

NVMe-oF Storage for Classified Agencies

Accelerating digital transformation

New insights with speed, agility, and no-compromise

Facial recognition, event correlation, and high-fidelity surveillance demand performance storage that simply cannot be delivered by traditional SAN or NAS storage. Artificial intelligence and big data analytics are game-changing technologies for your agency, yet achieving real-time insights across petabytes of images, video, text, and social media feeds in a shared file system is often described as impossible.

Pavilion shatters expectations with seamless integration into architectures like IBM's Spectrum Scale™ by leveraging NVMe-oF into a system designed for NVMe from the ground up. Previously, attempting to use 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 classified agency transformation.

Traditional SAN or NAS architectures based on monolithic architectures have reached their limits for big data and fast data analytics. In fact, it is rare to find leading internet and consumer-facing for-profit firms using anything except modern, scale-out storage solutions for massively parallel applications. Not only do the applications and systems not scale for correlation of images, videos, and structured queries, the sheer throughput requirements cannot be delivered by shared storage systems.

Pavilion's technology and partnership with leading classified solutions providers offer a better way. NVMe-oF has crossed the chasm and offers the reliability, security, and manageability that organizations have trusted with SAN for modern, rack-scale applications, enabling new technologies with low risk and high return.

Look forward — Look to Pavilion

Pavilion is leading the way in government transformation. Whether it is a logical evolution of a global file system improve storage utilization, shatter backup windows and achieve compliance, or a wholesale replacement of your infrastructure with Spectrum Scale, Apache HDFS with Spark™, our NVMe-oF Storage Platform provides unprecedented performance, availability and management features to future-proof your storage infrastructure as you bring insights to your big data analytics.

Performance

The Pavilion HyperParallel Data Platform™ is the most performant, dense, scalable, and flexible data storage platform. With up to 120GB/s of throughput, 20M IOPS of performance, and as little as 25µs latency from a host, over RDMA-based fabrics and through 20 parallel storage controllers to a distributed RAID-6 volume, Pavilion's performance is unparalleled. To achieve similar results in just 4 Rack Units (RU) of space, competing alternatives require at least 80 RU, or two full racks and as much as 14TB of DRAM at 10 times the acquisition cost.

Benefits

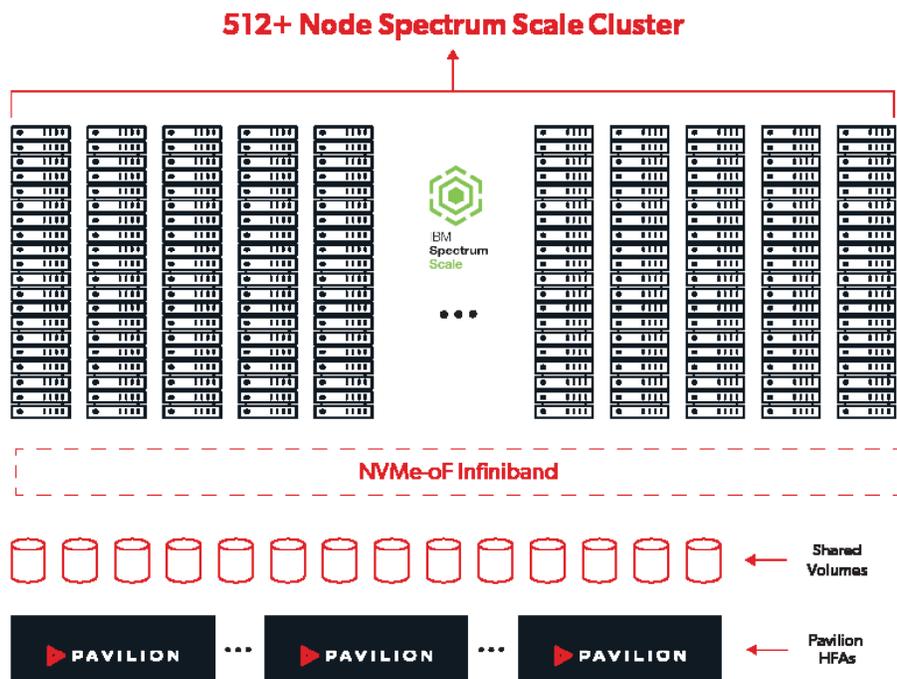
- Consolidate racks into a single 4U system
- Petabyte scalability, high-performance, and low-latency
- iSCSI and NFS protocols allow for easy migration from legacy systems to modern, scale-out applications
- 2X increase compute density
- Operate with nearline storage as a complete solution to existing backup and archive systems
- 67% faster indexing
- Enhance asset integrity and security
- Optimize the performance of mixed file sizes
- Continuous operations

Availability

Agencies demand no single points of failure, standards-based hardware and protocols as well as redundancy throughout a storage array. The Pavilion HyperParallel Data Platform features a completely fault tolerant design from controllers, power supplies, fans, management controllers, even dual PCIe switching fabrics. Using standard distribution NVMe-oF operating system drivers, multi-pathing to our 20 independent controllers assures fail-over in the event a network link or storage controller is unavailable.

Versatility

With up to 20 storage controllers and 40 Ethernet or InfiniBand fabric connections fully non-blocking at 100Gb/sec. the Pavilion array can serve as the ideal next wave of NVMe storage deployment for massive rack-scale workloads. Deploy the largest and fastest NVMe drives without concern for application performance impact in the event of a node recovery. Use a combination of read-intensive and high endurance drives for a bottom-of-rack configuration that services multiple workloads across a cluster.



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>