pavilion data

Solution Brief

Accelerating Apache Cassandra™

Solution Overview

Ultra-Low Latency Networked Storage For Today's Modern Cloud-Scale Applications

Key Benefits

- Disaggregate physical resources to independently scale compute, storage and network resources
- Low latency Flash storage provisioned to many servers from a single 4U appliance
- Reduce the cost of deploying large Cassandra databases in both CAPEX and OPEX
- Deliver orders of magnitude better application performance in high-scale Cassandra environments

Pavilion Benefits

- Fastest block storage for Cassandra databases
- Latency of direct-attached SSDs
- Up to 460 TB in 4U
- Frictionless deployment
- Data resiliency & high availability
- Up to 20 active-active storage controllers
- Multi-path IO support
- Space-efficient instant snapshots and clones
- Thin provisioning
- Pay As You Grow scalability
- Expand for capacity or performance, independently
- Increase storage utilization up to 10x or more

IT Infrastructure Demands are Changing

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 flexibly deployed so that ever-changing requirements can be satisfied on a day-by-day or even minute-by-minute basis. Compute, network, and storage resources all need to be able to scale independently to meet an ever-increasing and diverse set of application requirements.

This shift is creating a demand for disaggregated storage resources that can scale to new levels, where performance tuning can be eliminated so any workload can be satisfied at any time.

The Pavilion Memory Array

The Pavilion array 4U appliance delivers performance and density at a level that allows multiple racks of application servers to be provisioned with logical Flash storage resources over a low latency network. As a result, it is now possible to deploy shared storage in place of direct-attached SSDs in cloud-scale Cassandra environments.

Besides performance and capacity, the Pavilion array offers several important data management and enterprise availability features, including thin provisioning, instant zero-space snapshots and clones, and no single point of failure in an enterprise class chassis and platform. Active-Active controllers and support for multi-path IO ensures no single point of failure out to the host. Finally, the Pavilion array requires no proprietary software to be installed on application servers, freeing up host resources for applications and eliminating deployment complexity.

New Storage Challenges

In order to satisfy the performance requirements in today's cloud-scale environments, low latency storage is deployed as direct-attached SSDs in individual nodes of a high-scale latency-sensitive Cassandra environment. This leads to several problems that administrators must deal with. The storage cannot be shared effectively since it is stranded in a given server, so it tends to be under-utilized. In many cases, storage utilization is as low as 25%. In addition, storage provisioning decisions need to be made at server procurement time, meaning that determining the size of the storage in each server is done up front when hardware is purchased. This leads to inflexibility and higher costs. In addition, if storage needs to be scaled for either performance or capacity reasons, more server nodes may need to be deployed to accommodate more direct-attached SSDs. Finally, while some of the applications deployed offer robust data protection mechanisms, they rely on making copies of data on other server nodes, leading to more capacity being required and driving up costs.

Lower TCO With Low-Latency Networked Storage Using Pavilion and Cassandra

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 Over Fabric, 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):

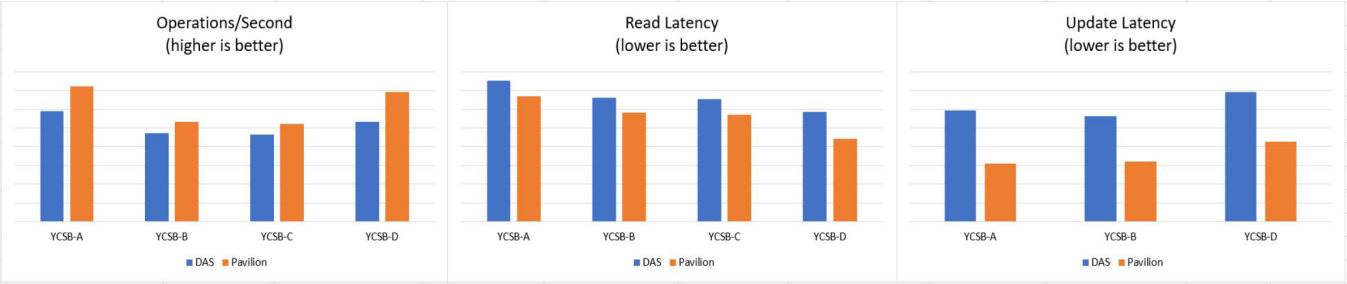

The Pavilion array can offer 100s of Terabytes of low-latency logical flash storage from a disaggregated 4U storage appliance. Racks of Cassandra nodes can be supplied low-latency storage capacity from a central storage appliance that can deliver up to 120 GB/second bandwidth and 20 million 4K read IOPS. In addition, the Pavilion array offers important data management features that lower the cost of deploying large clustered applications significantly.

Up to 4X+ less flash deployed in most environments

With the Pavilion array, you decide at deployment time how much storage to provision to each Cassandra instance. You are no longer constrained by the size of the SSDs that were purchased and installed in any given server. With thin provisioning, you only use what the database actually requires, regardless of how much capacity has been advertised to that specific Cassandra environment. This will greatly reduce the amount of raw Flash storage required to deploy Cassandra databases in these cloud-scale environments.

Simplify data protection and reduce server overhead

Space-efficient snapshots and clones allow an entire database to be backed up, or copied for test/dev purposes instantly without any performance impact. In addition, the Pavilion array is an enterprise platform with no single point of failure, ensuring maximum application uptime and data availability. This alleviates the need to distribute copies of each node's data to other nodes, lowering storage capacity requirements and reducing application and network processing overhead required to distribute that data.

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 can accommodate SSDs. This allows you to increase the compute density of a rack by leveraging 1U servers instead of 2U servers with front-loading drive bays. In addition, the Pavilion array requires no custom software to be installed on server nodes, allowing applications to take full advantage of host processing resources as well as simplifying deployment complexity.

All of the above combines to reduce your costs of storage and computing

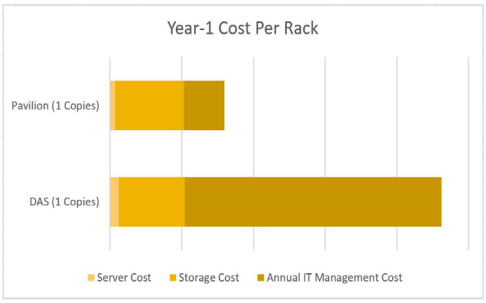

The "DAS" configuration consisted of 16 2U servers, each with 2x6.4TByte SSDs. The above YCSB test demonstrated that using Pavilion allows us to reduce the number of servers. Eliminating direct-attached SSDs allows us to use 1U servers, saving a large amount of acquisition costs, rack space and power and cooling. Through the use of thin provisioning and space-efficient snapshots, the total capacity can be reduced for the same usable capacity.

Using the Pavilion array 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 Cassandra environment dramatically.

As a result of the power, simplicity, and density offered by the Pavilion array, it is now possible to deploy flash as a flexible service for scale-out application deployments, allowing for cloud-scale agility and flexibility, and lowering TCO in your Cassandra environment.

pavilion data

Pavilion Data Systems, Inc.
2560 N First St., Suite 220, San Jose, CA 95131
E-mail: sales@paviliondata.com