PERFORMANCE
SOLUTION BRIEF

PAVILION

Benefits

- Universally unmatched performance
- Dramatically cut data processing time while simultaneously analyzing larger data sets
- Cut AI model training and inference time in half
- Ingest data faster than you ever thought possible
- Run twice as many VMs on existing infrastructure

Features

- High performance for multiple workloads, simultaneously
- Up to 120GB/s throughput
- Up to 20M IOPS
- Latency as low as 25µs

Pavilion HyperParallel Data Platform™

The Universally Unmatched Data Storage Platform

Part one of a four-part Solution Brief Series: **Performance**, Density, Scalable, and Flexibility

Modern applications such as AI/ML, Analytics, and HPC workloads are fueling the growth of the leading organizations of tomorrow. These organizations are leveraging the power of data to extract meaningful insights and gain a competitive advantage across a wide range of use cases, including:

- Analytics used for fraud detection,
- Retailer trend analysis,
- Facial recognition,
- 3D volumetric capture,
- Genomic sequencing,
- Tick data analysis for high frequency trading,
- and many more use cases requiring high throughput and ultra low latency.

Deployments using these applications are accelerating and storage performance now serves an even greater impact as organizations must move and process more data, at a faster rate, than ever before. Companies realize failure to embrace this new reality puts them at a true competitive disadvantage.

"I think Pavilion has given DDN, VAST Data and WekaIO a hard problem as your system's superior hardware architecture accelerates file IO beyond their limits."

Chris Mellor
Blocks and Files

These applications require a storage platform that can meet extreme performance requirements. Uniquely capable of meeting this need, the Pavilion HyperParallel Data Platform delivers universally unmatched high performance, with consistent, predictable ultra-low latency to power the most demanding applications.

The HyperParallel Architecture

Powered by up to 20 independent controllers per system, each Pavilion HyperParallel Data Platform delivers astonishing performance of up to 120GB/s of throughput, 20M IOPS, and latency as low as 25µs. Other solutions either cannot achieve this level of performance, or require multiple arrays aggregated together, dramatically driving up costs, to do so. The Pavilion HyperParallel Data Platform, which consists of the Pavilion HyperParallel Flash Array™, powered by Pavilion Hyper OS™, leverages the power of a unique switch-based design, that delivers order of magnitude performance improvements over legacy architectures.

Only Pavilion delivers this unparalleled performance, in a compact 4RU system.

Universally Unmatched High Performance for Every Data Type

The Pavilion HyperParallel Data Platform delivers industry leading performance across workloads. Building on the multi-controller architecture of the Pavilion HyperParallel Flash Array, only Pavilion delivers best-in-class performance for block, file, and object workloads simultaneously.

Best-in-Class Performance for Each Workload, From A Single System

Block		File		Object	
Read Up to 120GB/s Up to 20M IOPS 100µs	Write Up to 90GB/s Up to 5M IOPS 25µs	Read Up to 90 GB/s	Write Up to 56GB/s	Read Up to 80GB/s	Write Up to 35GB/s

Legacy architecture systems are typically designed to support one data type, and then have a secondary data type that runs on top of the first, dramatically impacting performance. This inefficient design is a function of the dual controller architecture and inherently limits the performance of the second data type.

Pavilion breaks free from the limitations of the past with up to 20 independent controllers per system and is uniquely capable of running multiple workloads natively, on independent controllers. Every workload gets the full performance of each controller that it is running on.

Multi-Controller Performance and Reliability

The Pavilion HyperParallel Data Platform, with up to 20 controllers per system, delivers universally unrivaled performance. Each controller is independent and delivers consistent, predictable performance. Customers can start with as few as four controllers. Each controller can deliver up to 6GB/s of performance, for a total of 24GB/s. They can then scale up performance as needed, up to 20 controllers for up to 120GB/s of throughput per system.

High Performance for Read and Write

Application performance is critical for both read and write operations, Pavilion delivers extreme high performance for both. Other vendors typically are able to show better read performance by adding expensive cache to their array. This can improve sequential read performance but does little for writes. As a result, many vendors do not publish write speeds, or if they do, it is usually far less than their read performance.

Part of the unique design, Pavilion uses a cacheless architecture which delivers extreme high performance for both reads and writes, so every application, including those that rely on high data ingest rates, can benefit. Plus, without all that expensive DRAM or SCM for cache, costs are dramatically reduced.

Unrivaled Performance at Scale

Supporting any size data set and delivering linear performance and unlimited scalability across any number of arrays uniquely sets Pavilion HyperParallel Data Platform apart from all competitors

Each Pavilion HyperParallel Flash Array can support up to 2PB of usable capacity. Customers that need to scale beyond that can leverage the power of the Pavilion HyperParallel File System™, which is part of Pavilion HyperOS. Customers that leverage the power Pavilion HyperOS enjoy unrivalled performance with a global namespace for file or object workloads to scale in any combination to any capacity.

File	Object
Up to 75+GB/s Read Up to 50+GB/s Write	Up to 50+GB/s Read Up to 30+GB/s Write

Performance is based on a single system and scales linearly across systems

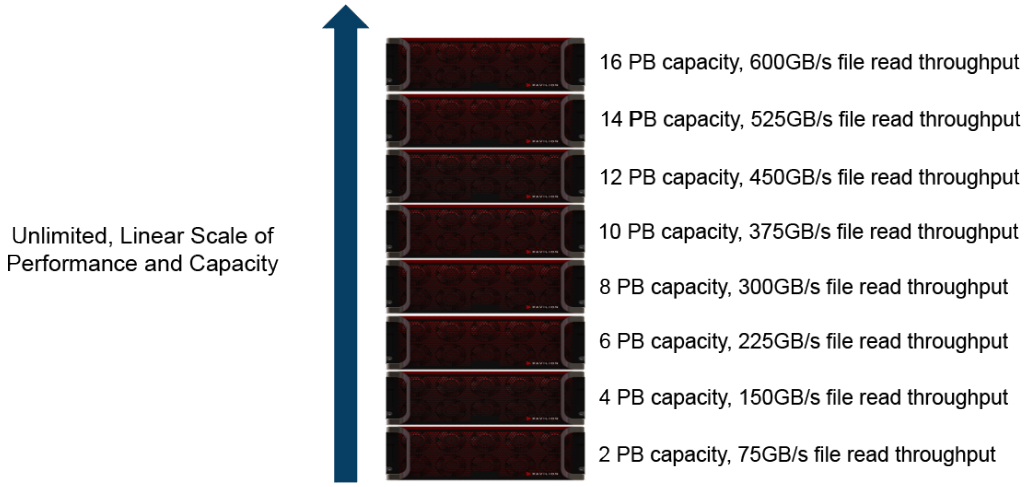

Performance shown is with Pavilion HyperOS 3.0 and global namespace for NFS.
Performance is maximum per system, individual results may vary.

Customers can also use their own external file system such as IBM Spectrum Scale, Lustre, or BeeGFS. Regardless of how customers scale, they always enjoy industry leading performance so they can get the most out of their applications.

Spectrum Scale		Lustre	
100GB/s Read	75GB/s Write	80GB/s Read	50Gb/s Write

Performance is from a single system and scales linearly across multiple systems.

The Pavilion HyperParallel Data Platform

Customers use Pavilion systems to power the only universally unmatched data platform with simultaneous, best-in-class performance for block, file, and object workloads, to process more data, faster than they ever thought possible. This in turn provides customers with measurable competitive advantage via speed to insights, lower operating costs, and greater flexibility to name just a few.. Applications including AI/ML, analytics, and HPC are running faster, doing more, and providing significantly more material benefits to the organization, making the innovators of today the leaders of tomorrow.