

## Features

**Key Benefits**

- Next generation disaggregated NVMe-oF platform disrupts current architectures
- iSCSI-based interfaces allow for easy migration from legacy systems to modern, scale-out applications
- Thin provisioning of NVMe volumes for unpredictable data surges
- Use encryption at rest to assure the privacy and security of fast data
- Use RDMA-based Infiniband or Ethernet and NVMe-oF drivers, including multipathing
- Operate with Ceph™ or other object stores as a complete solution to big data analytics

**Pavilion Benefits**

- Up to 920TB in 4U fully shared or partitioned
- 20m IOPS, 120GB/sec read and 90GB/sec write bandwidth @ 40µsec latency
- Create independent storage zones for DevOps, Test and backup
- Use zero-space snapshots and clones minimize network traffic and maximize productivity
- OPENCHOICE Storage lowers procurement costs and future-proofs

# NVME-oF STORAGE FOR RESEARCH FACILITIES

## Accelerating Discoveries

**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 to achieve precision no longer has to take days or weeks. Pavilion Data's NVMe-oF Storage Platform is redefining analytics at the junction of big data and fast data. Achieving rapid insights without compromise for outliers in standard deviations was previously considered impossible.

Pavilion Data makes the impossible, possible with seamless integration into architectures like IBM's Spectrum Scale™ and Apache Spark™ 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 Data is at the forefront of Massively Parallel Computing (MPC) and analytics.

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 data correlation, but the sheer throughput requirements cannot be delivered by shared storage systems.

Pavilion Data's technology and partnership with leading research facilities offer a better way. NVMe-Over-Fabrics 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 Data is leading the way in MPC analytics. Whether it is a logical evolution of a global file system improve throughput, make snapshots and clones without network impact, 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 and bring insights to your toughest analytics problems.

**Availability**

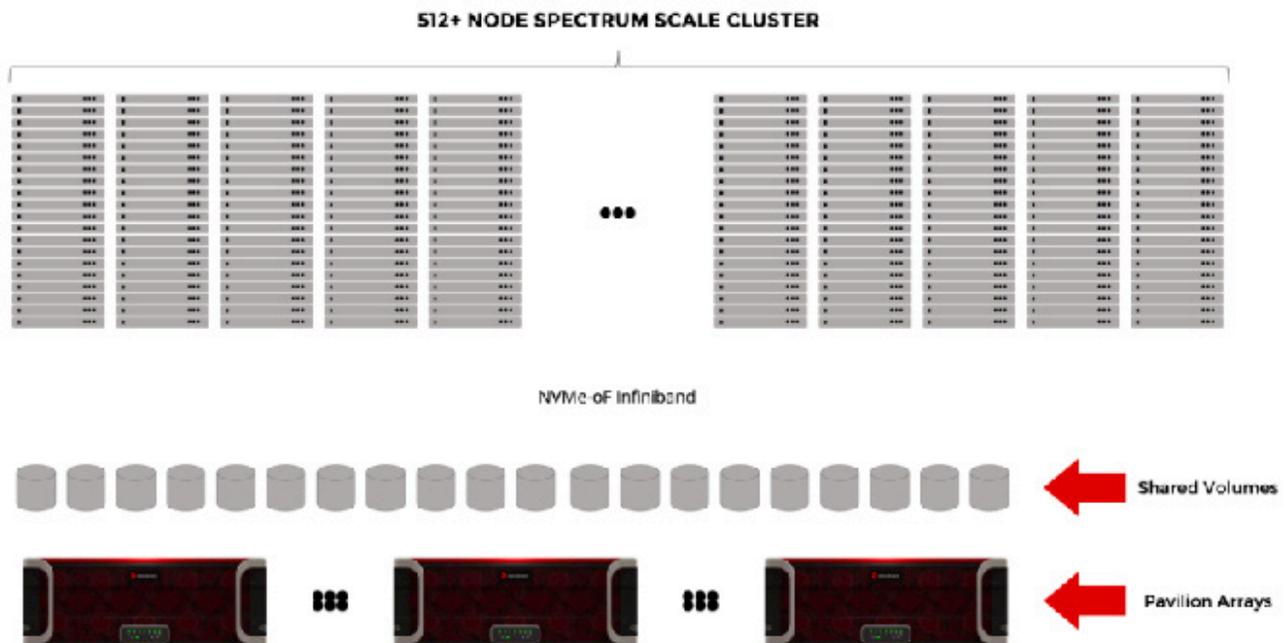
Pavilion Data has the industry's lowest latency disaggregated storage array. At 40 microseconds from a host, over RDMA-based fabrics and through 20 parallel storage controllers to a RAID-6 volume of OPENCHOICE NVMe SSDs, Pavilion Data's performance is unparalleled. To achieve similar results in just 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.

## Availability

Pavilion Data has the industry's lowest latency disaggregated storage array. At 40 microseconds from host, over RDMA-based fabrics and through 20 parallel storage controllers to a RAID-6 volume of OPENCHOICE NVMe SSDs, Pavilion Data'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. This radical performance density is an ideal fit for colocation facilities near trading exchanges.

## Management Features

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. With OPENCHOICE future-proofing your storage investment has never been easier.



In a large cluster, setting up shared volumes is easy. Use the GUI, CLI or framework integration (Redfish™ or Swordfish®) to assign NSD host clients or servers to a controller. Assign that same set of hosts to another controller and voila! You have a shared NVMe SAN. Use standard NVMe-over Fabrics multipathing for failover and you have resiliency. Since each Pavilion controller supports 4096 NVMe queues, it is possible to define each NSD node with 8 queues - up to 512 nodes to any given controller. Since the Pavilion platform has processing power associated with each of the 20 storage controllers, it is also possible to eliminate the need for NSD servers altogether.

With built-in encryption for data at rest, meeting security and privacy requirements is a fundamental part of the system design. Use consistent snapshots, encrypt those snapshots, then use standard backup and restore utilities to achieve high-fidelity compliance at a level of granularity that you define. Easily export clones to another cluster when export options in parallel file systems are limited.

Artificial intelligence, machine learning, and deep learning bring new opportunities to create new hypotheses and prove new theories. Pavilion Data offers a storage platform enabling you to deliver greater insights, faster using proven technologies systems and trusted storage management techniques. Learn more today at [www.paviliondata.com](http://www.paviliondata.com)