PAVILION

Statistics Netherlands

Customer

- Statistics Netherlands/CBS

Industry

- Government

Applications/Use Case

- Cloud-scale applications using Greenplum VMs for self-service

Business Problem

- Data growth & ingestion time
- Aging infrastructure and bloat
- Migration of Greenplum VMs to K8s containers
- Encryption of all data
- Non-invasive backup

Solution

- 2 Pavilion HFAs that are VMware certified
- Moved from hyperconverged to disaggregated performance storage and commodity servers
- Consolidated thousands of VMs running Greenplum into a seamless fabric
- Infrastructure designed to move forward with stateful K8s
- Early adopter of VMware NVMe-oF with proven results

Containers and VMs at Statistics Netherlands (CBS)

CBS or Statistics Netherlands, founded in 1899, is a Dutch governmental institution that gathers statistical information about the Netherlands and its populace. In Dutch, it is known as the Centraal Bureau voor de Statistiek (Central Agency for Statistics), often abbreviated to CBS. It is located in The Hague and Heerlen. CBS is responsible for collecting, processing, and publishing statistical data on the Netherlands population, economy, and society. The United States equivalent is the United States Census Bureau.

CBS's infrastructure was aging, creating significant challenges and infrastructure bloat that had resulted in IT users no longer being able to ingest, digest and serve the data at the speed and scale that the government required.

Legacy storage arrays with a dual-controller architecture restricted the ability to accommodate CBS's growth of data and optimize time to insight. They needed a standards-based (no proprietary hardware or software) high-performance infrastructure with linear scaling of CPU, Memory, and Storage. This solution must also support Kubernetes (K8s), vSphere 6.7, and be certified by VMware with plans to support VMware's proposed NVMe-oF driver for vSphere 7.

Planning migration of some VMs to K8s-based persistent containers, CBS needed support for K8s and the Container Storage Interface (CSI). They were also looking for a solution that solved the performance and scalability needs of today and was future-proofed for tomorrow's infrastructure.

CBS's data center architect, Kasper von Beem, and lead infrastructure product manager, Ian Schets, knew that their company needed to a new storage array that could better meet CBS' growing requirements of today and was future-proofed for tomorrow. They looked into various storage technologies but had difficulty finding a storage array that met all of CBS's needs.

Kasper and Ian looked at HCI but disqualified it because of the inability to meet the building blocks approach to infrastructure that CBS had in place. (linear scaling of CPU, memory, and storage). They also disqualified an 'NVMe-Based' or '100% SAN' storage array because it did not deliver the performance at a density that was required for the infrastructure.

Composed By Pavilion

Kasper and Ian reached out to Pavilion and worked with the sales team to design a solution that would fit their performance needs and scale easily with their plans for future growth. They found the Pavilion HFA's use of NVMe-oF gave them the same parallel performance as local NVMe SSDs. The Pavilion HFA also provides advanced SAN-like features including thin-provisioning, RAID support, snapshots, framework integration, encryption, performance monitoring, multi-pathing, and 24/7 proactive cloud-based support.

CBS found that the Pavilion HFA delivers up to 120 GB/s throughput, 40µs of latency, and 1.1 PB of storage in a compact 4U form factor. It is the industry's first

Business Value

- 4X reduction in data center footprint
- 70% OPEX reduction due to lower IT infrastructure costs (such as power and cooling)
- High-performance, low-latency, linear scaling, and petabyte scalability maximizes the efficiency of NVMe SSDs and cost 25X less (by $/IOPS) the incumbent vendor's proposal
- Support for persistent K8s containers
- Future-proofed with built-in headroom for compute and storage growth

You can't predict the future, so overkill is underrated.

Kasper von Beem, Data Center Architect — Statistics Netherlands (CBS)

Pavilion^{OS} is the perfect feature set. The bells and whistles of HCI bog down NVMe performance.

Ian Schets, Product Manager, Infrastructure — Statistics Netherlands (CBS)

hyperparallel flash array that unlocks the parallel performance of NVMe. CBS was a beta tester of vSphere 7, and Pavilion is one of the first vendors to support the VMware NVMe-oF driver included with vSphere 7.

Pavilion^{OS} allows CBS to make clones directly from snapshots, accelerating backup operations and easing backup completion headaches. CBS also found that getting equivalent performance from the incumbant's newest NVMe technology cost 25 times more than the Pavilion HFA solution. They also found that the two Pavilion HFAs they required took 90% less rack space than the configuration proposed by the incumbant.

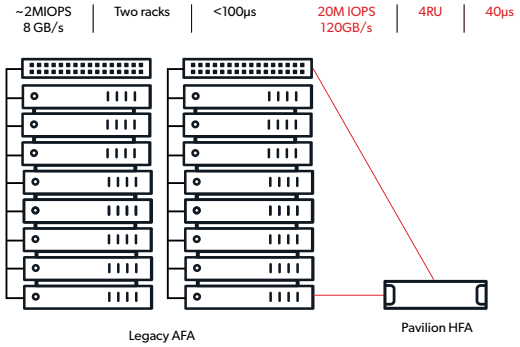

Pavilion was selected over the latest NVMe technology from the incumbent supplier. To meet the infrastructure requirements, the incumbent proposed 16 x All-Flash Arrays, while Pavilion proposed two HFAs. The Pavilion HFA was selected because it delivered a standards-based NVMe-oF in a simple, dense, and high-performance storage solution.

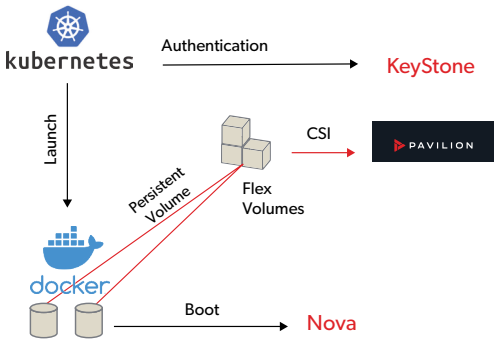

The Pavilion HFA gave CBS the bandwidth needed to process assets from additional sources, improving their production workflow. They also had an initiative to deploy services using containers, so the Pavilion HFA's CSI support of persistent K8s containers was also crucial in their decision. This would enable them to transition their infrastructure from old to new, automagically migrating VM's to containers with no impact to infrastructure configuration and no drama for IT.

The CBS team is now able to serve data with more speed and at a never-before-seen scale, delivering more accurate data on social issues based on reliable statistical information.

Summary

Pavilion worked closely with CBS to assure a smooth transition from their legacy environment and provides them with high performance, low latency, and gave them a linear price and capacity scalability. They found that the Pavilion HFA is ideal for

operations while its OpenChoice Storage™ feature future-proofed their storage solution and saved them production cost. It lets them redeploy existing or purchase new NVMe SSDs as their expansion needs arise.

If CBS has a problem down the road, Pavilion has a specialized team of support technicians waiting to help resolve any issues that may arise. The Pavilion HFA's VMware certification and support for persistent K8s clusters allowed CBS to sleep well at night.

Find Out More

The Pavilion HFA is defining the future of composable disaggregated NVMe-oF. 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 professional software support of a Pavilion HFA. We design, implement, and deliver a complete solution tailored to the environment. Contact us today to get in touch with our talented extended teams of professionals.

©2019 Pavilion Data Systems, Inc. All rights reserved.