GridGain

In-Memory Computing for Financial Services eBook

Part 3: Improving the Performance of Asset and Wealth Management, Spread Betting, Fintech, and Banking Applications

A GridGain Systems In-Memory Computing eBook

November 2017

Table of Contents

Chapter 1 – Transforming Asset and Wealth Management with In-Memory Computing

The asset and wealth management sector of the financial services industry is undergoing rapid changes. Competition among investment-services providers is growing, with a wider range of providers courting a younger, more tech-savvy client base. These young investors bring with them new expectations and new challenges. To court them successfully, providers must offer the services they demand, including a 24/7 online investment environment with access to a wider range of investment vehicles than the traditional equity and bond funds.

Technology is key for maintaining a competitive edge in this situation. Investors will no longer wait for slow software or browser refreshes, so fast performance against big datasets and streaming data is crucial. Providers also need a scalable solution that interoperates with other systems, so they can offer the full range of digital channels and investment vehicles. Sophisticated analytics capabilities are essential as well, both for meeting increased regulatory requirements and for predicting advantageous investment scenarios.

To achieve this level of performance, scalability, and analytical sophistication, many financial-services providers are turning to in-memory computing solutions. This chapter will discuss the increased expectations of investors, the new challenges providers are facing, and how providers can gain the edge they need with solutions such as the GridGain in-memory computing platform.

What Today's Investors Want

In the past, investors typically made their investments through large, institutional investment firms with advisors who would help them choose appropriate investment vehicles – generally funds oriented toward equities or bonds. Today's investors are different. Overall, they are younger and more technology-driven, and these demographic trends are creating a new set of expectations for investment providers.

Today's investors want investment services that offer the following characteristics:

- **Multiple sources of personalized investment advice:** Rather than relying completely on an advisor to guide them, today's investors want to take a more active approach to finding an investment strategy tailored to their needs. They're consulting peers, comparative recommendations, and reports of investment performance, looking for transparent explanations and a sense of personal involvement.

- **24/7 access to investments:** Today's investors want 24/7 online access to their investment environment, including from mobile platforms, with timely alerts of changes in their holdings and their value. They want to be able to make rapid decisions based on personal preferences from anywhere at any time.

- **Individualized, wider-ranging investment strategies:** Rather than limiting their investments to funds, today's investors are interested in a wider range of asset classes and investment vehicles, including socially conscious investments, startups with innovative ideas, and opportunities in other countries. They are less interested in traditional risk-management models and performance benchmarks than in choosing a strategy aligned with their personal goals and trying to come up with an acceptable risk level for that strategy.

These expectations – along with the willingness of these young investors to move their money from firm to firm if their expectations aren't being met – have encouraged a fast pace of change among providers of investment services.

Changes Faced by Investment Providers

As the client base skews toward younger, more tech-savvy investors with higher expectations, the world of investment-services providers is also evolving, seeing the following types of changes and increased pressures:

- **More non-traditional providers**: New investment services providers – not just banks and financial institutions – are getting involved in this area, offering new technology platforms and a wider range of investment opportunities, such as crowdfunding and startups.

- **Increased competition:** Large investment houses are competing against smaller, more technology-focused firms, as providers see increasing investor movement among investments and firms.

- **New distribution channels:** Providers are using newer channels (online, mobile, social media) to reach and influence people choosing between financial services firms.

- **Customer-centric, 24/7 business model:** In today's environment, providers need to put customers first and make sure all available digital channels are meeting customer demands, providing round-the-clock access.

- **Increased cost pressures:** With more competition to reach potential investors and greater transparency about fees, there is pressure on providers to reduce their charges to clients. In addition to reducing fees, some providers are using common platforms to consolidating functions (investment management, risk management, and so on) and increase efficiency.

- **Increased market-volatility pressures:** In today's post-financial-crisis world, market volatility is 20 to 25 percent greater than it was previously. Providers need to make sure they are offering investment strategies that adequately address this volatility and mitigate its risk.

- **Increased regulatory pressures**. Providers today must meet increased reporting requirements and implement new risk-management. Providers must do capital planning and stress testing to keep their balance sheets healthy. They must also provide high-quality data and reporting related to trades and transactions, as well as establish robust anti-fraud and cyber-security

programs. In addition, they must comply with strengthened requirements in the areas of consumer protection, liquidity, foreign banking, crime prevention, and data governance.

In this environment of increased competition, customer expectations, and regulation, as well as other pressures, providers are finding that technology is the key to maintaining a successful edge.

Meeting New Challenges with Cutting-Edge Technology

As investment providers seek ways to cope with the increased expectations and pressures of today's environment, they are focusing on how technology can help with the following three challenges:

- Providing a real-time, 24/7 individualized environment with a full range of investment services
- Facilitating customer access through multiple channels and platforms, including mobile devices and social media
- Ensuring that all provided advice and investment actions meet rigorous regulatory requirements

Let's take an in-depth look at the technologies and techniques that providers are using to meet these challenges.

To satisfy customers who want 24/7 access to individualized investment advice, trading, and management of their portfolios, providers are using the following technologies:

- **Content automation:** Automating communications between the customer and the financial advisor as much as possible

- **Robo-advisors:** Using robo-advisors as needed to select individualized investments for customers based on the information they provide

- **Virtual reality:** Using virtual reality experiences to help customers visualize potential performance of various investment strategies

- **Big-data management and analytics:** Analyzing large quantities of data to provide analysis for customers

- **Advanced analytics:** Employing analysis mechanisms that allow customers to source inventories of various assets and decide what they need in their portfolios

- **Scenario analysis:** Modeling market movements in particular directions to make sure investors are prepared for potential changes

- **Machine learning:** Analyzing historical actions and performance to predict current behavior, using techniques such as behavior analysis, neural networks, heuristics, and "what-if" analysis

- **Complex Event Processing (CEP):** Creating algorithms and models that can automatically place orders for customers based on pre-configured ordering parameters and a real-time analysis of current market conditions

- **Blockchain capabilities:** Using blockchain digital-ledger technology for messaging, instructions, and payment options to reduce transaction costs and processing time (For a more in-depth look at blockchain capabilities and requirements, see "Enable Bitcoin and Blockchain Technology with In-Memory Computing")

In addition to creating an investment environment that is available 24/7, providers face the challenge of enabling customers to access their investment services seamlessly from any communications channel or platform.

To meet these challenges, providers are using the following technologies and techniques:

- **Multi-channel integration**: Making the available capabilities (including advice, account setup, quote requests, real-time notifications of portfolio changes, and so on) similar across all digital communication channels and platforms, including social media and mobile devices

- **Session-state tracking**: Tracking the user session state across multiple platforms, to enable a consistent experience

- **Cloud solutions**: Putting data and software in the cloud to enable convenient access from any physical or virtual environment

- **Open source and open architecture solutions:** Using tools that are open source or open architecture to facilitate interoperability among platforms

A third important challenge for providers is making sure the investment advice they provide – and the actions they facilitate, such as trades – meet regulatory requirements that have increased substantially in recent years. These requirements include time-stamped orders, audit trails, and pre- and post-trade compliance checks to ensure that orders are executed with the best price and lowest cost for customers. (For an in-depth look at the regulatory challenges faced by financial-services providers, see "Achieving Real-Time Financial Regulatory Compliance with In-Memory Computing")

To meet these requirements, providers are using the following technologies:

- **Big-data analytics:** Real-time monitoring of large quantities of data (such as current prices of assets across all investment platforms) to ensure compliance with regulations

- **Fast-data analytics:** Performing fraud detection and prevention and making automatic decisions about whether to approve trades and other actions such as asset transfers

- **Fully integrated security:** Implementing rigorous security standards to ensure privacy and compliance with cybersecurity regulations

Fortunately, there is a technology that is available with these features and is well-suited to this type of high-speed, big-data use case: in-memory computing, as implemented in the GridGain in-memory computing platform. Let's look at why in-memory computing makes sense for providers of investment services.

Customer Case Study: Sberbank. One of the most noteworthy GridGain Systems financial services customers is Sberbank, the largest bank in Russia. Sberbank was faced with a similar problem to the one currently facing companies who are transitioning from traditional, human-advisor-based investment services to real-time, online investment environments. The bank was switching from a more traditional, brick-and-mortar setup – one in which people would come into their offices and manually process a limited number of financial transactions each day, during a limited time period – to a new world with online and mobile customers transacting with them 24/7. Their transactions were increasing from 30-40 per second to 3,000-4,000 per second as a result.

The company forecasted future throughput requirements and determined that it needed to move to a next-generation data-processing platform to handle the expected transaction volume. Sberbank analyzed more than ten potential solutions from vendors in the in-memory computing space and found that the GridGain in-memory computing platform was the most comprehensive solution. The bank concluded that GridGain would provide the next-generation platform with a significant improvement in performance and scalability.

The GridGain in-memory computing platform provided several other important capabilities that Sberbank's next-generation platform would require such as machine-learning and analytics, flexible pricing, artificial intelligence, ease of deployment, hardware independence of cluster components, and a rigorous level of transactional consistency. Of particular importance was the ability to conduct integrity checking and rollback on financial transactions. Sberbank could not find that level of consistency with other in-memory computing solutions.

In a January 2016 article in RBC, Herman Gref, the CEO of Sberbank, said that the bank selected the GridGain Systems technology to build "a platform that will enable the bank to introduce new products within hours, not weeks." He went on to state that the GridGain in-memory computing platform enables Sberbank to provide "unlimited performance and very high reliability" while being "much cheaper" than the technology used previously. Sberbank is using GridGain's in-memory computing platform to implement capabilities that could not be provided by the other vendors evaluated – a group that included Oracle®, IBM® and others.

Today's Best Investment for Asset and Wealth Managers

As the field of asset and wealth management widens to include both traditional and non-traditional investment-services providers, more and more providers are competing to attract today's young, tech-savvy investors. In this environment, providing state-of-the-art 24/7 services at real-time speeds is a must, and providers must up their game with greater transactional speed and analytic power to beat the competition.

Fortunately, in-memory computing solutions provide the level of performance that asset and wealth management companies need. The GridGain in-memory computing platform provides asset and wealth management services with the performance, scale, and flexible interoperability that they require to distinguish themselves in a competitive market. The GridGain in-memory data grid can be inserted between existing application and data layers with no rip-and-replace to improve the speed and scalability of existing applications while supporting ACID transactions in an always-on individualized investment environment. Sophisticated streaming capabilities enable the integration of market and news feeds with portfolio information. The GridGain in-memory database provides parallel processing capabilities to run ANSI-99 SQL analytics and other business processes in real time.

Combining an open source framework with enterprise-level features, the GridGain in-memory computing platform offers a high-performance, scalable, comprehensive, secure, and affordable solution – an elegant and efficient way to give investment providers, and the investors they serve, the high-performance edge they need.

Chapter 2 – Spread Betting Pays Off with In-Memory Computing

During the past decade, financial spread betting has become a major growth market globally. Tradefair, one of the largest betting companies in the UK, has seen 20 to 30 percent growth per year for the past 10 years, with more than a million people in Great Britain opening spread betting accounts.

Spread betting offers some compelling advantages, including low entry and transaction costs, preferential tax treatment, and a diverse array of products and options. Traders can bet on any type of event for which there is a measurable outcome that might go in either of two directions – for example, housing prices, the value of a stock-market index, or the difference in the scores of two teams in a sporting event. Spread betting can be conducted through brokerages, asset managers, online gambling firms, and any other player who is willing to provide a *spread*, or range of outcomes, upon which traders can bet that the measured value will rise or fall. The upfront cost to place a bet is minimal and rewards can be substantial.

However, spread betting is also a high-volatility, minimally regulated market with significant risks. To limit these risks – and increase the rewards – financial institutions involved in spread betting are using advanced mathematical models to analyze large amounts of data, predict outcomes, and devise optimal strategies. These computationally intense actions must be performed at very high speeds to take advantage of current market conditions.

Fortunately, there are technologies today that provide the real-time speed needed for such strategies. In-memory computing platforms such as the GridGain in-memory computing platform, built on Apache® Ignite™, can provide both exceptional performance and other valuable features, such as scalability, high availability, and fully ACID-compliant transactions.

This chapter will discuss the advantages and risks of spread betting, the technologies being used for it, and the reasons why in-memory computing is becoming the technology of choice for brokerages, asset managers, online gambling firms, and other players who want to succeed at spread betting.

Advantages of Spread Betting

Spread betting continues to grow in popularity because it offers numerous advantages, including:

- **Tax preference**. In the UK, traders pay no capital gains on the profits received. This tax advantage makes spread betting very lucrative compared to other strategies, such as buying futures or stocks.

- **Low cost of entry**. Unlike more regulated markets, in which a trader might need to buy at least 100 contracts or 100 shares, there are no minimum stakes with spread betting. A trader can bet as little as one dollar, which means there are relatively low margins.

- **Low transaction costs**. With spread betting, there is no direct commission. All commissions and risks are included in the spread.

- **Large upside potential.** Because of the low cost of entry, traders can potentially make big profits with relatively little investment, if they understand the risks.

- **24/7 operations with flexible entry and exit**. Unlike the stock market, where trades occur between 9:30 a.m. and 4:00 p.m., spread betting is a 24-hour market. Traders can enter, trade, and exit at any time, as well as see prices in real time.

- **Less regulation**. The products involved in spread betting are not technically considered to be financial products, so they are scrutinized less and have fewer regulatory reporting requirements. This lack of regulation helps to keep transaction costs low.

- **Diversity of products, markets, and strategies**. It is possible to make the outcome of almost any type of event (financial, sporting, weather-related, and so on) into a product that can be traded in spread betting markets. As a result, brokers or bookmakers can potentially create a wide variety of products. Also, those who can discern correlations between disparate assets can be creative in developing their own strategies for betting, as well as for hedging and risk management in general.

- **Single-account access**. Traders can access all types of spread betting products from a single account.

- **Fully electronic platforms**. Spread betting does not require going to see a broker or calling anybody. Everything happens online. Traders can receive their profits or pay their losses electronically.

- **Mobile trading**. Most spread betting platforms allow traders to use their smartphones or other mobile devices to connect and trade.

- **Automatic stop losses**. This option lets traders specify certain limits and thus limit their losses – saying, for example, "If that index falls below a certain level, I want to sell and get out of this position." There is no need to wait until the end of the day or whatever is the duration of the bet.

Of course, individuals and institutions that want to get involved in spread betting also need to consider the risks and how to manage them.

Increased Risk and Other Disadvantages

Spread betting can also have some disadvantages and increased risk compared to other forms of financial speculation, due to the following characteristics:

- **No "market maker" to assume risk**. In traditional financial markets, some of the risk is assumed by market makers. They hold large amounts of an asset and offer both sides of a position in response to market demands, buying when investors want to sell and selling when investors

want to buy. So, investors know there is always somebody on the other side who is willing to intervene in case of an asset falling sharply or any other condition. With spread betting, there are bookmakers instead of market makers. Bookmakers have no obligation to step in and help traders who bet erroneously or find the market moving sharply against them.

- **Less consumer-protection oversight**. With spread betting, as noted earlier, there is no financial authority providing oversight. While this lack of regulation provides some advantages, it also means there is less consumer protection. Traders are covered only by standard consumer-goods regulations. Fraud is not allowed, and there are age restrictions on certain activities, but traders are otherwise on their own with no financial authority providing backup.

- **Wider bid/offer spreads**. Bid/offer spreads exist both in the traditional financial-services business and in spread betting. In traditional markets, which are more liquid and more controlled, the bid/offer spread is typically narrower than in riskier transactions like spread betting. Wider bid/offer spreads provide more profit to brokers and bookmakers and less profit to traders.

- **Leveraged transactions**. Brokers in traditional financial services businesses and in spread betting businesses both typically provide margins – that is, loans to investors or traders – but the leverage is typically much more substantial with spread betting. With more leverage, the potential for loss is higher.

- **High interest rates for overnight financing**. Traders are not supposed to have open transactions at the end of the day. If they are losing at that point because the market is turning against them and they need to replenish their accounts, brokers or bookmakers will provide overnight financing – with significantly higher interest rates than would be available elsewhere.

- **High volatility**. Volatility in spread betting is much higher than in some other markets, as prices can fluctuate greatly depending on related events and how soon the event being bet on will draw to a close.

- **Low transparency**. With spread betting, traders can see the bookmaker's or venue's prices, but they do not know what other people are betting. Unlike in a more traditional financial services business, there are no clear reporting requirements. Traders or those representing them must make calculations using their own models or data analysis.

To manage the increased risks, it is important for anyone involved in spread betting to incorporate risk-management strategies.

Risk-Management Strategies

Participants in spread betting can limit its risks with the following combination of strategies:

- **Hedging**. Traders or those representing them can build a hedging framework by evaluating how one product offered by the spread betting business is related to others. Hedges are alternate

bets or investments that provide some protection in case of adverse movements in the original position. Developing a hedging framework typically involves mathematical or statistical models.

- **Modeling outcomes**. Being able to predict outcomes with a high degree of certainty is another good way to limit risks in spread betting. Successful predictions often result from building mathematical or statistical models that predict the outcome of a bet based on outcomes of related events. Such models are complex and rely on feeding accurate data inputs into probability regression analysis.

- **Subscribing to data services**. To get accurate and timely data as input for their models, traders or those representing them may subscribe to data services that provide information about relevant events. This information might include news feeds and market data from financial services markets or data from other spread betting venues that provide similar or related products.

- **Analyzing news**. To use incoming news effectively in spread betting, it is important to know how the news will play into people's opinions – and how opinion changes are likely to translate into the buying or selling of bets in a spread betting environment. A strategy called *sentiment analysis* involves analyzing historical patterns in people's opinions about certain types of news, then incorporating these patterns in the spread betting models for predictive purposes.

- **Using state-of-the-art technology**. Effective technology is the key to making all of these risk management strategies work together successfully. With the right technology, traders can stream data inputs into models that quickly compute event relationships and changing outcome probabilities while the events are occurring – so traders or those representing them can act fast with well-informed bets and hedging strategies.

The next section looks at the technology requirements of a state-of-the-art system for spread betting in more detail.

Essential Technologies for Spread Betting Systems

Both traders and spread betting venues need high performance and highly scalable systems to provide them with information, to perform complex modeling and develop hedging strategies, and to allow them to rapidly access and trade as many products as possible.

To create these systems, the following technologies are employed:

Big data. To know how to interpret events outside of the bid-ask spreads of the spread betting venue, traders or those representing them need to analyze large amounts of data from subscription services. Their systems need to have data repositories that get updated with each tick of new incoming data. Big data technologies provide ways to organize these large datasets into multiple pools and connect them in real time for immediate analysis.

Apache Hadoop® with MapReduce. Analyzing large datasets in real time requires speed and efficiency.

Complex event processing (CEP). This technology involves looking at multiple streams of incoming data and using artificial intelligence (AI) to identify meaningful events. It is particularly useful in spread betting systems involving mathematical models built to receive ticks from multiple streams, from multiple event lines. CEP helps models to update themselves, changing ratios and other parameters in response to meaningful information in the streams of incoming news and market data.

Robots. When there are many people receiving similar information and betting on the same assets, traders or those representing them can gain advantages by building their own algorithms or "robots" to automatically buy or sell depending on certain levels of information or outcomes of related events received in an electronic format.

Real-time systems. Speed and performance are critical when systems need to interpret incoming data in real time and take immediate action based on that real-time information. For this reason, spread betting venues and traders rely heavily on real-time systems.

Online and cloud platforms. The spread betting industry typically uses mobile and Internet-based platforms that must be highly scalable and performant.

Mobile trading. Most spread betting platforms not only allow for mobile trading access, they also enable the mobile trading to operate in a high-performance mode by accessing data and information within the parent platform.

Data partitioning and parallel processing clusters. Because there are many different assets to trade, as well as a wide variety of products and a diverse product base, spread betting systems need to use some additional performance improvement mechanisms, such as data partitioning and parallel processing clusters – that is, clusters of connected computers processing the data in parallel. These mechanisms are important for supporting 24/7 data access and transactions.

Scalable data architecture. When increasing numbers of users are coming into a spread betting system, it is important to be able to increase the processing power and increase the amount of memory being used. If the architecture is based on a cluster of computers, adding another computer into the cluster provides an easy way to scale up.

In-memory computing. Spread betting is an analytically intense process that uses performance-hungry models, and it is most effective when performed in the fastest possible way: using in-memory computing. In-memory computing involves keeping data in RAM for extremely fast access with no disk-related slowdowns. This makes it faster than any other storage-based computing method.

Chapter 3 – Faster Fintech with In-Memory Computing

Financial institutions have been running on computers for more than 50 years, but these businesses are currently going through rapid transformation. Fintech, the new and disruptive combination of finance and technology, is quickly changing the way banks, financial institutions, and market makers do business. Looking for every possible edge in speed, security, cost savings, customer service, and data analysis, they are turning to nimble outside developers and open source software to create applications that they used to create in house. The result is a vibrant burst of tools and capabilities that they, and their customers, have never had before.

Fintech Is on the Rise

The Fintech market grew from $1.8 billion in 2010 to more than $19 billion in 2015. According to a 2016 Accenture report, venture capitalists, private equity firms, and other players have invested $50 billion in almost 2,500 Fintech companies since 2010. There are many reasons why.

Driven by all that VC investment and an accompanying startup mentality, Fintech development is happening internationally, and the applications are easily reaching a worldwide market of global banks. Many barriers to entry have been removed as developers take advantage of lower-cost technologies and open source software to innovate and then implement and iterate their ideas quickly.

Banks are rapidly adopting cloud-based applications and outsourcing some of their data processing and analysis to Fintech companies that can do it faster, better, and within very specific regulatory strictures that they specialize in addressing. By offloading a significant percentage of their traditional back-end work, banks can focus on their core specialties and improve their customer service.

Today's Top Fintech Trends

According to January 2017 data from Venture Scanner, funding for Fintech development is directed most at consumer lending applications, followed by consumer payments, payments backend, and business lending. That means banks are currently most interested in using Fintech to **address consumer-oriented finance**, with 24/7 availability, digital banking, payment through social media connections, biometrics, contactless spending, and other trending consumer-facing technologies.

At the same time, financial institutions are eager to **automate investments and trade**, relying more on algorithmic and data-driven conclusions than on human guesswork. As what were formerly institutional strategies start to go mainstream, traders can implement intelligent high-frequency trading driven by machine learning, instant risk analysis, and sentiment analysis even as they lower fees.

Lending is also growing rapidly as concepts such as peer-to-peer lending and crowdfunding come to market and disrupt banking's traditional lending models. Even **insurance is being transformed** as everything from Internet-of-Things (IoT) sensors to social media analysis deliver more and better data at

the individual customer level, letting insurers customize policies and offer more relevant products to their customers.

Fintech Faces Challenges

Although there is no stopping the Fintech revolution, there are a few challenges in its path that must be addressed going forward:

- **Global turmoil** is always a hazard in financial industries, especially at a moment when a wave of isolationist thinking introduces uncertainty in the global banking sector.
- **Name recognition** can be a problem for Fintech companies hoping to attract consumers who feel safe and familiar with banks they have used for generations.
- **Reputational risk** is a concern when new and perhaps outsourced online technologies stumble.
- **Regulatory risk** confronts both Fintech and traditional financial firms trying to adopt new and untested practices.
- **Legacy technology** can slow Fintech acceptance. Banks may, for example, be totally unprepared for innovations like Bitcoin and blockchains.
- **Security** is always a concern when new kinds of data are collected, especially from consumers.

One way banks have addressed such concerns is to invest in or partner directly with Fintech innovators, signing cooperative agreements and working together to move forward more quickly. Financial institutions are getting a real jolt from Fintech innovations, and as it turns out, one real game changer is the implementation of cutting-edge in-memory computing technology.

In-Memory Computing: An Enabling Technology for Fintech

In-memory technology has been around for decades, ever since developers realized that working with data stored in memory was faster than working with data stored on disk. Today, that speed means everything, especially in split-second financial marketplaces. Digital transformation has pushed companies closer to customers who demand real-time interactions, and the amounts of Internet traffic, connected devices, and data—all of which require instant analysis—continue to grow. It is estimated that total global data will grow from eight zettabytes in 2015 to 35 or more zettabytes in 2020.

Luckily, the cost of memory continues to tumble—down 30 percent annually over time—so moving to in-memory platforms has gotten more cost-effective with each passing year. Thus, Gartner has reported, the in-memory technology market will grow to $10 billion by the end of 2019, representing 22 percent compound annual growth over time.

The combination of these two trends means that the time is right to leverage the improved performance and scale provided by in-memory computing to make Fintech implementations run optimally.

GridGain and Apache® Ignite™ Address Fintech's Needs

The growth of the in-memory computing market includes in-memory databases, in-memory analytics, in-memory data grids, in-memory app servers, in-memory messaging, event stream processing, and analytic servers. Apache Ignite, the open source version of GridGain Systems' in-memory computing platform, provides complete support for all of these components.

Apache Ignite:

- Supports data caching, massive parallel processing, in-memory distributed SQL, streaming, and much more.
- Slides in between the existing application and data layers easily, with minimal changes to code.
- Offers ACID-compliant transactions as well as analytics support.
- Works with all popular RDBMS, NoSQL, and Hadoop databases and offers a Unified API with support for a wide range of languages.
- Can be deployed on premise, in the cloud, or in hybrid environments.
- Can be configured for high availability with automatic load balancing and failover.
- Remains highly customizable for a multitude of uses as it handles computing and transacting on large-scale data sets in real time.

The GridGain and Apache Ignite in-memory computing platforms run everything in memory to yield the performance and scalability gains that Fintech developers need. The GridGain Enterprise Edition adds expected enterprise features such as better security, better network segmentation protection, data center replication, and rolling updates that allow updating of multiple nodes on a cluster without taking the whole cluster down.

Fintech Benefits from In-Memory Advantages

In a recent survey, GridGain asked its customers in the Fintech space for what they were considering using in- memory computing. The top three answers were high-speed transactions (more than 85 percent of respondents), real-time streaming, and database caching. Also on the list: building a hybrid transaction and analytics (HTAP) platform, and doing application scaling. Which data stores are Fintech developers most likely to connect to? Apache Cassandra™ was the clear leader, followed by Oracle® and MySQL®.

The survey also found that Fintech developers were most likely to run GridGain and/or Apache Ignite on premise, followed by AWS® and a private cloud. No one mentioned Microsoft Azure, SoftLayer, or Google Cloud Platform. SQL was by far the language Fintech developers prefer to use to access their data, followed by Java and C++. This was surprising given that C++ is used heavily in the Fintech space.

It is interesting to understand how GridGain's Fintech customers choose to deploy the in-memory computing platform. A powerful starting point is the **In-Memory Data Grid**, inserted between the application and database layers to cache the disk-based data from the RDBMS, NoSQL, or Hadoop databases in RAM, delivering a 1,000x performance increase. This in-memory key value store automatically replicates and partitions data caches across multiple nodes and deliver elastic on-demand scalability to add new nodes. Distributed in-memory transactions can also be ACID-compliant. The data

grid offers support for all popular RDBMS, with read-through and write-through and support for write behind. Setup remains completely flexible to address even the most unique use case.

Deploying In-Memory Computing for Velocity, Volume and Scalability in Fintech Implementations

Many Fintech users have adopted GridGain's in-memory computing solution to improve performance and functionality by taking advantage of its distributed computation and in-memory computing speed. It has strong use cases for:

- Core Banking and Trading Platforms
- Risk Management
- Real-Time Financial Analytics
- Big Data Analytics
- Compliance and Monitoring
- Financial SaaS Platforms

For Finastra, Improvements in Data Velocity and Volume

Finastra, a financial services software provider with over 2,000 clients including 48 of the world's 50 largest banks, needed to eliminate data processing bottlenecks. Its clients deal with huge amounts of trading and accounting data, and they need to manage high-speed transactions and conduct real-time reporting, something that Finastra's software architecture could not provide. There was a true need for speed. Nighttime batch processing was no longer acceptable.

Finastra's solution: To move to a Java-based IT stack that taps into a data lake as opposed to a traditional data warehouse, and to do that while creating an environment that can address the individual regulations of up to 130 countries where their clients conduct business. The new Finastra commodity servers each contain 256GB RAM to store and process transaction and market data in-memory and manage parallel processing across GridGain clusters.

Finastra has also used in-memory computing tactics to develop and launch its new FusionFabric Connect product, a cloud-based SaaS collection of modules that integrates many trading systems.

Tomorrow's Banking Today

Banks and financial institutions want—and need—it all: speed, high availability, security, scalability, and flexibility to give them a competitive edge today and tomorrow. Fintech developers have many of the answers, not just for the banks but also for themselves as they seek out their own consumers for new kinds of banking, investment, and insurance services. The common denominator is in-memory computing, the technology platform that can handle the unprecedented amount of data now available for analysis and process the billions of transactions that these institutions need to manage in near real-time. Moving money around the world has always been challenging, but Fintech is finally making it easier with the assistance of GridGain's in-memory computing platform.

Chapter 4 – Banking on In-Memory Computing with GridGain and Apache Ignite

Addressing the challenges highlighted in this eBook requires speed, scalability, availability, security and flexibility. In short, they need distributed in-memory computing. In-memory computing eliminates the disk access bottleneck that slows down applications built on disk-based databases. An in-memory computing platform enables users to process transactions 1,000 times faster than disk-based solutions and enables scale out to terabytes of in-memory data by adding new nodes to the cluster.

In-memory technology has been around for decades. Until recently, however, the cost of RAM made in-memory computing practical only for the highest value applications. The cost of memory continues to fall, dropping an average of 30 percent per year, which makes in-memory computing platforms economical for a wider range of use cases each year. Gartner projects that the in-memory technology market will grow to $10 billion by the end of 2019, a 22 percent compound annual growth rate.

In-memory computing platforms include key features that are now essential for many financial applications. These features, available in Apache Ignite and the **GridGain in-memory computing platform**, include an in-memory data grids with strong distributed SQL capabilities, in-memory database capabilities, streaming analytics, and native integrations with a variety of other open source projects including Apache® Kafka™, Spark™, Cassandra and Hadoop.

GridGain and Apache Ignite are deployed as an in-memory computing layer between the application and data layers. The products work with any RDBMS, NoSQL or Hadoop database. The **In-Memory Data Grid** with strong distributed SQL capabilities is a key-value store which can replicate and partition data caches across multiple nodes and deliver elastic on-demand scalability. Distributed in-memory ACID transactions are also supported. The Data Grid offers support for all popular RDBMSs, with read-through and write-through and support for write behind. Setup is flexible to address unique use cases.

The **In-Memory Database** leverages the systems distributed SQL capabilities. It is horizontally scalable, fault tolerant and ANSI-99 SQL compliant. It also supports all SQL, DDL and DML commands including SELECT, UPDATE, INSERT, MERGE and DELETE queries and CREATE and DROP table. The SQL syntax is ANSI SQL-99 compliant. GridGain can use any SQL function, aggregation, or grouping. GridGain supports distributed SQL joins and allows for cross-cache joins. Joins between partitioned and replicated caches work without limitations while joins between partitioned data sets require that the keys are collocated. GridGain supports the concept of fields queries as well to help minimize network and serialization overhead. The in-memory distributed SQL capabilities allow users to interact with the GridGain platform not only with the usage of natively developed APIs for Java, .NET and C++ but also using standard SQL commands through the GridGain JDBC or ODBC APIs. This provides a true cross-platform connectivity from languages such as PHP, Ruby and more.

The **In-Memory Compute Grid** enables distributed parallel processing of resource-intensive compute tasks. It offers adaptive load balancing, automatic fault tolerance, linear scalability, and custom scheduling. Built around a pluggable SPI design, it offers a direct API for Fork-Join and MapReduce processing.

The optional **Persistent Store** feature in the memory-centric Apache Ignite architecture is a distributed disk store that transparently integrates with GridGain as an optional disk layer. It may be deployed on spinning disks, solid state drives (SSDs), Flash, 3D XPoint or other similar storage technologies. Persistent Store allows organizations to maximize their return on investment by establishing the optimal tradeoff between infrastructure costs and application performance by adjusting the amount of data that is kept in-memory.

The **In-Memory Service Grid** provides control over services deployed on each cluster node and guarantees continuous availability of all deployed services in case of node failures. It can automatically deploy services on node startup, deploy multiple instances of a service, and terminate any deployed service. It is a load-balanced and fault-tolerant way of running and managing services across the grid.

Stream analytics establish windows for processing and run either one-time or continuous queries against these windows. The event workflow is customizable and often used for real-time analytics. Data can be indexed as it is being streamed to make it possible to run extremely fast distributed SQL queries against the streaming data.

In-memory Hadoop acceleration provides easy-to-use extensions to disk-based HDFS and traditional MapReduce, delivering up to 10 times faster performance. GridGain and/or Ignite can be layered on top of an existing disk-based HDFS and used as a caching layer offering read-through and write-through while the GridGain Compute Grid can run in-memory MapReduce.

Many enterprise financial services firms have adopted GridGain's in-memory computing solution to improve performance and functionality. Clients including Apollo Global Management, Barclays, Cambridge Associates, Citi, Fidelity, and ING have all taken advantage of GridGain's distributed computation and in-memory computing speed.

Conclusion

With the tight regulatory environment, competition from traditional and non-traditional industries, customer demands, and cost pressures that companies are facing today, financial services and fintech companies require big data technologies that make processes and transactions much faster and more efficient. Large companies accumulating massive amounts of data need to be able to perform analytics on that data in real time in a cost-conscious manner to ensure a good user experience. Many are finding in-memory computing platforms such as GridGain and Apache Ignite to be a key strategy for meeting these challenges.

Additional Resources

For more information about topics covered in this eBook, the following resources are available from the GridGain website:

eBook: In-Memory Computing for Financial Services: Part One

eBook: In-Memory Computing for Financial Services: Part Two

Datasheet: The GridGain In-Memory Computing Platform

Case Study: Misys (Finastra) Uses GridGain to Enable High Performance, Real-Time Data Processing

Article: **How Russia's Oldest Bank Found Itself on the Edge of In-Memory Computing in CIO Magazine**

Webinar Recording: Powering Financial Spread Betting with In-Memory Computing

Webinar Recording: Transforming Wealth and Asset Management with In-Memory Computing

Contact GridGain Systems

To learn more about how GridGain can help your business, please email our sales team at sales@gridgain.com, call us at +1 (650) 241-2281 (US) or +44 (0) 7775 835 770 (Europe), or complete our contact form to have us contact you.

About GridGain Systems

GridGain Systems is revolutionizing real-time data access and processing by offering an in-memory computing platform built on Apache® Ignite™. GridGain solutions are used by global enterprises in financial, software, e-commerce, retail, online business services, healthcare, telecom and other major sectors, with a client list that includes Barclays, ING, Sberbank, Finastra, IHS Markit, Workday, and Huawei. GridGain delivers unprecedented speed and massive scalability to both legacy and greenfield applications. Deployed on a distributed cluster of commodity servers, GridGain software can reside between the application and data layers (RDBMS, NoSQL and Apache Hadoop), requiring no rip-and-replace of the existing databases, or it can be deployed as an in-memory transactional SQL database. GridGain is the most comprehensive in-memory computing platform for high-volume ACID transactions, real-time analytics, web-scale applications and HTAP. For more information, visit gridgain.com.