

Pavilion HyperParallel Data Platform™ for kdb+

Accelerate Kx kdb+ and Lower TCO

Financial organizations deal with massive amounts of data that is growing and has to be evaluated in real-time as a guide to decision making. For example, 20 years ago the NYSE processed half a million trades and quotes per day. Today it's grown by over an order of magnitude. The trading firm that analyzes this data faster has an edge over their competition.

Many organizations use high-performance databases, such as kdb+ from Kx Systems, to obtain the fastest analysis. kdb+ takes a high-performance historical time-series columnar database called a historical database (HDB), which has traditionally been stored on direct access storage (DAS) in a kdb+ server and combined it with a real-time database (RDB) that is stored on memory in the same server. Now, the Pavilion HyperParallel Data Platform can be used to provide ultra-high performance with ultra-low latency that boosts kdb+ application performance, while reducing costs.

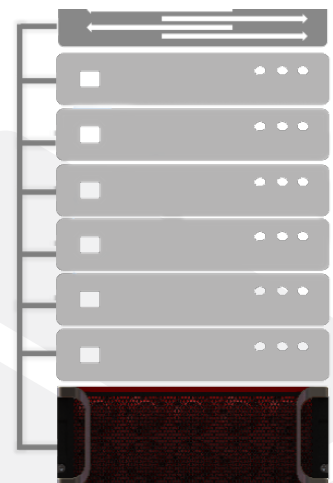
The kdb+ workload is parallelized across multiple clustered kdb+ servers and the results of the analysis are written in parallel to the kdb+ server's DAS. At the end of the day, the RDB is written in parallel to the HDB on the server's internal disk.

The performance of kdb+ is dependent on the performance of the storage infrastructure. Low I/O latency and high I/O transactions per second are essential for mission-critical kdb+ clusters.



Benefits

- Increases rack density 2X
- Petabyte scalability and linear scaling with high performance and low latency maximizes data center availability
- Protect your data and meet evolving security and compliance requirements
- Reduces storage TCO 50%
- Proven faster than many AFAs
- Concurrent block, file & object protocols give deployment flexibility
- Enterprise design and data integrity validation ensures reliable access to data
- Accelerates kdb+ searches 13X
- Simplifies rack scale management via Web GUI & CLI, vCenter, REST API, OpenStack, Kubernetes, Redfish, and Swordfish



Server-Side Storage Challenges

kdb+ offers the flexibility of deploying distributed resources in a scale-out fashion, but it is expensive when it comes to storage and storage management. Typically, storage is deployed with SSD-based DAS in the kdb+ servers, but this leads to significant problems that admins must deal with, including:

- HDB storage is often overprovisioned, as it is sized based on the initial deployment. Determining the size of the storage in each server is done before the requirements of the kdb+ the solution is known, which leads to storage bloat and increases the total cost.
- HDB storage needs to be replicated in each server in the cluster, resulting in the wasteful multiplication of storage requirements.
- When scaling for either performance or capacity reasons, more server nodes need to be deployed, which require more direct-attached SSDs, effectively expanding the infrastructure and adding unnecessary cost.
- The kdb+ data protection mechanism relies on making copies of data and storing them on other servers, leading to more capacity being required, which increases the cost of the kdb+ solution.

Admins often look at replacing the local SSDs in their kdb+ servers with NVMe-based shared storage from an all-flash array (AFA). However, most AFAs don't unlock the full performance of the flash technology due to their legacy, dual-controller architecture, leaving them with worse performance than a kdb+ system with DAS-based NVMe drives.

Pavilion changes all this. With fast-parallel access to flash, 100 or 200 Gb/s Ethernet or InfiniBand connections, and usage of the NVMe-oF protocol, the Pavilion HyperParallel Data Platform provides kdb+ with the benefits of fabric attached storage and the same or better performance than DAS. With Pavilion, kdb+ can now extend the benefits of parallelization to storage in addition to the CPUs.

The Pavilion HyperParallel Data Platform

The Pavilion HyperParallel Data Platform delivers never before seen NVMe performance (up to 120 GB/s throughput, 25µs of latency, and 20M IOPS) and density (2.2 Petabyte) in a compact 4U form factor. It provides kdb+ with the performance of locally attached NVMe SSDs by using the NVMe-oF protocol. It provides the ultra-low latency and ultra-high performance that kdb+ applications need, which is better than an AFA costing 25X more per IOPS. This enables customers to improve analysis and real-time decision making.

The value that a Pavilion HyperParallel Data Platform provides to kdb+ was proven through independent certification by the Securities Technology Analysis Center (STAC®). They analyzed the results of running the STAC-M3™ benchmarks with a solution comprised of kdb+ and the Pavilion HyperParallel Data Platform. The solution set records for four of the 17 benchmarks in the Antuco suite, beating all DAS and AFA solutions tested by STAC, and was the best of all the flash arrays they tested in another four benchmarks. The combination benefited from the parallel design of the Pavilion HyperParallel Data Platform, and the testing shows that the solution soared above the competition for several read-I/O intensive queries like volume-weighted average bid, as well as balanced compute and I/O workloads like statistical calculations.

STAC found that the Pavilion HyperParallel Data Platform was faster in four of 17 mean-response time benchmarks in the Antuco suite compared to all other solutions using kdb+:

- **STAC-M3. β1.100T.STATS-UI.TIME**
- **STAC-M3. β1.50T.STATS-UI.TIME**
- **STAC-M3. β1.10T.STATS-UI.TIME**
- **STAC-M3. β1.10T.STATS-AGG.TIME**

They also found that it was faster in eight of the 17 mean-response time benchmarks compared to all other solutions using kdb+ and flash arrays, including the tests above plus:

- **STAC-. β1.100T.VWAB-12D-NO.TIME**
- **STAC-M3. β1.10T.VOLCURV.TIME**
- **STAC-M3. β1.1T.MOHIBID.TIME**
- **STAC-M3. β1.1T.STATS-UI.TIME**

The Pavilion HyperParallel Data Platform requires no proprietary software to be installed on the kdb+ nodes and uses standard NVMe, NVMe-oF, Ethernet, and InfiniBand drivers, freeing up host resources for processing and reducing deployment risk.

The Pavilion HyperParallel Data Platform Delivers High Performance and Low Latency to kdb+

Fast & Dense



Enjoy the benefits of processing information faster using an ultra-high performance and ultra-low latency storage array. Search data 25 times quicker than with two local NVMe SSDs in a kdb+ server node, speeding time to results. Simultaneously use NVMe/RDMA, NVMe/TCP, iSCSI, NFS (v3, v4, pNFS, RDMA) and S3 for all kdb+ needs.

Disaggregating flash storage from each server in a rack no longer requires dedicating resources to managing local SSDs. Using kdb+ servers that don't have to support DAS storage can double the compute density of a rack. No custom software needs to be installed on kdb+ servers, enabling them to take full advantage of host processing resources and reduce risk.

Safe & Secure



Protect the SSDs with RAID-6 and the Pavilion SwarmController™, which enables multiple controllers to work together to rebuild a failed SSD at the rate of less than 5 minutes per TB. Get self-healing bit-rot support for data, to assure every process gets uncorrupted data.

Security is a must for kdb+. The Pavilion HyperParallel Data Platform uses a government standard data at rest encryption based on a 256-bit XTS-AES algorithm. It leverages the Intel AES-NI instruction set to ensure its always-on encryption does not impact performance. Meeting compliance requirements is a fundamental part of the system design. Use consistent snapshots, encrypt those snapshots, then use standard backup and restore utilities to achieve high-fidelity compliance at a level of granularity that you define.

Enterprise Strength & Enterprise Support



Pavilion adds value to kdb+ with enterprise features, end-to-end data integrity, a robust and modular chassis, and redundancy throughout the storage array. All features come in-the box, including thin provisioning, snapshots, clones, data at rest encryption, and more.

Don't let support concerns prevent the deployment of web-scale applications. The Pavilion global support team can act as an extension of your IT organization for all banking applications.

Economical & Flexible Management



An NVMe SSD is one of the most expensive components in a kdb+ configuration. With the Pavilion HyperParallel Data Platform, wasted capacity can be eliminated by scaling down flash storage deployed by 2-3X, while saving 50% over DAS in terms of \$/GB/ sec. Management overhead can be reduced by eliminating the need for volume managers or file systems.

Reduce storage management costs by using the Pavilion HyperParallel Data Platform's Web GUI or integrate with other management frameworks, including vCenter, Kubernetes, OpenStack, DTMF/Redfish, Swordfish, RESTful API, and more.

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 scalability and flexibility, providing customers unprecedented choice and control. Learn why Fortune 500 companies and federal government agencies choose Pavilion. Visit www.pavilion.io or follow the company twitter at <https://twitter.com/PavilionData>