

**Kinetica is a GPU-accelerated, in-memory analytics database that delivers truly real-time response to queries on large, complex and streaming data sets: 100x faster performance at 1/10 of the hardware of traditional databases. Kinetica’s fully distributed architecture, and simplified data structures lead to more predictable scale out on industry-standard hardware. In-database analytics via user-defined functions (UDFs) open the way for converged AI and BI workloads accelerated by GPUs. Kinetica comes with native geospatial and visualization pipeline for interactive location-based analytics. Organizations use Kinetica to simultaneously ingest, explore, analyze and visualize data within milliseconds to make critical decisions and find efficiencies, lower cost, generate new revenue, and improve customer experience.**

## GPU Acceleration Overcomes Processing Bottlenecks

GPUs are designed around thousands of small, efficient cores that are well suited to performing repeated similar instructions in parallel. This makes them well-suited to the compute-intensive workloads required of large data sets such as machine learning, deep learning, NLP, and OLAP.

- High performance GPUs from NVIDIA feature over 4,000 cores, versus 16 to 32 cores per typical CPU-based device.
- Parallel processing is ideal for scanning entire dataset & brute force compute.
- High performance computing trend to using GPUs to solve massive processing challenges.
- GPU acceleration brings high performance compute to commodity hardware.

**Kinetica** was initially built from the ground up for the US Army Intelligence and Security Command to track terrorist and other national security threats in real time—producing instant results while visualizing insights across >200 sources of streaming datasets.

With the growth of data from IoT, transactions and other sources, business users are running up against similar computational bottlenecks with the challenge of streaming and analyzing data in truly real time.

Kinetica is already battle tested, and has a strong partnership with NVIDIA for accelerated hardware. Additionally, Kinetica is certified on IBM, HP, Dell, SuperMicro and Cisco servers.

The US Postal Service relies on Kinetica to ingest, analyze, and visualize large and complex streaming data for real-time route optimization and on-time mail delivery.

## The Kinetica Advantage

### Performance

Makes real time a reality by ingesting massive-scale, streaming data, while delivering analytic results and producing visualizations in milliseconds. With Kinetica’s distributed architecture, the compute power is unmatched. Ingest streaming data—billions of records per minute—so applications and users can access “up to the moment” analytics as data is streaming through your enterprise.

- Realize 100x performance improvement on queries compared to CPU-based in-memory solutions
- Produces analytic results and visualizations in milliseconds
- Holds 100s of terabytes of data in-memory to enable an extremely low-latency analytics layer for your most critical decision-making across your enterprise

### Advanced Analytics and Machine Learning with In-Database Processing


- User-defined functions (UDFs) enable compute as well as data-processing, within the database
- Machine learning/AI libraries such as TensorFlow, Caffe, and Torch can run in-database alongside, and converged with, BI workloads.
- TensorFlow bundled with database for simpler deployment and faster time to value
- Fuzzy Logix integration for rich in-database analytics

### Enterprise-Grade Database

Kinetica’s enterprise-grade database is designed for simplicity, manageability, scalability, availability and security so enterprises can minimize complexity and quickly derive business value.

- **Simplicity:** No typical tuning or indexing required; ask and answer any question in real time. Familiar relational database seamlessly plugs into existing data architectures for easy integration with databases, data lakes, business applications, BI and ETL tools.

## REFERENCE ARCHITECTURE



## ADVANCED IN-DATABASE ANALYTICS ARCHITECTURE



- **Manageability:** Dynamic resource provisioning (add/remove nodes), enhanced process management, and improved resource utilization (compression, dictionary encoding) simplify administration and minimize costs
- **Scalability:** Easily scale up or out. Data written to Kinetica is automatically distributed across the cluster for scalable data management. Distributed parallel processing across the cluster improve performance for compute-intensive analytics and machine learning workloads.
- **High availability:** Kinetica supports inter-cluster active-active configuration and data replication for five 9's uptime
- **Security:** Kinetica features LDAP authentication with support for Active Directory, authorization with role-based access control, full encryption for data in movement and at rest with SSL, PLS, AES-256 and more.
- **Deploy on-premises or in the public cloud** for faster time to value.

### Easy APIs and Integration

Plug into your existing infrastructure, complementing your current business applications, EDW and Hadoop environment. Kinetica ships with industry standard connectors to make it easy to integrate with your existing infrastructure.

- Open source integration components include:
  - Apache NiFi
  - Apache Spark and Spark Streaming
  - Apache Storm
  - Apache Kafka
  - Apache Hadoop
- Kinetica's APIs are fully supported in REST, Java, Python, C++, JavaScript and Node.js
- Kinetica ships ODBC and JDBC drivers for integration with industry standard BI, ETL and SQL tools
- Full SQL-92 query support

### Complete Native Visualization and Geospatial Capabilities

Kinetica delivers realtime geospatial processing and visualization, ideal for fast moving, location-based IoT data, with significantly faster performance and at a fraction of the cost of legacy geospatial solutions.

Kinetica's geospatial and visualization features include:

- Support for geospatial data types such as points, shapes, tracks, labels
- A fully GPU-accelerated distributed rendering pipeline
- Visualizations based on billions of points in seconds
- Realtime drill-down all the way to individual points
- A full range of geospatial analytics
- Updates in realtime as data changes
- Geospatial visualizations are exposed as Open Geospatial Consortium (OGC) compliant Web Map Services (WMS) allowing for interoperability with other industry-standard geospatial applications

### Savings

Embedding GPUs into Kinetica's architecture means there are 4,000-plus cores per device, versus 8 to 32 cores per CPU-based device.

- Smaller hardware footprint with tangible savings: 1/10 the hardware costs on average, and 1/20 the power and cooling.
- Offload expensive relational databases to Kinetica and consolidate to a single data product with seamless integration to Hadoop for long-term storage

### Runs on Industry-Standard Hardware and Public Cloud

#### Example servers include:

Cisco UCS C240M4  
Dell PowerEdge R730  
HPE ProLiant DL380 Gen9 Servers, Apollo 2000  
IBM S822LC 'Minsky' machine  
NVIDIA DGX Station

#### Available on public cloud:

Amazon Web Services  
Microsoft Azure  
Google Cloud Platform  
Nimbix



For more information  
visit [qcmtech.com](http://qcmtech.com)



For more information on Kinetica  
and GPU-accelerated databases,  
visit [kinetica.com](http://kinetica.com)