

NVMe-oF Storage for Spark

Accelerate Spark with no-compromises



Modern applications that use Spark to analyze structured and unstructured data need to be future proofed. Reduce storage over provisioning, accommodate storage growth, consolidate applications and their data, speed time to results, and reduce TCO by disaggregating Spark storage with the Pavilion Hyperparallel Flash Array (HFA). It is the industry's first NVMe oF storage array. It scales to over a petabyte of data, provides ultra-high performance, ultra-low latency, linear scalability, and uses NVMe SSDs.

A DAS-based Spark infrastructure results in 2-3X overprovisioning and islands of storage—while using software-defined storage takes CPU resources away from Spark, impacting time to results and increasing TCO. Pavilion changes all that with its HFA array.

The Pavilion Hyperparallel Flash Array

The Pavilion HFA delivers never before seen NVMe performance (90 GB/s throughput, 40µs of latency, and 20M IOPS) & density (1.1 Petabyte) in a compact 4U storage array. It gives Spark applications the performance of locally attached NVMe SSDs enabling organizations to move to a Composable, Disaggregated Infrastructure (CDI) infrastructure, where application resources are readily available.

The Pavilion HFA requires no proprietary software to be installed on a server farm and uses standard NVMe, NVMe-oF, Ethernet, and InfiniBand drivers, freeing up host resources for processing, reducing deployment risk and increasing Spark density.

Benefits

- Increases Spark density 2X
- Petabyte scalability, high-performance, low-latency, and linear scaling maximizes data center efficiency
- Protection for your data and your business. Meet evolving requirements for data security and compliance
- Search 20X faster than DAS
- Deployment flexibility using concurrent protocols
- Enterprise design and data integrity validation ensure reliable access to data
- Rack scale/CDI management via Web GUI, vCenter, Kubernetes, RESTful API, OpenStack, DTMF/Redfish, and Swordfish



Benefits of Disaggregation

The Pavilion HFA provides high performance and low latency to Spark applications.

Scalable & Flexible



Provide up to 1.1 Petabytes to Spark deployments and simultaneously use the NVMe /RDMA (Ethernet, IB), NVMe/TCP, iSCSI, or NFS interfaces. Grow performance and capacity linearly without impacting on-going operations. Pavilion HFA's thin provisioning feature provides Spark with significantly better utilization than DAS, since the Pavilion array will only allocate physical space when the Spark application needs it. This results in physical space savings vs. DAS of up to 75% per server. No wasted capacity, no time spent messing with volume managers or file systems. Just set and forget.

Fast & Dense



The ultra-high performance, extreme low latency, and multiple storage controllers of the Pavilion HFA accelerates Spark workflows and boosts time to results. The Pavilion HFA lets Spark applications search terabytes of structured and unstructured data over 20 times faster than a DAS NVMe SSD.

Get 2X more density of a Spark cluster by using storage-less server nodes and moving all storage and storage management to the Pavilion HFA. No custom software needs to be installed on the Spark cluster, enabling it to take full advantage of host processing resources as well as simplifying deployment complexity.

Economical



Spark performs parallel computing across nodes and uses local SSDs and RAM to reduce the I/O and execution times of tasks. SSDs and RAM are some of the most expensive components of the Spark cluster, why not reduce them both? The Pavilion HFA disaggregates storage while providing Spark applications with ultra-high performance and ultra-low latency. This enables DevOps to eliminate DAS NVMe SSDs and remove RAM in every Spark server, reducing TCO.

With Pavilion, you no longer are constrained by the size of the SSDs. Thin provisioning allows the application to use the required amount of storage at any given time, regardless of how much capacity has been advertised to that specific Spark server. No more extra copies of data, the Pavilion HFA reduces the amount of raw flash storage deployed in a Spark environment, by up to 3X, reducing TCO. At 50% less expensive than a DAS SSD in terms of GB/s, the Pavilion HFA is designed for all your Spark environments.

Safe & Secure



Protect the SSDs with Pavilion's RAID-6 erasure coding and its 12% overhead. Its "Swarm Recovery" rebuilds a failed SSD ten times quicker than a DAS or an AFA. The Pavilion HFA ensures that the failure of an SSD does not impact Spark operations.

Keep data and snapshots safe with FIPS-compliant data at rest encryption that ensures its always-on encryption keeps data secure without impacting Spark performance.

The Pavilion Data Assurance feature works with the RAID feature to provide self-healing bit-rot support for data, assuring every Spark workflow gets uncorrupted data. Take consistent, instant, zero-footprint, and uncorrupted snapshots, encrypt them and provide them to backup and disaster recovery processes to speed operations and ensure consistency.

Enterprise Strength & Support



Get 24/7 proactive support, end-to-end data integrity, a modular chassis, and redundancy throughout the storage array to protect your Spark applications as infrastructures grow. With the Pavilion HFA, all the features come in-the-box, including thin provisioning, snapshots, clones, data at rest encryption and more.

Manage the Pavilion HFA via its Web GUI or use the management framework of your choice, including: vCenter, Kubernetes, RESTful API, OpenStack, DTMF/Redfish, Swordfish, and more.

Find Out More

Pavilion is defining the future of disaggregated NVMe-oF. Spark applications can be disrupted by the Pavilion HFA and its unprecedented availability, performance and versatility to future-proof the storage infrastructure.

Our expertise is in simplifying and optimizing NVMe to make the impossible, possible. When storage is business-critical, there's no substitute for the guaranteed performance, functionality, high availability, and OpenChoice support of a Pavilion NVMe-oF storage array. Use the Pavilion Hyperparallel Flash Array with its NVMe over fabrics support and make better decisions, faster! Contact us to learn more.