

Benefits

Expectation Shattering Storage

- Unrivaled performance, density, scalability, and flexibility
- Unlimited, linear scalability
- Industry leading performance for block, file, and object workloads
- Distributed, clustered file system
- Global namespace for NFS and S3
- High Availability and High Performance

Enterprise Features

- Thin Provisioning saves up to 2X on SSDs
- Native Block, File & Object Protocols
- Distributed DP Raid with 12X Faster Rebuild Than DAS and Legacy AFAs
- Instant and space-saving snapshots and clones
- Continuous Operations with Multi-Pathing
- Data at Rest Encryption
- Non-Disruptive Software Upgrades

Economical

- 2-3X Lower Storage TCO
- Linear Scaling of Storage & Compute
- 24/7 Proactive Support
- High density 4RU system
- Support for 40 x 100GbE ports per system

Pavilion HyperOS™ 3.0

Pavilion HyperOS™ 3.0 powers the **Pavilion HyperParallel Flash Array™** and together, the **Pavilion HyperParallel Data Platform™** delivers organizations with the highest performance, density, scalability, and flexibility in the industry. With support for block, file, and object protocols, **Pavilion HyperOS** shatters performance expectations to enable your applications to do more for your organization than ever before.

Built to unlock the power and performance of NVMe and NVMe-oF, **Pavilion HyperOS** enables the **Pavilion HyperParallel Flash Array** to deliver consistent, predictable, high performance, ultra low latency, and linear scalability. Designed for the most demanding environments, **Pavilion HyperOS** combines enterprise class management, security, and data protection features with a highly intuitive GUI to ensure maximum availability and ease of use. An API driven approach ensures that the **Pavilion HyperParallel Flash Array** easily integrates with your ecosystem for simple management.

Industry Leading Performance

The **Pavilion HyperParallel Data Platform** delivers industry leading performance density on a single system, with performance of up to 120GB/s read and 90GB/s write performance. The **Pavilion HyperOS** expands upon this with a global namespace for NFS and S3. Distributed across any number of arrays, the **Pavilion HyperOS** enables consistent predictable performance of 60GB/s read and 40GB/s write for both NFS and S3 from each array.

Native Multi-Protocol

Pavilion HyperOS supports block, file, and object protocols to run natively on any combination of controllers, across any number of **Pavilion HyperParallel Flash Arrays**. Each **Pavilion HyperParallel Flash Array** supports up to 20 independent controllers, each of which runs its own instance of **Pavilion HyperOS**.

Unlike traditional arrays, that enable multiprotocol by running one protocol on top of another, the **Pavilion HyperOS** runs each protocol natively on any combination of controllers natively for maximum performance.

Pavilion HyperParallel File System™

Part of the **Pavilion HyperOS**, the **Pavilion HyperParallel File System** is a multi-chassis clustered, distributed file system. The **Pavilion HyperParallel File System** expands upon that industry leading performance to deliver a high performance NFS and S3 object store that provides unrivaled levels of flexibility in performance, affordability, ecosystem integration, data services, and control across a global namespace.

Uniquely capable of providing independent, linear scalability of both capacity and performance, the **Pavilion HyperParallel File System** provides global namespace support for both NFS and S3 across multiple **Pavilion HyperParallel Flash Array** systems, enabling unlimited, linear scale.

Flexible Client Connectivity

Pavilion HyperOS provides flexible client connectivity across a number of clients for maximum ease of use and connectivity. Client plug-ins include NFS v3/v4, pNFS, NFS RDMA, a Gluster native client, Hadoop and Apache Spark.

Pavilion HyperOS is 100% standards-compliant. It uses standard Ethernet/InfiniBand networking interfaces, standard NVMe SSDs, standard host software stacks, and more. Pavilion does not require any custom software to be installed on application hosts or network switches. It supports standard 40Gb, 100Gb. Or 200Gb Ethernet/InfiniBand adapters for NVMe/Ethernet (RoCE v2), NVMe/TCP, NVMe/InfiniBand (RDMA), any v1.x compliant NVMe-Over-Fabrics driver. iSCSI, NFS v3 and v4, pNFS, and S3 Any or all of these protocols and fabrics can be operational in a single system at the same time.

Enterprise Grade Data Services

All storage services, APIs, and advanced data services are built-in and included with every array. Enterprise data services such as tiering, replication, snapshots, clones, security, encryption, thin provisioning, and proactive support are included in every instance of **Pavilion HyperOS 3.0**.

S3 data services include a single or global namespace, access controls, tiering, replication, snapshots, security, encryption, applicaiton plugins, and more.

Integrations for S3 include the most popular solutions for identity providers, monitoring and alerting, notificaiton targets, federation, orchestration, load balancing, and backup. Machine learning and Big Data integrations include Apache Spark, Hadoop, TensorFlow, kafka, splunk, Greenplum, Teradata, and more.

Performance and Scale

The **Pavilion HyperOS** enable linear, independent scaling of performance across controllers and arrays. Scale up to 10 systems per rack for an astonishing 1.2TB/s of performance per rack.

Enhanced Quality of Service

Four tiers of QoS support enable organizaitons to manage data access so that your most critical applications are always able to get the data they need. Cap based controls ensure that data is always available to those that need it.

Pavilion HyperOS is 100% standards-compliant. It uses standard Ethernet/InfiniBand networking interfaces, standard NVMe SSDs, standard host software stacks, and more. Pavilion does not require any custom software to be installed on application hosts or network switches. It uses standard 40Gb, 100Gb, or 200Gb Ethernet/InfiniBand adapters for NVMe/Ethernet (RoCE v2), NVMe/TCP, NVMe/InfiniBand(RDMA), NFS, S3, iSCSI and any v1.x compliant NVMe-Over-Fabrics driver. All of these protocols and fabrics can be operational in a single system or across an unlimited number of systems.

Block, File, and Object

Multi-Protocol

- Applications need access to lots of data, but much of it resides in silos. Unify the silos, driving collaboration across block, file, and object data. Simultaneously use NVMe/RDMA, NVMe/TCP, iSCSI, NFS and S3 for all your application needs.
- Enjoy industry leading performance with the native support for any supported protocol on any combination of controllers across an unlimited number of arrays for total choice and control

Ensuring Data is Safe & Secure



Distributed Dual Parity RAID

- Fully populated, the array organizes NVMe SSDs into 4 groups of 18 drives. Pavilion implements a distributed RAID 6 within a group, resulting in 16+2 RAID 6 protection. The overhead for RAID 6 is less than 12%. Optionally a hot spare can be defined, resulting in 15+2+1 RAID 6 protection.

Automatic "SWARM" Rebuild

- In the event of a drive failure, multiple controllers swarm the replacement drive in parallel to ensure fast rebuild. A 2TB drive is recovered in less than 5 minutes. With the **Pavilion HyperParallel Flash Array**, an application's SSDs are fully rebuilt with RAID 6 protection 12X faster than using DAS or an all-flash array.

Snapshot/Clones

- Multiple logical crash-consistent copies can be made of any volume and served out to different applications, such as a backup process, 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.

Continuous Operations with Multi-Pathing

- Achieving the necessary level of availability and reliability for massively parallel modern applications is costly. Underutilized and stranded capacity trapped in servers and the operational overhead of managing these isolated servers is an ongoing challenge for IT. There is no single point of failure and, I/O can be performed by multiple storage controllers, increasing data availability.

Data at Rest Encryption

- Security is a must, so Pavilion implements FIPS-compliant data at rest encryption and ensures the always-on encryption does not impact performance.

Non-Disruptive Software Upgrades

- Meeting compliance requirements is a fundamental part of the system design. All **Pavilion HyperOS** updates can be applied without disruption to ongoing operations.

CRC and Versioning

- **Pavilion HyperOS** includes a T10 dif cyclic redundancy check to protect against bit-rot. Building upon this, Pavilion also uses a patent pending technology to incorporate a version number for every write to protect against failed media writes. This protects organizations against rare, but occasional and significant issues that become more common as drives grow larger.



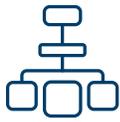
Controlling data growth & keeping cost in check



Thin Provisioning

- Volumes are logical NVMe disks and used as a regular block device. The volume is thin provisioned from the media group. The user receives the provisioned size of the volume, yet space is only allocated as-needed to maximize utilization. Volumes can range from under one TB to 100s of TBs and can be re-assigned or shared between systems as needed.
- Manage your data from a single system for total control with proactive monitoring by Pavilion Data

Economical & Flexible



Linear Scalability

- The **Pavilion HyperParallel Flash Array** supports the demands of multiple application with its cloud-like scalability. When an application's capacity needs to increase, DevOps can add an NVMe SSD without impacting performance. Conversely, when an application needs more I/O performance, additional storage controllers can be added without impacting capacity.

Multi-Fabric

- **Pavilion HyperParallel Flash Array** supports a variety of NVMe-oF Physical Fabric Ports, eliminating unnecessary protocol translations while enabling NVMe semantics for low-latency and high IOPS across a range of topologies including NVMe/Ethernet, NVMe/TCP, NFS, iSCSI, S3, and NVMe/InfiniBand.

Integration

- The **Pavilion HyperParallel Flash Array** lets you manage the array as you prefer. Whether you like to use a full-featured Web GUI and CLI that delivers deep insight into performance metrics at the volume/application, controller, port, or system level. You can also utilize its REST-based API to integrate with DTMF/Redfish and Swordfish or use the **Pavilion HyperParallel Flash Array** OpenStack and vCenter plug-in, or integrate with Kubernetes with its FlexVolume plug-in and CSI support for persistent volumes (PVs).

Proactive Support

- 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, performing problem-solving operations remotely.