



elasticsearch

Benefits

- Cut server & storage CapEx
- 50% reduction of networking cost due to future-proof design
- Accelerate searches and discovery
 - Search 9% more Docs per sec
 - Run 47% more Ops per Sec
- Save 50% of rack space compared to the nearest competitor
- Avoid BC/DR impact to Elasticsearch operations

Features

- Data integrity validation for Elasticsearch's distributed document store
- Multi-Platform
 - VMware certified NVMe/RDMA (RoCE v2), NVMe/TCP drivers, iSCSI, and NFS drivers for Elasticsearch on VMware 7
 - NVMe/RoCE and NVMe/TCP drivers for Elasticsearch on Linux
 - iSCSI, NFS, S3, SMB, and Kerberos CSI protocols for other Linux, Windows, and VMware applications
- Unified Block, NFS, and S3 drivers
- NFS and S3 Global Namespace
- Intuitive management via Web GUI, vCenter, Kubernetes, REST API, OpenStack, DTMF/Redfish, and Swordfish
- Cloud-based call-home support

NVMe-oF Storage for Elasticsearch

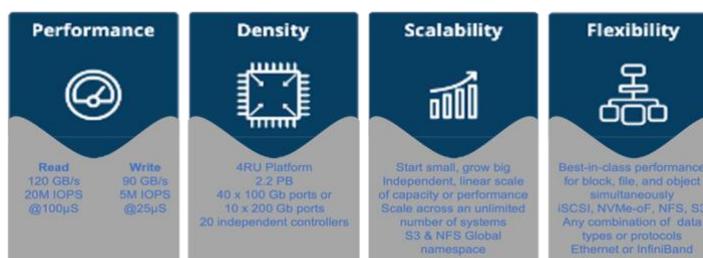
Accelerate Searching and Data Discovery

Your environment is diverse. Many organizations use a Direct Attached Storage (DAS) configuration for Elasticsearch that is easy to deploy but is complex, slow, and hard to scale. Analyzing data across hot, warm, and slow tiers using a DAS-based design will be slow and is lacking in enterprise features and centralized storage management. It is very easy for the cost to escalate when processing must be used for enterprise features or IT has to perform storage management for each Elasticsearch node.

Using a distributed DAS model was ideal for slow storage and network configurations, but today it just leads to slow operations. A server is a single point of failure, requiring IT to reserve at least half the storage space for copies of data from other nodes, this increases storage cost and network traffic and leads to expensive data/IO rebalancing. If a server runs out of space, Elasticsearch will fail, requiring IT to add additional servers & storage.

Focus on outcomes by relying on an NVMe-oF architecture

The Pavilion HyperParallel Flash Array (HFA) is an AFA that uses NVMe-over-Fabrics (NVMe-oF). Pavilion leads the industry in its support of [NVMe-oF](#) and is built for NVMe from the ground up. NVMe is not shoe horned-in like many other vendors. Pavilion boosts Elasticsearch operations on Linux and VMware, enabling customers to derive actionable insights in a short amount of time from big and fast data workloads.



The patented architecture of the Pavilion HFA provides its high-performance and ultra-low latency in a compact form factor that is more than twice as small as the nearest competitor. Its ultra-low latency and ultra-high performance let you better use server resources, reducing rack space, power, and HVAC costs. Its multi-controller design enables Elasticsearch customers to architect more complex searches, support more concurrent operations, achieve continuous operations, and shatter backup windows. This revolutionary storage system offers the most performant, dense, scalable, and flexible storage in the universe.

Elasticsearch Delivers More with Pavilion

The Pavilion HFA provides Elasticsearch with performance that is better than locally attached NVMe SSDs. It avoids customers having to create storage islands with its unified block, file, and object support. It supports NVMe/RDMA (RoCE v2), NVMe/TCP, iSCSI, NFS (v3, v4, pNFS, and RDMA), the Kubernetes CSI driver), and S3. VMware 7U3 has certified native drivers for NVMe/TCP, NVMe/RDMA, iSCSI, and NFS.

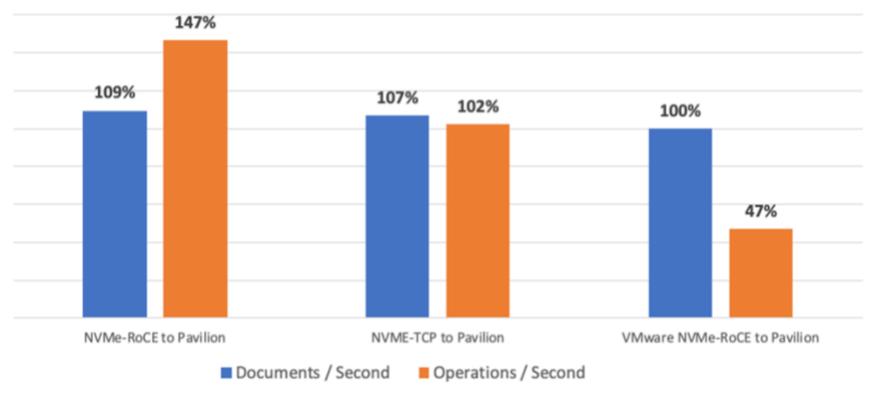
Two Elasticsearch servers working in a cluster and a single Elasticsearch Rally server to simulate the workload was used. All scenarios used the same server, operating system, Elasticsearch version, and dual parity RAID 6. Linux was CentOS 8.5 Stream and VMware was 7.0U2. The bare metal configuration had 6 NVMe SSDs. It was tested against a Pavilion HFA that used NVMe/RoCE.

The Pavilion HFA's NVMe/RoCE support was tested with the server running Linux, it delivered better results than the same server using INTERNAL NVMe drives. Testing showed that the Pavilion HFA had 109% of documents per second and 140% of the operations per second then the same server with INTERNAL NVMe drives.

SOLUTION BRIEF

The Pavilion HFA's NVMe/TCP support was tested with the server running Linux, it delivered better results than the same server using INTERNAL NVMe drives. Testing showed that the Pavilion HFA had 107% of documents per second and 102% of the operations per second then the same server with INTERNAL NVMe drives.

The Pavilion HFA's NVMe/RoCE support was tested with the server running VMware. Testing showed that the Pavilion HFA had 100% of documents per second and 47% of the operations per second then the same server with INTERNAL NVMe drives. The Pavilion HFA has the same documents per second, but with all the flexibility and functionality of the VMware ecosystem.



For full results contact Pavilion at info@pavilion.io.

Using Elasticsearch with servers that don't have DAS storage can double the density of a rack. In addition, the Pavilion HFA reduces storage management costs by 50% over DAS and saves money by reducing server hardware and paring data center space reducing the cost of deploying Elasticsearch clusters significantly. No custom software needs to be installed on the nodes, enabling them to take full advantage of host processing resources and reduce risk. Protect Elasticsearch SSDs with RAID-6 and the Pavilion Swarm capability, which will rebuild a failed SSD at the rate of 5 minutes per TB. It also provides Elasticsearch with self-healing bit-rot support, to assure that the distributed document store stays uncorrupted.

Security is a must for Elasticsearch. Pavilion uses a FIPS-compliant data at rest encryption and ensures its always-on encryption does not impact performance. Use consistent snapshots, encrypt those snapshots, and use standard backup and restore utilities to ensure that Elasticsearch data stays safe and secure. Instant snapshots and clones with a consistency group feature allow an entire Elasticsearch node cluster to be backed up or copied for test/dev purposes and on the fly without any performance impact to running Elasticsearch.

Get end-to-end data integrity, a robust and modular chassis, and redundancy throughout the storage array to protect your Elasticsearch document store as its infrastructure grows. All features come in-the-box, including thin provisioning, snapshots, clones, data at rest encryption and more. Support concerns won't delay the deployment of Elasticsearch. Pavilion provides 24/7 proactive support, a cloud-based call home service, and can act as an extension of your IT organization for Elasticsearch.

Zero-footprint snapshot capability provides an instant replica of an Elasticsearch cluster. Developers, testers, and other organizations that need access to real-time data can utilize this replica. The snapshot also removes the need to quiesce Elasticsearch for operations like backup, replication, business continuance, and DR, accelerating time-to-results.

Pavilion eliminates the need for a separate storage appliance for block, file, and object applications. It supports all three simultaneously and provides the highest performance per data center rack unit to every application. It enables the deployment of scalable clusters of Elasticsearch, SQL Server, and HDFS containers running on Kubernetes. Now users can combine and analyze relational, unstructured data, and S3 object data for searches, AI/ML, and other data analysis.

Learn More

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.