KX

STREAMING ANALYTICS

CAPTURE ANY DATA
DRIVE DECISIONS

Introducing KX Streaming Analytics

ABOUT KX

Right now, continuous intelligence is at the forefront of digital transformation – but it's always been at the heart of our business and our technology. In addition to powering the financial markets for over 20 years, KX is also the driving force behind the transformation of some of the world's most valuable businesses looking to their data to help them work smarter, react faster and perform better.

KX is the only technology that can capture any data – any time – and run anywhere. And no one else comes close to our powerful streaming analytics capabilities, with businesses now able to make the best-informed decisions. In-the-moment. In-context. And, in real-time.

SEAMUS KEATING
CEO - KX

Introduction

KX Streaming Analytics is used across a range of industries – from manufacturing to finance – to solve complex problems that other data platforms cannot address.

It is used to derive insights and improve outcomes in areas like performance, quality, yield, and design in industries ranging from automotive and semiconductors to energy and telecommunications, where the volumes of data from sensors, machines, and edge devices are especially challenging.

KX is similarly used by financial institutions to trade, manage risk and detect fraud through analysis of the enormous volume and velocity of data they accumulate across prices, quotes, and transactions.

In this overview, we explain how they do it using KX Streaming Analytics.

One platform to instantly capture, analyze, and visualize real-time and historical data simultaneously

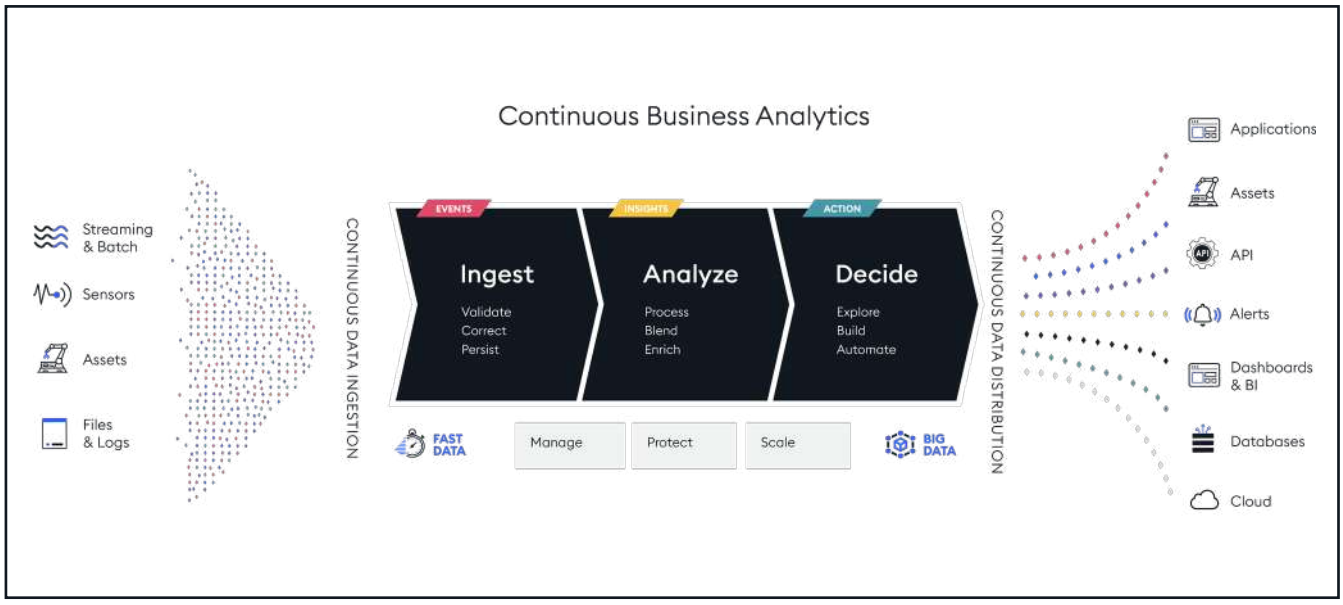

Knowing that something is likely to happen is a better insight than knowing that it has already happened.

Data, in all its facets and formats, is the new force fueling innovation and guiding successful organizations in informing their decisions, reforming their processes, and outperforming their competition. It's all about data: digital transformation is premised on it, machine learning is built upon it, operational efficiency is measured by it. Data is, without doubt, the new power and KX Streaming Analytics can help you harness it.

The Problem – Knowing too Late

At one time the humble abacus provided a competitive advantage by enabling faster and more accurate calculation than by hand. But speed is relative, and some of today's analytics platforms remain unfortunately abacus-like in their mode of operation. Many are based on an old-style polled model where data is persisted into a repository by certain processes, retrieved in queries by different ones only to evaluated in yet another set before the results are finally returned to the user.
The problem is that in the interim the data has changed and with it the market has moved, the fault has occurred, or the opportunity has passed. In short, they tell what has happened rather than what is happening, or even better, what might happen. For that, you need Streaming Analytics.

The Opportunity – Knowing Now, and Knowing in Advance

Streaming Analytics enables instantaneous decision-making based not only on data as it arrives but crucially on factoring in historical data as well. It moves beyond simple threshold analytics like "Has it exceed a limit?" on real-time data to include more contextual analysis that identifies patterns and insights based on moving averages, correlations, and machine learning over historical data. In short, it goes beyond telling you what has happened to telling you not only what is happening, but more importantly, what may be about to happen.

Overview

The Analysts' Views – Knowing How

In his paper on "The Five Levels of Streaming Analytics", Nick Heudecker from Gartner outlines the difficulties organization face in achieving streaming analytics from disparate interconnected technologies:

> **"Streaming analytics is a cross-functional discipline integrating technology, business processes, information governance and business alignment. It's the difficulty integrating these areas that keeps many organizations from realizing the value of their data in real-time."**

Mike Gualtieri from Forrester agrees on the complexity of the problem, the opportunity from getting it right, and the components necessary to do so. Visit www.KX.com to come to hear him outline the "7 Must-Haves for a Streaming Analytics Platform" which, like KX, covers the full lifecycle from data capture and processing through to development flexibility and fault-tolerance.

The KX Solution – Providing the Answers

KX has been a pioneer and trusted leader in in-memory computing, streaming analytics, and operational intelligence for over two decades. In independent tests, KX technology has been shown to be many orders of magnitude faster than traditional technologies; retrieving data in microseconds, aggregating data and performing complex queries in seconds to provide anomaly detection and predictive analytics that competing technologies may take hours or even weeks to deliver.

The KX Difference

ABOUT KX

KX is the fastest, best-informed, real-time decision-making engine in the world. Elegantly simple. Connect to the data you need. KX can capture any data – run anywhere – and is the highly efficient way to do incredibly difficult things. Our unrivaled streaming analytics platform drives the most demanding business decisions with real-time continuous intelligence.

Continuous Realtime Intelligence

KX enables real-time analysis of any data, whether at speed or at rest, providing firms with dynamic insights into what's happening, right now. Continuous updates enable them to make decisions instantly, and automatically.

World's Fastest Time Series Database

Independently verified as outperforming competing technologies, KX delivers powerful sub-microsecond latency for stream event processing with the ability to ingest and store 30 million updates per second and over 10 TB per day - in the moment, in-context and in real-time.

Simple and Unified

With interfaces to a wide range of third-party technologies and deployable on commodity hardware, KX integrates easily into existing environments whether in the cloud, on premises or at the edge. And for simplicity, it's the same platform in each case.

What Some of Our Clients say

WHY KX?

"Through the power of KX what we found was that our data flow would be 10 times faster than what it used to and would use only a 10th of the CPU. Which is fantastic "

BRIAN JONES,
HEAD OF SOFTWARE DEVELOPMENT,
ASTON MARTIN RED BULL RACING

"Within 5 hours of us bringing our transaction core online in KX we had detected and were able to stop a network of accounts that were being used for fraudulent activity."

TONY BRAUND,
PADDY POWER BETFAIR

KX Streaming Analytics

WHY KX?

Simplify, Consolidate, Visualize.

CAPTURE ANY DATA

KX Streaming Analytics is a cloud-based integrated platform to enable organizations to reduce costs and improve productivity through anomaly detection, analytics, predictions, corrective action, and automation based on both real-time and historical data. It can also run on-premises, on edge devices or in hybrid configurations of all as required.

RUN ANYWHERE

KX simplifies data analysis because it interprets streaming and historic data simultaneously within the same solution. This approach, coupled with open interfaces to other technologies, enables it to capture data from multiple sources, and run across multiple environments in order to analyze data and automate decision making in real time, in any environment

ANALYZE AND AUTOMATE

KX Streaming Analytics provides full life-cycle data ingestion, processing, analytics, and data management enabling mission-critical operations, 24x7, with resilient failover capabilities to ensure no downtime and no data loss.

KX Streaming Analytics
Core Components

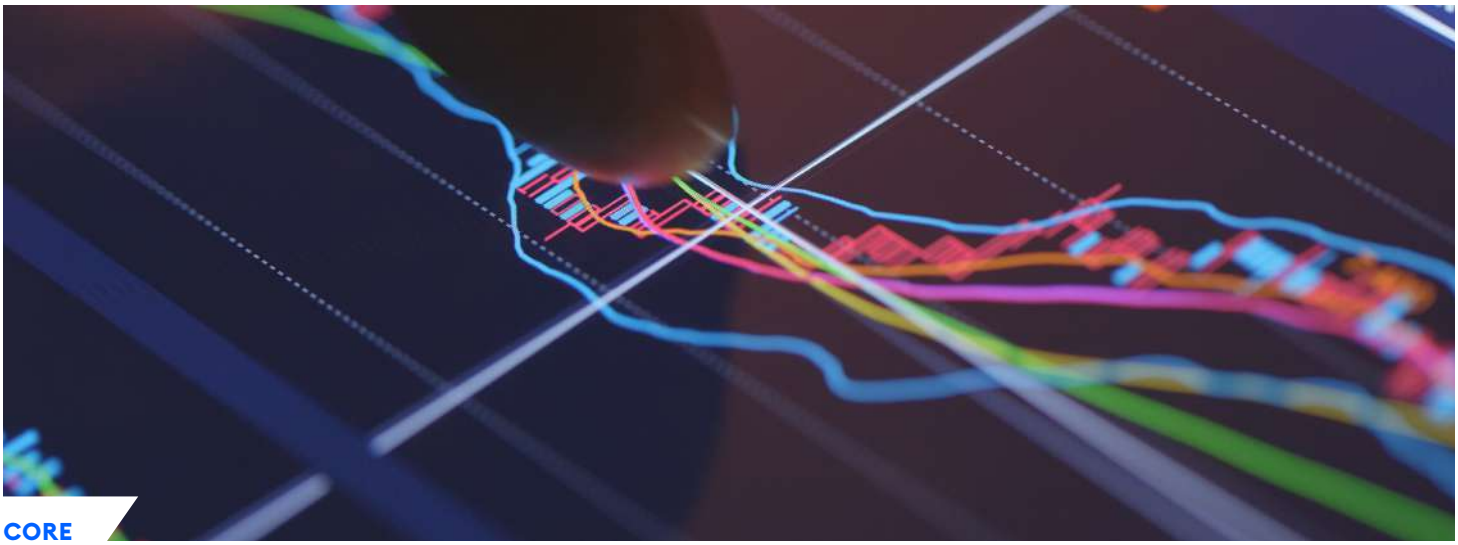

CORE

Kdb+ Time Series Database
The world's fastest time-series database for high-performance in-memory streaming analytics

KX Analyst
For managing, transforming, and interpreting large data sets at development time

KX Dashboards
Real time visualization, insights, and data exploration at run-time

KX Fusion APIs
For integrating with other technologies, programming languages, and APIs

Kdb+ Time Series Database – The World's Fastest Time Series Database

WHY KX?

Fast, Powerful, Scalable.

At the core of KX Streaming Analytics is kdb+, a powerful time-series database with a very low footprint executable of below 800K that includes the following components:

- a high-performance time-series columnar database
- an in-memory compute engine
- a real-time streaming processor
- an expressive query and programming language, q

Integral to kdb+ are its time-series analytics with temporal datatypes including nanosecond precision timestamps, a hybrid in-memory/on-disk architecture with a vector-oriented analytics engine all supported by robust message handling that enables multi-process and distributed solutions. Most significantly it operates on commodity hardware to run Cloud running Linux, Windows, and Mac and on edge devices where local processing may be required or optimal for specific use cases.

Kdb+ Time Series Database – The World's Fastest Time Series Database

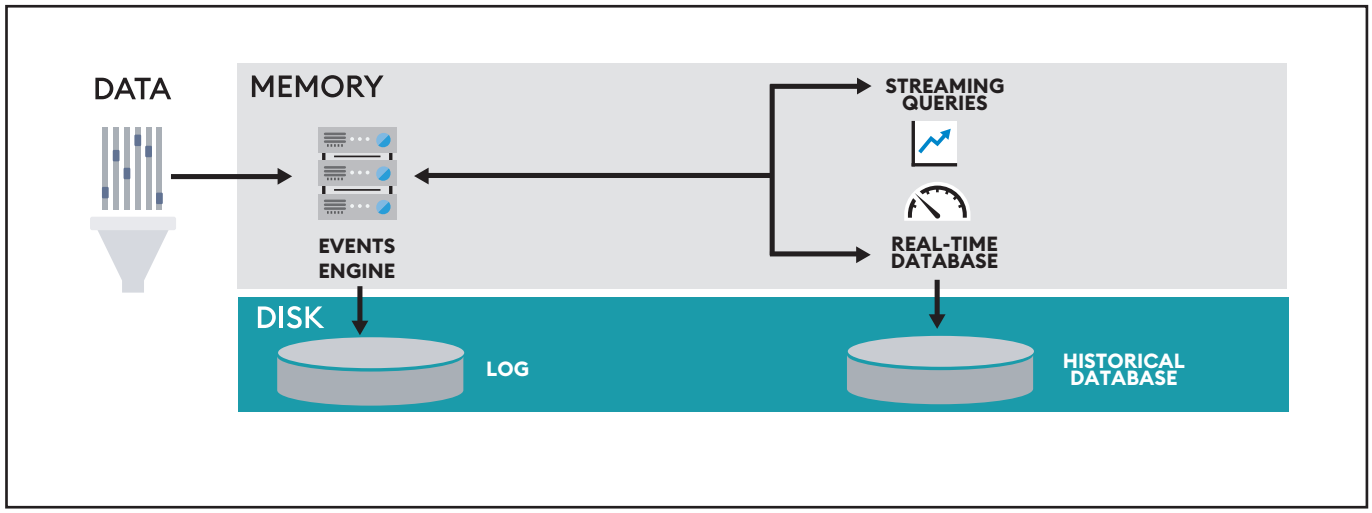

High-Performance, In-Memory Processing.

The diagram above illustrates a high-level reference architecture for a typical KX solution

- Data is collected by the events engine which logs messages to a transaction log file for resilience purposes
- The event engine persists the incoming data to the in-memory Real Time Database (RDB) via IPC.
- Dedicated streaming query engines subscribe to the event engine (for a subset of the data) to compute real-time analytics (e.g. weighted averages, standard deviations, interpolations) as data arrives
- At the end of a given time period (typically once per day) the RDB persists data from memory to disk.
- This data is immediately available in the Historical Database (HDB), allowing the RDB to purge the data from memory

Under this architecture, the most recent data lives in memory, which is the fastest to access, while older data lives on persistent disk, which is cheaper to maintain. An additional advantage of the architecture is that it allows users to query on-disk tables as if they were in-memory and vice versa, with no change to query syntax, hence reducing development and maintenance overhead.

As kdb+ processes can easily be connected or chained via IPC, it is also very simple to move processes from one server onto multiple servers to enable horizontal scalability and distributed processing for performance and fault tolerance requirements.

KX Analyst – Making Big Data Accessible, Usable and Valuable

KX Analyst provides an environment to manage, manipulate and explore massive datasets in real-time giving users the ability to easily load, transform, query, and visualize massive datasets in near real-time with minimal or no programming.

It includes a powerful transformation (ETL) tool that enables users to import, transform, join, and export any sized dataset without programming. Operations such as type casting, column renaming and reordering, data filling, and complex query filtering are all supported. It also provides a complete test-driven development framework with support for automated test creation using a behavioral data-driven (BDD) testing framework and a property library for automatically generating test cases. Code coverage is also supported.

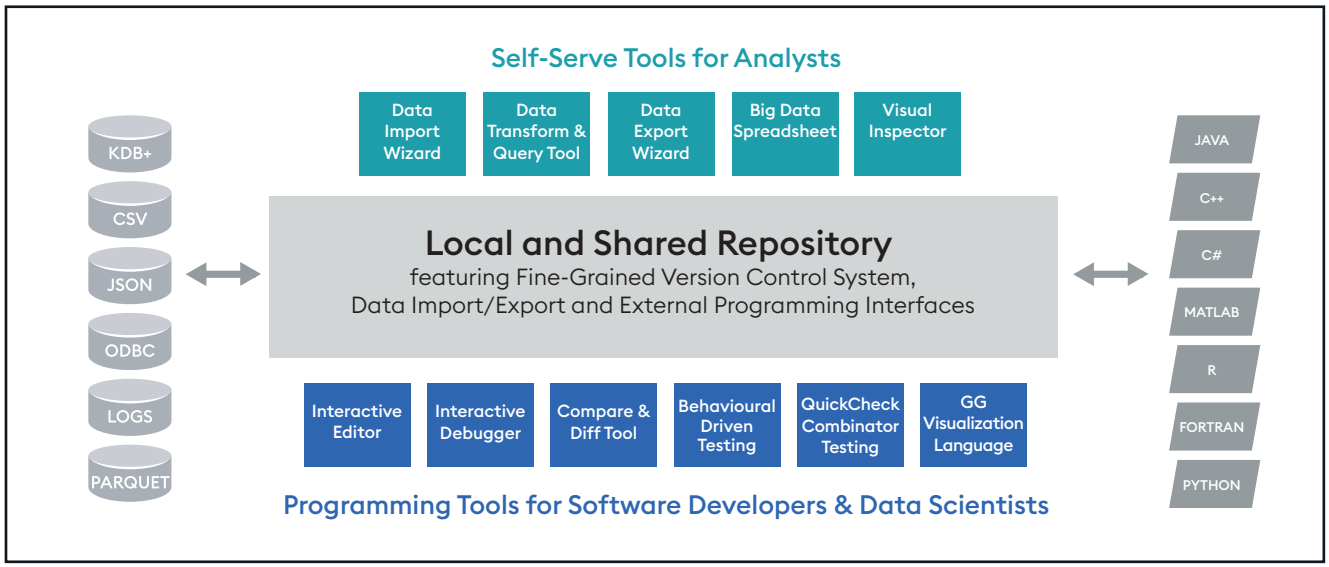

KX Analyst – Making Big Data Accessible, Usable and Valuable

WHY KX?

Capture, Analyse, Visualize.

By leveraging KX's server-side rendering technology, analysts can query and visualize tens of millions of records in a few seconds. This means that analysts can now search for and find patterns in large datasets in seconds or minutes, rather than days or weeks. With this functionality, it enables experienced developers and data scientists to quickly and easily create robust analytic libraries using a complete edit-debug-test development environment.

Simple	Fast
Most operations require no programming,	**Immediate answers, even for massive datasets,**
Complex transformations, joins and AND/OR Where queries without coding	100M record queries complete in milliseconds
Wizard interface for importing, exporting, and transforming data	100M record binned visualizations render in 2-4 seconds
Point-and-click visualization of massive datasets	100M record unbinned visualizations render in 6 minutes

KX Dashboards – Visualize, Stream and Share Big Data Insights

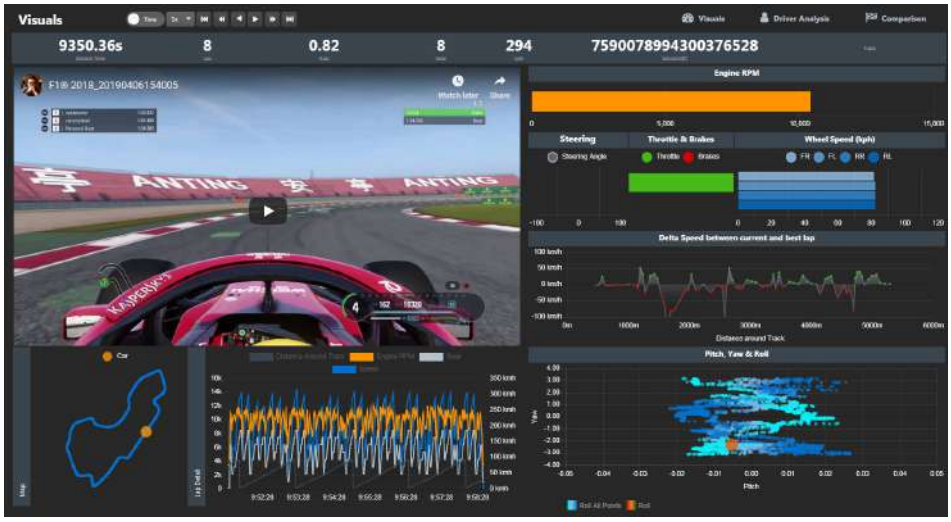

KX Dashboards provides a fast, accessible way to visualize big data, making it easy to query, transform, share and present and include comprehensive security and entitlement frameworks to meet your security, audit, and regulatory requirements.

By combining hardware-accelerated rendering with virtual scrolling to reduce display time, and binary transfer to optimize data throughput, KX Dashboards can render millions of records per second. Throttling, conflation by time intervals and server caching provide further support in maintaining high-frequency updates across multiple users.

Built for today's data-driven world, KX Dashboards makes it easy to uncover actionable and timely insights by visualizing both historical and real-time data, dynamically and interactively. In contrast to standard, largely static business intelligence visualization tools, KX Dashboards support streaming analytics and provide the flexibility to amend views and explore data in real time, as it arrives, eliminating the need to pre-process and store aggregated data. As a result, KX Dashboards let you to know now what your data is saying right now, and act upon it instantly.

KX Dashboards – Visualize, Stream and Share Big Data Insights

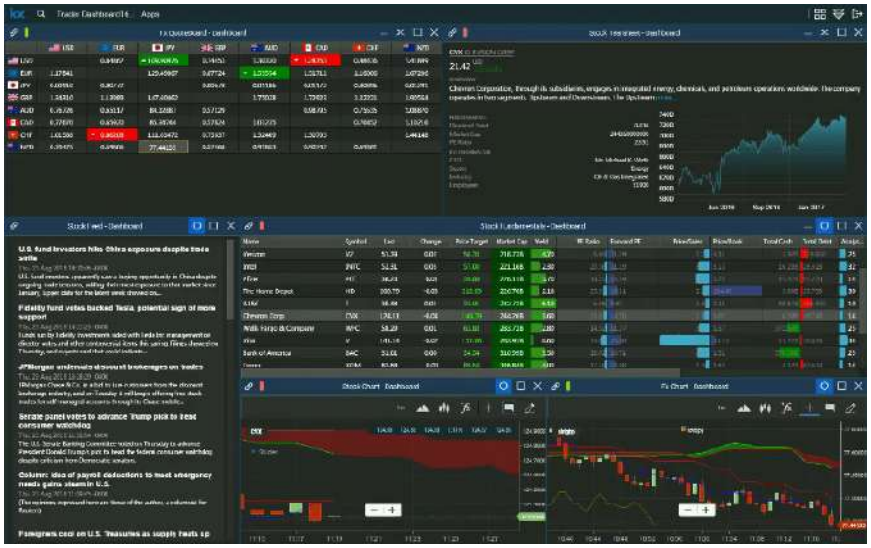

In addition to a wide range of native display options including tables, charts, 3D graphs, maps it includes filter, join, aggregate, and drill-down functionality to gain insights with real-time OLAP.

KX Dashboards also includes an extensible visualization layer for 3rd party integration and embedding custom visualizations via an SDK and a simple framework-agnostic API.

KX Fusion APIs – Simplifying Integration

While KX offers a single-stack solution to support the complete data analytics life-cycle from ingestion to processing and visualization, it also provides interfaces to other technologies and languages.

This approach enables organizations to preserve the value of existing investments and technical skills in other areas and accelerate solution development through code reuse and industry-standard interconnectivity. The functionality is delivered as Fusion Interfaces to enable non-KX programmers to interoperate with kdb+ data structures via a range of APIs including:

EmbedPy loads Python into kdb+, allowing access to a rich ecosystem of libraries such as scikit-learn, TensorFlow, and PyTorch. Python variables and objects become q variables – and either language can act upon them. Python code and files can be embedded within q code. Python functions can be called as q functions.

R can connect to kdb+ and extract data. Alternatively, you can Embed R inside q and invoke R routines. Q can connect to a remote instance of R and Q can load the R maths library

KX Fusion also includes publish and subscribe interfaces to commercial and open-source message buses like Kafka and Solace for further flexibility and ease of integration.

KX Streaming Analytics, given its capacity for storing and processing vast amounts of both streaming and historical data, is a particularly powerful platform for machine learning. There are additional technical reasons too including:

- array and times-series operations are perfect for the feature-engineering steps of sampling aggregating and joining datasets
- streaming analytic capabilities enable online training of models and real-time prediction
- the ability to manage, ingest, store and analyze huge datasets make KX the ideal engine to feed deep neural networks, where enormous volumes of data are required for effective training

Further interfaces include HDF5, Jupyter, Java, LDAP, MQTT, Prometheus, Protobuf, and FFI, an extension to kdb+ for loading and calling dynamic libraries using pure q.
Please visit code.KX.com/q/interfaces/ for further information.

KX in the Cloud

WHY KX?

Cloud-Ready, Certified and Optimized

KX is built for the cloud. Our streaming analytics platform takes advantage of native cloud architectures to deliver superior scalability and flexibility with cloud economies of scale. No longer limited to physical hardware constraints, KX can scale up or down instantly, as needed, depending on storage and compute requirements. Best of all, running KX in the cloud enables users to focus on delivering results rather than maintaining hardware infrastructure.

No matter what kind of cloud you use - public, hybrid, or private – KX is interoperable across all of them. KX in the cloud even integrates with existing on-premises KX installations. This level of flexibility allows workloads to be distributed accordingly to preferred platforms of choice while future-proofing for any future cloud migration needs.

KX is a certified Amazon Solutions Partner, Google Cloud Partner, and Microsoft Azure partner. Our platform is available via KX or marketplace purchase from AWS, GCP, or Azure directly. Customers using KX via the marketplace can onboard quickly, get automatic updates, and leverage unified billing all from their cloud vendor of choice. KX is optimized for native technologies and offers first-class connectivity to each cloud provider's native services. Whether you want to leverage object storage in AWS S3 or container deployments via GKE, KX provides a seamless elegant way to utilize your cloud investment.

Google Cloud Platform | aws marketplace | Microsoft | Azure Marketplace

Global Leader in Materials Engineering Solutions Saves Over 10,000 Hours a Year using KX

Global Product Manager

"We've revised our database architecture to include a [KX] database, a more performance based database to provide for that high speed data querying, as well as the data storage"

Organization: **Leader in Materials Engineering Solutions**
Geography: **Global**
Industry: **Semiconductors**
Employees: **22,000+**

This organization is the global leader in materials engineering solutions used to produce virtually every new chip and advanced display in the world. It's application is the most comprehensive control and productivity suite in manufacturing. It is powered by technologies that are widely proven and deeply mature in high volume factories worldwide.

Replaced legacy system at significant lower cost, including a 10x improvement in required hardware

KX

Leader in Materials Engineering Solutions

THE CHALLENGE

For years, the semiconductor and electronics industries has looked for solutions to solve complex manufacturing problems. A key challenge is the need for faster speed-to-data. A common customers complaint is that with the explosion of data, query speeds are too slow and want see improvements in speed to data and total cost of ownership. As an industry leader with a vast customer base, this organization proactively searched for a solution to meet these customer needs.

WHY KX

Over a year, 7 vendors were evaluated across 10 different objectives to enhance its application. This was an opportunity to look beyond the scope of traditional databases and found timeseries databases as a preferred solution that has had huge growth in adoption over the years. KX proprietary database is built for real-time data use cases taking advantage of a timeseries architecture coupled with relational database technology, a programming language to enable real-time calculations in the database, and the ability to manage both streaming and historical data in the same database. KX was selected for its superior performance results on ultra-fast speed to data, lower total cost of ownership, its ability to scale, and new user experience.

Saved 10,000+ hrs a year from false alerting and improved detection.

THE BENEFITS

- Over **100 million readings** per day
- Delivers **high performance** historian and streaming analytics
- **Supports +600 utilities** in electricity, water, gas, and transportation
- **Detect and predict** conditions faster

Improved Query Speeds by over 30x for an AI based Fault Detection Application.

Przemek Tomczak,
Sr VP IoT and Utilities, KX

"New tools are needed for this new era of digital information overload. Rich real-time analytics presents a significant opportunity for optimizing business operations, customer engagement, and offering new services and products."

Organization: **Provider of Advanced Data Analysis Solutions for Manufacturing**
Geography: **Global**
Industry: **Manufacturing**
Employees: **350+**

This global manufacturer is a leading provider of real-time, engineering and automation solutions for the global semiconductor and manufacturing industries for 20+ years. Its AI solutions for smart manufacturing are shaping the factories of the future, by connecting data driven manufacturing organizations with the knowledge to act.

Supporting next generation fault detection within its AI application, delivering full sensor trace analysis in real-time, 10x reduction in false alarms, and limitless modeling.

KX

Provider of Advanced Data Analysis Solutions for Manufacturing

THE CHALLENGE

As a leading supplier of smart manufacturing applications, it was mission critical to find a solution to manage massive volumes of sensor data closer to real-time, and that can be easily integrated within its ecosystem of applications. They needed a solution to replace their traditional database solution, to keep up with the requirements around smart manufacturing and can integrate with its innovative AI solution. To stay competitive and expand their capabilities for its customer base they needed a database system that manages timeseries data providing performance with ingestion and queries, fast analytics in real-time, the ability to work with both streaming and historical data in real-time, and a scalable and interoperable architecture.

WHY KX

After extensive testing against their legacy solution, KX outperformed for queries and ingestion, while greatly enhancing the ability to scale at lower TCO. In addition, with KX they exceeded their performance requirements, by orders of magnitude faster than the alternative – saving significant capex and opex costs.

Over 30x speed improvements in queries enabling real-time action

THE BENEFITS

- Easily able to **ingest 17 Million records per second**
- **Over 30x speed** improvements in queries enabling real-time action
- Supporting **next generation fault detection**
- New capabilities added for **enhanced service**

KX Clients

CLIENTS

CME Group · Standard Chartered · JPMorgan Chase & Co. · Aston Martin Red Bull Racing kx Innovation Partner

ANZ · Citadel · BIVA · Quantile · Bank of America Merrill Lynch

SEC · Singapore Exchange · Credit Suisse · Morgan Stanley · FDL Frontier Development Lab

CTC Chicago Trading Company · utilismart Corporation · Survalent · Barclays · Goldman Sachs

Mizuho · RxDataScience · SMBC Nikko · HSBC · Thomson Reuters

iex · ieso · Fidelity Investments · 深圳证券交易所 Shenzhen Stock Exchange · Square Point

UBS · republic wireless · Euronext · Total Committed to Better Energy

KX

STREAMING ANALYTICS

CAPTURE ANY DATA
DRIVE DECISIONS