



DDN Storage Solutions

For KX Systems

**Fully-Integrated and Optimized Data
Platforms for Financial Analytics**

EXECUTIVE SUMMARY

DDN Storage Solutions are proven to be the most effective and performant solutions for data intensive financial workflows. DDN parallel storage improves on traditional approaches to simultaneously reduce complexity, drive down cost and improve performance as data datasets grow.

This paper describes how DDN Storage Solutions are commonly used for KDB+ database applications. DDN enables finance institutions to optimize their storage infrastructure and serve data fast and efficiently with the highest parallelism, reducing time to results for institutions to gain insight of their data. The recent STAC M3 benchmark results ratify how our solutions accelerate the full range of financial analytics workflows.

DDN SOLUTIONS FOR FSI AND KDB APPLICATIONS

DDN Storage solutions for FSI are designed for performance and scale. Financial Analytics makes a particular demand on the storage system in order to deliver performance that is realized at the application level. Firstly, large multi Petabyte datasets must be accessible with consistent low latency. Secondly, the memory map functions often used in time series database analytics must be optimized. Thirdly, single thread and single client performance is often a critical piece of a workflow. DDN's parallel architecture, RDMA accelerated data and intelligent client all work together to deliver the performance whilst keeping hardware infrastructure small.

KDB+, a core technology of KX Systems, is a high-performance, streaming analytics and operations intelligence platform that allows firms to make better business decisions, solving complex problems. KDB+ is widely adopted by finance institutions. DDN has worked with KX for over ten years to build highly optimized storage with tunings specific for financial customers.

THE DDN SHARED PARALLEL ARCHITECTURE

The DDN shared parallel architecture includes network optimization and an intelligent client to provide superior performance, scalability, security, and reliability. Multiple parallel data paths extend from the drives all the way to customer applications, including KDB+. With DDN's true end-to-end parallelism, data is delivered with high-throughput, low-latency, and massive concurrency in transactions. This ensures applications complete queries in minimum time providing a faster response. Optimized parallel data-delivery directly translates to increased application performance and faster completion times.

The DDN Storage shared parallel architecture also contains redundancy and automatic failover capability to ensure high reliability, resiliency, and data availability in case a network connection or server becomes unavailable. The integrated architecture simplifies network deployment and reduce latency consolidating resources.

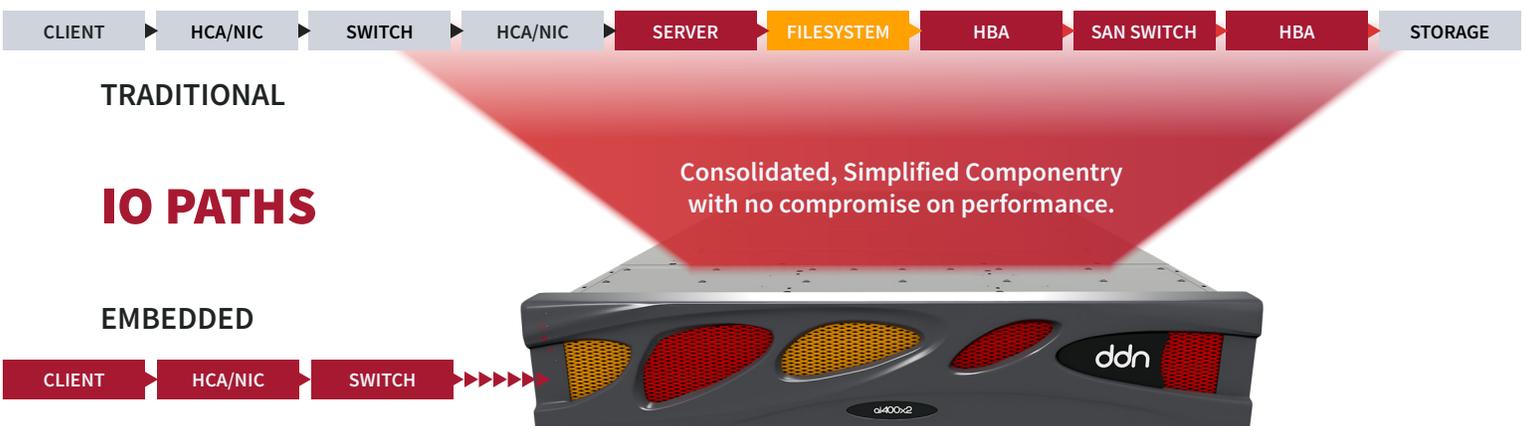


Figure 1.
Optimal IO Path with
DDN EXAScaler Appliances

DESIGN OF A KDB ENVIRONMENT WITH DDN EXASCALER

KDB+ (KX) SYSTEMS

The KDB+ database, developed by KX systems, is designed to be used for financial analytics to store time series data and to scale up/out when data increases. It is an ultra-fast cross-platform time-series columnar database with an in-memory compute engine, a real-time streaming processor and an expressive query and programming language called q.

Tick data analysis is generally done on three distinct data sets: real-time data, near real-time data and historical data. Real time data is kept in-memory, being moved to an intraday Database (near-realtime data), and from there to the historical database. Historical data require highly scalable and high-performance systems, as the volume of data produced per day is growing exponentially.

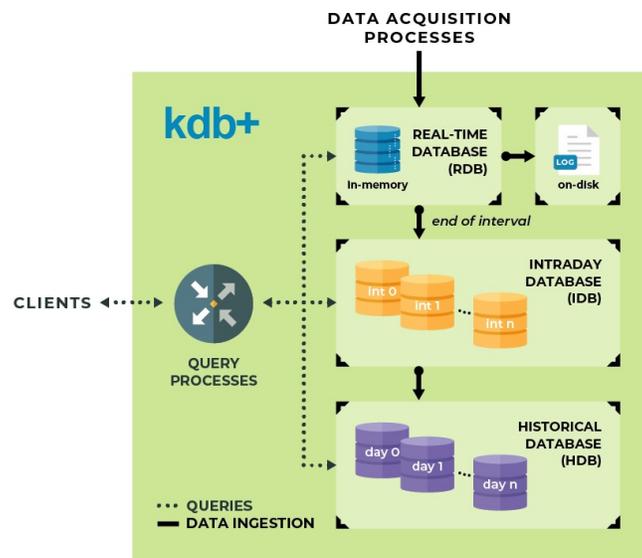


Figure 2.
Standard Data
processing with KDB+

THE DDN SHARED PARALLEL ARCHITECTURE

KDB+ Databases are stored as a series of files and directories on-disk, which facilitates the handling of the database on a filesystem, the database layout is partitioned by date. Besides partitioning, KDB+ allows segmentation. With segmentation a partitioned table is spread across multiple directories that have the same structure as the root directory. Each segment is a directory that contains a collection of partition directories, and it is configured through a par.txt file. The purpose of segmentation is to allow operations against tables to take advantage of parallel IO and concurrent processing.

Accessing data on disk KDB+ uses memory mapping (MMAP), which avoids the overhead of copying data between user and kernel space. DDN's EXAScaler provides a robust support for MMAP, additionally DDN has enhanced prefetch algorithms to optimize data access for mmap based workflows.

STANDARD CONFIGURATION WITH DDN EXASCALER

DDN EXAScaler guarantees the fastest speed and minimal operational latency on accessing historical data at the time that simplifies the architecture and setup of the whole solution.

The storage network between the KDB+ servers and the filesystem is by design a high-performance and low latency InfiniBand network, although other options like 100GbE are supported. KDB+ servers use the EXAScaler network protocol (Inet) to access the filesystem for maximum performance. This network layer simplifies deployments, particularly where compute systems have multiple network connections, and implements Remote direct memory Access between the client computers and the storage servers. EXAScaler also provides additional protocols like NFS, SMB or S3 to satisfy customer needs on heterogeneous environments.

The DDN EXAScaler filesystem can expand over one or multiple storage appliances whilst offering a single mount point to client computers. EXAScaler is regularly deployed with filesystems exceeding 100 Petabytes and 1 Terabyte per second of performance and it's scaling properties are well proven in AI and High performance Computing environments.

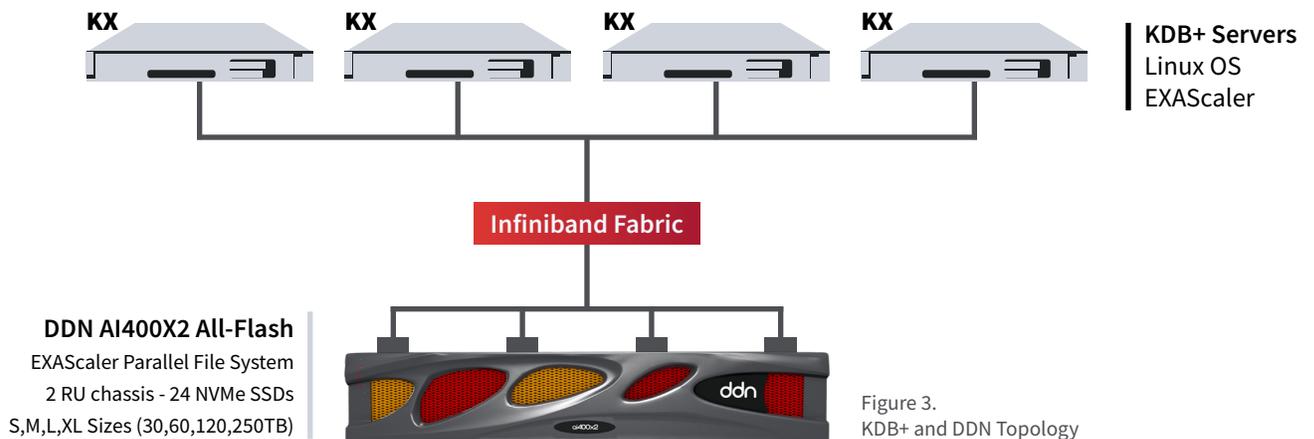


Figure 3.
KDB+ and DDN Topology

While other solutions require the user to partition the time series database space carefully across multiple mount points and directories to balance the usage of the resources and provide optimal access to data. DDN EXAScaler single namespace and internal intelligent automatic data balance drop that need, simplifying the configuration and administration, and allowing to scale seamlessly and transparently.

KDB+ Database can be segmented on multiple directories to allow for a higher level of parallelism using the par.txt file. All the directories will be allocated in the same single namespace not requiring extra management from filesystem configuration point of view. The KDB servers can share access to a common data set with optimal performance and without the need of extra management and data movement to be initiated by administrators.

DDN EXASCALER OPTIMIZATIONS FOR KDB APPLICATIONS

FSI applications using KDB+ although share a common database and follow similar principles on their architecture design, might highly differ in the way they handle their data and on the specific workflows they implement. Divergence on the usage of KDB+ translates in difference access patterns to the shared parallel filesystem. As well, variations on the architectural design of the environment, like the usage of fat nodes as KDB+ servers equipped with a vast amount of memory or standard servers with modest amount of resources mark differences on how optimally access and handle the data.

DDN EXAScaler is a mature and rich-feature filesystem which provides intelligent and very flexible algorithms to optimize the filesystem behavior and provide faster response to applications based on their IO data access pattern.

TAKING ADVANTAGE OF THE ADVANCE CACHE CONTROL

DDN EXAScaler features advance cache control mechanisms, between those mechanisms lays the option to prefetch full files up to a certain size into memory or decide how long keep data in memory once it has been accessed.

KDB+ servers with large memory capacity running workflows which access recurrently the same data benefit enormously from this cacheing approach. The corresponding tuning settings are:

```
ldml.namespaces.*.lru_max_age to increase timeout to flush data from cache  
llite.*.max_cache_mb to control maximum memory to be used by EXAScaler cache
```

KX applications save tick data in directory structures organized by dates and contain files for each table in the DB. Keep the smaller and most frequent referenced tables in memory help to reduce latency, accelerating global response to queries.

INTELLIGENT AND ADAPTATIVE PREFETCH ALGORITHMS ENHANCED FOR KDB APPLICATIONS

When loading the whole or most of the database into memory is not possible, due insufficient memory or large datasets, the behavior of the prefetch algorithm accessing data becomes critical. DDN has engaged very actively with KDB customers in improving the read-ahead support to address their specific needs.

Between other optimizations DDN has developed automatic recognition of semi-random patterns that result from queries along an index. Also, reverse table scans are successfully detected by the advance prefetch algorithm today as result of this work. A generalized ability to prefetch in reverse has been added to the EXAScaler prefetch algorithm, which allows EXAScaler to cover a wider spectrum of customer use cases with tuning.

DDN continues actively improving further these algorithms pursuing to provide our customers with the most adaptable and efficient storage solutions.

PROVEN PERFORMANCE WITH THE DDN AI400X/AI400X2 APPLIANCES

The DDN AI400X2 is a turnkey appliance, fully integrated and optimized for the most intensive AI, Analytics and HPC workloads. The appliance is proven and well-recognized to deliver highest performance, optimal efficiency, and flexible growth for deployments at all scale. A single appliance can deliver up to 90 GB/s of throughput and well over 3 million IOPS to clients via a HDR100 or 100 GbE network, and can scale predictably in performance, capacity and capability. The AI400X2 appliances are available in all-NVMe and hybrid NVMe/HDD configurations for maximum efficiency and best economics. The unified namespace simplifies end-to-end analytics workflows with integrated secure data ingest, management, and retention capabilities.



Figure 4. The DDN AI400X2 Appliance

The STAC report published recently (STAC/kdb211014.) demonstrated how a DDN AI400X2 delivers the highest performance with the minimal footprint for the STAC-M3 Antuco and Kanaga benchmarks.

STAC-M3 is the set of industry standard enterprise tick-analytics benchmarks for database software/hardware stacks that manage large time series of market data (“tick data”).

The DDN AI400X2 solution beat results of a traditional NAS Storage System, reducing dramatically the execution time of most of the metrics while using only a third of the NVMe devices and consuming 33% of the rack space, with the consequent power and heat savings.

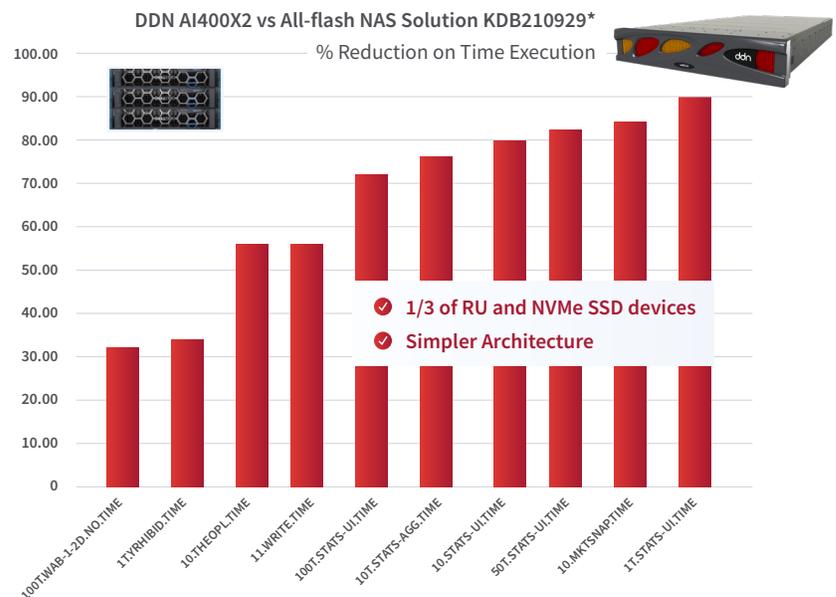


Figure 5. STAC Results DDN AI400X2 vs All-Flash NAS Solution

CONTACT DDN TO OPTIMIZE YOUR ANALYTICS WORKFLOWS

DDN has long been a partner of choice for organizations pursuing at-scale data-driven projects. Beyond technology platforms with proven capability, DDN provides significant technical expertise through its global research and development and field technical organizations.

A worldwide team with hundreds of engineers and technical experts can be called upon to optimize every phase of a customer project: initial inception, solution architecture, systems deployment, customer support and future scaling needs.

Strong customer focus coupled with technical excellence and deep field experience ensures that DDN delivers the best possible solution to any challenge. Taking a consultative approach, DDN experts will perform an in-depth evaluation of requirements and provide application-level optimization of data workflows for a project. They will then design and propose an optimized, highly reliable and easy to use solution that best enables and accelerates the customer effort.

REFERENCES

<https://stacresearch.com/KDB211014>

<https://stacresearch.com/KDB210929>

<https://kx.com/blog/what-makes-time-series-database-kdb-so-fast/>

<https://kx.com/blog/memory-mapping-in-kdb/>

ABOUT DDN

DataDirect Networks (DDN) is the world's leading big data storage supplier to data-intensive, global organizations. DDN has designed, developed, deployed, and optimized systems, software, and solutions that enable enterprises, service providers, research facilities, and government agencies to generate more value and to accelerate time to insight from their data and information, on premise and in the cloud.