

NVMe-oF Storage for Cassandra

Accelerate Cassandra with no-compromises

Cloud, big-data analytics, mobile, and cloud-delivered applications using Apache Cassandra are driving a new paradigm in IT infrastructure design.

Resources need to be freed. They need to be available and deployed so that the ever-changing requirements can be satisfied on a minute-by-minute basis. This means that compute, network, and storage resources all need to scale independently to meet an ever-increasing and diverse set of application requirements.

Direct-Attached Storage 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's NVMe-oF Storage Array

Pavilion 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 Cassandra environments.

The platform 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

- Increases density 2X
- Petabyte scalability, high-performance, low-latency, and linear scaling maximizes data center efficiency
- Protect your data and your business. Meet evolving requirements for data security and compliance
- Search 20X faster than DAS
- Deployment flexibility using concurrent protocols
- Enterprise design and data integrity validation ensure reliable access to data
- Rack scale/CDI management via Web GUI, vCenter, Kubernetes, RESTful API, OpenStack, DTMF/Redfish, and Swordfish

Benefits of Disaggregation

The Pavilion HFA provides high performance and low latency to Cassandra applications.

Scalable & Flexible

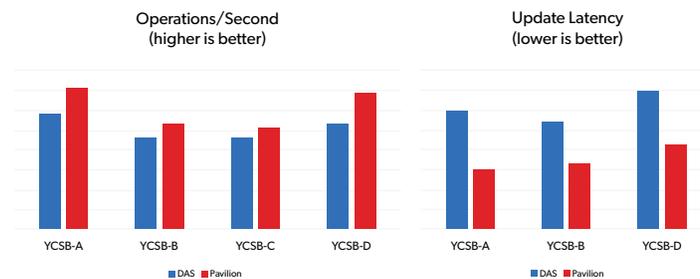


Provide up to 1.1 Petabytes, 90GB/s throughput, 20M IOPS, and 40µs of latency to Cassandra deployments and simultaneously use NVMe/RDMA (Ethernet, IB), NVMe/TCP, iSCSI, or NFS interfaces. Grow performance and capacity linearly without impacting on-going operations. Pavilion HFA's thin provisioning feature provides Cassandra with significantly better utilization than DAS, since the Pavilion array will only allocate physical space when the Cassandra application needs it. This results in physical space savings vs. DAS of up to 75% per server. No wasted capacity, no time spent messing with volume managers or file systems. Just set and forget.

Economical



The drivers behind deploying direct-attached SSDs as the primary storage in Cassandra environments were performance, scalability, fault isolation and agility. Many of the applications require the absolute lowest latency and flexibility to scale on demand, and thus the best performance available came from direct-attached SSDs installed in individual nodes. However, with the advent of high speed RDMA-capable networking and efficient block storage protocols like NVMe-oF, it is now possible to get the same performance advantages of direct-attached SSDs with flexible shared storage. Below is a performance comparison of direct-attached SSDs and a Pavilion HFA using the Yahoo! Cloud Services Benchmark (YCSB):



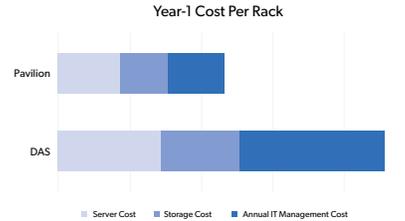
With Pavilion, you are no longer constrained by the size of the SSDs. Thin provisioning allows the application to use the required amount of storage at any given time, regardless of how much capacity has been advertised to that specific Cassandra server. No more extra copies of data, the Pavilion HFA reduces the amount of raw flash storage deployed in a Cassandra environment, by up to 3X, reducing TCO. At 50% less expensive than a DAS SSD in terms of %/GB/s, the Pavilion HFA is designed not just for exotic applications, but is the one platform for all your Cassandra environments.

Fast & Dense



The ultra-high performance, extreme low latency, and multiple storage controllers of the Pavilion HFA accelerates Cassandra workflows and boosts time to results. The Pavilion HFA lets Cassandra applications search terabytes of structured and unstructured data twice as fast as a DAS NVMe SSD.

The “DAS” configuration consisted of 16 2U servers, each with 2x6.4TB NVMe SSDs. Eliminating local SSDs allows customers to use 1U servers or smaller, providing up to 2-3X savings for rack space, power and cooling, increasing Cassandra’s density per rack. Not to mention the usage of thin provisioning and space-efficient snapshots greatly reduces the total raw capacity required for the same usable capacity.



The result? The power, simplicity, and density offered by the Pavilion HFA provides a flexible service for scale-out Cassandra deployments, increasing agility and flexibility and lowering TCO.

Safe & Secure



Protect the SSDs with Pavilion’s RAID-6 erasure coding and its 12% overhead. Its “Swarm Recovery” rebuilds a failed SSD ten times quicker than a DAS or an AFA. The Pavilion HFA ensures that the failure of an SSD does not impact Cassandra operations.

Keep data and snapshots safe with a FIPS-compliant data at rest encryption, the Pavilion HFA’s always-on encryption keeps data secure without impacting Cassandra performance.

The Pavilion Data Assurance feature works with the RAID feature to provide self-healing bit-rot support for data, assuring every Cassandra workflow gets uncorrupted data. Take consistent, instant, zero-footprint, and uncorrupted snapshots, encrypt them and provide them to backup and disaster recovery processes to speed operations and ensure consistency.

Enterprise Strength & Enterprise Support



Get 24/7 proactive support, end-to-end data integrity, a modular chassis, and redundancy throughout the storage array to protect your Cassandra applications as infrastructures grow. With the Pavilion HFA, all the features come in-the-box, including thin provisioning, snapshots, clones, data at rest encryption and more.

Manage the Pavilion HFA via its Web GUI or use the management framework of your choice, including: vCenter, Kubernetes, RESTful API, OpenStack, DTMF/ Redfish, Swordfish, and more.

Find Out More

Pavilion is defining the future of disaggregated NVMe-oF. Cassandra applications can be disrupted by the Pavilion HFA and its unprecedented availability, performance and versatility to future-proof the storage infrastructure.

Our expertise is in simplifying and optimizing NVMe to make the impossible, possible. When storage is business-critical, there’s no substitute for the guaranteed performance, functionality, high availability, and OpenChoice Storage™ support of a Pavilion NVMe-oF storage array. Use the Pavilion Hyperparallel Flash Array with its NVMe over fabrics support and make better decisions, faster! Contact us to learn more.