WHITE PAPER

GridGain

Powering Financial Fraud Detection and Prevention with In-Memory Computing

A GridGain Systems In-Memory Computing White Paper

POWERING FINANCIAL FRAUD PREVENTION WITH IN-MEMORY COMPUTING

Financial fraud is now a multi-billion-dollar business and growing rapidly, with Juniper Research predicting that online fraud alone will climb from $10.7 billion in 2015 to 25.6 billion in 2020. Failure to detect and prevent fraud can harm the reputations of financial firms and reduce confidence in the industry as a whole.

Protecting their customers from fraud and protecting themselves from fraud-related losses are high priorities for financial institutions. However, fraud prevention is not a simple task, and firms must tackle it simultaneously with other crucial tasks such as ensuring regulatory compliance. To accomplish these data-intensive tasks in a timely manner, financial firms need solutions that are flexible, scalable, reliable, and fast enough to analyze extremely large datasets in real-time.

Fortunately, today's in-memory technologies provide powerful tools for combating fraud – tools that perform complex processing, modeling, and analysis of big data in real-time. This white paper will discuss what financial fraud is, how firms are addressing the problem, and why in-memory computing technologies such as the GridGain in-memory computing platform are perfectly suited to the task of detecting and stopping fraud wherever it occurs.

WHERE FINANCIAL FRAUD OCCURS

The term "financial fraud" encompasses a wide variety of illegal practices. Financial fraud typically shows up in the following arenas:

- Checks that are written by unauthorized institutions or officers
- Credit cards that are stolen and used illegally
- Mortgages that are illegally manipulated
- Corporate financial statements that are changed or illegally manipulated
- Securities that are traded using illegal techniques
- Payments that are requested fraudulently or rerouted to improper destinations
- Identity theft in which thieves steal financial information or impersonate others in order to make money
- Forgery of documents, signatures, banknotes, or works of art to produce financial gain
- Computerized banking and computer-based financial transactions employed improperly to produce financial gain
- Tax evasion in which corporations or people avoid paying taxes that they owe

The Financial Fraud Research Center at Stanford University estimates that Americans currently lose $50 billion dollars a year to fraudulent practices such as these. The magnitude of this loss provides strong incentives for financial firms to find more effective techniques for fraud prevention.

EVOLVING TECHNIQUES FOR FRAUD PREVENTION

Traditional approaches to identifying fraud relied on manual verification and analysis. These approaches involved people auditing transactions directly, looking at who was providing passwords and other information. There were few apparent ways to automate this process. However, technology has progressed to the point where firms are now able to collect much more information and employ sophisticated techniques to automatically process and analyze data to detect fraud as it happens.

The techniques that banks and financial services firms currently use to detect fraud include the following:

- Statistical and multi-channel analysis: Calculating parameters (such as averages and performance metrics), linking together data from multiple sources (channels), and analyzing correlations between different data measurements to find patterns that help with fraud detection
- Models and probability distributions: Calculating models and probability distributions that predict how financial data will behave, so actual data can be compared against predictions and variations can be flagged as potential indications of fraud
- User profiles: Computing and maintaining user profiles associating personal data with specific users to help identify atypical behavior or attributes that may be fraudulent or indicate identify theft
- Real-time algorithmic analysis: Using algorithms to identify and validate user actions as they occur, in real-time
- Data clustering and classification: Analyzing known patterns and profiles and classifying them for use in algorithms and models—essentially creating a data repository for fraud detection

GridGain

- Artificial intelligence and machine learning: Using machine learning techniques such as neural networks to refine automated fraud detection, reducing false positives (false alarms) and improving behavior-based predictions for current transactions and users

Performing these knowledge-intense activities in real-time on extremely large datasets requires high performance and highly scalable technologies, as the next section discusses.

TECHNOLOGIES USED FOR FAST DATA ANALYSIS

To apply the processing and analysis techniques needed for fraud detection in real-time or near real-time on a large scale, financial firms combine these techniques with technologies that can provide fast data analytics in a high intensity, transactional environment.

These technologies include the following:

Big Data

The first step in using financial data for fraud detection is to prepare data and make it available for analysis. Big data technologies provide ways to organize large datasets into multiple pools and connect them in real-time for immediate fraud detection and additional analysis.

Apache Hadoop® with MapReduce

Stopping fraudulent transactions in large datasets in real-time requires speed and efficiency. The average speed of executing a transaction may be only milliseconds, and within those single-digit milliseconds, the processing system must analyze the transaction, validate it, and check all available data pools without affecting the performance of processing the transaction. Hadoop with MapReduce is designed to help in situations exactly like these. It organizes hierarchical data to improve performance, allowing quick conclusions as to whether a transaction should be stopped.

Complex Event Processing (CEP)

This technology, used in many financial institutions today, involves looking at multiple streams of incoming data and using artificial intelligence (AI) to identify meaningful events, such as potential fraud. It uses neural networks and other AI paradigms to decide how incoming data elements affect the behavior of the system as a whole as transactions are processed.

Near Real-Time Systems

Trying to solve fraud issues after they occur is an expensive strategy and it poses risks to a company's reputation. A much better approach is to stop those transactions while they are happening. This approach requires extremely fast and efficient processing so financial firms are turning to near real-time systems.

Data Partitioning and Parallel Processing Clusters

When there are many transactions coming in at the same time, lining them up one by one to check them for fraud is not an option. To operate in real-time and maintain acceptable performance, the system must include multiple processors operating on the data simultaneously – that is, clusters of connected computers processing the data in parallel. It is also important to have the distribution available on the clusters to process those transactions regardless of where they occur, while maintaining data consistency. A system with data partitioning and parallel processing clusters is essential to meet these needs.

Scalable Data Architecture

We are operating in the world of constantly growing data. Large financial institutions are experiencing 20 to 30 percent data growth year over year, and they cannot risk running out of space. They must be able to add more storage while not losing performance, which means they need a scalable data architecture.

In-Memory Computing

Combating fraud is an analytically intense process that uses performance hungry models and it must be performed in the fastest possible way: using in-memory computing. Because in-memory computing involves keeping data in RAM for extremely fast access, with no disk-related slowdowns, it is faster than any other storage-based computing method. In the next sections, we will discuss how in-memory solutions such as the GridGain in-memory computing platform have evolved to be fast, affordable, and comprehensive in their ability to combine all of the technologies listed above.

THE MOVE TO IN-MEMORY GRID COMPUTING: FASTER, BETTER ROI

For applications that require heavy analytics and real-time (or near real-time) transaction processing of hundreds or even millions of transactions per second, the market is now moving from disk to in-memory computing. The reasons for this trend involve both performance and Return On Investment (ROI).

1,000x Faster. The move from disk-based to memory-based storage is a key factor in improving performance. However, simply moving to memory is not sufficient to guarantee the extremely high memory processing speeds needed when many people are connecting to place transactions or when

risk management strategies require fast, data-intensive analysis. This level of speed requires cluster computing, with multiple machines performing analyses at the same time, and parallel distribution of data. These capabilities, which also guarantee the high availability levels required in financial transactions, are fully provided in the GridGain in-memory computing platform. Clients who have implemented GridGain have found that they can process their transactions at least 1,000 times faster.

10x ROI Improvement. The cost of memory has dropped roughly 30% per year since the 1960s, so memory has become much more affordable in recent years. While it may still be slightly more expensive than disk, the performance is so much better that it improves ROI significantly. Clients who have implemented the GridGain in-memory computing platform have seen a tenfold or more improvement in their ROI.

With these substantial improvements in speed and ROI, it is not surprising that many financial institutions are turning to the GridGain in-memory computing platform for big-data applications such as fraud detection and prevention.

FINANCIAL INSTITUTIONS USING IN-MEMORY COMPUTING

Financial institutions use GridGain for a variety of fraud detection use cases involving high-volume transaction processing and big-data analytics, such as checking for compliance with anti-money-laundering (AML) and "know your customer" (KYC) regulations, looking for market manipulation, or monitoring other regulated areas. They are using complex event processing for real-time or near real-time customer views and analysis of positions, so they require ultra-low latency in real-time or near real-time data processing and analytics.

Banks who have implemented the GridGain in-memory computing platform – including Barclays, Citi, Sberbank, and others – are seeing a measurable difference at the transactional level. They no longer need to export the data to another system for analysis and approval (or disapproval) of the transaction. That model too often involved performance degradation and post-transaction processing delays, with clients unable to complete transactions until all required steps were performed. In contrast, because the GridGain in-memory computing platform does most of the required computing in a distributed and inmemory fashion, it can process transactions with no noticeable slowdown to the clients. Because GridGain verifies the integrity of each transaction before allowing it to go through, the result is a safer environment for clients.

Customer Case Study: Sberbank

One of the most noteworthy GridGain Systems financial services customers is Sberbank, the largest bank in Russia and the third largest in Europe. Sberbank was faced with a similar problem to the one currently facing spread betting. They were 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.

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 and provided the best results, allowing the bank to significantly improve performance. With GridGain, the company was able to generate one billion transactions per second in a test environment using only 10 Dell® blades with a combined memory of one terabyte. This system cost about $25,000, which is a significant reduction compared to the days when using in-memory technology cost millions of dollars.

The GridGain in-memory computing platform provided several other important capabilities that Sberbank needed, including machine-learning and analytics, flexible pricing, artificial intelligence, scalability, 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. Sbebank 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.

GRIDGAIN SYSTEMS: A LEADER IN IN-MEMORY COMPUTING

With companies facing tremendous data growth and the need for real-time and near real-time fraud detection, demand for the GridGain in-memory computing platform is growing dramatically. This comprehensive platform contains a complete feature set that surpasses the capabilities of mere in-memory databases, making it well suited to financial use cases involving machine learning, risk analysis, real-time analytics, complex event processing, and other capabilities oriented toward financial fraud prevention.

As a complete in-memory computing platform, GridGain helps users consolidate onto a single high performance and highly scalable big data solution for transactions and analytics, resulting in lowered TCO. Advanced SQL functionality and API-based support for common programming languages enable rapid deployment. In addition, the rapidly decreasing cost of memory boosts ROI for in-memory computing initiatives. Financial services companies can now build less expensive systems that perform thousands of times better. Sberbank, Barclay's, and Citi realized such benefits with the GridGain in-memory computing platform.

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.

MEETING THE CHALLENGES OF REAL-TIME FRAUD PREVENTION

As financial institutions and other companies are inundated with ever-increasing amounts of data to process and analyze for potential fraud, they are looking for high performance and highly scalable ways do so in real-time and near real-time in order to stop fraud before it affects their finances and reputations. Fortunately, in-memory computing solutions can now provide the level of performance and scale these companies need. The GridGain in-memory computing platform offers a scalable, comprehensive, and affordable solution – an elegant and efficient way to stop fraud in its tracks.

GridGain

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 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.