

December 29, 2023



OFFERING OVERVIEW

Kinetica Drives Breakthrough Spatial and Time-Series Analysis Aided by GenAI

Kinetica Database and Cloud Service Tap GPU Power
On Demand to Analyze Data on a Massive Scale



Doug Henschen

VICE PRESIDENT AND PRINCIPAL ANALYST



Produced exclusively for Constellation Research clients

TABLE OF CONTENTS

- Executive Summary3
- Kinetica Company Overview4
- About Kinetica5
- Functional Capabilities.....7
- Use Cases 15
- Pricing 16
- Analysis and Observations 17
- Recommendations 19
- Analyst Bio..... 21
- About Constellation Research 22



EXECUTIVE SUMMARY

Kinetica is a real-time database that is purpose-built to enable analytics on time-series and spatial data on a massive scale. The product was developed in 2016, just as data from sensors and machines began to explode. Spatiotemporal data from smartphones, vehicles, aircraft, cyberattacks, networks, satellites, weather tracking, and everything else that moves and is measured over time is growing faster than any other data type. Kinetica's technical advantages enable it to quickly handle complex, high-scale analyses that would cripple conventional data warehouse platforms and even specialized databases lacking Kinetica's combined time-series and spatial analytical capabilities.

This Offering Overview report recaps the differentiating features and capabilities of Kinetica, including comprehensive native vectorized processing, distributed and columnar architecture, on-demand graphical processing unit (GPU) acceleration, and industry-leading temporal and spatial analytical capabilities. The report also details Kinetica's new SQL-GPT feature, which uses generative AI (GenAI) for natural-language querying by way of a tunable large language model (LLM) built directly into the database to ensure accuracy, consistency, and data security. Technology buyers should use this report to evaluate Kinetica for implementation.

BUSINESS THEMES



Data to Decisions



Future of Work



Technology Optimization

KINETICA COMPANY OVERVIEW

Kinetica DB Inc. is the developer of a real-time database that is purpose-built for enabling analytics on time-series and spatial data at a massive scale. The company was launched in 2016 by co-founders Nima Negahban and Amit Vij. Negahban, CEO and the original developer of the Kinetica database, previously developed GPU-powered systems for big data applied in use cases ranging from biotechnology to high-speed financial trading. Vij, Kinetica's president, previously developed software for analyzing and visualizing high-scale data. He worked in both the commercial and federal markets, helping to develop GPU-based geospatial intelligence systems for multiple U.S. defense and security agencies.

Kinetica's distributed, vectorized database was designed to take advantage of GPU power to analyze the increasingly massive scale of data streaming from connected devices and myriad digital channels. In 2020, Kinetica extended its database to run on modern, many-core x86 central processing units (CPUs) supporting AVX-512 (Intel's 512-bit, Advanced Vector Extension instruction set). CPU support greatly expanded deployment options—including to a database-as-a-service (DBaaS) offering, which is available on multiple public clouds, and to a freely downloadable Kinetica Developer Edition—while retaining superior vectorized performance and the option to use GPUs for the most demanding workloads.

In 2023 Kinetica became one of the earliest database vendors to innovate with GenAI, adding an integration with OpenAI's ChatGPT in May, a database-native LLM in September, and a real-time Vector Similarity Search feature in December.

KINETICA DB INC.

Company: Kinetica DB Inc.

Headquarters: Arlington, VA

Founded: 2016

Type: Private

No. of Employees: More than 50

Website: [Kinetica.com](https://www.kinetica.com)

ABOUT KINETICA

Overview

The Kinetica database (Kinetica) is designed to handle time-series and spatial analytics with super low latency and on a massive scale. It was developed in 2016 as data from sensors and machines—data combining both time-series and spatial meaning—was coming to the fore. Such data is generated at high volumes by automobiles, aircraft, cyberattacks, inventories, networks, buses and trains, satellites, smartphones, weather fronts, and anything else that moves and is measured over time. This spatiotemporal data and related analyses are the fastest growing of any type and are poised to surpass all other forms of data and analysis in terms of scale and value.

Kinetica's support for time-based analysis spans from milliseconds and seconds to minutes and hours, whereas spatial movement might range from the movement of goods between shelves in a warehouse to trucks between delivery destinations to distances between satellites in Earth orbits. Automakers are using Kinetica to study vehicle performance for predictive maintenance applications as well as driver behavior and traffic patterns for autonomous driving solutions. Utilities are studying the potential for outages to optimize their networks. Financial services companies managing complex portfolios are studying billions of market transactions each day to assess fast-changing risk and regulatory exposures. Insurers are predicting building and property damage ahead of and in the wake of storms, floods, or fires to optimize the deployment of adjusters and to spot fraud.

What sets Kinetica apart is its ability to handle the combination of time-series and spatial analysis at scale and with speed, returning results within less than a second that might take hours, days, or weeks for conventional databases, if not causing them to crash. Before this report gets into the specific technical capabilities of Kinetica, here's a brief survey of the high-scale analytical database market.

Market Segment: High-Scale Analytical Databases

The high-scale analytical data platforms market has flourished and drastically evolved over the last two decades. Whereas the market once centered on a dozen mainstream database management systems (DBMSs) that were deployed on-premises, today most of the attention has turned to data warehouse platforms, blended lake/warehouse platforms, and specialized databases that are offered either as services on public clouds, as software, or both.

Software that can be deployed on-premises or in private clouds is still very relevant, because this option satisfies data-residency and other regulatory or policy requirements that limit or prohibit the use of public-cloud deployments and/or services.

General-purpose relational DBMSs that use structured query language (SQL) are used for the vast majority of data warehouses and data marts that support business intelligence (BI) and analytics. Options such as Oracle Database and Amazon Redshift, for example, offer distributed architectures for scalability, columnar storage, and/or in-memory processing options for query acceleration as well as support for advanced analytics not expressed in SQL. Nonetheless, DBMSs used for warehouses and marts are general purpose in nature, using conventional x86 infrastructure and focused almost entirely on consistent and highly structured data organized into columns and rows.

Combined lake/warehouse (aka lakehouse) offerings emerging in recent years support data engineering, data science, and data warehouse/BI workloads against a shared storage environment underpinned by object storage. To address analysis of both highly structured and less-structured and variable data, lakehouse platforms, such as Snowflake and Databricks, harness two or more query/processing engines, with the most prominent combination being a SQL engine and a Spark or Spark-like engine. Here, too, these engines are general purpose in nature and are not designed for specific types of analysis.

Some of the database and lakehouse types noted above are sometimes described as “multimodal” databases because they address both transactional and analytical use cases (e.g., Oracle Database,

Microsoft SQL Server, and IBM Db2) and/or offer a few SQL-based or extended SQL functions to address specific types of analysis, such as JSON document analysis, graph analysis, or time-series analysis. But their one-size-fits-all, generalized capabilities are not up to the rigors of data-type-specific, high-scale analysis with the most demanding performance requirements.

Enter specialized analytical DBMSs and services, of which there are many types, including document databases, graph databases, time-series databases, spatial databases, in-memory databases, and more. In almost all cases, specialized databases are used alongside, rather than in place of, conventional data warehouse or lakehouse platforms. As the category name suggests, specialized databases typically address use-case-specific needs and requirements on a subset of an organization's data.

The downside of specialized DBMSs/services is that they can be too narrowly focused for organizations requiring a combination of analysis capabilities. A specialized database focused solely on time-series analysis, for example, might do well at temporal analyses while having glaring weaknesses in handling spatial analyses. What's more, the workaround of using two or more specialized databases/services to address a combination of needs introduces administrative and cost burdens, complexity, and latency. Constellation sees Kinetica as a specialized database, but one that addresses two types of analysis—time-series and spatial—that are frequently and increasingly required in tandem.

FUNCTIONAL CAPABILITIES

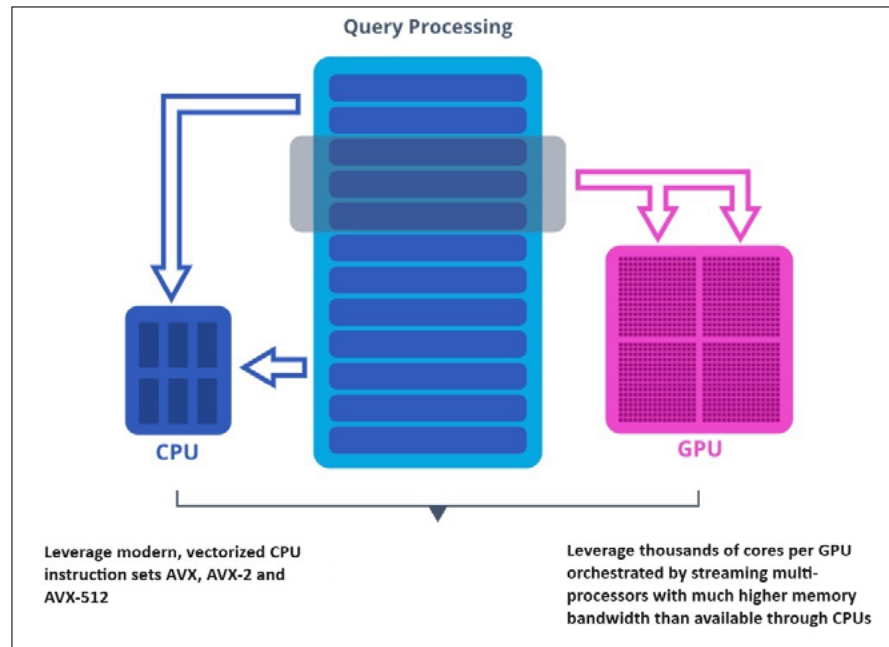
Kinetica is a vectorized, memory-first columnar database that harnesses GPU acceleration, as needed, for optimized speed and performance in combined time-series and spatial analyses. The database is designed for near real-time analysis on a massive scale. Kinetica supports subsecond (change data capture speed) ingestion and analysis, but it is not a streaming platform such as Confluent, Kafka, or Amazon Kinesis, nor a streaming data analytics product such as Azure Stream Analytics or Google Cloud Dataflow. Whereas streaming products focus on data only within a limited window of visibility, Kinetica is focused on enriching the latest data by way of complex queries and analyses, adding context discerned from all of the historical data available within the database.

Performance-enhancing design elements and features of the product include the following:

- **Distributed.** Kinetica’s distributed, massively parallel processing (MPP) architecture addresses the requirements of scale. Deployments can be scaled up, by increasing the storage capacity and processing power of each cluster, or scaled out, by adding nodes or clusters to the deployment. Sharding of data across nodes can be done automatically or specified and optimized by administrators.
- **Columnar.** Kinetica’s columnar design is crucial to performance, supporting ultralow-latency reads of data. Column-type options include standard base types—int, long, float, double, string, and bytes—as well as subtypes supporting date/time, geospatial, and other data forms. The database can be queried using SQL (ANSI SQL-92), JDBC, APIs, or applications that already support connectivity to PostgreSQL.
- **Vectorized.** Vectorization is the key differentiator behind Kinetica’s breakthrough performance (see Figure 1). Whereas most databases were designed and evolved with the limitations of conventional CPUs and sequential processing, Kinetica was designed from its inception to leverage the vectorization capabilities of GPUs and, as of 2020, modern CPUs. Most databases are now partially vectorized by supporting single instruction, multiple data (SIMD) instruction sets for CPUs (such as AVX, AVX-2, and AVX-512). The question is whether the product was rewritten to take full advantage of vectorization. In Kinetica’s case, *all* analytical functions (see the “Time-series and spatial analytics” bullet below) as well as core database functions were written from the start to take advantage of vectorization. This yields dramatic performance advantages, particularly on spatial and temporal queries at scale.

With the brute force of vectorized processing combined with vector-native analytical functions, Kinetica handles aggregations, predicate joins, windowing functions, graph solvers, and other complex analyses with far greater efficiency than conventional databases. Customers attest that massive multiway joins and full table scans that slow conventional databases to a crawl or that fail outright on those platforms can be handled within seconds by Kinetica.

Figure 1. Kinetica's Query Engine Automatically Assigns Vectorized Tasks to CPUs or, as Needed, GPUs, for Optimized Processing



Source: Kinetica

- **Memory-first tiered storage.** Tiered storage helps with optimization of performance and cost. Recent data from the last hour, day, or week, for example, can be stored in RAM or even GPU memory, while data from last week, last month, or last year can be stored more affordably on local disks or cloud storage. Kinetica intelligently manages data across GPU memory, system memory, SIMD, disk/solid state drive (SSD), and object storage such as Amazon S3.

Kinetica can also query external data warehouses, lakehouses, and data lakes, joining that data with data managed by Kinetica via highly parallelized queries. The database offers prebuilt connectors for a vast array of data sources, including Snowflake, Databricks, Salesforce, Kafka, BigQuery, Redshift, and many others.

Kinetica can query two types of external tables: materialized and logical. Materialized external tables pull data from external sources and cache that data in a persisted table within Kinetica. Data can be refreshed on demand or, configurably, on database startup. This mode ensures fast

response times at the cost of the data being only as current as the last refresh. Logical external tables pull data from external sources upon each query against the external table. This mode ensures that queries will always return the most current source data, but with performance penalties for reparsing and reloading data from source files with each query.

- **Time-series and spatial analytics.** Kinetica is purpose-built to support low-latency time-series and spatial analysis on large, fast-moving datasets. The combination offers obvious speed advantages over moving data or results between separate systems optimized only for time-series analysis or only spatial analysis.

Differentiation starts with support for specialized, low-latency data joins. On the temporal front, for example, organizations often need to combine data from multiple tables wherein time stamps don't exactly match. Kinetica's ASOF joins let users set an interval within which to combine data on values that are close to each other "as of" a specified time.

Spatial data contains geometry (points, lines, polygons) representing features such as locations, boundaries, or regions. Spatial joins match these geometries based on conditions such as containment, intersection, or proximity. Kinetica was purpose-built to handle the extreme efficiency challenges faced in spatial analysis, including:

- **Complexity.** Determining spatial relationships, such as containment, intersection, and proximity, requires complex algorithms, making these joins computationally intensive.
- **Indexing.** Traditional database indexing methods are less effective for spatial data. Spatial indexes such as R-trees or quadtrees are used, but optimizing these indexes efficiently for all possible spatial scenarios can be challenging.
- **Volume.** Spatial datasets can be large and complex. Processing huge amounts of spatial data efficiently is a significant challenge for most databases.

- **Optimization.** Query optimization for spatial joins is complex, due to the variety of spatial relationships and query approaches. It's difficult for most databases to efficiently evaluate all of these possibilities.

Kinetica was also purpose-built to accommodate sophisticated spatial joins that address complex operations that go beyond simple geometric relationships. Here are a few examples:

- **Distance joins.** Joining based on distance relationships, such as finding all points within a certain radius of other points, involves complex distance calculations that can be computationally intensive.
- **Spatial clustering.** Identifying clusters of spatial data points based on proximity involves algorithms such as density-based spatial clustering of applications with noise (DBSCAN) or K-means, both of which group points that are located closely together.
- **Spatial aggregation.** Aggregate functions based on spatial relationships are used to calculate the average value of points within a certain area or to find the centroid of a set of polygons.
- **Topological relationships.** Topological relationships, such as connectivity, adjacency, or containment, are analyzed to, for example, find polygons that share a boundary or points that lie within a polygon.
- **Fuzzy spatial joins.** Matching spatial data based on approximate or fuzzy relationships accommodates imprecise or uncertain spatial data.
- **Temporal-spatial joins.** Joining datasets based on both spatial and temporal relationships is done to, for example, find all the points within a specified distance of a certain location within a specific time frame.

Kinetica is also differentiated by its time-series, geospatial, and graph functions. On the time-series front, Kinetica supports windowing functions for aggregation and ranking over time as well as date/time functions and advanced date-time conversion functions. On the spatial front, more than

130 high-performance geospatial functions are supported through SQL or Kinetica's API. What's more, points and shapes are native geospatial objects in Kinetica, and tracks can be generated from sequential data to represent the paths of moving objects. On the graph front, Kinetica's graph API supports modeling of spatial data as graphs. SQL queries and built-in graph solvers, such as shortest-path analysis, can then be piped directly to maps—for example, to show the shortest path from point A to point B, factoring in speed limits, traffic, and other restrictions.

- **SQL-GPT.** Plenty of database vendors have introduced natural-language-to-SQL (NL2SQL) query features based on GenAI, but far fewer have delivered features offering the accuracy and performance that customers expect while meeting demanding security requirements. Kinetica introduced one of the database market's first integrations with OpenAI's ChatGPT with its first SQL-GPT release in May 2023.

In September 2023, Kinetica advanced SQL-GPT by adding the option of using a database-native LLM fine-tuned to be fully aware of Kinetica's syntax as well as conventional industry data definition languages (DDLs). This combined understanding enables users to harness Kinetica's specialized analytical functions—such as time-series, graph, geospatial, and vector search—via natural language queries.

In December 2023, Kinetica added SQL-GPT for Telecom, a database-native version of the LLM fine-tuned for the telecommunications industry. Customers can now choose to use a public LLM to power SQL-GPT (such as OpenAI ChatGPT or another option), or they can use Kinetica's database-native LLMs, with more industry-specific tunings expected to follow.

SQL-GPT and Kinetica's native LLM options meet customer expectations for accuracy, performance, privacy, and security, as detailed below.

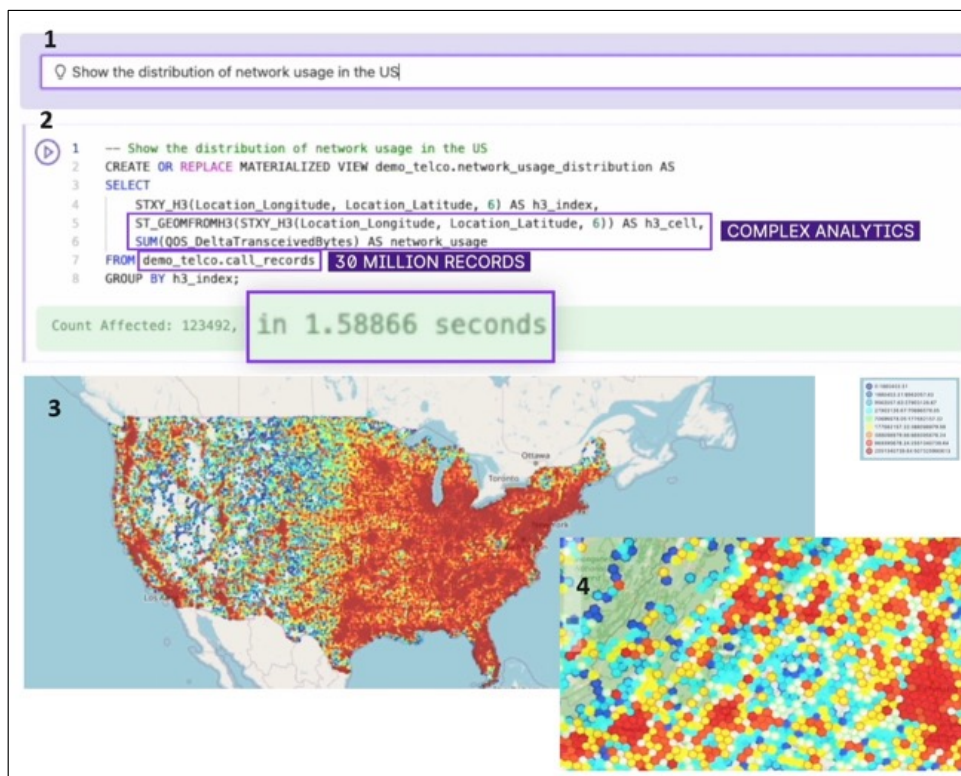
- **Accuracy.** Most public LLMs, such as OpenAI, are tuned for creativity, so they can respond to queries inconsistently. Kinetica's LLMs, in contrast, are tuned to generate accurate, consistent, and optimized SQL and Kinetica-specific expressions for functions that go beyond

standard ANSI SQL. The result, says Kinetica, is unwavering consistency and a reliable guarantee that queries will remain functional over time.

- **Performance.** Natural-language querying is, by its very nature, an iterative conversation with data, but the conversation stops when the underlying database can't quickly respond to a query such as "Show the distribution of network usage in the U.S." The native LLM generates a query that triggers a sophisticated, Kinetica-supported geospatial join (a capability not supported in standard SQL) that unites the shape of the U.S. with the precise, granular details of network usage, by location, at that point in time (see Figure 2).

As a next step, the Kinetica database harnesses vectorized CPU processing or, as needed, GPU processing, to analyze relevant data and deliver results with extraordinary speed. Finally, Kinetica can then quickly deliver interactive visualizations, such as multilayered maps, that

Figure 2. Kinetica SQL-GPT Turns (1) Natural Language Into (2) Complex SQL Queries for Analyses Such as (3 & 4) Interactive Visualizations of Spatial Data



Source: Kinetica

can be explored down to precise underlying details, such as the state of individual network nodes. This might trigger additional queries, natively translated by Kinetica LLMs, seeking more detail on specific cities or specific types of network activity.

- **Privacy.** While some enterprises and public agencies are banning the use of public LLMs that may use customer-submitted data to train their models, SQL-GPT uses only metadata (e.g., table and column names and types) to generate accurate SQL, never sending customer data to public LLMs. And, with Kinetica's native LLM option, all inferencing can take place within the customer's on-premises environment or within the customer's cloud perimeter. SQL-GPT also offers visual workflows to enable customers to fine-tune the model on their own data, but no external API calls are required, and data never leaves the customer's environment.
- **Security.** The database controls the entire interaction with the LLM, so it applies all the same data-access control rules it enforces for the underlying data tables to the SQL result an LLM may return. This prevents unwanted leakage of metadata to users who should not have access to that metadata.

- **Vector Similarity Search.** Kinetica's latest GenAI-enabled feature is Vector Similarity Search, introduced in beta in October 2023 and, at this writing, expected to become generally available in Q1 2024. The feature creates and stores vector embeddings and supports searching of these high-dimensional vectors in real time. Add in the power of Kinetica's vectorized CPU and GPU processing, and millions of vector embeddings can be loaded faster, search indexes can be built faster, and queries will be handled at higher volumes than is possible with conventional databases or rival specialized databases.

Many database vendors have introduced vector embedding features, but Kinetica's key differentiation is its ability to ingest a high volume of embeddings in near real time. The product's GPU backbone is well suited for this new breed of workload: fast ingest and query along with spatial correlations between vectors.

The concept of real-time vector embedding creation and search represents a visionary leap in using temporal and spatial attributes from sensor and device data. Kinetica's approach involves transforming dynamic, time-sensitive information into multidimensional vectors that encapsulate both temporal and spatial contexts.

Industries such as automotive, financial services, telecom, defense, and logistics stand to benefit from this innovation. For instance, in the automotive industry, Vector Similarity Search could enable instant analysis of real-time sensor data from vehicles, facilitating predictive maintenance and enhancing safety measures. In financial services, the ability to swiftly process and analyze temporal and spatial data could revolutionize risk assessment and fraud detection. Similarly, in telecom, defense, and logistics, the technology could power real-time tracking, anomaly detection, and strategic decision-making.

USE CASES

Kinetica addresses use cases across a range of vertical industries, as detailed below:

- **Automotive** use cases include near real-time analysis of vehicle performance, predictive maintenance, and autonomous vehicle development.
- **Financial services companies** use Kinetica for near real-time monitoring and detection, advanced analytics and prediction, and risk analysis spanning fraud, money laundering, insider trading, and other illicit activity.
- **Logistics** use cases are focused on dynamic route optimization across warehouses, distribution centers, cities, states, regions, and countries.
- **Oil and gas** use cases span exploration, extraction, and production and include drilling and well analytics, predictive maintenance, and supply chain optimization.

- **Property and casualty insurance** use cases include tracking and predicting the timing, severity, and impact of storms, fires, and floods to direct claims adjusters to the places most likely to be impacted, to detect potential fraud, and to reduce settlement delays.
- **Public sector** use cases range from smart city monitoring, disaster management, and sustainability planning to threat intelligence and cybersecurity.
- **Telecommunications** firms use Kinetica for near real-time performance tracking, network optimization, IoT-enabled predictive maintenance, and smarter commercial decision-making.
- **Utilities** use Kinetica for grid optimization, predictive maintenance, and near real-time customer insight by location and segment.

PRICING

Kinetica offers five deployment options, as follows:

- **Kinetica Developer Edition** is free for individual users but does not support multiuser deployments, advanced features (detailed below), or production deployment.
- **Kinetica Cloud** is a free cloud DBaaS solution that offers up to 10GB of storage but lacks support for advanced features (detailed below) and production deployment.
- **Kinetica Cloud Dedicated**, set for beta launch in January 2024, offers dedicated instances of Kinetica managed by the vendor on its Kinetica Cloud DBaaS running on AWS. The service uses a common control plane, but provides dedicated Kinetica clusters to each customer to ensure consistent performance and robust security. This dedicated offering adds advanced features including Reveal dashboarding, user defined functions (UDFs), GPU acceleration options, and support for multiple users and multinode clusters. Pricing is pay-as-you-go, starting at \$4.50 per hour of cluster usage, which includes a Kinetica license fee and the cost of AWS (compute and storage) infrastructure.

- **Kinetica Managed Service on AWS/Azure** is a Kinetica-managed cloud offering deployed from the AWS Marketplace or Azure Marketplace and provisioned in the customer’s virtual private cloud (VPC) environment. Kinetica does not have access to the customer VPC unless the customer opts in to a remote support option, in which case designated support personnel from Kinetica are granted specific roles to monitor and manage the infrastructure and the database services. Kinetica never has access to the data within the database itself or any user accounts to log into the database.

Kinetica Managed Service on AWS/Azure includes all of the advanced features available on Kinetica Cloud Dedicated but is available pay-as-you-go (described above) or as annual subscriptions based on a (discounted) Kinetica license fee. Customers pay the costs of cloud (compute and storage) infrastructure. Additional costs apply for premium support.

- **Kinetica Enterprise Edition** is available for on-premises deployment or customer self-managed deployments on the cloud(s) of their choice. Enterprise Edition supports all of the features supported in Kinetica Cloud Dedicated and Managed Service offerings, but adds support for external authentication (LDAP, Active Directory, Kerberos) and high availability (HA) replication. This edition is sold through annual subscriptions based on the size of the cluster. Additional costs apply for GPU acceleration, HA, and premium support.

ANALYSIS AND OBSERVATIONS

Kinetica is a specialized analytical DBMS and DBaaS service that is specifically designed for the rigors of ultralow-latency temporal and spatial analysis on high-scale data, from tens or hundreds of terabytes into the petabytes. Buyers should consider the product’s strengths and weaknesses, detailed below and summarized in Figure 3.

Strengths

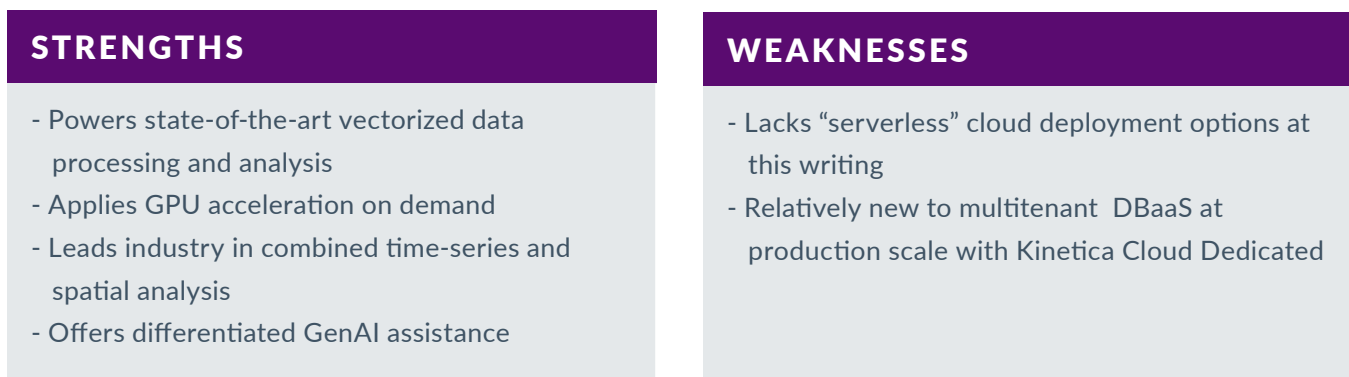
- **Powers state-of-the-art vectorized processing and analysis.** This capability tackles analyses not possible with conventional databases or even specialized databases lacking vector-native analytical functions.

- **Applies GPU acceleration on demand.** GPU acceleration on demand goes beyond the limitations of CPU-based processing, applying brute force to the most taxing, high-scale analyses.
- **Leads industry in combined time-series and spatial analysis.** Differentiation includes support for specialized, low-latency joins and functions for combined time-series and spatial analysis.
- **Offers differentiated GenAI assistance.** SQL-GPT, a database-native LLM developed by Kinetica, uses GenAI to generate consistent, optimized, and customer-tunable SQL without exposing data to security risks.

Weaknesses

- **Lacks “serverless” cloud deployment options.** Kinetica Cloud Dedicated and Managed Service deployments must be appropriately sized in advance and then subsequently monitored and manually scaled up or down, as needed, to balance performance and utilization. Elastic scalability features are on the roadmap.
- **Relatively new to multitenant cloud delivery.** Kinetica has years of experience hosting instances of its database in the cloud and has run the multitenant Kinetica Cloud service since 2022. Kinetica Cloud Dedicated is set to become generally available in January 2024, so it’s early days for delivering Kinetica as a multitenant service for production deployments (versus Kinetica Cloud, which supports small-scale, nonproduction deployments).

Figure 3. SWOT Diagram



Source: Constellation Research

RECOMMENDATIONS

Any organization dealing with, or expecting to grow into, high-scale analytical workloads primarily focused on the combination of time-series and spatial data should consider Kinetica. “High-scale” ranges from the tens of terabytes into the petabytes. Organizations managing data at this scale invariably have incumbent data warehouse deployments.

Specialized analytical databases such as Kinetica are most often considered as options to take on the most demanding workloads that are either performing poorly or that are impossible to address, practically speaking, on incumbent data warehouse platforms or by using general-purpose compute platforms, such as Spark. Less-demanding workloads, such as BI reporting and routine ad hoc analysis, continue to run on the incumbent platform.

Triggers for considering Kinetica might include:

- Performance constraints on incumbent platforms tied to growing volumes of time-series and spatial data and related analyses.
- Functional limitations of incumbent platforms (or other specialized database options being considered) due to a lack of time-series and/or spatial data management and data analysis capabilities.
- Performance constraints of specialized analytical database options that do not fully embrace vectorized processing and/or GPU acceleration.

Would-be buyers should begin their assessment with organizational and tech strategy considerations before considering specific vendor offerings. Organizational considerations include existing budgets; existing technology skills; incumbent technology dependencies; and the desire (and executive and budgetary commitment) to innovate with data, new sources of data, and more-advanced analytics and data science. Tech strategy considerations include cloud strategy; on-premises requirements; data lake and data science strategy; and BI, analytics, and AI/machine learning (ML) ambitions.

It all starts with a clear understanding of where the organization is coming from and where it wants to go—and on which clouds and with what level of commitment to technology spending, skills building, and analytical innovation. Based on conversations with dozens of firms that have deployed analytical data platforms and specialized analytical databases, Constellation offers the following cautions and suggested best practices:

- **Think big and long-term.** It's all too common for organizations to outgrow deployments within just a few years, through either unanticipated organic growth or business-changing acquisitions. Don't ignore history, but look beyond it to consider future possibilities and plan deployments that will stand the test of time and emerging requirements.
- **Look for deployment flexibility.** Does the analytical platform you are considering support on-premises deployment as well as cloud and/or multicloud deployment? Are licenses or subscriptions portable, so you can leverage training and financial investments? Is there flexibility to mix and change deployment modes?
- **Consider available skills and training resources.** Evaluate your existing talent, the availability of training, and the cost and availability of professionals experienced with the platforms you are considering. There are plenty of SQL-savvy data management professionals out there, but how many have experience deploying, managing, and/or working with the specific platforms you are considering? Take into account the size of each vendor's customer community and its level of activity.
- **Seek out reference customers.** Look for reference customers with similar data, data scales, analytical needs, and workload requirements. Talk to them at length about the strengths and weaknesses of the platform and supporting vendor you are considering. Do all of the above before mounting pilot projects with each short-listed vendor to test your own data and key workloads.

ANALYST BIO

Doug Henschen

Vice President and Principal Analyst

Doug Henschen is vice president and principal analyst at Constellation Research focusing on data-driven decision-making. His Data to Decisions research examines how organizations employ data analysis to reimagine their business models and gain a deeper understanding of their customers. Data insights also figure into tech optimization and innovation in human-to-machine and machine-to-machine business processes in manufacturing, retailing, and services industries.

Henschen's research acknowledges the fact that innovative applications of data analysis require a multidisciplinary approach, starting with information and orchestration technologies; continuing through business intelligence, data visualization, and analytics; and moving into NoSQL and big data analysis, third-party data enrichment, and decision management technologies. Insight-driven business models and innovations are of interest to the entire C-suite.

Previously Henschen led analytics, big data, business intelligence, optimization, and smart applications research and news coverage at InformationWeek. His experiences include leadership in analytics, business intelligence, database, data warehousing, and decision support research and analysis for Intelligent Enterprise. Further, Henschen led business process management and enterprise content management research and analysis at Transform magazine. At DM News, he led the coverage of database marketing and digital marketing trends and news.

[🐦 @DHenschen](https://twitter.com/DHenschen) [📄 constellationr.com/users/doug-henschen](https://www.constellationr.com/users/doug-henschen) [in linkedin.com/in/doughenschen](https://www.linkedin.com/in/doughenschen)

ABOUT CONSTELLATION RESEARCH

Constellation Research is an award-winning, Silicon Valley–based research and advisory firm that helps organizations navigate the challenges of digital disruption through business model transformation and the judicious application of disruptive technologies. Unlike the legacy analyst firms, Constellation Research is disrupting how research is accessed, what topics are covered, and how clients can partner with a research firm to achieve success. Over 350 clients have joined from an ecosystem of buyers, partners, solution providers, C-suite, boards of directors, and vendor clients. Our mission is to identify, validate, and share insights with our clients.

Organizational Highlights

- Named Institute of Industry Analyst Relations (IIAR) New Analyst Firm of the Year in 2011 and #1 Independent Analyst Firm for 2014 and 2015.
- Experienced research team with an average of 25 years of practitioner, management, and industry experience.
- Organizers of the Constellation Connected Enterprise—an innovation summit and best practices knowledge-sharing retreat for business leaders.
- Founders of Constellation Executive Network, a membership organization for digital leaders seeking to learn from market leaders and fast followers.

 www.ConstellationR.com

 [@ConstellationR](https://twitter.com/ConstellationR)

 info@ConstellationR.com

 sales@ConstellationR.com

Unauthorized reproduction or distribution in whole or in part in any form, including photocopying, faxing, image scanning, emailing, digitization, or making available for electronic downloading is prohibited without written permission from Constellation Research Inc. Prior to photocopying, scanning, and digitizing items for internal or personal use, please contact Constellation Research Inc. All trade names, trademarks, or registered trademarks are trade names, trademarks, or registered trademarks of their respective owners.

Information contained in this publication has been compiled from sources believed to be reliable, but the accuracy of this information is not guaranteed. Constellation Research Inc. disclaims all warranties and conditions with regard to the content, express or implied, including warranties of merchantability and fitness for a particular purpose, nor assumes any legal liability for the accuracy, completeness, or usefulness of any information contained herein. Any reference to a commercial product, process, or service does not imply or constitute an endorsement of the same by Constellation Research Inc.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold or distributed with the understanding that Constellation Research Inc. is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. Constellation Research Inc. assumes no liability for how this information is used or applied nor makes any express warranties on outcomes. (Modified from the Declaration of Principles jointly adopted by the American Bar Association and a committee of publishers and associations.)

Your trust is important to us, and as such, we believe in being open and transparent about our financial relationships. With our clients' permission, we publish their names on our website.

San Francisco Bay Area | Boston | Colorado Springs | Denver | Ft. Lauderdale | New York Metro
Northern Virginia | Portland | Pune | San Diego | Sydney | Washington, D.C.