

The Next-Generation Analytics

Database for IoT

Kinetica is the world's most performant and scalable analytics database for IoT workloads, which demand real-time insights on location and time series data.



Real-Time Analytics Database

Leverage simultaneous streaming data ingest and analysis, in the context of historical and integrated data. Realize the full benefits of your real-time data by fusing it with other data, providing full context and allowing you to use more history than you could with a streaming platform alone.

Real-Time Ingest

Kinetica's lockless, distributed, key-value database enables real-time ingestion at scale by minimizing overhead. Native Kafka integration accelerates the ingest of new streams.

Context

Real-time data benefits from the additional context provided by a database to fuse multiple data sources and leverage deep history. ANSI SQL makes it easy to use and integrate with popular front end tools and applications.

Analytics

Perform advanced analytics in-database, using hundreds of fully vectorized in-database functions or bring-your-own models with ONNX and PMML for real-time machine learning inference.



U.S. AIR FORCE



BBP
Highest Building Performance



Best-in-Class Location Intelligence

Analyze many data points with time and space components simultaneously, using a single vectorized database instruction. Getting value from sensor data requires blending spatial, time series and graph analytics.

Spatial Functions

Sensor data is typically geo-tagged, necessitating new geo-spatial capabilities and competencies. Kinetica has over 100 spatial functions such as, geo-joins, point in polygon, map matching,, and many more. Kinetica can also create visualizations from geospatial data like heatmaps, choropleths, contours, and others.

Time Series

Sensor data is temporal. Kinetica has a multitude of advanced time series functions such as aggregations, windowing functions, inexact joins, and others.

Graph

Seamlessly use your relational data in a native graph context for understanding geospatial and non-geospatial relationships. Perform real-time route optimization and other graph use cases using hyper efficient vectorized graph algorithms.

Visualize at Scale

Visualize and interact with billions of data points, instantly leveraging powerful server-side rendering.



Greater Performance on Fewer Nodes

Achieve exponential performance gains by combining traditional parallel processing with vectorization. Maintain performance SLAs while slashing data infrastructure spending. Reduce environmental impact with better computing efficiency for better resource efficiency.



A top U.S. bank, a 700 node Spark cluster running queries in hours took seconds on 16 nodes of Kinetica.



A top U.S. retailer consolidated 100 nodes of Cassandra and Spark into 8 Kinetica nodes.



A top U.K. pharma achieved identical performance between a 88 node Impala cluster and a 6 node Kinetica cluster.