IDC ANALYZE THE FUTURE

IDC WORKBOOK

Sponsored by: Kinetica

To keep up with the pace of digital transformation, organizations are turning to active analytics powered by GPU-accelerated data platforms to make real-time decisions and uncover new insights.

Building Active Analytical Applications with GPU-Accelerated Data Platforms

September 2019

Written by: Carl W. Olofson, Research Vice President, Data Management Software

Introduction

As the analytics world is increasingly driven by streaming data, Internet of Things (IoT) data, and ever-changing complex combinations of data reflecting real-time business conditions, pressure has been building for an analytic database technology that can keep up. Although database management systems (DBMSs) are becoming more effective at ingesting and organizing such data, executing complex queries on the data with results that are usable in the moment remains a daunting challenge.

AT A GLANCE

WHAT'S IMPORTANT

Active analytical applications require the ability to analyze both streaming and historical data to make in-the-moment decisions. They need to be able to drive machine learning and other artificial intelligence elements with the data and may include the need for locational awareness. Such applications are key elements in digital transformation.

In response, a new kind of database query processing technology has emerged — one that powers analytic calculations using graphics processing units (GPUs). Unlike conventional technology that limits its processing to CPUs, these new products can crunch large sets of numbers in parallel in a fraction of the time that the same queries require on regular systems. This is because GPUs, which were originally developed to perform the calculations necessary to drive real-time graphics, can concentrate terrific processing power on many problems simultaneously. GPU databases promise to be a key to handling very complex queries, where answers are needed to drive ongoing processing in an increasingly dynamic, stream-driven, smart processing environment.

The use of the CPU remains relevant for certain types of data analytics. A data platform that is able to leverage both GPU and CPU has a significant advantage when it comes to performing different kinds of analytics as well as applying machine learning to gain insight from massive volumes of complex data.

A complete GPU-accelerated data platform should combine streaming data analytics with the ability to analyze historical and location data and also deploy machine learning algorithms. This allows organizations to shift from passive analytics to active analytics that run continuously, always updating with the latest data and information. Businesses and governments are able to use GPU-accelerated data platforms as the foundation on which to build active analytical applications that react dynamically to real-time data and events. This can help organizations develop new data-driven products and services that are critical for continued innovation and digital transformation.

Implementation or Selection Recommendations

To build active analytical applications that leverage real-time data and react instantly to users and context, organizations require a unified, GPU-accelerated data platform that offers simple ways for developers to leverage streaming, historical, and location data as well as machine learning to deliver new insights. Such a platform can help organizations transform their business models and keep pace with digital transformation. In particular, the following verticals and use cases are well-suited for active analytical applications:

- Logistics: Enabling responsive routing and same day delivery
- Financial services: Moving from batch to continuous risk analysis
- Telecommunications: Accelerating networking optimization at massive scale to reduce the amount of time needed for analysis from years to hours
- Automotive: Making self-driving smarter as well as enabling in-car, edge analytics
- Retail: Powering customer 360 initiatives to enable targeted marketing offers based on context
- Healthcare: Bringing new medicines to market faster

Optimization Strategies

Anyone considering adopting a data platform that can support real-time analytics on a large scale, complex historical data analytics, analytics that involve location intelligence in real time, and data handling for the purposes of training machine learning systems should take a look at a GPU-accelerated solution.

In particular, such investigation should include the following:

- Look at one's own shifting data environment and future data analytics needs.
- Examine the various GPU-accelerated database systems available on the market, looking not just for a DBMS but also for a unified data platform.
- Review business use cases that involve streaming, historical, and location analytics as well as machine learning.
- Consider the role of active, data-driven analytical applications in powering business transformation and innovation.

Worksheet Section

See Table 1 for questions prospective customers should ask themselves to determine if it makes sense to pursue GPU-based database software.

TABLE 1: ***Is GPU-Based Database Software Right for You?***

Question	Answer
Do you require large volumes of data queries?	
Is your existing system overwhelmed with the analytic requirements of the database and its users?	
Have you been expanding the processing capacity of your database servers, but they still can't satisfy the demand?	
Does your business depend on immediate analysis of streaming data?	
Are you pursuing Internet of Things data analytics strategies that demand immediate answers?	
Are you seeking to add machine learning to the business and operationalize artificial intelligence?	
Do you have a need to do location analytics at massive scale?	
How important is the ability to address a breadth of analytics in a single platform (e.g., streaming, location, historical, graph, machine learning)?	Rate on a scale of 1 to 5, with 5 being most important
How important is speed to problem resolution, and are your current database systems lacking in the performance needed to achieve that speed?	Rate on a scale of 1 to 5, with 5 being most important

If you answered "Yes" to four or more of the questions above, and/or answered 3 or higher to the last questions, you should consider a GPU-based data management solution such as that offered by Kinetica.

About the Analyst

Carl W. Olofson, Research Vice President, Data Management Software

Carl Olofson has performed research and analysis for IDC since 1997. He manages IDC's Database Management Software service and supports the Data Integration Software service. Mr. Olofson's research involves following sales and technical developments in the structured data management (SDM) markets, including database management systems (DBMSs), dynamic data management systems, database development and management software, and dynamic data grid managers, including the vendors of related tools and software systems. Mr. Olofson also contributes to Big Data research and provides specialized coverage of Hadoop and other Big Data technologies. Mr. Olofson advises clients on market and technology directions and performs supply- and demand-side primary research to size, forecast, and segment the database and related software markets.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Corporate USA

5 Speen Street
Framingham, MA 01701, USA

T 508.872.8200

F 508.935.4015

Twitter @IDC

idc-insights-community.com

www.idc.com