocity and varieties of data to be analyzed with precision. But precision no longer has to take days or weeks with the Pavilion HyperParallel Data Platform™. It has industry leading performance, density, scalability, and flexibility that redefines analytics and AI/ML applications for big data. Achieving rapid insights for outliers in standard deviations was previously considered impossible.

Pavilion seamlessly integrates into architectures like IBM's Spectrum Scale™, Apache Spark™, and Ceph by leveraging NVMe-Over-Fabrics into a system designed for NVMe SSDs from the ground up. Previously, using NVMe technology for anything other than metadata access was unheard of. However, data processing and storage architectures have evolved rapidly, and Pavilion is at the forefront of Massively Parallel Computing (MPC) and analytics.

Reduce storage over provisioning, accommodate storage growth, consolidate applications, accelerate time to results, and reduce TCO by disaggregating a research application's storage with the Pavilion HyperParallel Data Platform. It scales to over 2 PB in each array, provides ultra-high performance, and ultra-low latency.

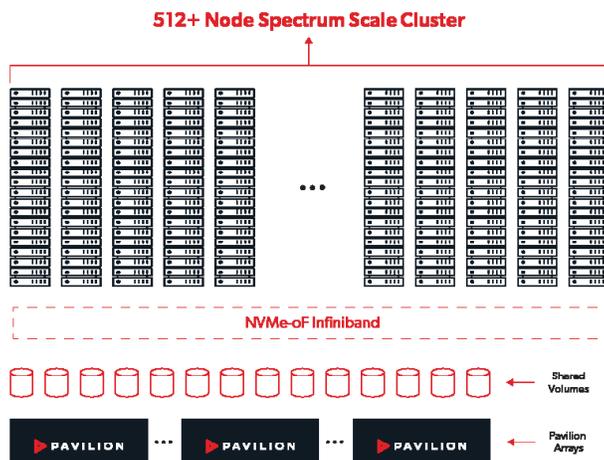
Direct attached storage and legacy all-flash arrays can't analyze this data quickly enough to provide an unrivaled customer experience. Pavilion offers the performance, reliability, security, and manageability, and user experience that today's leading research agencies need.

The Pavilion HyperParallel Data Platform

The Pavilion HyperParallel Data Platform delivers never before seen NVMe performance (120 GB/s throughput, 25µs of latency, and 20M IOPS) and density (2.2 Petabyte) in a compact 4U form factor. It provides applications with the performance of locally attached NVMe SSDs, enabling the move to a Composable, Disaggregated Infrastructure (CDI) infrastructure, where application resources are

Benefits

- Increases density 2X
- Reduces storage deployed by 75% and storage TCO by half
- Petabyte scalability, high-performance, low-latency, and linear scaling maximizes data center efficiency
- Protects your data and your business. Meets evolving requirements for data security and compliance
- Concurrent block, file, & object protocols for deployment flexibility
- Enterprise design and data integrity validation ensure reliable access to data
- Rack scale/CDI management via Web GUI, vCenter, Kubernetes, RESTful API, OpenStack, DTMF/Redfish, and Swordfish
- Operates with Ceph™ or other object stores as a complete solution for big data analytics



readily available. The Pavilion HyperParallel Data Platform uses NVMe SSDs and supports multiple block and file protocols with NVMe-oF/RDMA, NVMe-oF/TCP, iSCSI, NFS (v3, v4, pNFS, and RDMA) and S3.

The Pavilion HyperParallel Data Platform requires no proprietary software to be installed on a server farm and uses standard NVMe, NVMe-oF, Ethernet, and InfiniBand drivers, freeing up host resources for processing and reducing deployment risk.

Customer Choice and Control

The Pavilion HyperParallel Data Platform delivers high performance and low latency to applications.

Scalable & Flexible



Provide up to 2.2 Petabytes using NVMe-oF, iSCSI, NFS, or S3 in any combination. Grow performance and capacity linearly without impacting on-going operations. Scale across multiple arrays to support any capacity without limits.

Fast & Dense



Get an unrivaled experience by processing information faster using an ultra-high performance and ultra-low latency storage array. Search data 25 times quicker than with a DAS NVMe SSD, speeding time to results.

Disaggregating flash storage from each server in a rack no longer requires dedicating resources to managing local SSDs. Using servers that don't have to support DAS storage can double the compute density of a rack. No custom software needs to be installed on application servers, enabling them to take full advantage of host processing resources and reduce risk.

Safe & Secure



Protect the SSDs with RAID-6 and Pavilion SwarmController, which rebuilds a failed SSD at the rate of less than 5 minutes per TB. Get self-healing bit-rot support for data, to assure every process gets uncorrupted data.

Security is a must for every organization. Pavilion uses a FIPS-compliant data at rest encryption based on a 256-bit XTS-AES algorithm. It leverages AES-NI instructions to ensure the always-on encryption does not impact performance. Meeting compliance requirements is a fundamental part of the system design. Use consistent snapshots, encrypt those snapshots, and use standard backup and restore utilities with the confidence that data is safe and secure.

The failure of an SSD, data corruption, or data exposure does not impact banking operations.

Enterprise Strength & Enterprise Support



Get 24/7 proactive support, end-to-end data integrity, a robust and modular chassis, and redundancy throughout the storage array to protect your applications over infrastructures growth. All features come in-the-box, including thin provisioning, snapshots, clones, data at rest encryption and more.

Don't let support concerns prevent the deployment of applications. Pavilion provides 24/7 proactive support and can act as an extension of your IT organization for all applications.

Block, File & Object



Lots of data, resides in silos, unify access and drive collaboration across block and file data. Simultaneously use Ethernet, InfiniBand, NVMe/RDMA, NVMe/TCP, iSCSI, NFS (v3, v4, pNFS, RDMA) and S3 for all application needs.

In a large cluster, setting up shared volumes is easy. Assign NSD hosts to a controller and assign that same set of hosts to another controller and voila! You have a shared NVMe storage. Use standard NVMe-over Fabrics multipathing for failover and you have resiliency. The Pavilion HyperParallel Data Platform has the processing power to eliminate the need for NSD servers altogether.

Economical & Flexible Management



NVMe SSDs are the most expensive components of an application cluster, why not reduce them? Scale down flash storage deployed by 2-3X, and save 50% over DAS in terms of \$/GB/sec. No wasted capacity, no time spent messing with volume managers or file systems. Just set and forget.

With the Pavilion HyperParallel Data Platform you won't waste SSD space. It only allocates physical space as the application requires/consumes it. Save TCO by having no wasted capacity and no time spent messing with volume managers or file systems.

Reduce storage management costs by using the Pavilion HyperParallel Data Platform Web GUI or integrating with other management frameworks, including: vCenter, Kubernetes, RESTful API, OpenStack, DTMF/Redfish, Swordfish, and more.

Find Out More

Pavilion shatters customer expectations and resulting organizational outcomes by revolutionizing data processing for modern AI/ML, HPC, Analytics, Enterprise Edge and other data-driven applications. The Pavilion HyperParallel Data Platform, powered by Pavilion HyperOS, delivers unmatched performance and density, ultra-low latency, unlimited scalability and flexibility, providing customers unprecedented choice and control. Learn why Fortune 500 companies and federal government agencies choose Pavilion. Visit www.pavilion.io or follow the company twitter at <https://twitter.com/PavilionData>