kinetica

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 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.

REFERENCE ARCHITECTURE

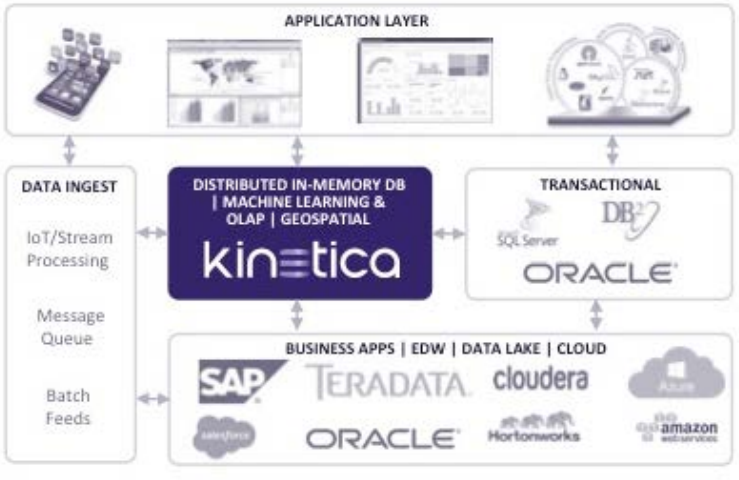

The Kinetica Advantage

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

Simplicity

Kinetica is a cohesive, unified suite in one product. Our patented database design redefines how data is stored and processed to maximize the massive parallel compute power of GPUs.

- Plug into existing data architecture; deep integration with open source and commercial frameworks
- None of the typical tuning or indexing required by traditional systems is needed, which frees business users to ask and answer any question in real time
- Connect with common BI tools like Tableau, Kibana and Caravel
- A converged solution, replacing complex solutions that are made up of multiple disparate components

Kinetica 6.0 Features

User Defined Functions (UDFs) converge AI and BI workloads accelerated by GPUs Advanced in-database analytics converge AI, BI, ML, NLP, etc. into one powerful DB for business users.

NVIDIA NVLink™ support. Accelerates database performance as data moves between GPU and CPU 3x faster on average compared to the traditional PCI Express.

'Reveal' Visualization Framework for interactive real-time data exploration. Location-based analytics integrate with all major map providers: Bing, ESRI, Google, Mapbox, open street maps.

VRAM Boost Mode. Prioritize certain data or use cases for ultra low latency. "Pinning" data in very fast GPU Video RAM (VRAM) for lightning fast query performance, while also still being able to leverage cluster-wide system RAM to both scale up and scale out to multi-terabyte in-memory processing.

Full SQL-92 query support. GPU-accelerated SQL-92 query support through certified JDBC and ODBC connectors.

New visual installer. Allows for easy click button installation of Kinetica across hundreds of nodes.

Runs on Industry-Standard Hardware

Example servers include:

Cisco UCS C240M4

Dell PowerEdge R730

HPE ProLiant DL380 Gen9 Servers, Apollo 2000

IBM S822LC 'Minsky' machine

NVIDIA DGX-1

ADVANCED IN-DATABASE ANALYTICS ARCHITECTURE

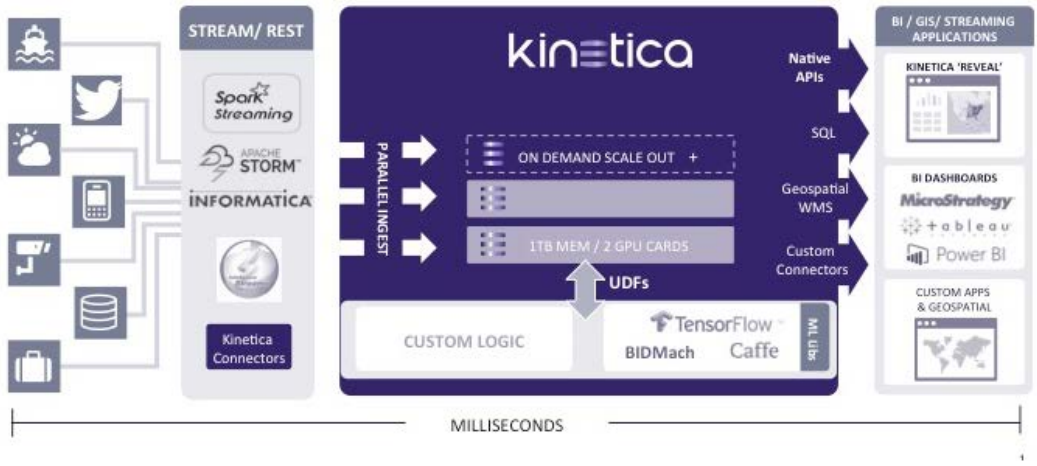

Predictably Scalable

Built from the ground up as a distributed database, Kinetica can be easily scaled, either horizontally or vertically, to meet requirements and SLAs for both data ingestion and OLAP query performance.

- Clients or systems writing data to Kinetica are automatically routed to parallel connections across the cluster
- OLAP queries are executed using fully distributed GPU-acclerated processing across the cluster

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 and SQL tools
- Integration with industrystandard ETL tools such as Informatica for data ingest and publishing

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 middleware.

Kinetica's geospatial and visualization features include:

- 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

kinetica

For more information on Kinetica and GPU-accelerated databases, visit kinetica.com