

Benefits

- Simultaneous SMB data ingest with NFS processing and file sharing
- Superior scale and uptime for non-stop discovery
- NFS performance of up to 75GB/sec Read and 50GB/sec Write in a 4 RU footprint
- Native support 3rd party file systems including Spectrum Scale, Lustre, BeeGFS with unmatched performance and density
- Superior GPU and GPU Direct solution

Features

- Hybrid cloud storage with simultaneous block, file, and object (S3) protocols
- Ultra-low latency for uninterrupted data ingest, processing and sharing
- Capacity to handle multiple petabytes in a highly dense 4U footprint
- Scale with near-linear performance across chassis as demands increase
- Client connectivity for NFS v3, v4, NFS RDMA, pNFS, Gluster, Spark, Hadoop

Cryo-EM Workflows at Scale with Pavilion HyperParallel Storage

Securely Accelerate Discoveries with Unmatched Performance and Flexibility

An essential part of drug discovery is the ability to accurately depict the intricate interactions between a drug and the receptor. Cryo-EM is well-suited for providing vital information about molecular binding sites that influence the discovery of vaccines and therapies. Many pharmaceutical research organizations now rely on cryo-EM to accelerate structure-based drug design.

Consequently, cryo-EM has seen rapid advances in microscopes, detectors, and algorithms, with a singular goal of resolving near-atomic resolution of single protein particles. Accelerating sub 3 Angstrom resolution reconstructions routinely requires larger GPU/CPU resources to process hundreds of thousands of raw particle images.

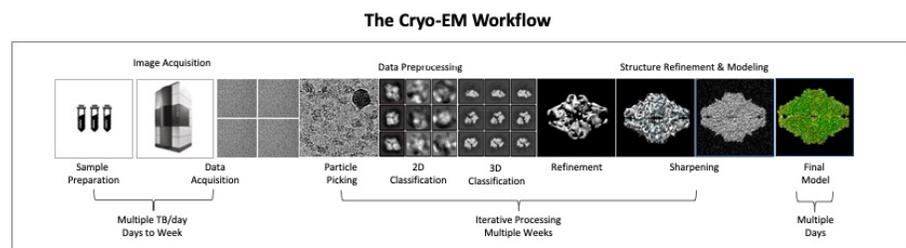
Cryo-EM pipelines, beginning at data acquisition and followed by complex data processing stages, continuously pushes hardware and software requirements to unprecedented levels.

This solution brief highlights how Pavilion Data Systems uniquely meets the challenges associated with ingest and accessibility of raw data, performance, and availability of data for 3D construction, with a keen eye on security and optimization throughout the process.

The Cryo-EM Workflow

The cryo-EM workflow can be broken into three distinct and critical parts:

- 1) Sample preparation and data acquisition
- 2) Data pre-processing
- 3) Image analysis



Sample Preparation and Data Acquisition

Typically stretching over multiple days to a week, this stage of the workflow is dominated by advancements in direct electron detectors and their ever-increasing capabilities. These detectors now routinely capture movies instead of single integrated images. This technique produces multiple (30-100) frames per imaging area.

The fields of view of these detectors are becoming increasingly larger (8k by 8k is already available). Consequently, individual frame sizes are substantially larger than they were a couple of years ago.

A high degree of automation is now possible in this stage of raw data creation, and researchers can now acquire over 1,000 movies in a day.

Using high-resolution, movie-generating cameras to continuously capture images at higher frame rates requires a highly performant, available, and scalable data repository to effectively meet multiple research projects' needs. Researcher Challenge: Cryo-EM is the fastest growing research area at many organizations. There is a vast demand for microscope times. Lack of performance storage and data services can halt research activities. More instantaneous transformation of raw data from movies to micrographs and accurately storing all associated metadata enables quick identification of errors and issues with sample preparation, instrument calibration, and operation.

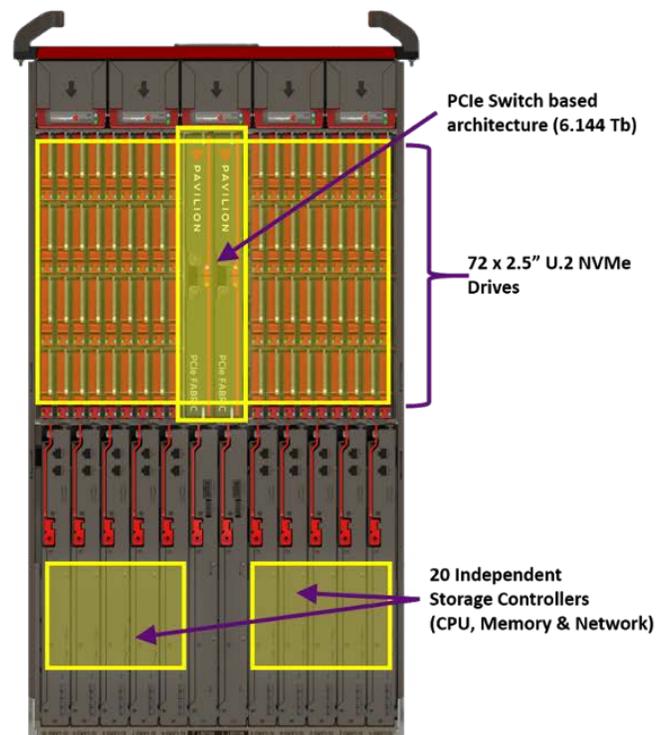
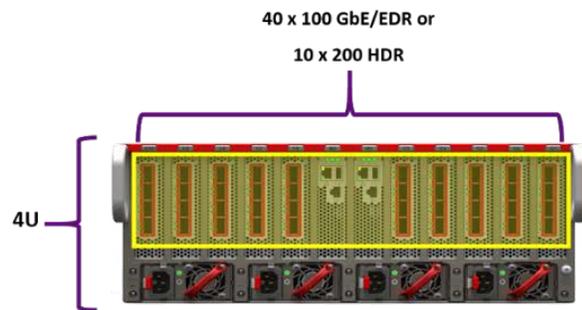
Pavilion's unique architecture and HyperParallel file system allows it to ingest from multiple microscope cameras simultaneously via SMB while being able to share the data via SMB or NFS seamlessly for visualization, processing and analysis. Also, with Pavilion's density allowing up to 2 PB of capacity in 4RU of space, experiments can run longer or use more cameras to capture the angles needed without having to stop to move data. Also, Pavilion's PHFS allows seamless automated data migration of older data to any standard NFS tier for archiving and longer-term retention, freeing up the NVMe capacity for more experiments.

With unmatched data ingest speeds of 90GB/sec write for block data and 56GB/sec write for file data Pavilion is the ideal platform for rapid movement from microscopes to shared file systems for researchers. As data is processed and analytics performed, Pavilion offers up to 2PB of capacity and read speeds of up to 120GB/sec block and 90GB/sec file access. With up to 20 storage controllers, Pavilion offers RAID-6 along with high availability controller fail-over to avoid any unnecessary downtime. Pavilion SwarmController™ capabilities can harness multiple controllers in unison to rebuild a failed drive faster than any alternative on the market.

Pavilion HyperParallel Flash Array

Unique Architecture:

- Built like a network switch – PCIe Backplane
- Built from the ground up to support NVMe and NVMe-oF
- Multi/Many Controllers
- Any controller to Any Drive Connectivity
- Cache-less and Tier-less
- DMA from any drive and any controller (nanoseconds)
- RDMA from hosts to Pavilion (microseconds)
- Compact form factor



Data Pre-Processing

This essential first phase of data processing also spans multiple days to weeks. Motion Correction is the first step, right after data acquisition. Any single or multi-particle data processing involves movie frame alignment that converts each 2-GB movie into a single 50MB integrated frame. With installations capable of producing multiple TB of data/day, this crucial first step relies on GPUs for faster completion. The resulting micrographs are foundational to all subsequent processing.

CTF Estimation, to adjust for contrast variations in the aligned micrographs, followed by particle picking and sorting are the essential next steps in data pre-processing. Once particles are picked within a micrograph, they are excised into individual, smaller images to form a particle stack. Stacks of one to several million particles are common, resulting in more intermediate data rapidly produced by multi-core servers deployed in this stage. New ML/AI-based techniques are being deployed to automate particle picking to reduce or eliminate manual intervention. The goal is to accelerate this manual process, make it more iterative, and improve accuracy.

Researcher Challenge: Researchers desire to reduce time to get to actionable insight incrementally. Researchers are often forced with the need to copy (or move) large amounts of raw data before data pre-processing can begin. Researchers need an in-situ approach to data processing instead of being burdened by data movement processes. This is harder in a multi-research environment, with multiple concurrent projects underway.

Pavilion empowers researchers with a scalable, centralized shared file system that dramatically reduces time to insights. This file system provides direct client access for SMB, NFS v3, v4, NFS RDMA, pNFS, Gluster, Spark, and Hadoop across Ethernet or InfiniBand connections. With centralized, shared storage, researchers do not need to make multiple copies of data or move files between storage systems. Scaling from a single 4RU chassis and up to 2.2PB of capacity across an unlimited number of systems is linear, such that performance is not impacted. In addition, the shared global namespace spans multiple chassis to allow growth without impact to operations.

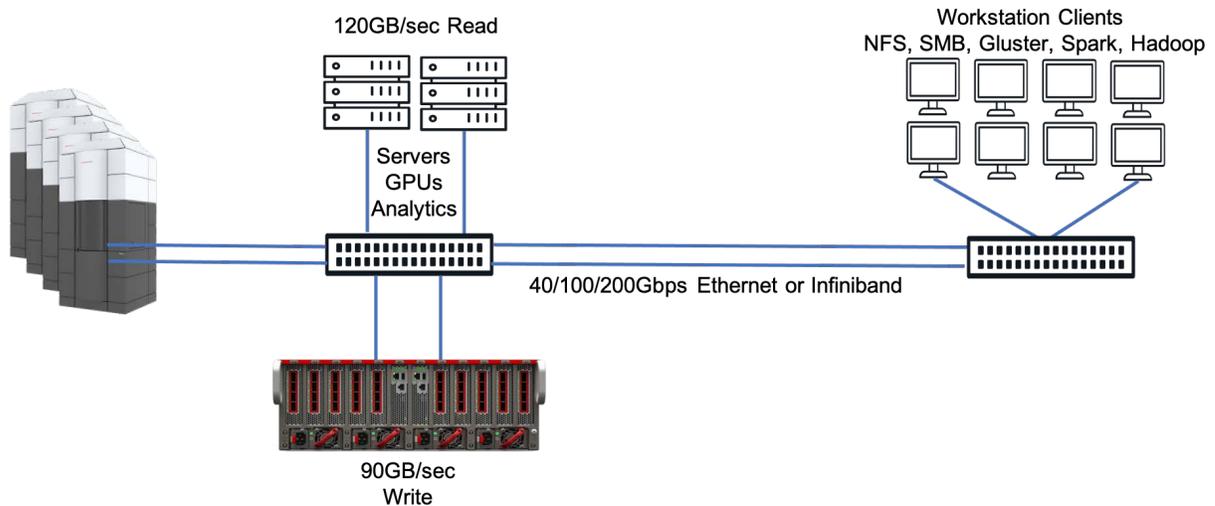


Image Analysis

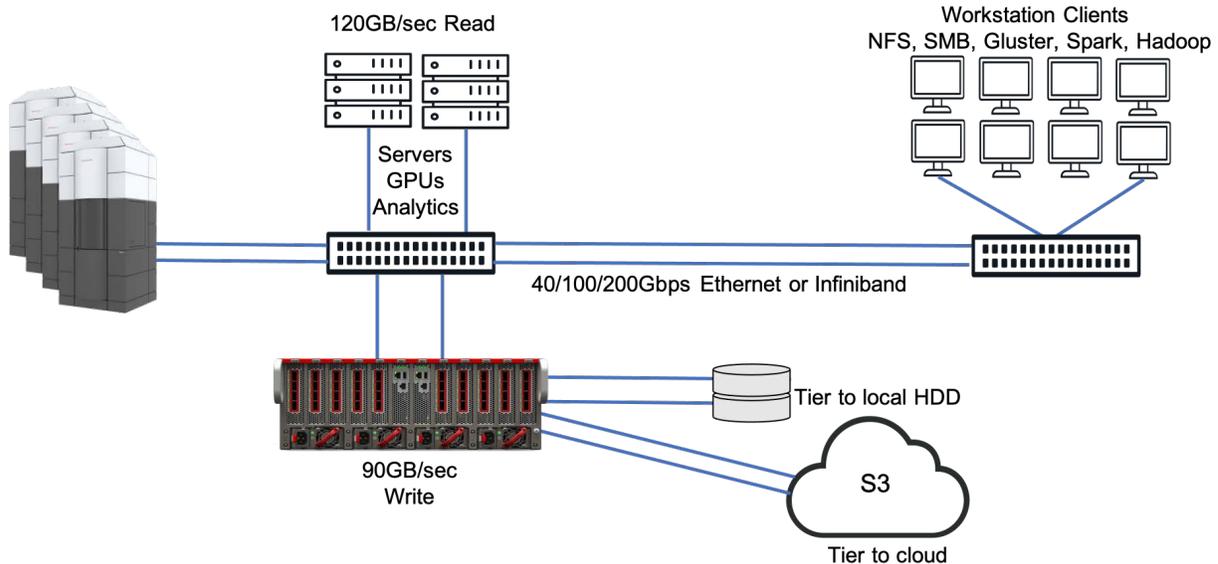
2D and 3D classification form the bulk of the processing-intensive stages of the workflow. With the addition of image refinement and sharpening processes, this is the most iterative and time-consuming part of the workflow and can easily stretch over multiple weeks. Traditional HPC-based compute complexes and newer GPU-based clusters are equally deployed to meet computational requirements.

Researchers are constantly interested in the latest open source and commercial software tools and platforms to improve accuracy and automate individual steps. Reliance on ML/AI is increasing by the day. There is a growing need to process data using multiple tools and software versions.

High throughput data access is paramount in these stages to ensure compute needs are not I/O bound. Stages of structural modeling, refinement, and visualization also need low latency data access.

Researcher Challenge: There is a lot of intermediate data produced during the various stages of a workflow that does not need to be stored on primary data storage systems. Automated mechanisms to offload data to more cost-effective storage tiers are a key optimization opportunity to reduce cost. The availability of tiered data for recall and iterative processing is also desired. In addition, data protection (Backup, DR, Deletion protection) becomes crucial.

The Pavilion HyperParallel Data Platform delivers the highest performance, lowest latency solution on the market for primary storage. As data goes through the stages of a workflow it can be easily and programmatically offloaded to lower cost storage tiers and public, private, or hybrid S3 object stores. Pavilion's shared global namespace is transparent to backup, replication, and business continuity applications widely used for data protection.



Data Sharing: Scientific experimentation is seldom done in isolation. There is a strong need to share data, results, and techniques with others. In many cases, raw and supplementary files are shared to enable processing via newer cutting-edge algorithms. These datasets are typically multiple TB in size. Centralized, secure access and data sharing eliminates the need to make copies.

Several organizations make their analysis results available to other researchers or teams. EMDB and EMPIAR are examples of publicly available repositories.

Researcher Challenge: As Cryo-EM becomes more mainstream for structural biology, this data is only expected to grow. Drug discovery and related activities create irreplaceable intellectual property. Securely sharing information is very important.

Pavilion implements FIPS-compliant data at rest encryption and ensures the always-on encryption does not impact performance. Data encryption can be applied across the entire namespace to assure privacy and protection for valuable research. Pavilion also supports T10 DIF (Data Integrity Field) to ensure no data is corrupted when writing to NVMe SSDs. With client access from standard operating systems plugins, users can be on Windows, Linux, or Mac securely sharing the global namespace in a simple and extremely performant manner. For example, a Gluster client can read at speeds up to 75GB/sec and write at speeds up to 50GB/sec.

Conclusion

To truly unlock the potential of cryo-EM in drug discovery, accelerated digital image analysis techniques must keep up with developments of direct electron detectors. High-performance compute, and storage are the underpinnings of these critical IT processes. Pavilion uniquely provides data storage solutions to meet the most stringent of requirements in this arena. With partnerships with cutting-edge computational technology providers, Pavilion is truly helping revolutionize the utility of cryo-EM in many organizations, with increased research productivity.

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 scale, and flexibility, providing customers unmatched choice and control. Learn why Fortune 500 companies and federal government agencies choose Pavilion. To schedule a demo, visit www.pavilion.io Follow the company on LinkedIn.