

## REPORT REPRINT

# After \$50m funding round, Kinetica adds geospatial muscle to its GPU database

**JAMES CURTIS**

**16 NOV 2017**

The company has rolled out an updated version of its GPU-powered database with enhanced geospatial processing. With \$50m in funding secured earlier this year, Kinetica is pouring money into its development efforts.

---

THIS REPORT, LICENSED TO KINETICA, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT 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 RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2017 451 Research, LLC | [WWW.451RESEARCH.COM](http://WWW.451RESEARCH.COM)

Kinetica has rolled out a new version of its GPU-powered database that adds a number of new updates. Already ideal for geospatial processing, the company is adding support for OpenGL/EGL, along with a collection of built-in functions to further drive geospatial processing. Enterprise management features and increased cloud support have been included as part of the latest 6.1 release. In June the company landed a \$50m series A round, which Kinetica noted would help beef up its development efforts. This latest release is evidence of that.

---

## THE 451 TAKE

Kinetica has done well to continue to mature its database. For instance, in our previous coverage, Kinetica reported adding in-database analytics with a tie-in using UDFs; with its most recent 6.1 release, the company has added enhanced geospatial support, administration tooling and broader cloud support. While organizations are still figuring out ways to leverage GPUs for analytics, we anticipate good growth and adoption of GPUs. However, the data-warehousing space is also changing in that many vendors are building more advanced analytical capabilities into their database platforms - the competition is rising, which means Kinetica will need to continue to drive its differentiation within this market.

---

## CONTEXT

Amit Vij and Nima Negahban founded Kinetica in 2009. The pair was managing an IT consultancy firm working for the US government, and 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. The idea proved useful, and the company was able to secure a contract, as well as other government contracts, including the US Post Postal Service (USPS), which then enabled the pair to launch a company around a GPU-powered database.

We initiated coverage of Kinetica in 2016, and the company has grown a good deal since then. In our previous report, we had the company with 10 paying customers, a number that has since grown to 'dozens.' Total employee count was last reported at 50, and is expected to hit 95 by December.

The company reported securing a series A round of \$50m in June that was co-led by Canvas Ventures and Meritech Capital Partners. Other participating investors included Citi Ventures, as well as existing investors Ray Lane and GreatPoint Ventures. Total funding to date, which includes an early \$13m seed round, comes in at \$63m.

## PRODUCTS

At its core, Kinetica is a vectorized columnar database designed for analytical workloads. The company points to its hybrid architecture that leverages both CPUs and GPUs. When data lands in Kinetica, it lands first in memory, and then is moved to GPU at the time the query executes. Data, however, can also be persisted to disk. Moreover, in a previous release, Kinetica introduced the ability to 'pin' data in GPU memory for certain workloads and queries that may require it. Kinetica also touts the product's massively parallel processing (MPP) architecture and its scaling capabilities.

For its 6.1 release, Kinetica added a number of updates, most notably around geospatial processing. Kinetica was already using GPUs for rendering maps, but certain types of calculations could not be run on the GPUs, and instead had to be processed on the CPU. To address this, Kinetica added support for OpenGL/EGL, which enables all the rendering calculations to run at greater performance. Additionally, the company has added more than 80 geospatial operations that can be exposed via SQL command, which Kinetica claims provides a similar experience as ArcGIS and PostGIS.

For cloud enhancements, Kinetica is offering what it calls 'one-click' deployments on the public cloud, specifically on AWS and Microsoft Azure. Users have the option to bring their own license or grab the image from the marketplace. Users also have the choice of metered pricing, allowing for rapid spin-up or spin-down of a cluster. While not directly related to cloud, Kinetica is also simplifying its general deployment process, driven through a new GUI. The new deployment process allows for less technical users to deploy Kinetica without having a deep knowhow of Linux or having to set up a configuration file.

A handful of new enterprise features have been added, such as compression. With some internal testing, the company is claiming 100TB of data can be compressed to 1TB, realizing, however, that the compression results will vary by data type. Enhanced monitoring has been added, as well as the ability to add or delete nodes on the fly without database disruption. Security has likewise been improved. Along with LDAP, authentication, authorization and encryption, Kinetica has added auditing functionality so that a security administrator can monitor security changes.

## USE CASES

We have noted the USPS as a Kinetica customer. The use case for USPS leverages geospatial data for logistics and route optimization in delivering mail packages. However, other verticals and use cases include deep learning, specifically 'deepchem' within the health and life sciences vertical. This specific use case involves Kinetica's UDF capabilities to enable in-database analytics. Retail is another vertical where Kinetica has been used to carry out predictive modeling for sales within a specific store or region. In the energy sector, Kinetica is used to map energy grids, such as gas and power lines, and to help drive productive analytics on wells. There are also companies using Kinetica to drive faster BI efforts that involve Kinetica's recent partnership with Tableau to drive greater BI performance.

## COMPETITION

As a GPU-powered database vendor, Kinetica will compete with the other vendors in this space, specifically MapD, BlazingDB, SQream Technologies and Brytlyt.

MapD is likely Kinetica's primary competitor, although the architectures and strategies are slightly different. MapD employs a GPU-memory-first approach, although the company notes that it leverages (and can run in) CPU memory as well, whereas Kinetica points to its hybrid in-memory CPU and GPU approach. BlazingDB and SQream Technologies both employ a disk-based approach with their GPU-powered database offerings, and likewise cater to the largest of large datasets.

Along with the core GPU database vendors, we see competition with the analytic-driven database vendors that have been expanding analytical capabilities, including built-in machine learning. For instance, IBM recently released the IBM Integrated Analytics System that blends Db2 and Netezza technologies, as well as the company's Data Science Experience offering and Apache Spark. Teradata recently announced the Teradata Analytics Platform, which combines the Aster Analytics engine with Teradata, providing more advanced analytics capabilities. Pivotal, with its Greenplum product, includes advanced analytics by way of the Apache MADlib open source library.

Others include Oracle with its updated Exadata X7 system, as well as the Oracle 18c Database. The company also offers Oracle Advanced Analytics, which enables some advanced analytics to be carried out in-database. SAP's in-memory database HANA boasts the ability to carry out both operational and analytical workloads. There is also SAP Vora, an in-memory compute engine built on top of Apache Spark to extend the analytical capabilities of the Spark framework, which can be also integrated with HANA.

## SWOT ANALYSIS

### STRENGTHS

Kinetica offers a hybrid CPU- and GPU-enabled database that is differentiated by the ability to do advanced in-database analytics via UDF functionality.

### WEAKNESSES

Kinetica has made strong gains since its founding days, but its GPU database, as well as the market generally, is still maturing.

### OPPORTUNITIES

Geospatial is already a good match with GPUs, so Kinetica should be able to drive a sort of land-and-expand strategy with existing customers that already have geospatial use cases in the works.

### THREATS

While there are only a handful of other GPU-powered database vendors, the larger competitive field is likely to come from the analytic database vendors that have recently begun to step up their analytics capabilities, especially with more advanced analytical capabilities such as machine learning.