PAVILION

## Benefits

### Shatter Expectations

- Industry leading performance, density, and ultra-low latency
- Unlimited scalability
- Unrivaled flexibility
- Unique multi-controller architecture

### Get more from your data

- Real-time performance at scale to get results faster
- Consistent and predictable performance for your applications
- Concurrent block, file and object protocols simplifies your data center
- Extreme high performance for all data types

### End-to-End NVMe

- 72 U.2 2.5" NVMe drives in 1-4 Media Groups
- Simultaneous NVMe-oF/TCP, NVMe-oF/Ethernet, NVMe-oF/InfiniBand

### Massive Performance & Capacity

- Up to 120/90 GB/s R/W B/W, 20M IOPS, and 25μs latency
- NFS RDMA for extreme low latency
- Up to 2.2PB of capacity per array
- Up to 22PB of capacity per rack
- Scale out across multiple racks

### Ease of Use

- Single interface to manage all arrays
- Use the Pavilion HyperParallel File System or use your own
- VMware vSphere 7 compatible for NVMe-oF and RDMA

### Economical and Enterprise Features

- Instant and space-saving snapshots and clones
- Distributed DP RAID with 12X Faster Rebuild Than DAS and Legacy AFAs
- Data Assurance for Integrity Validation
- Performant controller failover
- Thin Provisioning
- Always-On Data at Rest Encryption
- Tiering

# Pavilion Hyperparallel Flash Array™

## The Pavilion HyperParallel Flash Array is the most performant, dense, scalable, and flexible storage platform in the universe.

We now live in a data centric world, a world that requires a new type of storage designed for modern workloads. A storage solution that can deliver best in class performance for a variety of workloads, not just one. Storage that can scale easily and without limits. A solution that enables users to move more data, faster than ever before, enabling organizations to extract more actionable information than was ever thought possible and achieve a competitive advantage.

## Consistent, Predictable, High Performance Across Workloads

The Pavilion HyperParallel Flash Array delivers customers industry leading performance for all their workloads. Uniquely capable of delivering industry leading performance and ultra low latency for block, file, and object workloads, simultaneously and with unlimited scale, the Pavilion HyperParallel Flash Array has revolutionized what customers can do with their storage.

- Up to 20 independent storage controllers and parallel architecture boost performance
- Up to 40 Ethernet or InfiniBand ports for flexible connectivity
- Up to 120/90 GBs R/W throughput, 20M 4K Random Read IOPS[1], and latency as low as 25 μs
- 4RU Chassis
- Active/Active data access
- No single point of failure
- End-to-End NVMe with simultaneous NVMe/Ethernet, NVMe/InfiniBand, and NVMe/TCP protocols
- Native S3, iSCSI, and NFS (v3/v4/pNFS/RDMA) support
- Global namespace for NFS and S3
- Up to 72 U.2 2.5" NVMe NAND and SCM SSDs
- NVMe-oF/RDMA (Ethernet and InfiniBand)
- NVMe-oF/TCP
- Standard drivers



## Unlimited Scale, Maximum Performance

The Pavilion HyperParallel Flash Array empowers customers to scale without limits. The power and flexibility of the unique multi-controller architecture enables performance and capacity to scale linearly and independently. Start with as little as four controllers and 18 drives. Then scale up to 20 controllers and 72 drives drives per array for over 2PB of capacity. Then scale across multiple arrays using any combination of NFS, S3, or block protocols.

[1]Based on Pavilion Data Internal Analysis

PAVILION

## Features

### 100% Standards Compliant

- Simultaneous NVMe/Ethernet, NVMe/TCP, NVMe/InfiniBand, NFS v3/v4, pNFS, NFS RDMA, S3, and iSCSI with standard drivers
- Up to 40 100Gb Ethernet and InfiniBand ports
- Up to 10 200Gb Ethernet and InfiniBand ports

### Flexible Management

- Web GUI and CLI
- vCenter, Kubernetes, and OpenStack
- RESTful API, DTMF/Redfish, and Swordfish
- Broad ecosystem integration
- Hadoop and Spark plugins

### Endurance & Reliability

- Data Versioning
- Continuous Operations
- Modular Architecture With Hot-Plug Support
- 24/7 Proactive Support
- Telemetry to Cloud-based Portal
- Designed for High-Reliability with No Single Point of Failure
- Replication

## Truly Unified Storage

Eliminate the burden of using different platforms to get best of breed performance for different workloads and dramatically simplify your datacenter while reducing costs. Legacy solutions implemented multiple protocol support by running a primary protocol on a single controller and then a secondary protocol on top, dramatically limiting performance.

The Pavilion HyperParallel Flash Array is the first solution to truly support high performance and ultra-low latency for block, file, and object workloads, simultaneously, on a single platform. The Pavilion HyperParallel Flash Array leverages up to 20 independent controllers per array to run each protocol natively on one or more controllers. Capable of up to 120GB/s throughput, 20M IOPs, and as little as 25µs latency per system, the Pavilion HyperParallel Flash Array enables organizations to run all workloads on a single, easy to manage platform.

Then scale those workloads across multiple arrays, in any combination. Always get the performance and capacity you need for every workload.

## High Performance for Read and Write

Legacy vendors often claim to offer high performance, but their specifications often reflect read performance only. The Pavilion HyperParallel Flash Array delivers high speed, low latency for both reads and writes. Capable of delivering up to 90GB/s write performance, the Pavilion Platform always delivers the performance your applications need.

## Pavilion HyperOS™

The Pavilion HyperParallel Flash Array is powered by Pavilion HyperOS 3.0. Delivering a modern storage solution to solve the challenges of today, tomorrow, and beyond, Pavilion HyperOS 3.0 takes advantage of the multicontroller architecture of the Pavilion HyperParallel Flash Array to deliver unrivaled performance, ultra-low latency, density, scalability and flexibility. Advanced features of Pavilion HyperOS 3.0 include:

- Global namespace support for NFS and S3
- Distributed, clustered file system
- Native support for multiple protocols, including NFS (v3, v4, pNFS, RDMA), S3, and iSCSI
- RDMA support for block and file workloads
- Advanced data protection features not found anywhere else
- Enterprise grade features including replication, tiering, security, encryption, snapshots and clones, encryption, and compression

## Multi-Fabric

The Pavilion HyperParallel Flash Array supports up to 40 Ethernet and InfiniBand ports and simultaneous protocols that eliminate unnecessary protocol translations while enabling NVMe semantics for low-latency and high IOPS across a range of topologies including NVMe/Ethernet, NVMe/TCP, S3, NFS, pNFS, NFS RDMA, iSCSI, and NVMe/InfiniBand.

## Cacheless Design

The unique design of the Pavilion HyperParallel Flash Array allows for data to be written directly to storage without the need to go through cache first. This unique architectural advantage delivers applications with greatly reduced latencies, when compared to legacy designs. Total costs are greatly reduced as there is no longer a need for an expensive caching tier using SCM.

## Instant Space Saving Snapshots and Clones

Protect your data with multiple logical and crash-consistent copies. Then enable your users to do more with clones that can be made of any volume and served out to different applications individually. These copies are consistent, space-efficient, instant, and writeable. Created instantly without physical data copy activities occurring; blocks are then written as the copy or original is modified over time.

## Thin Provisioning

Always have the right capacity for your users with thin provisioned volumes. Volumes are logical NVMe disks and used as a regular block device. The volume is thin provisioned from the media group. The user sees the provisioned size of the volume, but space is only allocated as-needed to maximize utilization. Volumes can range to 100s of TBs and can be re-assigned or shared between systems as needed.

## Controller Failover without Sacrifice

Get the best of both worlds and never choose between performance and availability again. Legacy, dual controller systems require you to choose between performance, by using both controllers resources, or availability by using active-passive or active-active mode at lower utilization. The Pavilion HyperParallel Flash Array leverages the power of multiple controllers to eliminate this burden, uniquely delivering high performance, even in the event of a controller failure.

## Distributed Dual Parity RAID

Fully populated, the array organizes NVMe SSDs into 4 groups of 18 drives. Pavilion implements Dual-Parity RAID within a group, resulting in 16 RAID protection. The overhead for RAID 6 is less than 12%. Optionally a hot spare can be defined, resulting in 15+2 RAID 6 protection.

## Pavilion SwarmController™ Rebuild

In the event of a drive failure, multiple controllers swarm the replacement drive in parallel to ensure fast rebuild. A 1 TB drive is recovered in less than 5 minutes. With the Pavilion HFA, an application's SSDs are fully rebuilt with RAID 6 protection up to 12X faster than using DAS or an all-flash array.

## High Availability, Low Overhead

Achieving the necessary level of availability and reliability for massively parallel modern applications can be costly. Converged and hyperscale solutions result in underutilized and stranded capacity trapped in servers in the event of a failover. Plus, the operational overhead of managing these isolated servers is an ongoing challenge for IT. The Pavilion HyperParallel Flash Array disaggregates storage from compute resources, delivering a more efficient, simple, and affordable solution with the same or better performance and availability.

**PAVILION**

## Non-Disruptive Software Upgrade

Take advantage of maximum availability with compliance requirements that meet government standards. All Pavilion HyperOS updates can be applied without disruption to ongoing I/O operations, so you never have to worry about going offline for updates.

## Data Assurance

As drive capacities continue to get larger, there is a small, but growing chance that there will be an error on a drive. All data written to the Pavilion HyperParallel Flash Array is validated on subsequent read operations, and if corrupted it will be rebuilt using the distributed RAID 6 protection. This ensures that applications never get corrupted data.

The Pavilion HyperParallel Flash Array also uses a unique, patent pending technology to confirm that all your data is safe. On every write the Pavilion solution generates a version number to confirm that rare, but significant drive errors do not corrupt your data. This technology, found only on the Pavilion platform, provides an additional layer of data protection.

## Data at Rest Encryption

Trusted by governments, Pavilion implements FIPS-compliant data at rest encryption and ensures the always-on encryption does not impact performance.

## Framework Integration

A full-featured Web GUI and CLI deliver deep insight into performance metrics at the volume/application, controller, port, or system level. The Pavilion HyperParallel Flash Array can also be managed using vCenter, Kubernetes, RESTful API, OpenStack, DTMF/Redfish, Swordfish, and more.

## Support Telemetry

The Pavilion HyperParallel Flash Array reports issues to Pavilion's cloud-based support portal for analysis by Pavilion. The support portal enables Pavilion to proactively respond to any issues, often solving problems before you even know they occur.

# Technical Specifications

| Common Specs | |
|---|---|
| Power Consumption | 3200 Max Watts |
| Heat Dissipation | 10922 Max BTU |
| System Height | 4 Rack Units |
| Protocols & Transports | Simultaneous NVMe/Ethernet using RoCE v2, NVMe/InfiniBand, NVMe/TCP, NFS v3/v4, pNFS, NFS RDMA, S3, iSCSI |
| Software Services | Snapshots, Clones, Data at Rest Encryption, Thin Provisioning, RAID, Pavilion SwarmController RAID Rebuild, Data Assurance |
| Management | Web GUI, CLI,  vCenter, Kubernetes, RESTful API, OpenStack, DTMF/Redfish, Swordfish, and SNMP v2c |
| Support | 24/7 Proactive Support with Cloud-based Telemetry |
| Physical Interfaces | Up to 40 ports of 100 Gb, or up to 10 ports of 200 Gb, Ethernet or InfiniBand for Storage Services<br>Up to 2 ports of 10 GigE for Management |
| Dimensions | 17.5" W x 31.1" D x 6.9" H |
| Weight | 147 lbs. (66.7) (loaded with SSDs) |

*All performance measurements taken in fully sustained mode with RAID6 enabled using NVMe-oF block protocol over Ethernet.