



# Kinetica, the Database for Time & Space

## Enabling the Most Innovative IoT Analytics Use Cases in the World

**With 25 billion IoT devices now in use, the Internet of Things is moving from the initial phase where value came from connectedness, to the next phase where value comes from turning data exhaust into actionable intelligence that drives industry transformation. The focus over the next decade for IoT initiatives centers on improving business, life, and society, using the insights gained.**

Even though organizations using IoT are generating masses of data, making sense of this data remains a hurdle. Most businesses are not harnessing or analyzing IoT data and are thereby missing out on insights that could improve business decisions. According to McKinsey, the Internet of Things will generate up to \$11T in economic value by 2025.



Unlike transactional data and web data, IoT data is inherently geo-tagged and time series in nature. Machine generated observations record the longitude and latitude of things at continuous intervals. IoT data is also characterized by extreme volume and velocity with rapidly perishable value. Prior generation data warehouses and data lakes lack the sophisticated time series and

geospatial analytics required to derive valuable insights. Cobbling together various technologies for time series, geospatial, real-time analytics, and streaming results in latency, complexity, and cost issues.

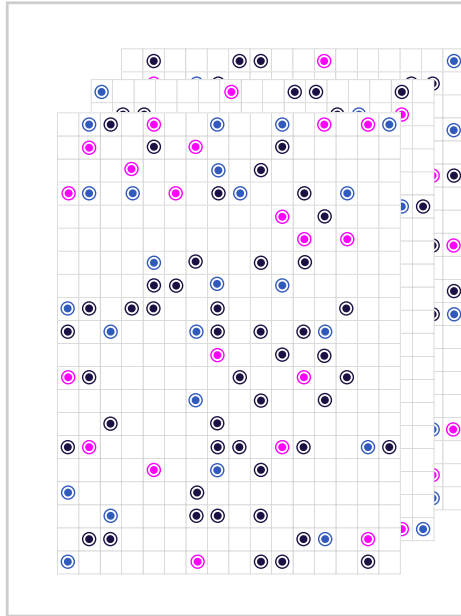
New times call for new capabilities. The breakthrough engineering behind Kinetica is data level parallelism. Data level parallelism leverages vectorized chips on Azure to operate on a set of values at one time. Prior generations of data analytics platforms exclusively rely on task level parallelism that operates on multiple tasks at one time. Task level parallelism requires burdensome indexes, materialized views, and other hints that restrict analytic agility. Data level parallelism is orders of magnitude faster than task level parallelism for certain workloads, including real-time processing, time series analysis, spatial calculations, and graph solving. Data level parallelism compute efficiencies radically reduces development effort.

Current generation databases require burdensome indexes, materialized views, and other hints to achieve performance that restrict analytic agility. Requirements must be known in advance and designed with specific use cases in mind. This impairs analytic agility. Through Kinetica's use of vectorization, that is exponentially more performant, no such tricks and workarounds are needed, thereby radically increasing analytic agility.

Ingest multiple data streams in real-time





Fuse data for full context



Innovative architecture enables analytic agility

Performance analytics across space and time



## 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.




**The United States Postal Service uses Kinetica to combine millions of streaming location events from vehicle transmitters with billions of historical events, all available for immediate analysis.**



### 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.


A graphic featuring three fighter jets flying in formation against a light blue sky. Below the jets is the U.S. Air Force logo, a stylized eagle with wings spread. Underneath the logo, the text 'U.S. AIR FORCE' is written in a bold, black, sans-serif font.

**The United States Air Force uses Kinetica to analyze billions of aircraft track points per day to obtain a fused operational picture across all of North American airspace and more rapidly detect anomalies that may be airborne threats.**



## 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.

Kinetica helps many of the world's largest companies solve some of the world's most complex problems, including the US Air Force, NORAD, USPS, Citibank, Telkomsel, MSI, OVO, and Softbank, among others. Kinetica is the first fully vectorized database in the cloud, enabling organizations to fuse historical and streaming data to perform real-time analysis of spatial and temporal data at unparalleled performance and scale.

**kinetica**

Organizations across the public sector, financial services, telecommunications, energy, healthcare, retail, automotive, and beyond can load and analyze fast-moving data simultaneously, delivering instant insight. Kinetica offers flexible deployment, pricing, and support models across private and public clouds.

Kinetica has a rich partner ecosystem, including Dell, HP, IBM, NVIDIA, and Oracle and is privately held, backed by leading global venture capital firms Canvas Ventures, Citi Ventures, GreatPoint Ventures, and Meritech Capital Partners. For more information and to try Kinetica, visit [kinetica.com](https://www.kinetica.com) or follow us on [LinkedIn](#) and [Twitter](#).