

Kinetica is promoting its in-database capabilities for a few reasons. One is that it effectively opens up the door for machine learning, deep learning and other advanced analytical workloads to be run alongside BI workloads, all within the same environment. There is also the benefit of less data movement because, theoretically, a data scientist can develop algorithms and models and then make those available to business analysts, for instance. As such, the company views this function as broadening its user profile – a type of data-democratization strategy.

In addition to in-database analytics, Kinetica has unveiled a front-end visualization tool called Reveal. Visualization tools make a lot of sense, especially when paired with GPUs, primarily since such pairing enables interactive or drill-down analytics, something traditional BI visualization tools often struggle with, particularly on extremely large multi-hundred, million-row tables. Spatial or location-based data is a particularly favorable use case with GPUs. As such, Kinetica processes the spatial data on the GPUs and then sends a small PNG file to Reveal, which allows users to interactively query a map. Reveal also enables users to create dashboards by dragging and dropping widgets, as well as integrate with several map providers such as Google, Bing, ESRI and others.

Other enhancements include VRAM Boost, which allows users to ‘pin’ data in GPU memory (VRAM) for some workloads and use cases that require low latencies. Kinetica runs in a hybrid CPU in-memory architecture that leverage GPUs, so with VRAM Boost, customers have the option to have certain workloads run in GPU memory (VRAM), which offers even greater performance.

COMPETITION

Kinetica is part of a small but growing collection of GPU-powered database vendors. And although there may not be as many players in the market currently, there does exist a certain level of competition among them. The company differentiates itself with a hybrid in-memory CPU- and GPU-enabled database, so the addition of VRAM Boost further enhances its hybrid story. That being the case, MapD is likely Kinetica’s closest rival, given that it runs primarily in VRAM, but can leverage CPUs as well.

Further, Kinetica’s Reveal visualization tool is comparable to MapD’s Immerse visualization offering. And while both firms enable third-party BI tool integration, each of their visualization tools is particularly suited to work with GPUs – especially well with drill-down analytics. BlazingDB is another contender, although its GPU database runs on disk. SQream Technologies deploys a NoSQL database under its covers, while brytlyt leverage PostgreSQL.

While we noted that Kinetica’s GPU-powered database can be competition for the data-warehouse suppliers, with the company’s newly released in-database analytics functionality, we believe it is even more so now. As such, we see Kinetica encroaching on a number of data-warehouse vendors, including Oracle, IBM, Teradata, Microsoft, SAP, HPE and Amazon Web Services, among others.

Lastly, some in-memory and NoSQL database providers also deserve a mention. These include SAP HANA, MemSQL and VoltDB, as well as NoSQL specialists such as Redis Labs, Cassandra, MapR-DB and Aerospike.

SWOT ANALYSIS

STRENGTHS

By enabling in-database analytics, Kinetica is demonstrating that GPUs can have a broader appeal for organizations, particularly for more advanced analytics.

WEAKNESSES

The GPU market is maturing in many ways as organizations continue to understand and leverage the technology. Consumption can also be challenging, particularly on the cloud front, where instances can be limited.

OPPORTUNITIES

GPUs shine for drill-down analytics, and when combined with the company's Reveal tool, this should give Kinetica additional opportunities. However, the notion of driving both BI and more advanced analytics such as machine learning within a single environment should, at a minimum, spur new customer conversations.

THREATS

While the GPU market is growing and does have an expanding roster of players, the overall sector is not as large as some of the other database segments we track - data warehousing, for instance. The biggest challenge lies in educating organizations on what GPUs can do, given that most of them still have a CPU-first mentality.