

REPORT REPRINT

Kinetica gets serious about performance, joins emerging GPU-enabled database market

JAMES CURTIS

17 OCT 2016

The company is an emerging player in the small but growing GPU database market. Kinetica just released an updated version of its in-memory database that it pairs with GPU processors for maximum performance.

THIS REPORT, LICENSED EXCLUSIVELY TO KINETICA, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR REPOSTED, IN WHOLE OR IN PART, BY THE RECIPIENT, WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2016 451 Research, LLC | WWW.451RESEARCH.COM

The GPU database market has seen some new entrants as of late. While Kinetica can trace its origins back to a project for US Army Intelligence, the company officially became an entity in 2015, when it branched out on its own. Kinetica is already at version 6.0 with its in-memory columnar database and reports a number of recent improvements in performance, scaling, visualization and SQL support, to name a few. The company has also reported early investments from former Oracle executive Ray Lane's investment firm GreatPoint Ventures totaling \$13m.

THE 451 TAKE

GPU analytical processing is starting to pick up. We have seen a number of new entrants to this market recently. Kinetica has a solid history that has allowed it to prove out its technology with the US government. The company has also zeroed in on use cases that are a good match for its GPU-driven database. The fact that Kinetica is also integrating well with Hadoop as a type of 'speed layer' should resonate well with companies, simply because performance is generally universally valued. While it might be early for GPU-driven databases, it's showing signs of growth.

CONTEXT

Amit Vij and Nima Negahban didn't necessarily envision starting a database company based on GPU technology, but that's how it turned out. In 2009, Vij and Negahban were driving an IT consultancy company working for the US government. The US Army Intelligence division was looking for a system capable of handling hundreds of real-time data feeds where analytics could then be run on the data in real time, allowing for quick decision-making capabilities. After some failed attempts at using some open source databases, the two decided to leverage GPUs as a basis to build out an analytics-based database architecture system to handle the high compute requirements. The idea proved useful and the two won the US Army Intelligence contract. Word got out on the US Army Intelligence database offering and two were able to secure a deal with the United States Postal Service, resulting in the company's first commercial contract. The USPS contract enabled the company to formally launch Kinetica and it began to market and sell its GPU-based in-memory database offering publicly.

In addition to counting USPS as a customer, Kinetica claims more than 10 paying customers (exact numbers not disclosed), with a number of proofs of concept under way. There are roughly 50 employees and the company is actively hiring. Kinetica is based in San Francisco.

PRODUCTS

Kinetica's core product offering consists of a columnar in-memory database that leverages high-performance computing hardware, specifically GPUs, for accelerated processing. Kinetica touts the product's massively parallel processing (MPP) architecture and that it runs on GPUs, which consist of thousands of small cores that can effectively process data.

Kinetica targets organizations with demanding analytical needs, particularly organizations that possess high cardinality data sets, which require significant compute power, given that the data is quite complex. In addressing this particular need, Kinetica identifies certain reasons for using its GPU-accelerated database. One is that Kinetica is able to ingest significant amounts of streaming data for which it can process and provide millisecond response times. The company also points out its ability to scale linearly based on the amount of data received. Upon ingest, Kinetica maintains a simple data structure – similar to a relational structure model consisting of a table in which rows can be added. This simple data model enables Kinetica to process the data quickly using the GPUs. Conversely, with relational databases, the data model is defined up-front in order to ensure optimized querying capabilities based on limited compute resources. With Kinetica, the inverse is true because processing is much more abundant and can be done via 'brute force' after the data has been ingested.

While Kinetica can handle operational and transactional data, the company primarily positions its offerings as what it calls the 'speed layer' in a broader data environment. Kinetica works to enable integration capabilities and cites its ability to ingest data from, for instance, Apache NiFi, Kafka, Spark and Storm, as well from messaging queues and ETL tools. There are APIs for programming languages (e.g., Java, C++, Python) and integration with some third-party BI and visualizations tools. However, Kinetica points out its native geospatial visualization front-end tool as a core part of its offering.

The company's release cadence occurs twice yearly. Version 5.0 was released in June and version 6.0 is expected to be generally available in Q4. Some noted enhancements include improved high-availability functionality. Previously, organizations had to set up two separate clusters for failover. With version 6, Kinetica is introducing intra-cluster failover support. Also, full SQL-92 query support has been added that enables greater compatibility with third-party BI tools using JDBC and ODBC connectors. Further, a number of performance enhancements have been added, such as optimization and tuning updates, along with NVIDIA NVLink support, which accelerates data movement between the CPUs and GPUs.

CUSTOMERS AND USE CASES

Kinetica claims about 30,000 users of its database spread across 10+ paying customers. Customers that the company can publicly identify include the USPS, IronNet and PG&E. USPS uses Kinetica to reallocate resources dynamically based on numerous data feeds to track mail trucks and delivered mail pieces delivered. IronNet, a cybersecurity firm, uses Kinetica for anomaly detection to identify and react to potential cyber threats. PG&E, an energy utility company, leverages Kinetica for its geospatial capabilities in optimizing energy usage from smart meters. Other use cases include customer 360, supply chain optimization, product and services recommendations, ride-sharing analytics and IoT scenarios for equipment failures.

Kinetica also envisions scenarios where Kinetica is integrated with Hadoop environments, particularly when those environments that might leverage Kafka, AWS Kinesis or other streaming tools. Data can then be moved to HDFS, NoSQL databases, object stores (AWS S3, for instance) or SANs. However, the streaming data can also be simultaneously sent via Spark Streaming, for instance, to Kinetica. However, other historical data from HDFS, object stores, NoSQL database and SANs can also be processed in Kinetica for analytics. In the Hadoop environment, Kinetica would again play the role of processing certain analytical workloads that may not make sense with Hadoop or NoSQL databases.

In terms of go-to-market, Kinetica drives a direct sales model for its US and EMEA sales efforts. Currently, the primary deployment avenue is on-premises, as the company is certified for IBM, Dell, Hewlett Packard Enterprise and Cisco hardware running NVIDIA GPUs. Kinetica is offered on AWS but not on Microsoft Azure because GPU servers aren't offered. The company also has an OEM relationship with IronNet.

COMPETITION

Kinetica notes that it sees SAP HANA in a number of competitive situations, and given how SAP positions HANA as an in-memory analytics product, this makes a lot of sense. However, while HANA runs in-memory, its architecture does not leverage GPUs.

The GPU database market may be small, but it is growing. There are a handful of GPU database vendors that Kinetica is likely to see. MapD, a newcomer to the GPU space, similarly targets real-time analytics workloads. MapD, however, runs in GPU video memory. SQream Technologies is another that offers an analytics-based offering that runs a NoSQL database under its covers, and Brytlyt leverages PostgreSQL.

In addition to SAP HANA, we expect Kinetica could possibly encounter other in-memory database vendors as well as database vendors targeting real-time analytical use cases. These include MemSQL, VoltDB, HPE Vertica and some of the Hadoop vendors with their SQL-on-Hadoop tools targeted at analytical use cases. Some NoSQL players may also arise such as Redis Labs, Cassandra, MapR-DB and Aerospike.

We also believe that Cray's Urika-GX is a possible competitor, which combines Hadoop, Spark and the Cray Graph Engine in a unified offering that targets analytical use cases that require heavy compute capabilities. The Cray XC series-supercomputing lineup does leverage GPUs, but the Cray Urika-GX is a CPU-based machine.

SWOT ANALYSIS

STRENGTHS

Kinetica strengths lie in its ability to provide a high-performance and highly scalable offering that makes great use of GPUs that work well for high cardinality workloads.

WEAKNESSES

While Kinetica has had a number of development years under its belt, the company is still a relative newcomer with a handful of customers. Certain functionality on the product is still maturing.

OPPORTUNITIES

Augmenting an existing environment with a so-called 'speed layer' can be quite compelling and intriguing to customers. Offloading workloads, particularly analytics that was previously addressed with Spark and NoSQL tools, are also possible opportunities for Kinetica.

THREATS

The unknown is perhaps hampering some organizations to adopt. GPUs are still gaining traction as a whole and organizations are trying to sort out how to leverage this technology that clearly has performance advantages.