

NVIDIA AND KINETICA FOR ACCELERATED ANALYTICS

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 accelerated analytics to transform their data-driven businesses into AI enterprises.

INTEGRATED SOLUTION

The Kinetica engine, running on NVIDIA® DGX™ Systems and NVIDIA GPUs, provides real-time analytics on data in motion and at rest with at least 10–100X faster performance at 1/10 the cost of traditional systems. NVIDIA and Kinetica together deliver unmatched performance, predictable scalability across multiple high-density nodes, and seamless integration with industry-standard connectors to data sources and applications. Kinetica's engine includes a distributed, in-memory GPU database and delivers advanced analytics, location-based visualization, and streamlined machine learning.

INDUSTRY CHALLENGES

- > **Data Warehouses:** Relational database management systems (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 or streaming data.

NVIDIA and Kinetica. Better Together.

Unmatched Performance

- > The NVIDIA GPU-accelerated Kinetica engine takes advantage of vector and matrix processing for real-time ingestion, processing, and analytics of both data in motion and data at rest with at least 10–100X faster performance.
- > NVIDIA NVLink™ is a high-bandwidth, energy-efficient interconnect that allows data sharing at rates 5 to 12X faster than traditional PCIe interconnects.

Scalable Across Multiple High-Density Nodes

- > NVIDIA GPUs deliver more throughput with less infrastructure: 1/10 the hardware costs on average and 1/20 the power and cooling.
- > Distributed architecture scales on demand linearly, reliably, and predictively for both ingestion, querying, and AI.

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

- > Kinetica integrates easily with open-source and commercial frameworks such as Apache Kafka, NiFi, Spark, Storm, and extract, transform, and load (ETL) tools.
- > Application programming interfaces (APIs) are fully supported in REST, Java, Python, C++, Javascript, and Node.js. Open database connectivity (ODBC) and Java database connectivity (JDBC) drivers integrate with industry-standard business intelligence (BI) and structured query language (SQL) tools.

Interactive Location-Based Analytics

- > The NVIDIA GPU-accelerated Kinetica “Reveal” interactive visualization framework enables real-time data exploration.
- > A distributed geospatial pipeline includes native geospatial object types, functions, WMS support, and map-based visualization for analytics on fast-moving, location-based data.
- > Server-side rendering leveraging the power of GPUs on billions of points of data to create a seamless and highly performant experience.

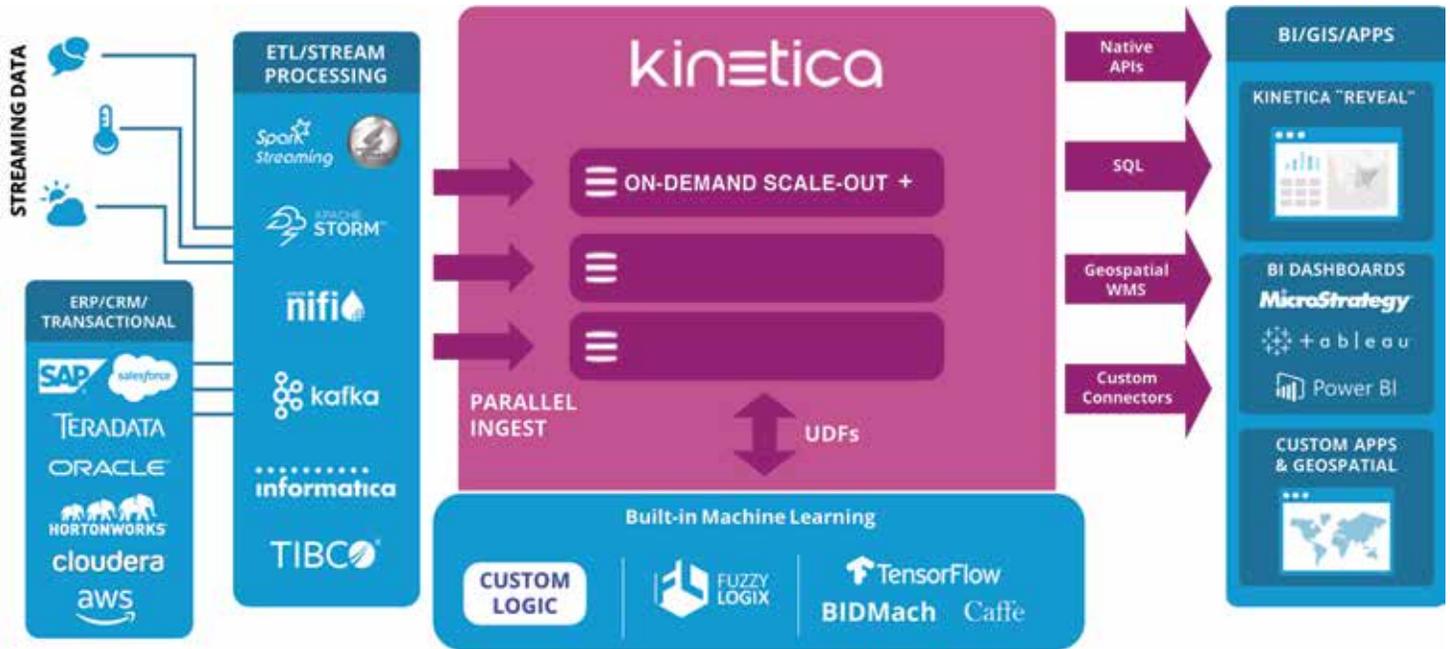
Machine Learning Augmented Analytics

- > Advanced in-database analytics with UDFs bring together AI, BI, machine learning, deep learning, natural language processing, and other data analytics into one powerful database platform.
- > The framework provides direct access to NVIDIA CUDA® APIs via UDFs deployed within Kinetica.
- > It seamlessly integrates with open-source and third-party machine learning frameworks and tools such as TensorFlow, Caffe, Torch, Fuzzy Logix, Deepchem, H2O, and Anaconda.

Enterprise-Grade Database Capabilities

- > The engine delivers simplified administration and dynamic resource provisioning, five nines availability, and easy deployment on premises or in the public cloud, including AWS, Azure, and Google.
- > Robust security includes LDAP authentication, authorization with role-based access control, and full encryption for data in movement and at rest with SSL, PLS, and AES-256.
- > Advanced compression techniques for data storage maximize memory and disk utilization.

Kinetica Architecture for AI and BI Workloads



Industry Insights

Category	Use Case
Telecommunications	Perform advanced customer churn analysis and monitor network performance in real time.
Finance	Leverage real-time risk analysis for stronger portfolio management.
Federal	Bring fast and actionable analytics to security and public services.
Retail	Manage supply chains in real time for replenishment and inventory management.
Automotive	Analyze large volumes of autonomous vehicle data in real time to improve systems.
Logistics	Manage and optimize traffic patterns and congestion in real time.
Energy	Implement real-time drilling and well analytics.
Healthcare	Accelerate drug discovery with machine learning and advanced analytics.

Recommended Hardware

NVIDIA data center GPUs are available in servers, supercomputers, and cloud services around the world. You can now get end-to-end accelerated analytics solutions powered by NVIDIA GPUs with supporting software technologies and support from NVIDIA experts.

NVIDIA TESLA®
SERVERS IN EVERY SHAPE
AND SIZE



Hewlett Packard Enterprise IBM CISCO DELL

DGX SYSTEMS
THE AI SUPERCOMPUTER FOR
INSTANT PRODUCTIVITY



 NVIDIA.

CLOUD
EVERYWHERE



amazon Microsoft Azure IBM Cloud Google Cloud Platform

Find Out More

NVIDIA GPUs for accelerated analytics help customers effectively analyze, visualize, and unleash the power of AI to transform their digital business into an AI enterprise.

Website: www.nvidia.com/analytics

Twitter: [@NVIDIAADC](https://twitter.com/NVIDIAADC)

Blog: blogs.nvidia.com

Kinetica's GPU-accelerated analytics database delivers fast in-database analytics, machine learning, and interactive location-based analytics.

Website: www.kinetica.com

Contact: info@kinetica.com

Partner Webpage: www.kinetica.com/partner/nvidia

Twitter: [@KineticaHQ](https://twitter.com/KineticaHQ)

Blog: www.kinetica.com/blog