



## Features

### Spark Benefits

- Reduce the cost of deploying large cloud-scale Spark deployments in terms of both CAPEX and OPEX
- Disaggregate physical resources to independently scale compute, storage and network resources
- Simplify database provisioning, deployment and protection
- Deliver orders of magnitude better application performance in high-scale cloud, web, and SAAS environments

### Pavilion Benefits

- 40  $\mu$ s Latency
- 14TB - 1 PB in 4U
- Frictionless Deployment
- Data Resiliency & High Availability
- Up to 20 Active/Active Controllers
- Multi-path IO support
- Independently scale compute and storage.
- Space-Efficient, Instant Snapshots and Clones
- Thin Provisioning
- Standard Ethernet
- **OPENCHOICE** Storage

# NVME-OF STORAGE FOR SPARK

## Deploy Disaggregated Storage to Spark

Big-data analytics, mobile, and cloud-delivered applications using Apache Spark are driving a new paradigm in IT infrastructure design.

Resources need to be freed. They need to be available and deployed so that the ever-changing requirements can be satisfied on a minute-by-minute basis. This means that compute, network, and storage resources all need to scale independently to meet an ever-increasing and diverse set of application requirements.

## Direct-Attached Storage Challenges

Whilst DAS offers the flexibility of deploying distributed resources in a scale-out fashion, it's expensive when it comes to storage and storage management. Typically, storage is deployed as direct-attached SSDs in individual servers but this leads to significant problems.

- Storage is not shared effectively since it is stranded in a single server and results in underutilization of NVMe; in some cases as low as 25%.
- Storage provisioning decisions are made at procurement time. Determining the size of the storage in each server is done before the requirements of the application are known. This leads to inflexibility and higher costs over time.
- When scaling for either performance or capacity reasons, more server nodes need to be deployed to accommodate more direct-attached SSDs, effectively expanding the infrastructure unnecessarily.
- And whilst some applications offer data protection mechanisms, they rely on making copies of data on other database server nodes, leading to more capacity being required, bloating the cost of infrastructure even more.

## Pavilion's NVMe-oF Storage Array

Pavilion delivers never before seen NVMe performance and density that allows customers to provision logical flash storage resources over a low latency network. As a result, you can now deploy shared storage in place of direct-attached SSDs in cloud-scale Spark environments.

The array requires no custom software to be installed on application servers and includes important data management and availability features, such as thin provisioning, instant zero-space snapshots and clones, and no single point of failure.

## Deliver Disaggregated NVMe-oF for Spark Deployments:

### Operational Benefits:

- Pavilion's instant, full-speed, zero-space Snapshots and Clones can be used to create additional copies of big datasets
- This enables multiple workloads to concurrently run on the same version of the data, without degradation or making copies of data over a network
- Reduce management complexity by moving storage to a shared appliance where a low latency pool of shared storage can be centrally managed
- Manage the array through a UI interface, and/or supported by REST API, which allows integration into existing management frameworks

### Infrastructure Benefits:

- Simplify the environment by leveraging a single, high-speed storage array to deliver large-scale Spark environments
- Deploy 'storage-less' servers to realize never-before-seen density per rack.
- Reduce IT expenditure on hardware acquisition, rack-space, power and cooling
- Use Thin provisioning for less raw flash capacity to be installed
- Transition to CDI infrastructure and achieve optimized levels of processing, storage and network bandwidth that can be scaled independently

The diagram represents the reference architecture for a modern Apache Spark implementation across complex persistent layers for modern applications, leveraging Pavilion's disaggregated storage array.

