

NVME-OF STORAGE FOR



Features

GPFS Benefits

- Orders of Magnitude Better Throughput
- Increase operational Flexibility
- Improve Performance at Scale of GPFS
- Deliver more density per rack
- Reduce costs by migrating to Composable, Disaggregated Infrastructure

Pavilion Benefits

- 120GB/s of Throughput
- 14TB - 1 PB in 4U
- Frictionless Deployment
- Scalability and Flexibility
- High Performance and High Availability.
- Management Ease
- Scale Compute and Storage Separately.
- Lower TCO by leveraging scale-out architecture
- Data Resiliency & High Availability
- Space-Efficient, Instant Snapshots and Clones
- Thin Provisioning
- Standard Ethernet
- **OPENCHOICE** Storage

GPFS environments require shared storage and until now, the only way to take advantage of NVMe's low latency and parallelism with Spectrum Scale GPFS was to use either of the following:

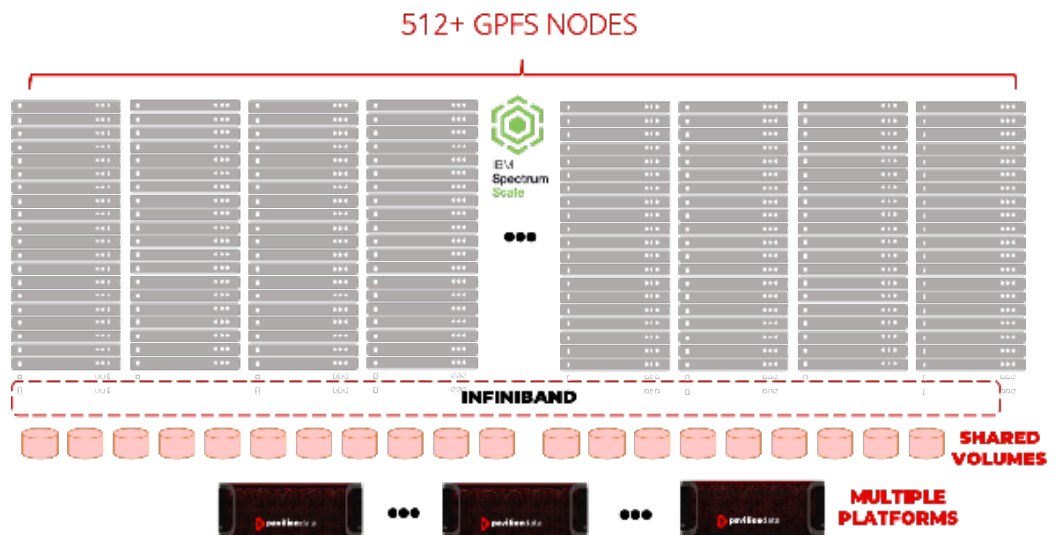
1. Legacy All-Flash Arrays limited to a pair of active-active controllers with NVMe drives.
2. Software Defined Storage (SDS) that adds processing overhead and management challenges to every server in the cluster.

However, due to recent technology advancements, Pavilion Data Systems now offers the world's fastest NVMe™-over-Fabrics (NVMe-oF) Platform for IBM Spectrum Scale™ GPFS. Loaded with manageability features you expect from an enterprise Storage Area Network (SAN), Pavilion provides end-to-end NVMe storage that delivers orders of magnitude better performance than AFA's or SDS in a category IDC defines as [Composable Disaggregated Infrastructure \(CDI\)](#).

Pavilion's NVMe-oF Storage Array

Our Platform delivers never before seen NVMe performance that allows customers to provision logical flash storage resources over a low latency network. As a result, we deliver shared storage with orders of magnitude more throughput for parallel file systems like GPFS.

The platform requires no custom software to be installed on application servers and includes important data management and availability features, including thin provisioning, instant zero-space snapshots and clones, and no single point of failure.



Performance-Optimized Platform For Gpfs

Pavilion's Platform provides resiliency, manageability, performance and flexibility functionality. All features are delivered by the array. It leverages a memory-speed distributed transaction log to optimize IO latency. An intra-cluster network allows components within the system to communicate at nanosecond latencies. Key features include:

Dual Parity RAID

Implemented within a media group. Each group has 18 drives resulting in a 16+2 RAID group. This results in only 12% space overhead in 4U using commodity 2.5" U.2 NVMe SSDs..

Management GUI, CLI and API


Our platform provides administrative access with a full-featured REST API and CLI for deep insight into performance metrics at the volume/ application, controller, port, or system level. It also includes a proactive telemetrics subsystem that reports issues to Pavilion's cloud-based support portal with plug-and-play support for controllers and drives to be powered down, hot-removed, or inserted at any time.

Rapid Rebuild with SWARM Recovery

One concern with larger drive capacities in RAID systems is the recovery from drive failure. The larger the drive the more time it takes to recover. Whilst generally true, Pavilion swarms multiple controllers in parallel to recover a failed drive resulting in a 12+TB SSD being recovered in well under an hour.

N+1 Active-Active Controllers


With up to 20 Active-Active Controllers, any controller can access any volume on the media (read and write) directly after a failover. Multiple controllers can perform IO for a volume, availability is increased in the event of failures across the network, controller and IO Line card. If the path between the controller and the array fails, a controller with a different access path takes over, maximizing application up time.



Linear Scalability As Controllers And Hosts Are Added

In the example below, an FIO-generated 70:30 R:W Mixed 4K Random IOPS workload was measured with one client connected to each controller in the platform. These results are achieved without installing any custom software in the application hosts and no host CPU cycles are required to deliver this type of performance.


System Scalability With Mixed R/W (70:30) 4K IOPS Workload



Why Pavilion For GPFS

The Scatter Diagram on the right compares ease to use with throughput for different available storage architectures. GPFS is about throughput, but requires shared storage. The typical SAN lacks significantly in throughput.

Pavilion’s Composable, Dissagregated Platform offers orders of magnitude more throughput for GPFS clusters. Our ease of use is also highly depedant upon the performance density we can provide from a single platform. A performance density that has never-before been in shared storage.



OPENCHOICE Storage™

Pavilion Data believes in an **OPENCHOICE** of SSD procurement. Negotiate the best possible price for the optimal SSD matching your workloads directly with the drive maker. Drop it into our NVMe-oF Storage Array. Redeploy existing DAS drives into our array. Run different workloads (bandwidth intensive volumes, endurance-sensitive volumes) within the same Array. When the latest NVMe technology is available, use it immediately – don't wait for the “next generation” array with planned obsolescence and forklift upgrades.

A Revolutionary Storage Array For Spectrum Scale

GPFS continues to shape the customer experience across many industries and sectors, notably Media & Entertainment and Financial Services. As the demand for parallel file system performance grows exponentially, there is an insatiable need to ingest, process, store and serve this data.

Pavilion Data is the industry’s first standard- based NVMe-oF Storage Array delivering NVMe performance that scales in a linear fashion while minimizing data center footprint and requiring no special server-side software.

To lean more, please visit www.paviliondata.com