



## Facial Recognition V2.0

### 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

### Solution

- 2 Pavilion HyperParallel Data Platforms providing over 150 GB/s b/w

### Results

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

*We thought it was not feasible to construct a GPFS cluster with thousands of nodes that has the density, performance, latency and manageability required by [this customer]. Pavilion is the only solution to achieve these expectations .*

#### Solution Consultant

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.

This client is a government agency that was looking for a storage solution for the world's fastest GPFS cluster with thousands of nodes to support real-time analytics and facial recognition. The agency demanded high throughput (>100GB/s), high performance (>10M IOPS), and low latency (<100 μs) Pavilion was the only solution that could meet these requirements, and actually exceeded them.

### 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 unlock the parallel performance provided by an NVMe SSD. Since the client wanted to build a several thousand-node GPFS cluster, Pavilion setup a GPFS environment and demonstrated that the Pavilion HyperParallel Data Platform™ combined with GPFS provided the performance their applications need.

The Pavilion HyperParallel Data Platform's use of NVMe-oF gave the client the same parallel performance as local NVMe SSDs. The Pavilion HyperParallel Data Platform 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.

The Pavilion HyperParallel Data Platform delivers up to 120 GB/s throughput, 25μs of latency, and 2.2 PB of storage in each compact 4U form factor. It is the most performant, dense, scalable, and flexible data storage platform in the industry.

Combining the Pavilion HyperParallel Data Platform with GPFS enabled the customer to optimize the ingest and retrieval of large and small files. It provides them with ultra-high performance, consistent video playback, improved storage utilization, enhanced global collaboration, and reduced production costs.

### Summary

Pavilion worked closely with the client to extend its initial footprint to support several hundred nodes, and that provided them with high performance, low latency, and linear price and capacity scalability. The client found that the Pavilion HyperParallel Data Platform is ideal for their operations.

### Find Out More

Pavilion shatters customer expectations and resulting organizational outcomes by revolutionizing data processing for modern AI/ML, HPC, Analytics, Enterprise Edge and other data-driven applications. The Pavilion Hyperparallel Data Platform, powered by Pavilion HyperOS, delivers unmatched performance and density, ultra-low latency, unlimited scalability and flexibility, providing customers unprecedented choice and control. Learn why Fortune 500 companies and federal government agencies choose Pavilion. Visit [www.pavilion.io](http://www.pavilion.io) or follow the company twitter at <https://twitter.com/PavilionData>