Proactively Building Competitive Advantage: Mastering Business with the Speed of In-Memory Computing

An ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) White Paper
Prepared for GridGain Systems

September 2014

EMA™ IT & DATA MANAGEMENT RESEARCH, INDUSTRY ANALYSIS & CONSULTING

Proactively Building Competitive Advantage: Mastering Business with the Speed of In-Memory Computing

Table of Contents

Proactively Building Competitive Advantage: Mastering Business with the Speed of In-Memory Computing

Introduction

The pace of business is increasing whether or not your organization wants to participate. Competitors and new business entrants are looking for ways to expand or establish their positions in your market. One of the key techniques of establishing competitive advantage is to decrease the time between when data is available and when action can be taken based on that information. Demanding decreased time between information and action, many organizations are increasing their use of low-latency operational and analytical processing.

By reducing the latency in operational and analytical processing, organizations can move from a traditional information strategy that often looked backward at the events that happened to a more proactive approach that allows organizations to capitalize on events as they occur and even predict what may happen next. This change to a more proactive operational mindset can improve their businesses by increasing their revenues, lowering their costs or improving their overall margins. Three industries where decreased latency can be effective are:

- **Financial Services:** Companies are moving beyond "simple" compliance associated with regulatory reporting on a quarterly or annual basis to using that same information to actively manage investment strategies.
- **Consumer Package Retailing:** Organizations are using faster decision making to change the nature of their businesses from the tactical management of supply chains to proactive and coordinated delivery of the right product to the right customer within a profit profile.
- **Utilities:** Using information from operational sources, the Utilities raise their level of service and profitability from meeting bare government mandated service level agreements to optimized generation and distribution of electricity.

The response required by the situations described above requires truly low latency or near-realtime processing to be effective. The amount of time afforded stakeholders to make decisions in these situations is not measured in days or even hours. With the influx of social media data sources, the window for product adoption can be impacted in hours. In each of these situations, the ability to provide actionable information for corporate decisions rests on the ability to provide information in near-realtime.

Recent advances in in-memory processing technologies offer the ability to meet these near-realtime processing challenges. There has been a wide range of options to date:

- In-memory accelerators for development and run-time environments to support low-latency application development initiatives and online applications.
- Databases that offer either partial or full utilization of in-memory strategies for speeding relational database performance.
- In-memory processing engines for use with near-realtime workloads on multi-structured data stores such as the Hadoop Distributed File System (HDFS).

Each of these approaches solves a point problem: application performance or database response or streaming applications. However, to be directional instead of reactive, an in-memory technology solution should support long-term architectures as well as near-term requirements. The in-memory solution should provide speed and scalability configured as a single architectural layer with monitoring and management capabilities as opposed to a "toolkit" of technologies that require assembly. This solution should support application development processing challenges as well as database focused workload

environments. Like a fabric, a comprehensive in-memory solution should act as a flexible layer for all types of applications, regardless of development language; and handle structured data, unstructured data as well as traditional and new data warehousing including Hadoop. The solution should also have the flexibility to support multi-structured data sources from Internet of Things (IoT) sources such as RFID tracking tags for consumer products and smart grid sensors for utilities. As such, the in-memory solution effectively becomes a data processing fabric that fits in between diverse applications and data stores, effectively becoming a platform that adds both speed and scale to a variety of performance-critical workloads. Finally, the solution should support a wide range of use cases including operational transactions, analytical workloads and operational analytical processes.

Regulatory Compliance to Anticipatory Investment in Financial Services

Regulatory compliance is one of the "necessary evils" of the world of the financial services industry. Since 1933 in the United States, the Security Exchange Act has dictated reporting and compliance requirements. In 2002, the Sarbanes-Oxley Act added further requirements. Moreover, in 2010, the Dodd–Frank Wall Street Reform and Consumer Protection Act established additional and enhanced standards. However, the United States is not alone in the establishment of compliance standards for financial services. Since 1988, the Basel Accords have also specified wide-ranging compliance requirements for organizations doing business in the European Financial Services industry. All of these compliance measures have placed reporting and compliance burdens on the industry. Some estimate that Dodd-Frank costs alone have amounted to at least $34 billion in the first two years of the Act's existence.

However, imagine if those compliance requirements provided, instead of a burden to financial services organizations, a competitive advantage. Literally, while some organizations would be flailing at compliance and costing the organization person-hours and market capitalization , others could be using the base results of compliance reporting to measure their performance and gain share in the marketplace. Using in-memory computation techniques, organizations can not only satisfy compliance requirements, but also manage their holdings, transactions and financial positions in near real-time. The ability to produce compliance reporting daily or even intraday allows financial services organizations to measure their business activities better than their less proactive competitors do. They can understand not only where they sit in the marketplace, but often project where their competitors are. This advantage is much like being able to "count cards" during blackjack or poker, understanding where you sit in comparison to the cards on the table and remaining in-deck enables you to plan better than your fellow players. Financial services organizations that have this insight cannot only understand which trades they should not make due to liquidity risk concerns, they can also understand which financial positions their competitors should not or could not make for the same considerations.

Access to this type of low-latency information provides a competitive advantage. Financial positions can be determined in a fraction of the time of competitors using traditional techniques and technologies. This will require an in-memory solution that can support the operational transactions required for compliance reporting and the analytical workloads required for identifying time series patterns and performing optimization queries. In addition, the technology platform should be able to meet future challenges. Many of the regulatory schemes being applied have sections that allow for relatively fast changes in requirements. Dodd-Franks is an example of regulatory laws that has components that can change over time. Also, the circumstances of the financial services industry is one with constantly changing product offerings and ways to package those offerings. An in-memory platform that is implemented needs to work across these multiple environments and data sources.

Active Asset Allocation in Consumer Package Goods

In the initial days of eCommerce and online consumer goods sales, there was a competitive advantage in an online presence. However, when "brick and mortar" retailers entered the online space and aligned their supply chains to provide not just the breadth of offerings that online retailers such as Amazon and Overstock.com provided, but the "instant gratification" of in-store pickup, the concept of omni-channel retailing was created. Feeding into the broad concept of "customer experience management", omni-channel retailing enables consumers to see all the products they could ever want online along with the ability to make a same day pickup at a physical location. However, omni-channel retailing is not without its challenges.

To properly operate an omni-channel retailing environment, consumer goods retailers must consider the costs associated with fulfillment, understand where the demand for goods and services is around a particular geographic region, and know how to respond with supply chain fulfillment to meet demands. The link between projected consumer demand and supply chain is critical. While the dual-fulfillment strategy associated with omni-channel retailing is the leading edge of the consumer retail industry, it may come at the expense of margins. Because the tradeoff between market share and profit margin can be a tricky adjustment, organizations that attempt to get involved with omni-channel strategies must have all their critical information available to them.

In July 2013, Sir Terry Leahy, former CEO of Tesco, the 3rd largest retailer in the world by revenue, said:

> *"Understand customers' wants and needs, as well as keep an eye on competitors, by using research that relies on bigger samples of data, preferably in real-time, in order to improve efficiency and create a window into the outside world that illuminates what's in front of you."*

However, Leahy indicated that failing to capitalize on these data sources can leave an organization "flying blind" and unable to meet the challenges of an omni-channel strategy.

With these types of hurdles, business stakeholders must have the correct information readily available. CFOs need to know the potential impacts to product and profit margin when decisions are made to properly modify pricing across delivery channels. This requires an in-memory platform to handle operational transactions to identify pricing opportunities. Traditional disk-based, or non-in-memory solutions simply cannot handle the processing latency or data transfer requirements. COOs and supply chain executives must understand not only where current demand, as provided by Point Of Sale (POS) systems, is located, but also recognize where potential demand is projected to occur. This type of processing requires an in-memory platform to handle low latency operational analytical workloads to insert real-time analytics directly within an operational process flow. This knowledge will enable adjustments to the supply chain to allow for optimization of inventory levels. Finally, it is essential for CEOs to know how their overall strategy is working and how the strategy of competitors in terms of offering and pricing alter the overall marketplace. These circumstances dictate an in-memory platform to handle the changing nature of analytical requirements across data domains and structures.

Moving from Service Level Agreements to Optimized Power Generation

The times are changing in the utilities industry. Utilities are facing situations in which understanding the nature of power demand and generation are key to the profitability of their organizations. These conditions come from disruptive and distributed energy sources such as wind and solar, along with regulatory requirements associated with both power Service Level Agreements (SLAs) and production targets. With government-prescribed profit margins and threats from "free" power sources such as solar power generation, utilities must grasp the "now" of their business as well have the ability to "peek" into the future to understand how to best manage their power generation operations.

Due to the operational constraints on power generation sources such as coal and nuclear plants and the variable nature of wind and solar, utilities are being pressured on multiple fronts to ensure that transmission networks are stable and that natural gas-fed power generation sources are utilized to meet demands without turning those gas-powered plants into non-revenue generating capital investments . Analytics allow organizations to understand the demands of customers based on temperature, season and other environmental conditions along with the ability of distributed power generation assets to provide inputs to the overall power grid. Often the demands of customers and the power producing ability of distributed assets such as solar and wind rely on similar inputs, but with inverted results.

For example, in Germany in June 2013, a relatively breezy (higher wind potential) and sunny (higher solar opportunity) Sunday (low customer consumption) brought together a set of circumstances that created a negative value per megawatt hour (MWh) for generated electrical power. With the minimum requirements of the power grid met by a combination of distributed wind and solar sources, a surplus of generated power made it necessary to pay partners to take on the excess electrical production.

Most businesses cannot endure long periods where they pay customers to take their products, let alone using $800m-$1b assets such as coal, nuclear and natural gas power plants to do so. The use of low-latency predictive analytical models to understand environmental conditions from localized weather stations, customer demands from smart grid power sensors and power producing capabilities combine for potent market opportunities. Understanding how various aspects of the power grid, both production and consumption, interact provide utilities with the ability to meet regulatory requirement and maximize margins associated with power production.

Generation of environmental/weather and customer-demand analytics in near-realtime provides the ability to understand how to best utilize a wide range of power generation assets. An in-memory platform that supports these analytical requirements must utilize not only structured data information from internal resources, but augment with Internet of Things (IOT) sources from distributed power production and smart grid sensors. With an in-memory platform, utilities can use this same information to look outside their local area to conditions in other grids to determine if they can leverage opportunities outside their grid. This greatly expands both the business opportunities and the importance of an in-memory architecture layer that can support inputs from multiple components and expand as needed.

EMA Perspective

Across many industries, the requirement to lower the latency of operational and analytical processing is increasing. This requirement is pressuring historically focused and batch oriented processing schemes to predict future events and conditions and to do so in much shorter timeframes. The concept of big data and the Internet of Things is pushing the "edge" of processing and latency in ways not conceived of just a few years ago. In recent EMA research, respondents indicated that the most important use case associated with their big data initiatives was around the speed of response and low latency processing:

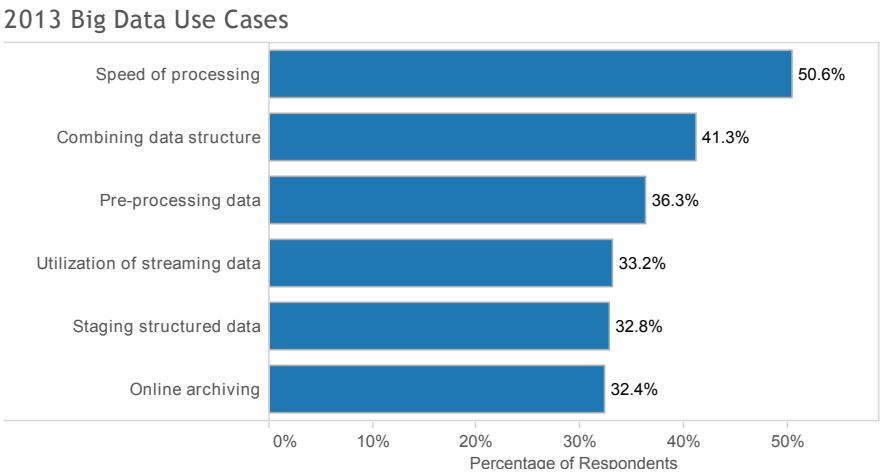

GridGain provides an "In-Memory Data Fabric" that meets the challenges of this leading use case of **speed of processing** across multiple workloads, data structures and implementation choices. This single architectural layer, which additionally offers out-of-the-box high availability, enterprise-level security and powerful monitoring and management, supports both application development and backend database implementations to provide near-realtime processing.

The GridGain In-Memory Data Fabric is also a comprehensive compute platform that supports multiple implementation methodologies and languages. From programming languages such as Java. NET and C++ to access protocols such as SQL and web-based REST interfaces, GridGain supports both legacy and future project challenges. GridGain also supports multiple data sources. From multi-structured information stored in Hadoop's HDFS, JSON documents in MongoDB and structured data managed in traditional relational database management systems (RDBMS), GridGain enables near real-time processing across the spectrum of data sources.

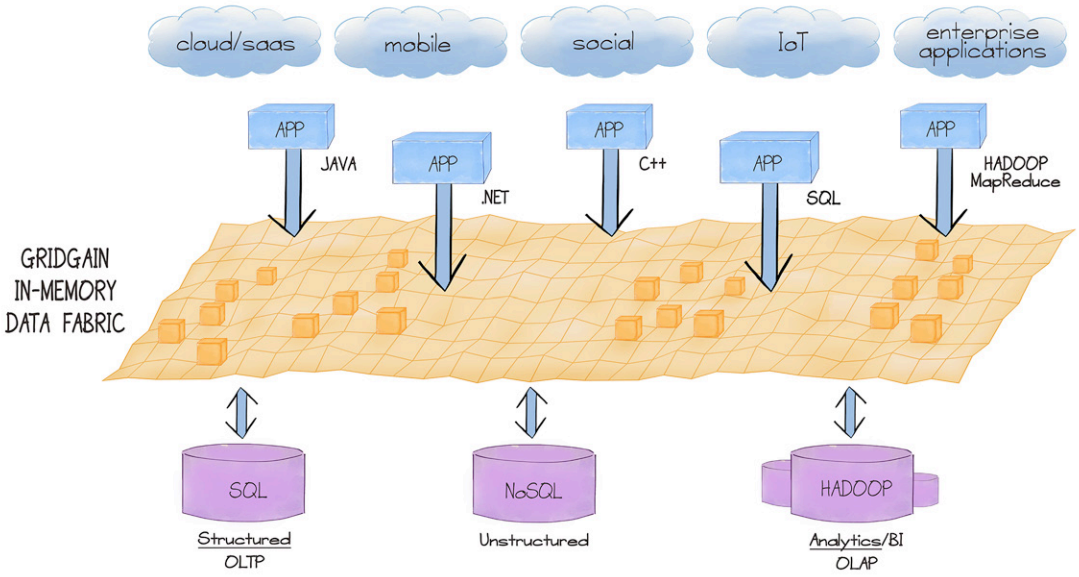

By removing the technical constraints of most in-memory technologies, GridGain provides a directional architecture that will support existing requirements to enhance processing performance and support the unseen challenges that will power new competitive advantage opportunities.

In financial services, organizations often need to make decisions on market information within the hour to fully capitalize on time-sensitive circumstances. Risk management and trading executives in these situations must comprehend not only the tactical costs (stock price, loan risk), but the overall strategic exposure that these situations bring to make not just faster, but better decisions than their competition. GridGain has both the technological capabilities and the experience to enable financial services organizations to make low-latency processing possible and financially viable.

In the consumer package retailing industry, latency requirements for analytical and operational processing were once measured in seasons with the release of fashions in Paris and New York or annually with the debut of electronics at the Consumer Electronics Show in Las Vegas. Now trends are measured using trending information on Twitter and Facebook. These influences can influence customer demand and "cool factor" of consumer products within days if not hours. Consumer retail organizations must align not only pricing but supply in these timeframes for both online and physical channels. Combining these structured sources from Point Of Sale (POS) and supply chain systems, and multi-structured data from social media sources, GridGain bridges the gap between traditional and big data processing requirements.

For the utilities industry, processing timeframes are just as intense. Power production and demand can change within hours based on weather fronts. Distributed solar assets do not generate power during snowstorms and customer demand can leap during these times, impacting the ability of utilities to balance their power generation requirements just as fast as a sunny morning can change into a cold and cloudy afternoon. GridGain makes possible the linkage between internal platforms to manage power generation with distributed power generation and IoT-based sensors for both generation and consumption.

In each of these situations, the ability to provide actionable information for corporate decisions rests on the providing information in near real-time. In-memory processing technologies that allow for a

flexible approach to the position and structure of data provide the highest probability of meeting the requirements presented by these new low-latency processing challenges. Limiting in-memory processing to a single database or data management architecture limits the ability of IT organizations to make decisions on how to best manage the processing of analytical, operational or operational analytical uses cases. Organizations such as GridGain offer the flexibility and breadth of processing options needed to make multiple use cases possible.

With its offering of an In-Memory Data Fabric, EMA considers GridGain a leading innovator of the type of in-memory processing platforms that helps the transition from traditional processing methodologies to in-memory processing across a wide set of use cases. GridGain is able to meet the challenges of operational, analytical or analytical operational workloads as well as a diverse set of implementation options. CIOs, Chief Data Officers (CDOs) and technology architects should regard GridGain as a strong addition to their data management architectures and a viable, and economically feasible, option to add competitive advantage via near-realtime processing requirements now and going forward.

Corporate Background

GridGain, a leading provider of an open source In-Memory Data Fabric, offers comprehensive in-memory computing solution to equip the real-time enterprise with an improved level of computing power. Enabling high-performance transactions, real-time streaming and analytics in a single, highly scalable data access and processing layer, GridGain enables customers to predict and innovate ahead of market changes. Fortune 500 companies, government agencies and innovative mobile and web companies use GridGain to achieve enhanced computing performance and business insights. GridGain is headquartered in Foster City, California.

Most of GridGain's customers are in the following markets:

- Financial services
- E-Commerce
- Media
- Risk management
- Healthcare
- Online gaming

Launched in 2010, GridGain Systems is privately held and headquartered in Foster City, California. To download the GridGain In-Memory Data Fabric, please visit http://www.gridgain.com/download/.

HIGHLIGHTS

Vendor name: GridGain Systems

Product(s) name: In-Memory Data Fabric

Product function: In-memory computing platform.

Vendor contact: info@gridgain.com

Availability: General Availability

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA's clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals and IT vendors at www.enterprisemanagement.com or blogs.enterprisemanagement.com. You can also follow EMA on Twitter, Facebook or LinkedIn.

Corporate Headquarters:
1995 North 57th Court, Suite 120
Boulder, CO 80301
Phone: +1 303.543.9500
Fax: +1 303.543.7687
www.enterprisemanagement.com

2968.092314