NVIDIA

kinetica

Solution Brief

CHALLENGES OF TODAY'S DATABASES

Data Warehouses. RDBMS and data warehouse technologies enable organizations to store and analyze growing volumes of data on high performance machines, but at high cost.

Distributed Computing. Hadoop and MapReduce enable distributed storage and processing across multiple machines. Storing massive volumes of data becomes more affordable, but performance is slow.

In-Memory Databases. Affordable memory allows for faster data read and write speeds, and enables faster analytics. At scale, compute processing now becomes the bottleneck. Response times seriously degrade for high-cardinality datasets, and systems struggle to ingest and query simultaneously; they can't deliver acceptable response times with real-time, streaming data.

Large, Complex, and Streaming Data tax the compute-intensive workloads of CPU-based solutions.

NVIDIA GPU Engine Powers Kinetica's Accelerated Data Processing

Businesses have spent the last decade determining how to store, manage, and query data to drive business decisions. Today's businesses must unleash the power of AI to transform their data driven businesses into AI enterprises.

NVIDIA GPUs provide Kinetica the power to perform brute-force queries on these large, complex, and streaming datasets. The outcome is remarkable performance increases and tangible savings on hardware. Benchmarks show that NVIDIA GPUs enable Kinetica to deliver 100 to 1000x faster analytic performance than other CPU-based in-memory databases.

Kinetica's advanced geospatial capabilities and its Reveal interactive visualization framework are particularly well suited for analytics on fast-moving, location-based IoT data. Kinetica has native geospatial object types and functions and a full rendering pipeline, enabling users to visualize and interact with data in real time.

Kinetica is further extendable through user-defined functions. UDFs have direct access to CUDA APIs, and can take full advantage of the distributed architecture of Kinetica. This is the first time in-database processing is available within a database that can fully utilize the parallel compute power of the GPU on a distributed platform.

Enterprises can now perform advanced analytics faster and converge AI and BI workloads on a single database platform with this industry-leading, end-to-end solution. Kinetica and NVIDIA provide unmatched performance, predictable scalability across multiple high-density nodes, and seamless integration with industry-standard connectors to data sources and applications.

Recommended Hardware

TESLA
Servers in every shape and size

DGX-1
The AI supercomputer for instant productivity

CLOUD
Everywhere

Unmatched Performance

Discover and act on business insights faster with a Kinetica and NVIDIA GPU-enabled solution.

NVIDIA® NVLink™ is a high-bandwidth, energy-efficient interconnect that allows data sharing at rates 5 to 12 times faster than the traditional PCIe Gen3 interconnect.

50-100X Faster on Queries with Large Datasets

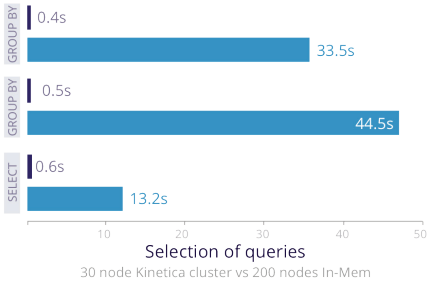

Scalable Across Multiple High-Density Nodes

Distributed architecture scales on demand linearly and predictively for both ingestion and querying.

Distribution across GPUs delivers more throughput with less infrastructure : 1/10 the hardware costs on average and 1/20 the power and cooling.

Real-Time Data Handlers for Structured and Unstructured Data

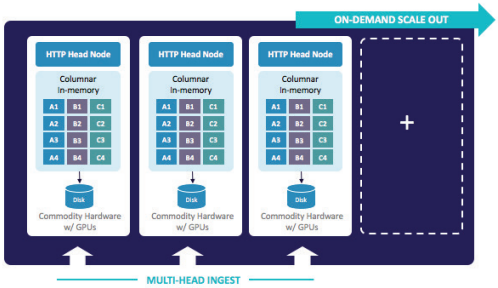

Seamless Integration with Industry-Standard Connectors to Data Sources and Apps

Integrate easily with open source and commercial frameworks, and with third-party BI applications.

APIs are fully supported in REST, Java, Python,C++, Javascript, and Node.js.

ODBC and JDBC drivers integrate with industry-standard BI and SQL tools.

User-Defined Functions (UDFs) Accelerated by GPUs Converge AI and BI Workloads

Advanced in-database analytics converges Artificial Intelligence, Business Intelligence, Machine Learning, natural language processing, and other data analytics into one powerful database platform.

Direct access to NVIDIA CUDA APIs via UDFs deployed within Kinetica.

Converged AI and BI

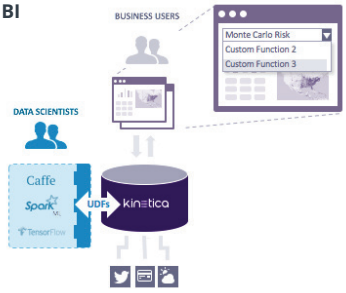

Interactive Location-Based Analytics

Distributed geospatial pipeline for analytics on fast-moving, location-based data with native geospatial object types.

Kineticas Reveal interactive visualization framework for real-time data exploration.

High-Speed Geospatial Pipeline

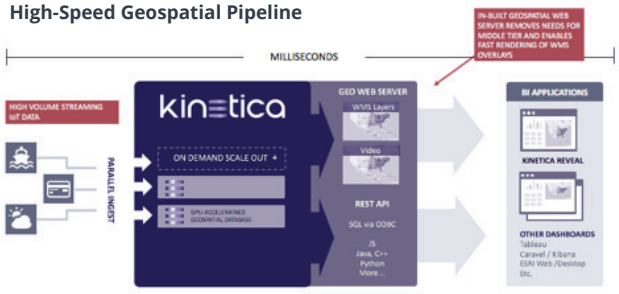

kinetica

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