



Empowering Fintech with In-Memory Computing

A GridGain Systems In-Memory Computing White Paper



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 SOME 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.

APACHE® IGNITE™ ADDRESSES 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, 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
- Works with all popular RDBMS, NoSQL, and Hadoop databases and offers a unified API with support for a wide range of languages
- Can be used as a memory-centric, distributed SQL database
- 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 2016 survey, GridGain asked its customers in the Fintech space what they were considering using in-memory computing for. 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? 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.

The In-Memory SQL Grid is the horizontally-scalable, fault-tolerant, ANSI SQL-99 compliant portion of GridGain and Ignite, with support for all SQL and DML commands such as SELECT, UPDATE, INSERT, MERGE, and DELETE queries. It is a mature, in-memory solution to supplement or replace a disk-based RDBMS. Geospatial support is built into the product, and all the communication to the SQL grid is done through ODBC and JDBC APIs without custom coding.

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 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.

In-memory streaming and continuous event processing 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.

With all these features and tools enabled and running optimally, financial services developers can vastly improve their applications by:

- Processing high-speed transactions faster than ever
- Speeding up search dramatically
- Deploying better systems for the management of IoT components

In the GridGain survey of Fintech customers, it was interesting to note that Compute Grid was deemed the most important product feature, above Data Grid. In many vertical industries GridGain surveys, the result is usually opposite. They were followed by the in-memory file system, Service Grid, and ANSI SQL-99 compliance.

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 Misys, Improvements in Data Velocity and Volume

Misys, 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 Misys's software architecture could not provide. There was a true need for speed. Nighttime batch processing was no longer acceptable.

Misys'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 Misys commodity servers each contain 256GB RAM to store and process transaction and market data in-memory and manage parallel processing across GridGain clusters.

Misys 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.

For Sberbank, A Fault-Tolerant and Scalable Open-Source Platform for Growth

Sberbank, the largest bank in Russia and Eastern Europe and the third largest in Europe, was expecting significant growth in its transaction volume and wanted to minimize its legacy dependency on Oracle by migrating to an open-source data grid architecture for its next-generation platform. The bank needed to introduce new products in hours, not weeks, and the platform had to have virtually unlimited performance and very high reliability while lowering costs and reducing human interaction.

The GridGain solution Sberbank ultimately implemented offers a very high level of redundancy and high availability and has the security levels Russian regulators demand. In testing, the bank was able to generate up to one billion transactions per second on an array of ten Dell R610 blades with one terabyte of memory, assembled at a cost of just \$25,000.

GRIDGAIN SYSTEMS: A LEADER IN IN-MEMORY COMPUTING

The GridGain in-memory computing platform helps Fintech developers consolidate onto a single high-performance and highly scalable big-data solution for transactions and analytics, resulting in lower TCO.

Advanced SQL functionality and API-based support for common programming languages enable rapid deployment. This, coupled with the rapidly decreasing cost of memory, boosts ROI for in-memory computing initiatives, enabling Fintech developers to create solutions that can perform thousands of times better. Any Fintech company seeking speed, security, and rapid deployment can yield the benefits of the GridGain in-memory computing platform.

Clients enjoy the following:

A Unified High-performance Architecture

The GridGain in-memory computing platform consists of multiple grids connected by a clustered in-memory file system. The In-Memory Data Grid, In-Memory Compute Grid, In-Memory SQL Grid and In-Memory Service Grid are interconnected. Computations occur as close as possible to the data used in the computation. Additional features such as high throughput, low latency, load balancing, caching, in-memory indexing, streaming, Hadoop acceleration and other performance improvements are crucial to success in real-time modeling, processing, and analytics.

Scalability

The GridGain in-memory computing platform excels in terms of scalability, allowing companies to add cluster nodes and memory in real-time with automatic data rebalancing. As a hardware-agnostic solution, clients can choose their preferred hardware for scaling up.

Full SQL Support

GridGain is ANSI SQL-99 compliant and the In-Memory SQL Grid supports DML users can leverage their existing SQL code using the GridGain JDBC and ODBC APIs. For users with existing code bases which are not based on SQL, they can leverage their existing code through supported APIs for Java, .NET, C++, and more.

High Availability

The GridGain in-memory computing platform provides essential high availability features such as data-center replication, automatic failover, fault tolerance, and quick recovery on an enterprise-level scale.

Transaction Processing

The GridGain in-memory computing platform supports ACID-compliant transactions in a number of user-configurable modes.

Security Features

The GridGain in-memory computing platform supports authentication, authorization, multiple encryption levels, tracing, and auditing.

Open Source Framework

GridGain is based on Apache® Ignite™, a popular open source project with many contributors that has been tested globally. GridGain Systems was the original creator of the code contributed to the Apache Software Foundation that became Apache Ignite and fully supports the technology behind Apache Ignite. The GridGain Enterprise Edition extends the features in Apache Ignite to provide enterprise-level capabilities and services, such as additional security, data center replication, auditing mechanisms, a GUI for management and monitoring, network segmentation, and a recoverable local store.

Production Support

[GridGain Systems Support](#) is available for [GridGain Community Edition](#), [GridGain Enterprise Edition](#) and [GridGain Ultimate Edition](#) users. The Enterprise and Ultimate editions include rolling updates, faster availability of all releases and patches, and 24/7 enterprise-level support.

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 realtime.

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.

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)208 610 0666 (Europe), or complete our [contact form at www.gridgain.com/contact](http://www.gridgain.com/contact) and we will contact you.

About GridGain Systems

GridGain Systems is revolutionizing real-time data access and processing with the GridGain in-memory computing platform built on Apache® Ignite™. GridGain and Apache Ignite are used by tens of thousands of global enterprises in financial services, fintech, software, e-commerce, retail, online business services, healthcare, telecom and other major sectors, with a client list that includes ING, Raymond James, American Express, Societe Generale, Finastra, IHS Markit, ServiceNow, Marketo, RingCentral, American Airlines, Agilent, and UnitedHealthcare. 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, continuous learning and hybrid transactional/analytical processing (HTAP). For more information on GridGain products and services, visit www.gridgain.com.