PAVILION

Benefits

- Continuous Operations with no Single Point of Failure
- Hot Swappable Components
- Redundant Fans and Power Supplies
- Distributed DP RAID (16+2 and 15+2+1 with Rapid "SWARM" Rebuild
- Redundant Internal High-Speed PCIe Fabric
- Redundant Out-of-Band Supervisor Modules
- Up to 20 Storage Controllers Enable Linear Scalability of Compute And Storage
- Redundant Network Interfaces for Each Controller
- NVMe-oF Multi-Path I/O
- Non-Disruptive Upgrades
- 40 Ports Of 100G Ethernet and InfiniBand with NVMe/RDMA, NVMe/TCP, NFS, and iSCSI

Pavilion Delivers High Availability, High-Performance, and Low-Latency

Key availability features in a purpose built NVMe-oF storage array with no compromises

Traditionally, applications that implement high availability place individual SSDs in separate servers. In the event a server or SSD fails, the application fails over to another node. Whilst this might seem simple, this has resulted in significant architectural drawbacks.

Organizations have huge amounts of stranded and under-utilized capacity that is either held captive in distributed servers or in additional copies distributed in the cluster in order to be able to fail over to another node. This need not be the case anymore.

The Pavilion Hyperparallel Flash Array

The Pavilion Hyperparallel Flash Array (HFA) delivers never before seen NVMe performance (120/90 GB/s R/W throughput, 40µs of latency, and 20M IOPS) and density (1.1 Petabyte) in a compact 4U form factor. It provides applications with the performance of locally attached NVMe SSDs, enabling organizations to move to a Composable, Disaggregated Infrastructure (CDI) infrastructure, where application resources are readily available. The Pavilion HFA is highly-available, uses NVMe SSDs, and simultaneously supports multiple block and file protocols with NVMe/RDMA, NVMe /TCP, NFS, and iSCSI.

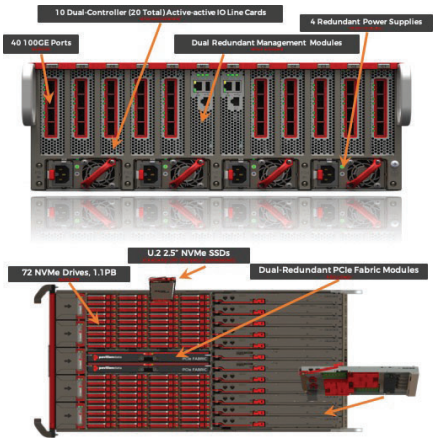

Pavilion allows multiple racks of application servers to simultaneously access storage over a high-speed low latency network at direct-attached SSD-class speeds. This is delivered to mission-critical applications with high availability and uptime. Just like other enterprise-class storage arrays, these availability features are self-contained within the array and do not require custom software on any application servers to achieve this level of reliability and availability.

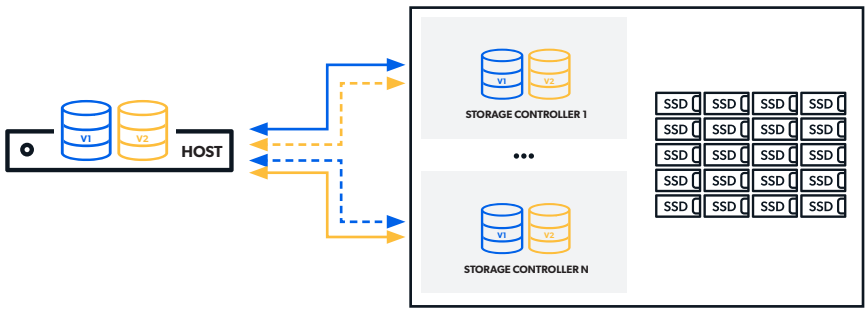

High Availability Features	The Pavilion HFA is designed from the ground up with key high availability features that improve application uptime especially in cloud-scale environments.
No Single Point of Failure	Every component is at least dual redundant, including network ports, internal PCIe fabric, I/O controllers, supervisor modules, power supplies, and fans.
Hot-Swappable Components	All components in the chassis are hot swappable for maximum serviceability, including SSDs, I/O line cards, supervisor modules, PCIe fabrics, fans and power supplies.
Distributed Dual-Parity RAID with Hot Spare Support and Rapid "SWARM" Rebuild	All user data volumes are provisioned from a drive group containing 18 NVMe SSDs in a distributed RAID-6 16+2 configuration. This ensures that up to two SSDs can fail without interrupting application access to data. The entire system contains up to 4 zones of SSDs, each with one or two RAID groups. Hot spares are also supported (15+2+1), enabling rebuilding to full redundancy, allowing more time to schedule drive replacement without increased data risk. In the event of a drive failure, multiple controllers swarm the replacement drive in parallel to ensure fast rebuild. A 1 TB SSD will be rebuilt in less than 5 minutes, which is over 12X faster than using DAS or an all-flash array.
Redundant Internal PCIe Fabric	Pavilion's patented architecture employs a high-speed PCIe network to connect all of the internal components, including I/O Line Cards, the NVMe Drive Array, and the Supervisor Modules. This fabric is fully-redundant and is implemented on dual-redundant swappable PCIe switch cards contained in the chassis.
Redundant Supervisor Modules	All components are managed by redundant out-of-band management controllers, or supervisors. Management of the array is done independently of the I/O controllers and the data paths, providing greater flexibility and consistent performance even during maintenance operations. If the supervisor interface is unavailable, I/O operations continue.
Multi-Path I/O Support	Pavilion's multi-path I/O support (dual paths) allows for uninterrupted data availability even under full line card, storage controller, NIC, or cabling failure. The use of industry standard multi-path NVMe-oF allows the operating system to transparently route around the failure to get to data, with the application needing no changes whatsoever.

Up to 20 Independent, Redundant I/O Controllers	All I/O controllers are active and serve I/O operations simultaneously, providing linear performance improvement as you add controllers. Each volume is available through multiple controllers, providing full availability even in the event of controller failure via multi-path I/O.
Non-Disruptive Software Upgrades	Meeting government compliance requirements is a fundamental part of the system design. All the Pavilion HFA's OS updates can be applied without disruption to ongoing I/O operations
Redundant Multi-Fabric Connectivity	The Pavilion HFA supports up to 40 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, NVMe/InfiniBand, NFS, and iSCSI.
Performance & Safety	Pavilion makes the only self-contained, disaggregated NVMe-oF array that delivers high availability features without requiring users to install proprietary software on application hosts. As a result, users can enjoy the simple and seamless deployment that they get with traditional shared storage arrays, but with the density and scalability to power multiple racks of servers with high performance and low latency enabling centralized storage with lower risk.

pavilion.io | Pavilion Data Systems, Inc. | 2550 N First St., Suite 250 San Jose, CA 95131 | E-mail: info@pavilion.io