

## Benefits

- Grow capacity as needed
- Scale out block, file, or object data sets in any combination
- Scale capacity or compute, independently

## Features

- Scale up compute or capacity independently within a system
- Start with as few as 18 drives and grow to 72 within a system
- Grow to over 2TB per system
- Begin with as few as 4 controllers and grow to 20 within a system
- Unlimited scale out across systems

# Pavilion HyperParallel Data Platform™

## The Universally Unmatched Data Storage Platform

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

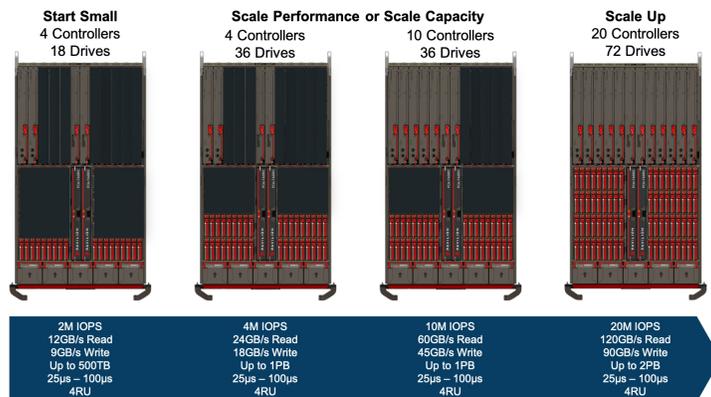
Unleash the power of compute without storage limitations. Doing so allows organizations to do more than ever before with their data. AI/ML, high performance computing (HPC), and analytics will continue driving demand for greater volumes of data.

As the need to process ever larger data sets becomes more critical, the need to expand storage capacity becomes critically important.

Pavilion gives customers unrivaled data storage choice and control with the ability to scale compute or capacity as needed. Customers can start small and grow to meet any capacity or performance requirements, without limits.

## Scale Up

The Pavilion HyperParallel Data System is the only storage platform that gives customers the ability to begin with as few as 4 independent controllers and 18 drives. Organizations can then scale either capacity or performance as needed, independently of each other. Customers can add more controllers, or they can add more drives, or both at the same time, up to 20 independent controllers and 2PB of usable capacity per array.



## Independent, Linear Scale

Pavilion customers can linearly scale up capacity or compute to meet their performance requirements.

Every controller is independent, with its own CPU, memory, networking, and OS instance, giving customers the full benefit of each one that is added to the system.

The unique design of the Pavilion HyperParallel Data Platform enables each controller to deliver linear performance growth. Every controller in a Pavilion HyperParallel Flash Array offers up to 6GB/s of throughput and 1M IOPs of performance. A system with 10 controllers delivers up to 60GB/s and 10M IOPs. A fully populated system provides up to 120GB/s of throughput and 20M IOPs.

With legacy systems, performance cannot typically be scaled independent of capacity and there is some amount of performance loss with every array that is added to a cluster.

Software defined solutions provide even worse performance gains, as every incremental node reduces overall performance, and the amount of performance reduction increases with every node added as additional memory and processing resources are consumed with managing the software.

## Scale Without Limits

Customers benefit from linear scale out across arrays, using block, file, or object data in any combination.

Organizations can leverage the power of parallel file systems such as Spectrum Scale, Lustre, BeeGFS, or the Pavilion HyperParallel File System™ to scale linearly across any number of systems to support the largest data sets.

Customer applications run faster, so users can be more productive and organizations can get more from their file-based workloads. On a single system, the Pavilion HyperParallel Data Platform delivers unmatched file performance from a 4RU system with up to 90GB/s of read performance and 56 GB/s of write performance. When scaled across multiple systems with Pavilion HyperOS, the Pavilion HyperParallel Data Platform gives customers up to 75GB/s of read and 50GB/s of write performance from each system.

Linear scale out performance extends to object workloads, with up to 80GB/s of read and 35GB/s of write performance from a single system. When scaled across arrays, the Pavilion HyperParallel Data Platform continues to give customers unmatched storage performance with up to 60GB/s of read and up to 40GB/s of write performance from each system.

# of Systems	Footprint (RU)	Block		File		Object	
		Read (GB/s)	Write (GB/s)	Read (GB/s)	Write (GB/s)	Read (GB/s)	Write (GB/s)
1	4	120	90	90	56	80	35
2	8	240	180	150	100	120	80
3	12	360	270	225	150	180	120
4	16	480	360	300	200	240	160
5	20	600	450	375	250	300	200
6	24	720	540	450	300	360	240
7	28	840	630	525	350	420	280
8	32	960	720	600	400	480	320
9	36	1080	810	675	450	540	360
10	40	1200	900	750	500	600	400

### About Pavilion

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](http://www.pavilion.io) and follow the company on LinkedIn