



NVME-OF STORAGE FOR CASSANDRA

Features

Cassandra Benefits

- Reduce the cost of deploying large cloud-scale Cassandra deployments in terms of both CAPEX and OPEX
- Disaggregate physical resources to independently scale compute, storage and network resources
- Simplify database provisioning, deployment and protection
- Deliver orders of magnitude better application performance in high-scale cloud, web, and SAAS environments

Pavilion Benefits

- 40 µs Latency
- 14TB - 1 PB in 4U
- Frictionless Deployment
- Data Resiliency & High Availability
- Up to 20 Active/Active Controllers
- Multi-path IO support
- Independently scale compute and storage.
- Space-Efficient, Instant Snapshots and Clones
- Thin Provisioning
- Standard Ethernet
- **OPENCHOICE** Storage

Deploy Disaggregated Storage and lower TCO in the process.

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.

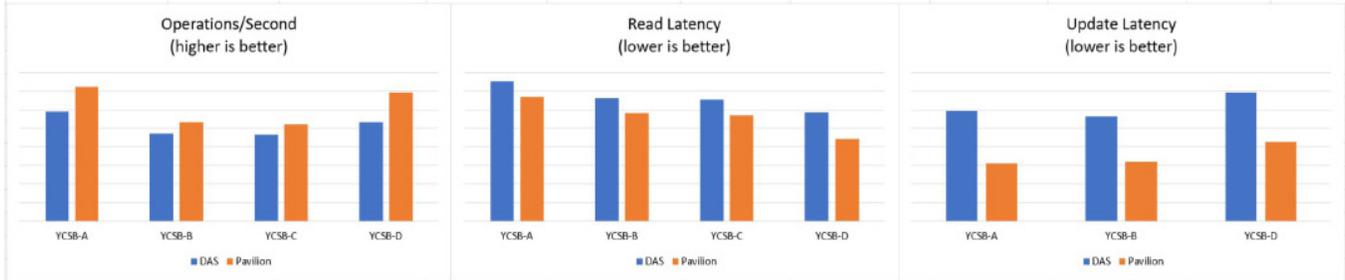
Pavilion's NVMe-of Storage Array

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.

Deliver Disaggregated NVMe-oF for Cassandra Deployments:

Up until now, 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 Memory Array using the Yahoo! Cloud Services Benchmark (YCSB):



Pavilion’s Storage platform provides the following key benefits to Cassandra deployments:

Deploy up to 4X+ less flash deployed:

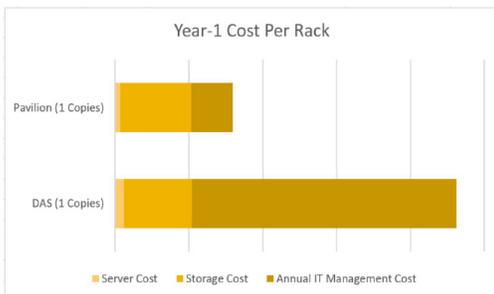
By leveraging thin-provisioned logical flash storage, our platform delivers the required needs of the applications. You can decide at application deployment time how much storage to provision to any given node and are no longer constrained by the size of the SSDs that were purchased and installed in 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. Our platform 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, you 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 Cassandra to take full advantage of the application host processing resources as well as simplifying deployment complexity.



The “DAS” configuration consisted of 16 2U servers, each with 2x6.4TB SSDs. Given the YCSB test demonstrated above, Pavilion allows you to reduce the number of servers. By eliminating servers and therefore SSDs allows customers to use 1U servers providing savings across rack space, power and cooling. Not to mention the use 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 NVMe-oF provides the first flexible service for scale-out Cassandra deployments, increasing agility and flexibility and lowering TCO in the process.