

## Benefits

- Leverage PB-scale GPU technology to visualize network resources and threat anomalies
- Share, aggregate, and automate GPUs to maximize efficiency and Investment
- Universally unmatched NVMe-oF and RoCE to reduce costs and increase time-to-insights

## Features

- Hybrid cloud storage for block, file, and object (S3)
- Unmatched performance, density, scalability, and flexibility
- Extreme read performance of up to 120GB/s
- Unrivaled write performance of up to 90GB/s
- 20M IOPs at less than 25µs of latency as measured from the host
- Proven with [OmniSci®](#) and validated by [NVIDIA®](#)

# Geospatial Intelligence with Pavilion, OmniSci and NVIDIA

## Government agencies and telecommunications firms realize superior analytics and visualization at massive scale and affordable investment levels

With an ever-increasing volume, velocity, and variety of data from sources like satellites, sensors, drones, IoT, and telematics, the ingest, processing, and rendering of spatiotemporal data quickly surpasses the capacity of local storage for CPU and GPU powered servers, such as NVIDIA DGX™ systems. This is especially true when months and years (not hours) of data are needed for visualization and real-time response.

To circumvent the limitations of local DGX storage, customers often filter or statistically sample data reducing the analytics fidelity. Innovative solutions are needed to enable precise insights at the speed of curiosity.

Pavilion and OmniSci address this problem with a new database architecture and storage infrastructure that delivers 100x faster queries, server-side rendering, and speed of thought visualization with supercomputing speed and scale at an affordable cost.

OmniSci leverages the power of GPUs to render interactive visualization of massive geospatial data sets with millisecond latency. Pavilion enables this with the performance of internal storage, but at multi-petabyte scale using proven NVMe-oF and RDMA (RoCE) networking.

The solution addresses network reliability and design optimization for telecommunications firms, customer experience, and capital planning.

Pavilion and OmniSci enable information dominance in the battlespace, geospatial intelligence, healthcare fraud, waste, and abuse for government agencies.

### How it works:

Pavilion storage is the data ingest engine for thousands of data sources. It has unmatched storage capacity, read, and write performance with extremely low latency connected by RoCE to ingest sources, servers performing heavy query and analytics, GPUs for rendering, and S3 object stores (cloud or on-prem) for collaboration. OmniSci manages the workflows, processing, and rendering in real-time with unprecedented analytics and visualization.



Ingest batch, stream, real-time



Process



Render

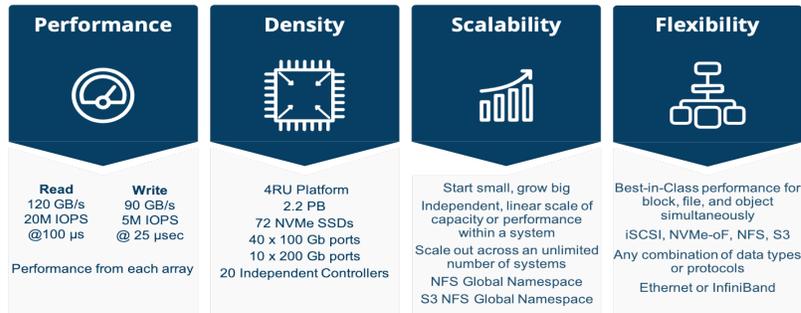


Collaborate



## The Pavilion HyperParallel Data Platform

The Pavilion HyperParallel Data Platform™, powered by Pavilion HyperOS™, is ideal for large-scale analytics and visualization. With up to 20 storage controllers, 2.2PB of NVMe capacity in just four rack units, this revolutionary storage system offers the most performant, dense, scalable, and flexible storage data platform in the universe.



### Unmatched Performance

In HPC analytics, storage performance is critical for both read and write operations. Pavilion delivers exceptionally high performance for both. Other vendors show acceptable read performance by adding expensive DRAM to their systems. This can improve sequential read performance but does little for writes. As a result, many vendors do not publish write speeds, or if they do, it is usually far less than their read performance.

To analyze and visualize insights in real-time across petabytes of data, OmniSci and Pavilion leverage the parallel processing capabilities of GPU technology from firms like NVIDIA. This can be in the form of aggregated GPUs using technology from VMware™ or using dedicated NVIDIA DGX systems.

A new technology from NVIDIA is Magnum IO GPUDirect™ Storage. Enabled by RDMA and RoCE, this technology is ideally suited for Pavilion HyperParallel Storage. GPUDirect eliminates the CPU from the data path, thereby feeding hungry GPUs faster.

In testing with NVIDIA's benchmarking tool [gdsio](#), Pavilion's [performance](#) dominates the competition with read bandwidth >110% higher than competitors, write bandwidth >169% higher than competitors, and latencies reduced by as much as 73%. All of this with 40% to 67% fewer rack units.

In short, Pavilion has the [pole position](#) according to *Blocks & Files* and others in GPU and GPUDirect storage performance.

### Unrivaled Scalability

The unique network-based architecture of the Pavilion HyperParallel Data Platform allows customers to add capacity or processing power independently of each other. Start with as few as 18 drives and then scale capacity up to 72 drives in each 4RU system for a total of up to 2.2PB.

Scale processing power and accelerate throughput and IOPS by increasing the number of controllers. Each controller is entirely independent and comes with its own processor, memory, networking, and OS instance. Begin with as few as four controllers and increase throughput and IOPS as needed by adding more controllers, up to a total of 20 in each system.

With Pavilion HyperOS™ 3.0, it is now possible to [scale-up, and scale-out](#) in a linear fashion across Pavilion systems using high-speed interconnects like 200Gb/s InfiniBand or Ethernet.

### Universal Flexibility

In an HPC big data environment (like OmniSci), block storage protocols are used to read/write data from/to OmniSci's SQL engine. In contrast, the data lake and collaborate functions are optimized on the Pavilion HyperParallel File System™ and S3 object store, respectively.

Most traditional storage arrays support either block or file and object data. None support the [flexible combination](#) of block, file, and object in a single system with high performance simultaneously.

### About Pavilion

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 on [LinkedIn](#).