

# Record-Setting Storage Performance— Without Breaking the Bank



## RECORD PERFORMANCE

Set 17 world records in performance for query-response times and aggregate throughput running STAC-M3 Benchmarks



## SUPERIOR TO ALTERNATIVES

16x faster than Optane DAS server, 4.5x faster than all-flash NAS and 2.1x faster than the NVMe-oF SAN with a file system layered on top



## SCALES LINERALLY

Run metadata intensive applications and share data sets concurrently with Monte Carlo simulations and back testing workloads without performance degradation



## EASY MANAGEMENT

Petabyte datasets with small and large files require no performance tuning. Seamless movement of data between hot and cold tiers

Financial services organizations have long been at the forefront of adopting leading-edge information technology. Massive data volumes, low-latency, and complex analysis are common attributes of high frequency trading and advanced analytics environments including fraud prevention, risk management, market simulations and more. There is a constant battle between researchers, scientists, and traders because more complex models will be more insightful and better predictors of business outcomes. The need to collect, store, and analyze vast amounts of real-time and historical data requires a revolutionary approach to storage.

## KX HIGH PERFORMANCE DATABASE

To obtain the fastest analysis possible, many Financial Services organizations use high-performance databases, such as kdb+ from Kx Systems, a First Derivatives Company. Kdb+ performs large scale simultaneous complex analytics on streaming, real-time, and historical data. The implementation consists of a Real Time Database (RTDB) that runs on streaming in-memory data, and a Historical Database (HDB) that is mapped to persistent storage. Incoming data streams are partitioned and load-balanced, and in parallel, written to persistent high-performance storage where it is analyzed to generate fast actionable insights. There is a heavy dependency on storage systems to deliver low latency and high transactions per second to kdb+ environments.

## CHALLENGES SCALING PERFORMANCE FOR MASSIVE DATA SETS

Scaling an existing in-memory analytics engine by adding memory and CPU cores to a single server allows more data to fit in memory for real-time analysis; however, it comes with some significant drawbacks. The data cannot be shared across multiple servers, and additional management effort is required when sharding data across multiple instances. Most importantly, this approach quickly becomes cost prohibitive. Lastly, today's state-of-the-art All Flash Arrays (AFAs) are a backend option for the HDB store, but their suboptimal architecture can leave a lot of the promised performance on the table.

## WekaFS™: RECORD-BREAKING PERFORMANCE AT SCALE TO KDB+

Pushing the limits of what's possible, WekaFS solves the storage performance challenges for the most demanding data-intensive environments. With a clean, modern design approach, WekaIO achieved the unimaginable. The Weka File System (WekaFS) is the only NVMe-optimized, shared, coherent, POSIX-compliant file system with a fully distributed architecture that delivers high throughput and ultra-low latency at Exabyte scale to thousands of kdb+ clients. Weka's parallel file system delivered record-breaking storage performance to kdb+, as shown in Figure 1.

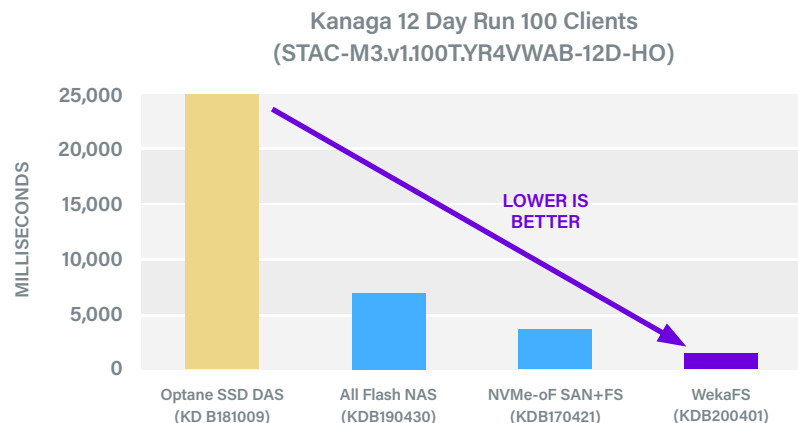


Figure 1 – Volume Weighted Bid for a random 1% of data for 12 Days over a 5-year data set

Running the STAC-M3 “Tick Analytics” Benchmarks, a standard for many financial organizations. WekaFS set 12 STAC-M3 world records for mean query-response times and 5 world records for aggregate throughput. The “Kanaga” suite of STAC-M3 benchmarks presents multiple years of market data to measure the volume-scalability of a database stack. The ability of a tick analytics stack to handle increasing volumes of historical data is essential to today’s trading organizations as market data volumes can quickly grow. The Kanaga results in Figure 1 show that WekaFS is the clear performance leader when compared to top storage alternatives.

The benchmark is a read-heavy workload that stressed the storage I/O and metadata capability. WekaFS was 16x faster than the Optane DAS server, 4.5x faster than all-flash NAS and 2.1x faster than the NVMe-oF SAN with a file system layered on top. The system under test used HPE servers and had a combined performance density of up to 113GB/sec available bandwidth to the client nodes, a maximum observed read throughput of 54.7GB/sec, and 8.8 million 4K IOPS in 10U, with the ability to fully distribute data, metadata, and system services. Distributing both data and metadata evenly across the storage cluster allows performance to scale linearly to fully saturate any CPU or GPU based compute clusters. This means users can run metadata intensive applications and share data sets concurrently with Monte Carlo simulations and algorithmic backtesting workloads without performance degradation. The full STAC Report can be found at STACresearch.com/KDB200401.

Petabyte scale datasets are handled effortlessly, including both small and large files with no need for performance tuning. WekaFS is easy to manage and is fully cloud-enabled for hybrid or public deployments. An optional integrated object-based data lake is possible, allowing for seamless movement of data between hot and cold tiers with the cost-benefit of hard disk drives for massive scalability. A rich enterprise feature set including snapshots, clones, encryption, authentication, and backup via its snap-to-object, cloud-bursting, and more. Figure 2 shows WekaFS deployed as a dedicated storage appliance deployed in a production environment with kdb+.




Figure 2 – Weka File System in a Production Environment

Find out more about Kx Systems and Weka offerings, and explore the possibilities of accelerating your critical financial services workloads. Go to <https://www.weka.io/solutions/financial-services>