

NVMe-oF Storage for Aerospike

Disaggregated Storage with No Compromises for Aerospike Environments

Fraud prevention, recommendation engines, payment processing, and IoT are just a few applications where Aerospike is driving a new paradigm in IT infrastructure design. With a hybrid architecture that enables in-memory and Flash tiering, Aerospike can handle hundreds of terabytes of data and trillions of transactions per month with exceptionally low Total Cost of Ownership.

Today's Storage Architectures lack the flexibility to accommodate ever-changing requirements. Ideally compute, network, and storage resources should scale independently and deliver optimal NVMe storage utilization without compromising performance.

New Store Challenges

While it offers the flexibility of deploying distributed resources in a scale-out fashion, it's expensive when it comes to storage and storage management. Typically, storage is deployed as direct-attached SSDs in individual servers but this leads to significant problems that admins must deal with.

- Storage is not shared effectively since it is stranded in a single server and results in underutilization of NVMe; in some cases as low as 25%.
- Storage provisioning decisions are made at procurement time, meaning that determining the size of the storage in each server is done before the requirements of the application are known. This leads to inflexibility and higher costs over time.
- When scaling for either performance or capacity reasons, more server nodes need to be deployed to accommodate more direct-attached SSDs, effectively expanding the infrastructure unnecessarily.

And whilst some applications offer data protection mechanisms, it relies on making copies of data on other database server nodes, leading to more capacity being required and bloating cost of infrastructure even more.

Pavilion Hyperparallel Flash Array

The Pavilion Hyperparallel Flash Array (HFA) delivers never before seen NVMe performance and density that allows customers to provision logical flash storage resources over a low latency network. As a result, you can now deploy shared storage in place of direct-attached SSDs in cloud- scale Aerospike environments.

The Pavilion HFA requires no custom software to be installed on application servers and includes important data management and availability features, including thin provisioning, instant zero-space snapshots and clones, and no single point of failure.



Benefits

- Reduce TCO with a seamless hybrid DRAM and NVMe storage architecture
- Disaggregate physical resources to independently scale compute, storage and network resources
- Choose performance or endurance NVMe SSDs as the Flash tier of the hybrid architecture
- Deliver performance similar to Direct-Attached Storage in high-scale Aerospike environments

Pavilion Benefits

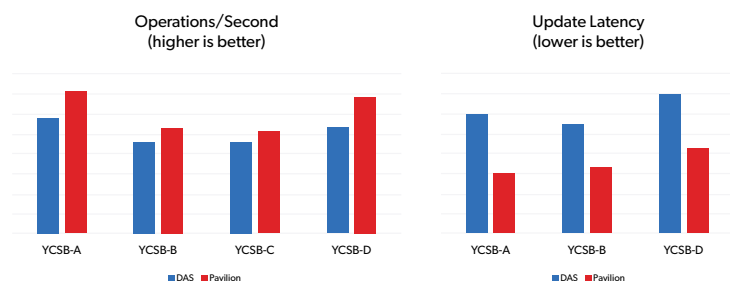
- Efficient block storage for Aerospike databases
- 40 μ s latency
- Up to 1.1 PB in 4U
- Frictionless deployment
- Data resiliency & high availability
- Up to 20 active-active storage controllers
- Multi-path I/O support
- Space-efficient instant snapshots and clones
- Thin provisioning
- Pay As You Grow scalability
- Standard Ethernet
- OpenChoice Storage™

Deliver Disaggregated NVMe-oF Storage for Aerospike Deployments:

Until now, Aerospike clusters were deployed on DAS SSDs because of the performance and fault isolation requirements of the environments. The absolute lowest latency requirement, in particular, drove the need for DAS due to its latency fulfillment.

However, with new high-speed RDMA-capable networking and efficient storage protocols like NVMe-oF, it is now possible to get the same performance advantages with shared storage.

The Pavilion Platform offers low-latency logical NVMe storage from a disaggregated 4RU appliance, allowing racks of database nodes to be supplied low-latency storage capacity from a central storage appliance. It delivers up to 90 GB/s bandwidth and 20 million 4K Read IOPS and provides the data management features that lower the cost of deploying Aerospike.



We can offer more than a petabyte of low-latency logical flash storage from a disaggregated 4U storage appliance. 10 Aerospike nodes can deliver up to 3.2 million 4K reads requests per second and 1.0 million 4K writes per second. In addition, the Pavilion HFA offers enterprise features that lower the cost of deploying large clustered applications:

Deploy up to 4X+ less flash:


By leveraging thin-provisioned logical flash storage, Pavilion HFA delivers the required needs of the applications. Decide at application deployment time how much storage to provision to any given node and longer be constrained by the size of the SSDs that was purchased and installed in any given server. Thin Provisioning allows the application to use the required amount of storage at any given time, regardless of how much capacity has been advertised to that specific database node. This greatly reduces the amount of raw flash storage deployed in these cloud-scale environments.

Simplify data protection and reduce server overhead:

Instant snapshots and clones allow an entire clustered database to be backed up or copied for test/dev purposes and on the fly without any performance impact. The Pavilion HFA provides no single point of failure, ensuring maximum application uptime and data availability. This removes the need for multiple copies of each node's data on other nodes, lowering the storage capacity requirements whilst reducing application and network processing overhead required to distribute that data to additional nodes.

Increased Compute Density per Rack by deploying Disk-less server nodes:

By provisioning high-speed logical flash storage volumes to each server in a rack, users no longer need to purchase servers that accommodate SSDs. This provides the ability to increase the compute density of a rack by leveraging 1U servers instead of 2U servers with front-loading drive bays. Pavilion also requires no custom software to be installed on database nodes, allowing Aerospike to take full advantage of the application host processing resources as well as simplifying deployment complexity.



NVMe-oF for Aerospike

The "DAS" configuration on the left consisted of 16 2U servers, each with 2x6.4TByte SSDs. Eliminating direct-attached SSDs allows us to use 1U servers, saving a large amount of acquisition costs, rack space, power and cooling. Through the use of thin provisioning and space-efficient snapshots, the total deployed capacity can also be reduced. Using the Pavilion HFA also reduced the operational costs, as a storage admin can manage approximately twice the SAN storage as DAS storage. Together, this reduces the operational costs of the Aerospike environment dramatically.