



## Facial Recognition

Facial recognition is a game-changer for government agencies concerned with safety, improvement of mass-transit security and citizen privacy. IBM Spectrum Scale (also known as GPFS) is an ideal platform for this type of big data analytics where integration with social media, geo-location and other real-time feeds deliver unprecedented accuracy and agility.

### Customer

- Government Agency via a Strategic Partner

### Industry

- Government

### Use Case

- AI/ML based analytics and facial recognition

### Challenge

- Data growth
- Limited rack space
- Encryption of all data
- Non-invasive backup

### Solution

- 2 Pavilion HFAs providing over 180 GB/s b/w
- Snapshots provide instant & space-saving copies of data for backup

### Results

- 10X lower TCO (CapEX/OpEX) than alternative
- 2X higher compute density than alternative
- Linear scalability accommodates growth

This client is a government agency that was looking for a storage solution for the world's fastest 512 node GPFS cluster to support real-time analytics and facial recognition. The agency demands 100% of native NVMe performance, with ½ the footprint of DAS with NVMe allows for 2x density compared to their current configuration. They expect scalability to double that of traditional DAS NVMe. They were looking for a solution that had four features:

- Encryption at rest
- On-demand reconfiguration to accommodate an unknown quantity of daily data collection and analytics
- Thin provisioning to assure ample capacity for unknown data growth
- Snapshots for zero-downtime backup

### Composed By Pavilion Powered By Spectrum Scale

The client looked at installing NVMe SSDs into their existing SAN, but found that doing so did not provide unlock the parallel performance provided by an NVMe SSD. Since the client wanted to build a 512-node GPFS cluster, Pavilion setup a GPFS environment and demonstrated to the customer that the Pavilion Hyperparallel Flash Array (HFA) combined with GPFS provided the performance their applications need.

They found that the Pavilion HFA's use of NVMe-oF gave them the same parallel performance as local NVMe SSDs. The Pavilion HFA also gave them advanced SAN-like features including thin-provisioning, snapshots/clones, framework integration, encryption, performance monitoring, multi-pathing, and 24/7 proactive cloud-based support.

A single 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 hyperparallel flash array that unlocks the parallel performance of NVMe.

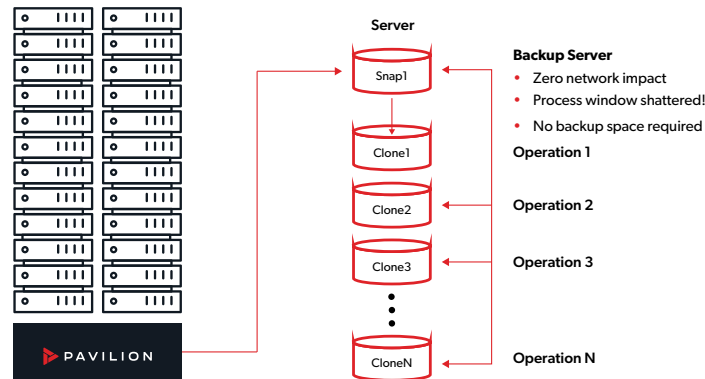
Combining the Pavilion HFA with Spectrum Scale enabled the customer scale out in a linear fashion with a single shared namespace to optimize the ingest and retrieval of large and small files as well as live stream input sources. It provides them with ultra-high performance, for analytics that drive real-time results with cost-effective technology.

The ability to make clones directly from snapshots accelerated operations and eased completion headaches. The client found that expanding their DAS for equivalent performance would cost 25 times more than the Pavilion HFA, they also found that the Pavilion HFA used 95% less rack space. So the customer replaced their DAS with the Pavilion HFA. It gave them the bandwidth needed to process assets from additional sources, improving their overall productivity.

Today, it is not feasible to construct a 512 node GPFS cluster that has the density, performance and manageability required by [this customer]. Pavilion is the only solution to achieve our performance goals without compromise of space, power or latency.

### Solution Consultant

The client used the Pavilion HFA's zero-footprint snapshots and clones to make multiple instant and transparent copies of real-time analytics and export these to other teams and clusters for further their insights. They also export the snapshots with their existing backup solution and presented clones of the snapshot to minimize traffic within the primary cluster.



### Summary

Pavilion worked closely with the client to assure a smooth transition from their legacy SAN environment and provides them with high performance, low latency, and gave them a linear price and capacity scalability. The client found that the Pavilion HFA is ideal for big data analytics. Its OpenChoice feature future-proofed their storage solution and saved them production cost by letting them redeploy existing or purchase new NVMe SSDs as their expansion needs arise.

Operations staff also found that the Pavilion HFA's SAN-like management for snapshots, clones, and thin provisioning aligned to the processes they had implemented for the SAN.

### Find Out More

The Pavilion HFA is defining the future of composable disaggregated NVMe-oF. Our system is an ideal part of a complete Media and Entertainment workflow. Our expertise is in simplifying and optimizing NVMe to make the impossible, possible. Taking the infrastructure to the next level, requires the "Midas touch" of local, proven experts. 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 partner with leading organizations to design, implement and deliver a complete solution tailored to the environment. Contact us today and get in touch with our talented extended teams of professionals.