

HIGH AVAILABILITY AT 40µs LATENCY

Highlights

Key HA Features

- No single point of failure
- All hot swappable components
- Dual parity RAID (16+2 or 15+2+1)
- Redundant internal high-speed PCIe Fabric
- Redundant out-of-band supervisor modules
- Multiple storage controllers
- Redundant fans and power supplies
- Redundant network interfaces for each controller
- NVMe-oF multi-path I/O
- Continuous Operations

Benefits

- 40 µs Latency
- Petabyte scalability in 4RU
- Frictionless Deployment
- Data resiliency & High Availability
- Space-efficient, instant snapshots and clones
- Thin Provisioning
- Linear scalability of compute and storage
- 40 ports of 100G Ethernet and Infiniband with RDMA, TCP/IP, or NFS
- **OPENCHOICE** Storage

Server-Side Architectural Limitations

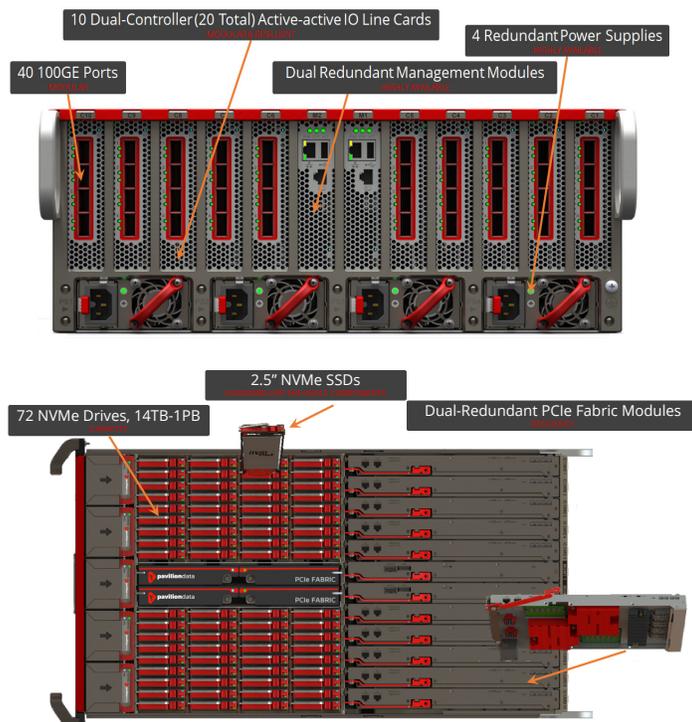
Today, modern applications rely on high availability by placing individual SSDs in separate servers. In the event an SSD fails, the application fails over to another node. Whilst this might seem simple, this has resulted in significant architectural drawbacks.

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.

Reliability at Scale with Pavilion's NVMe-oF Storage

Innovations in both networking and storage technology now make it possible to deliver shared storage with the same performance as locally-attached SSDs for large-scale distributed applications. However, in order to take advantage of these new capabilities, a shared storage system must have critical availability features built in.

Pavilion allows multiple racks of application servers to simultaneously access storage over a high-speed low latency network at direct-attached SSD-class speeds. We deliver this 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 storage array is designed from the ground up with key high availability features to application uptime in cloud-scale environments. Specific features are listed below.

No Single Point of Failure

Every component is at least dual redundant, including network ports, SSDs, internal PCIe fabric, I/O controllers, supervisor modules, power supplies and fans.

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.

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.

Dual-Parity RAID with Hot Spare Support

By default, all user data volumes are provisioned from a drive group containing up to 18 NVMe SSDs in a RAID-6 (16+2) configuration. This ensures that up to two drives can fail without interrupting application access to data. The entire system contains up to 4 zones of media, each with its own independent RAID group. Hot spares are also supported (15+2+1), enabling the array to rebuild to full redundancy, allowing more time to schedule drive replacement without increased data risk.

Redundant Supervisor Modules

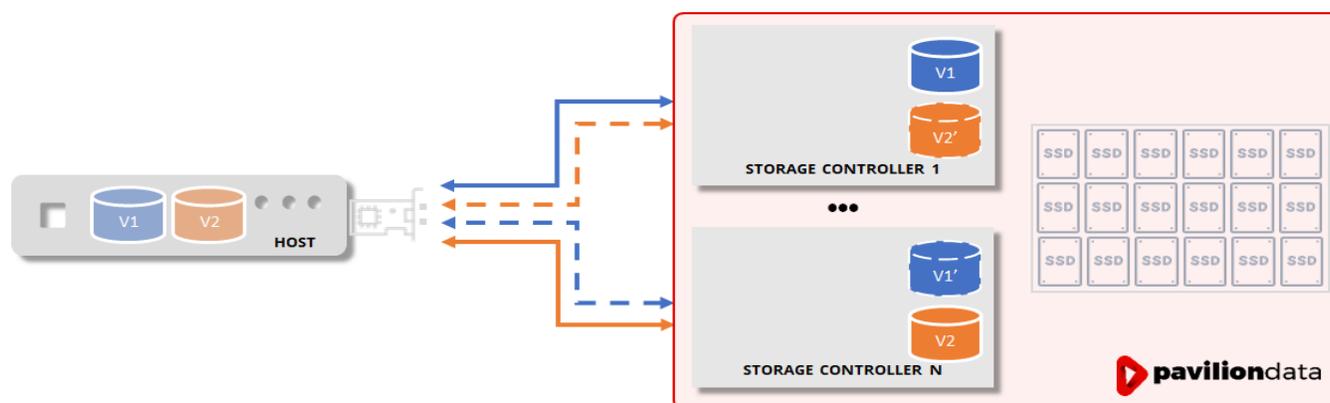
All components are managed by redundant out-of-band management controllers, or supervisors. Management of the array is done independently of the I/O controllers and the data paths, providing greater flexibility and consistent performance even during maintenance operations.

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.

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.



Performance and Safety with Pavilion's NVMe-of Storage Array

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.

