CYBERSECURITY SOLUTION BRIEF

Benefits

- Universally unmatched performance
- Eliminate overprovisioning, ingest bottlenecks, and filtering with 90GB/s write performance
- Get the performance and ultra-low latency of direct-attached NVMe SSDs
- Get unlimited, linear scalability of performance and capacity
- Reduce analytics storage expense by 2x with DAS consolidation
- Dramatically cut data processing time while simultaneously analyzing larger data sets
- Replace aging file/object data lakes with a linear scaling HyperParallel Block/File/Object data lake
- Train GPUs 3x faster with ½ the footprint of alternatives

Features

- Ingest, analyze, train, and visualize on a single scale-out and scale-up platform
- Up to 120GB/s read throughput
- Up to 90GB/s write throughput
- Up to 20M IOPS
- Latency as low as 25µs

PAVILION

Actionable Intelligence with Pavilion™

The Universally Unmatched Data Storage Platform

The threat of cyber-attacks is real for both private organizations and government agencies who face the daunting task of securing and protecting their data and resources. The challenges include threats such as supply-chain attacks, ransomware, and even ransomware-as-a-service (RAAS). That's in addition to insider threats and increasingly sophisticated social engineering attacks.

Complacency is the enemy of security. To strengthen and secure resources requires an ecosystem of innovation. Fortunately, there are many cyber tools being developed every day, along with new methods and capabilities to rapidly detect, prevent, and minimize the risks involved. More recent tools take advantage of artificial intelligence, which is a force multiplier for any cybersecurity team. These innovations are where cyber teams focus their efforts, but many do not realize that I/O has become a significant limitation for these next-generation cybersecurity tools.

There are 5 primary challenges in cybersecurity:

1) High-speed data ingest
2) Low-latency analytics and AI
3) Large scale analytic data lakes
4) AI/ML/DL Training
5) Visualization & Real-time response

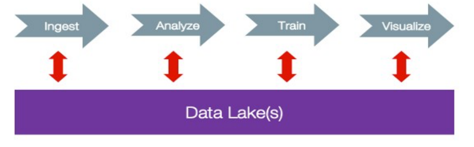

There are **five areas** where data I/O is the limiting factor in the success or failure of a cybersecurity initiative.

First is the **need to ingest data at high speed** (examples include Syslog, network tap, etc.). An increasing number of sources are generating ever-greater volumes of data, and that data needs to be collected in a central location for analysis. Check within your security organization and determine if your team has implemented filtering because of a data ingest problem. High levels of data filtering can hide a hacker's attack from detection.

The next area is a **requirement for in-depth analytics**. Analyzing the data is historically the biggest challenge of all. Adding to the already daunting task of processing growing volumes of data are new analysis methods leveraging artificial intelligence. Many of these tools consume a growing number of CPU cycles and take advantage of GPU capabilities, increasing the number of parallel I/O requests. With GPUs, the problem is less about overall throughput and more about the query/processing latency. Extremely low latency is needed for real-time and near-real-time actionable cybersecurity responses. This issue requires greater I/O capabilities and is often missed in the storage architecture design, resulting in a backlog of data, slow data processing, and ultimately limiting the power of cybersecurity tools.

The third issue is **storing all of the data** ingested in a single shared namespace for multiple analytics tools to process and predict potential threats. While processing and ingest are vital factors in many data center designs, the growing rate of data collection is often underestimated. The result is that the available storage capacity limits the volume of data collected. This requires the addition of more compute resources needed to store the data. The alternatives are to store less data, archive to low-speed resources like tape, have multiple separate data lakes that update each other through ETL processes, or statistically sample the data. The challenge is in the event of a breach, portions of the collection have already been deleted or archived to free up space, which means a reduction in the fidelity of real-time analytic results.

The fourth challenge is **AI model training**. To deliver quality results, AI models need to be trained on the largest possible datasets. GPU-based computing has revolutionized the rate at which AI models can be trained, delivering orders of magnitude improvements. To be effective, GPU-based compute systems require access to massive datasets with the lowest possible latency.

Finally, new **visualization technologies** are bringing better insights into potential cyber threats. These tools are both throughput and latency-sensitive as they comb through massive data lakes, analyze real-time streams and provide a portal for cyber operations centers, and flag potential threats with real-time responses. Like AI/ML/DL training, massively parallel architectures are needed to collect, synthesize, report, and automate real-time threat responses.

All five of these issues have the **same fundamental challenge**, which is the **need for high throughput, low latency data access that can scale seamlessly to support any size dataset.** CPU and GPU powered servers come with high speed, low latency NVMe SSDs, which provide the needed performance, but cannot scale and limit datasets to an unacceptably small size. Traditional scale-out storage offers capacity but cannot deliver the performance required.

The Pavilion HyperParallel Data Platform™ addresses this challenge with **unmatched performance** which enables organizations to accomplish outcomes previously thought impossible. Capable of delivering up to 120GB/s read and 90GB/s write performance, the Pavilion HyperParallel Data Platform solves the data ingest challenge with write performance that exceeds the read performance of most other solutions. With ultra-low latency of as little as 100µs for reads and 25µs for writes, measured at the host to include network latency, Pavilion delivers performance comparable to internal NVMe SSDs. Combined with the ability to linearly and independently scale performance and capacity across systems, the Pavilion HyperParallel Data Platform uniquely solves all of the performance, latency, and scalability challenges of the modern data center.

Pavilion's architecture is uniquely equipped to deliver an all-new level performance aligned to meet and exceed many cybersecurity workflows. The system provides the most performant, dense, scalable, and flexible storage platform in the universe.

Let's break down the cyber workflow and illustrate how Pavilion allows you to reimagine cybersecurity to shatter expectations, giving you choice and control over your infrastructure design.

1) Pavilion delivers up to 90GB/s write capabilities for data ingest in a single 4RU system for data ingest. That's more than most scale-out systems can deliver for just read performance. With up to 20 storage controllers, Pavilion eliminates the need to filter ingest data, assuring that all data from streaming, batch, and real-time sources can be delivered to the data lake while avoiding the need to over-provision resources for unpredictable bursts.

2) To deliver the lowest latency queries to the largest data lakes, Pavilion designed the HyperParallel Data Platform like a network switch. With a redundant 6.1Tbps non-blocking fabric inside the array, read throughput is 120GB/sec at 100-microsecond latency. **That's 50% lower latency than the nearest competitor in 25% of the data center footprint.**

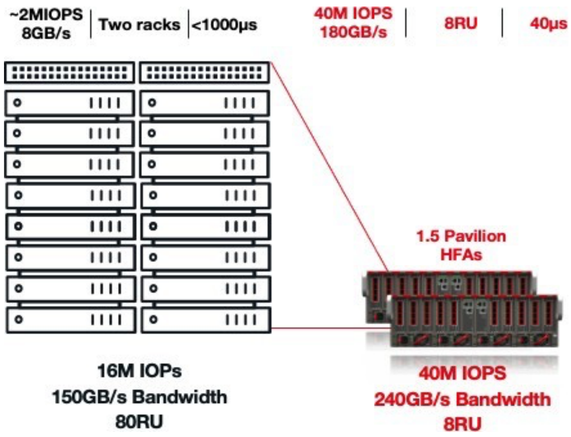

Pavilion delivers more bandwidth, more throughput, and lower latency than competitors in a smaller footprint to dramatically reduce costs.

3) To unify all data into a shared namespace, Pavilion packs 2.2PB in a 4U enclosure. That's before any compression. Scale-out in a linear fashion using Pavilion's HyperOS™ 3.0 with multi-chassis, a single shared namespace using NFS, or the most performant object store available on the market. The system can also work with scale-out file systems like IBM Spectrum Scale™, BeeGFS™, or Lustre™. In the case of Spectrum Scale, you can run the NSD Server function directly on Pavilion's storage controllers, dramatically reducing network traffic and latency.

   Of course, eliminating NSD Servers also reduces CapEx and OpEx expenses. There are no separate data lakes with a shared scale-out file system (Pavilion HyperOS or 3rd party), no messy ETL unifications, and all data is accessible for analytics, training, and visualization.

4) In AI/ML/DL training, Pavilion's massively parallel design and unmatched throughput can saturate an NVIDIA DGX-A100 with 2 HyperParallel Flash Arrays. With NVIDIA Magnum IO GPUDirect Storage, Pavilion can deliver a round trip to the client sub 5ms of latency for read and write. Whether you are using GPUDirect Storage or not, Pavilion has the capabilities to drive extremely low latencies so that AI/ML can be fully realized.

5) For visualization, reporting, and real-time response, Pavilion's hyperparallel design assures no-compromise ingest, a massively scalable data lake, and extremely low latency for tools like OmniSci™, Tableau™, and Graphistry™. Latencies are measured in microseconds regardless of whether the data is block, file, or object.

A cybersecurity workflow leveraging Pavilion can optimize data ingest streams, lower query response times, store more data and allow for rapid visualization and real-time threat response. Pavilion's storage platform offers universally unmatched storage enabling Cyber teams to reimagine what is possible.