WHITE PAPER

GridGain

Driving High-Frequency Trading and Compliance with In-Memory Computing

A GridGain Systems In-Memory Computing White Paper

With high-frequency, algorithmic, and quantitative trading becoming the norm for today's financial services companies, everyone is looking for a technical edge. Companies are racing to beat each other on latency, performance, and analytical complexity. At the same time, they need to maintain transactional-level compliance and risk-management controls. As a result of these requirements, firms engaging in high-frequency trading face unprecedented technical challenges—and they are looking to in-memory computing for answers.

This paper looks at the current state of high-frequency trading—why it's popular and what types of strategies and technologies are being used—and then explores how in-memory computing can meet the technological challenges and increase profits within this market segment.

WHAT IS HIGH-FREQUENCY TRADING?

High-frequency trading is a method of trading securities on the financial markets that involves high-speed, rules-based strategies, and multiple simultaneous trades—with all of the decisions driven by computerized, quantitative models. Basically, this method involves computer programs analyzing the situation in the market and making decisions on the right time to buy, sell, or perform other financial activities. The idea is to predict the market's movements and take actions that will cause you to benefit when those predictions come true—if they come true, and if your technology is fast enough to take advantage of the situation before other traders do.

High-frequency trading became very popular in the 2008 to 2009 timeframe, a period of market instability, market volatility, and financial crisis. High-frequency trading strategies work especially well in periods with a lot of price volatility. If prices move quickly and companies can be equally quick about getting market information and executing trades, they can make substantial profits—as Goldman Sachs did with their high-frequency trading in 2008-09.

While high-frequency trading encompasses a variety of distinct strategies, which we'll discuss in the next section, all of these strategies involve the following basic steps:

1. Obtaining market information—at the fastest speed possible
2. Processing the information through prediction algorithms—at the fastest speed possible
3. Executing trades based on the information—again, at the fastest speed possible
4. Fine-tuning the prediction algorithms based on how they perform

Having the fastest speed possible for the first three steps is crucial—it's what brings the most wins to traders and makes them more profitable than the rest of the market.

The fourth step, fine-tuning your algorithms based on transaction-cost analysis and back testing, is also extremely important. You need to analyze how well your algorithms performed, based on real-time data from the market, and ask these questions: Can the algorithms be improved? Should something have been done differently?

The first three steps are heavy transactional processes, requiring online transaction processing (OLTP), while the last one is primarily a heavily analytical process, requiring online analytical processing (OLAP). In high-frequency trading, both paradigms, OLTP and OLAP, need to work together for maximum profitability.

The next section takes a closer look at the types of strategies involved in high-speed trading to give a better idea of the type of analytics and transactional speeds that come into play.

Strategies Involved in High-Speed Trading

The general category of high-frequency trading includes a variety of individual strategies, including the following:

- Market making
- Market-taker
- Arbitrage
- Statistical arbitrage
- Pairs trading
- Momentum trading
- Pinging
- News and sentiment trading

Let's take a closer look at each of these strategies.

Market Making

Market-making strategists bet on both sides by placing limits on a sale order at slightly above the current market price, while at the same time placing limits on a buy order at slightly below the current market price.

This strategy is based on analyzing a security's bid-ask spread where the bid is the price a buyer is willing to pay, and the ask is the price a seller is willing to accept. The wider the spread between bid and ask at a given moment, the more profit can be made. In the world of high-frequency trading, computers are used to quickly identify and act on those advantageous spreads.

If you're an investor, the size of that spread is going to determine whether your bid or your ask is going to be the most successful. To determine what your spread should be, you need to analyze a lot of data. You need to analyze how other market makers are pricing the securities on the market in which you operate, as well as on other electronic markets. You also need to analyze the volatility of the instrument for which you are making markets, as well as how its price depends on other markets.

If your spreads are better than those of others, you're going to win most of the execution flow, which will bring most of the profit to you. But you need a great deal of analytical speed and transactional speed in order to perform all the necessary analysis and then act quickly in setting up your spread.

Market-Taker

These days, in most of the liquid markets, traders either pay a fee to the market or take in a rebate from the market based on how they are affecting liquidity (whether they're taking it or adding it). By analyzing the rebates and setting up a strategy accordingly, a firm can find some ways to benefit. For example, for a given security, they might buy on a market that provides a rebate and sell on one that does not charge a fee, creating a difference in prices that benefits them. We refer to this approach as the market-taker strategy.

Arbitrage

Arbitrage basically involves finding inefficiencies within the fair markets and taking advantage of those inefficiencies. This strategy heavily exploits the idea of latency, which describes the time that passes between the electronic sending of a signal and the receipt of the signal. For example, when you send an order out to the exchange, some latency occurs before the exchange receives the order and then fulfills the order.

Being fast in responding to market inefficiencies and placing your order can sometimes bring huge profits. However, with more and more players watching the market for inefficiencies and finding them quickly, there have become fewer arbitrage opportunities. By the time you send your order, it may be too late. As a result, algorithms must now be smarter and more complex in order to extract profitability from market inefficiencies. A successful strategy requires speed and analytics working together.

Statistical Arbitrage

This form of arbitrage is much more complex. It involves looking at short-term related securities and making determinations about what is happening in the short-term versus the long-term. Basically, it means trying to come up with a short-term strategy to extract value from knowing how different securities have traded historically versus how they are trading during this short spectrum of time.

This technique uses advanced mathematical strategies and requires both analytical speed and transactional speed, so a combination of OLAP and OLTP approaches is extremely useful. You need to observe the market quickly and make decisions quickly.

Pairs Trading

This trading strategy, which is considered to be a type of statistical arbitrage, involves pairing a long position stock with a short-position stock within the same economic sector or industry. For example, two airlines, such as Delta and United, would share the same cost structure and have similar expanses—oil, airport fees, and so on.

Suppose that, historically, the stock of these two airlines has traded more or less similarly, responding in much the same ways to factors such as oil prices, weather, and general economic prosperity. If you were to spot a difference between the normal stock trajectories of these two companies, you might project that this difference will be temporary—and act to take advantage of it. You could buy more stock in the airline that is currently performing less well, with the expectation that its value will move closer to that of the other airline. At the same time, you could sell the higher-performing airline's stock short, expecting it to decrease.

Sometimes pairs trading is also used to pair a derivative with an underlying asset, because derivatives are traded on one exchange and underlying assets are traded on a different exchange—for example, an options exchange versus an equities exchange or a futures exchange. You may see, historically, that certain derivatives and certain underlying assets typically move together: when the price of one goes up, the price of another goes down, according to a particular ratio. If you're seeing a discrepancy from that ratio in different markets, you can take advantage of that imbalance.

Momentum Trading

This is a strategy used to stimulate the market by placing a number of trades on orders for a certain security in a particular direction. Usually, those orders are canceled in less than a second, and this action triggers competing traders' comput-

ers to react by buying or selling. This reaction then creates an artificial price to change the market, and the momentum trader can act against this change.

In certain markets, momentum trading is considered a questionable practice. Critics say that momentum traders should pay for their activities because erroneous information is being sent to the market. However, this strategy has survived the criticism and continues to be used—most successfully in situations where the infrastructure supports speedy information access and fast transaction processing.

Pinging

This strategy is also quite controversial. It involves placing small orders, such as one lot (100 shares), in order to expose large hidden orders that exist on the exchange. Exchanges have traditionally hidden certain types of orders, such as discretionary orders and iceberg orders, which are large single orders divided into smaller lots in order to hide the full quantity. Pinging enables investors to find out if there are any hidden orders sitting on a particular security. By sending some small orders, they can see what happens —what they can detect in the way of liquidity—and how prices are affected.

This strategy requires heavy analytics because you need to make very quick decisions based on what happens after you place the small orders. You have to act very quickly to take advantage of what you've learned before other investors do so.

News and Sentiment Trading

This is a very recent phenomenon that has become quite popular on Wall Street nowadays. The strategy involves using historical knowledge to anticipate market reactions to news coming out about companies, sectors, countries, industries, or economic indicators. Investors employing this strategy will carefully monitor industries or companies they know to be susceptible to certain types of news or events.

There are feeds available from large market-data providers such as Bloomberg and Reuters that help investors interpret news from the market in ways that can digitally feed into algorithms. The algorithms analyze the news to determine decisions about what to buy or sell immediately based on the sentiment that this news is likely to cause in the market.

This strategy requires significant processing power and speed. You have to understand the sentiment of the news story, analyze how the prices of securities have changed historically based on that type of sentiment, and then make decisions and act upon them very quickly. You want to be the first to act on the news, before others act and change the pricing. Acting that quickly requires very heavy analytical capabilities, OLAP, as well as extremely fast transactional processing power, OLTP.

TECHNOLOGIES USED IN HIGH-SPEED TRADING

To implement the types of strategies we've been discussing for high-frequency trading, you have to have very strong, very fast infrastructure. Currently, the technologies that firms use to create this infrastructure include the following:

- **Dark Fiber Cabling**. To be very fast, you want to connect your decision engine directly to the backbone of the exchange with fiber-optic cable. That way, as soon as you detect the data, you can make a decision based on it and send your order very quickly to that exchange matching engine. The term dark fiber refers to a privately operated optical-fiber infrastructure that you can lease, rather than having to build your own.
- **Exchange Co-Location**. The closer you can get your trading computers to the exchange's servers, the better your high-frequency trading strategies will work. The ideal situation is to put your decision-making logic right on the same server or data center where your exchange is hosting its matching engine. That way, you will be closest to the liquidity, closest to the source of information, and closest to the processing engine of the orders that you will be sending out.
- **Hardware-Based Programming Logic (FPGAs and GPUs)**. Field programmable gate arrays (FPGAs) and graphics processing units (GPUs) are technologies that involve pre-programming your decision logic onto hardware boards, so that decisions are made using hardware instead of software. While this technology is somewhat more expensive than using software, it can also provide responses that are hundreds of times faster. The faster performance comes from doing most of the processing on the hardware. High-frequency traders have been among the first sectors in the financial services industry to adopt this hardware-based processing technology.
- **Parallel Processing Clusters**. Traders are also using parallel processing clusters to expedite high-frequency trading, especially when dealing with heavy analytics. You need to expedite the processing and analysis of the data, and a single server usually won't perform adequately for these types of analytics. To handle these situations, companies and firms are using clusters made up of multiple nodes, so that whatever they're processing can be run on multiple nodes at the same time to expedite the results.

- **In-Memory Computing**. This is one of the newest and most successful technologies being used for high-frequency trading, as well as other Big Data applications. It involves keeping data in memory, instead of on disk, to provide massive improvements in performance and extreme scalability to handle massive sets of data.

Full-featured, in-memory computing technologies, such as the GridGain in-memory computing platform, combine several of the technologies discussed in this section—in-memory computing, data acceleration using Hadoop and MapReduce, complex event processing of multiple data streams, and parallel processing clusters that harness the power of large numbers of computers in a grid.

The next section takes a closer look at how in-memory computing has evolved into a technology that works extremely well for high-speed Big Data use cases such as high-frequency trading.

IN-MEMORY COMPUTING: BENEFITS AND USE CASES

As noted earlier, blockchains can only grow, never shrink, since their integrity depends on no block being deleted. Large-scale blockchain subscribers must be prepared to quickly process, store, and analyze fast-growing amounts of transaction information arriving in real time. 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.

For applications that require heavy analytics and 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).

1000x Faster

The move from disk to memory is a key factor in improving performance. However, simply moving to memory is not sufficient to guarantee the extremely high memory-processing speeds needed at the enterprise level. Enterprise-level speed requires cluster computing, with multiple machines performing analyses at the same time, and parallel distribution of data. These capabilities are important for providing high availability, disaster recovery, and concurrency across systems—and they are all provided in the GridGain in-memory computing platform. Clients who have implemented the GridGain in-memory computing platform have found that they can process transactions about 1000 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.

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 companies who are transitioning from legacy systems and methods to real-time, synchronized transaction processing. 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.

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.

GRIDGAIN SYSTEMS: A LEADER IN IN-MEMORY COMPUTING

Demand for the GridGain in-memory computing solution is growing dramatically, for reasons such as:

Attractive Economics

After dropping 30% per year since the 1960s, the cost of memory has finally reached a price point that makes memory-based solutions affordable.

Proven Performance

As noted earlier, Sberbank became our client after testing GridGain against competing products and finding that GridGain could generate 1 billion transactions per second on a $25,000 system.

Similarly, we have found that most companies who give GridGain a chance—who do a proof of concept with us, or look at our technology—choose our technology. We deliver high throughput, low latency, load balancing, in-memory indexing, and other performance improvements that are crucial to success in running high-frequency trading or algorithmic trading strategies.

Scalability

As a company's data grows, its systems need to scale up to analyze the data quickly and provide better predictions and analytics. GridGain excels in terms of scalability, allowing you to add cluster nodes and memory in a very economical fashion in real-time with automatic data rebalancing.

High Availability

At GridGain, we pay a great deal of attention to high availability. We understand that when we're dealing with financial institutions, it's essential to provide such features as data-center replication, automatic failover, fault tolerance, and quick recovery on an enterprise-level scale.

Open Source Framework

More and more companies are embracing the vendor-agnostic flexibility of open source solutions. GridGain is built on Apache Ignite, one of the top 5 Apache Softeare Foundation (ASF) projects for the last two years. GridGain donated the original code to the project and continues to be the leading contributor. GridGain Systems provides global support and services for Ignite. GridGain adds integration, security, deployment, management and monitoring capabilities that help take business-critical systems into production.

A Comprehensive Platform

GridGain provides not just an in-memory data grid but also a much more comprehensive in-memory computing platform. Among the use cases are: high-performance computing (required for high-frequency trading), machine learning, risk analysis, real-time analytics, complex event processing, Hadoop acceleration, and other capabilities oriented toward Big Data. The product features allow us to provide faster speed without sacrificing high availability.

With all of these features, it's not surprising that GridGain has received positive attention from a number of analysts. Gartner named GridGain a 2014 "Cool Vendor" for in-memory computing. Forrester and many other industry analysts are also looking into us and say that GridGain is one of the leading companies in in-memory computing today.

HOW COMPANIES USE GRIDGAIN

GridGain is typically used to:

- Add speed and scalability to existing applications
- Build new, modern, highly performant and scalable transactional and/or analytical applications
- Build streaming analytics applications, often with Apache Spark, Apache Kafka™ and other streaming technologies
- Add continuous machine and deep learning to applications to improve decision automation

Companies start in any one of these areas. Over time, as GridGain is used with more projects, it becomes a common in-memory data access layer that can support the data, performance and scalability needs for any new workload:

- Data services and other APIs that help deliver an omnichannel digital business
- New customer-facing applications that support new products and services
- Real-time analytics that help improve operational visibility or compliance
- Streaming analytics, machine and deep learning that help improve the customer experience and business outcomes

GridGain helps deliver these types of projects faster while giving companies a foundation for a more real-time, responsive digital business model and the ability to be more flexible to change.

GridGain

IN-MEMORY DATA GRID (IMDG) FOR ADDING SPEED AND SCALABILITY TO APPLICATIONS

A core GridGain capability and most common use case is as an IMDG. GridGain can increase the performance and scalability of existing applications and databases by sliding in-between the application and data layer with no rip-and-replace of the database or application and no major architectural changes. GridGain supports all common RDBMSs including IBM DB2, Microsoft SQL Server, MySQL, Oracle and PostgreSQL, NoSQL databases such as Cassandra and MongoDB, and Hadoop.

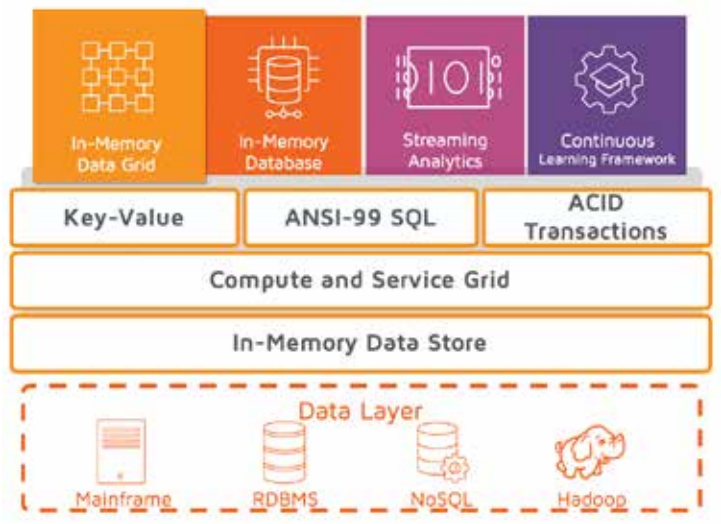

Figure 1. GridGain as an In-Memory Data Grid (IMDG)

GridGain generates the application domain model based on the schema definition of the underlying database. It then loads the data and acts as the data platform for the application. GridGain handles all reads and coordinates transactions with the underlying database in a way that ensures data consistency in both the database and GridGain. By utilizing RAM in place of disk, GridGain lowers latency by orders of magnitude compared to traditional disk-based databases.

The primary benefits and capabilities of the Ignite IMDG include:

- ANSI SQL-99 support including DML and DDL
- ACID transaction support
- In-memory performance orders of magnitude faster than disk-based RDBMSs
- Distributed in-memory caching that offloads queries from the existing database
- Elastic scalability to handle up to petabytes of in-memory data
- Distributed in-memory queue and other data structures
- Web session clustering
- Hibernate L2 cache integration
- Tiered off-heap storage
- Deadlock-free transactions for fast in-memory transaction processing
- JCache (JSR 107), Memcached and Redis client APIs that simplify migration from existing caches

HYBRID IN-MEMORY DATABASE (IMDB) FOR HIGH VOLUME, LOW LATENCY TRANSACTIONS AND DATA INGESTION

A GridGain cluster can also be used as a distributed, transactional IMDB to support high volume, low latency transactions and data ingestion, or for low cost storage.

The GridGain IMDB combines distributed, horizontally scalable ANSI-99 SQL and ACID transactions with the GridGain Persistent Store. It supports all SQL, DDL and DML commands including SELECT, UPDATE, INSERT, MERGE and DELETE queries and CREATE and DROP table. GridGain parallelizes commands whenever possible, such as distributed SQL joins. It allows for cross-cache joins across the entire cluster, which includes joins between data persisted in third party databases and the GridGain Persistent Store. It also allows companies to put 0-100% of data in RAM for the best combination of performance and cost.

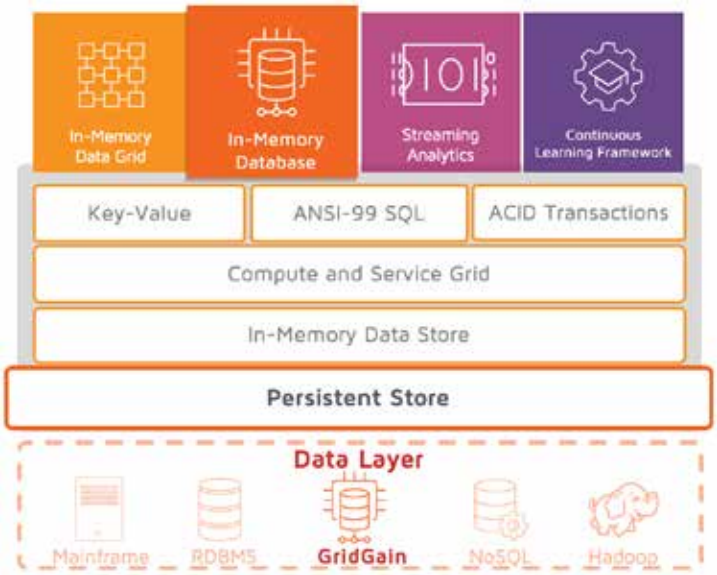

Figure 2. GridGain as an IMDB

The in-memory distributed SQL capabilities allow developers, administrators and analysts to interact with the GridGain platform using standard SQL commands through JDBC or ODBC or natively developed APIs across other languages as well.

The primary capabilities of the GridGain's hybrid IMDB include:

- ANSI SQL-99 compliance
- ACID transactions support
- Full support for SQL DML including SELECT, UPDATE, INSERT, MERGE and DELETE
- Support for DDL commands including CREATE and DROP table
- Support for distributed SQL joins, including cross-cache joins across the entire cluster
- SQL support through JDBC and ODBC without custom coding
- Geospatial support
- Hybrid memory support for RAM, HDD, SSD/Flash, 3D XPoint and other storage technologies
- Support for maintaining 0-100% of data in RAM with the full data set in non-volatile storage
- Immediate availability on restart without having to wait for RAM warmup

STREAM INGESTION, DATA MANAGEMENT, PROCESSING AND REAL-TIME ANALYTICS FOR STREAMING ANALYTICS

GridGain is used by the largest companies in the world to ingest, process, store and publish streaming data for large-scale, mission critical business applications. It is used by several of the largest banks in the world for trade processing, settlement and compliance; by telecommunications companies to deliver call services over telephone networks and the Internet; by retailers and e-commerce vendors to deliver an improved real-time experience; and by leading cloud infrastructure and SaaS vendors as the in-memory computing foundation of their offerings. Companies have been able to ingest and process streams with millions of events per second on a moderately-sized cluster.

Figure 3. GridGain for Stream Ingestion, Processing, and Analytics

GridGain is integrated and used with major streaming technologies including Apache Camel™, Kafka, Spark and Storm™, Java Message Service (JMS) and MQTT to ingest, process and publish streaming data. Once loaded into the cluster, companies can leverage GridGain's built-in MPP-style libraries for concurrent data processing, including concurrent SQL queries and continuous learning. Clients can then subscribe to continuous queries which execute and identify important events as streams are processed.

GridGain also provides the broadest in-memory computing integration with Apache Spark. The integration includes native support for Spark DataFrames, a GridGain RDD API for reading in and writing data to GridGain as mutable Spark RDDs, optimized SQL, and an in-memory implementation of HDFS with the GridGain File System (GGFS). When deployed together, Spark can access all of the in-memory data in GridGain, not just data streams; share data and state across all Spark jobs; and take advantage of all of GridGain's in-memory loading and processing capabilities including continuous learning to train models in near real-time to improve outcomes for in-process HTAP applications.

CONTINUOUS LEARNING FRAMEWORK FOR MACHINE LEARNING AND DEEP LEARNING

Ignite also provides Ignite Machine Learning (ML), MPP-style machine learning and deep learning with real-time performance on petabytes of data. Ignite provides several standard machine learning algorithms optimized for collocated processing including linear and multi-linear regression, k-means clustering, decision trees, k-NN classification and regression. It also includes a multilayer perceptron for deep learning along with TensorFlow integration. Developers can develop and deploy their own algorithms across any cluster as well as using the compute grid.

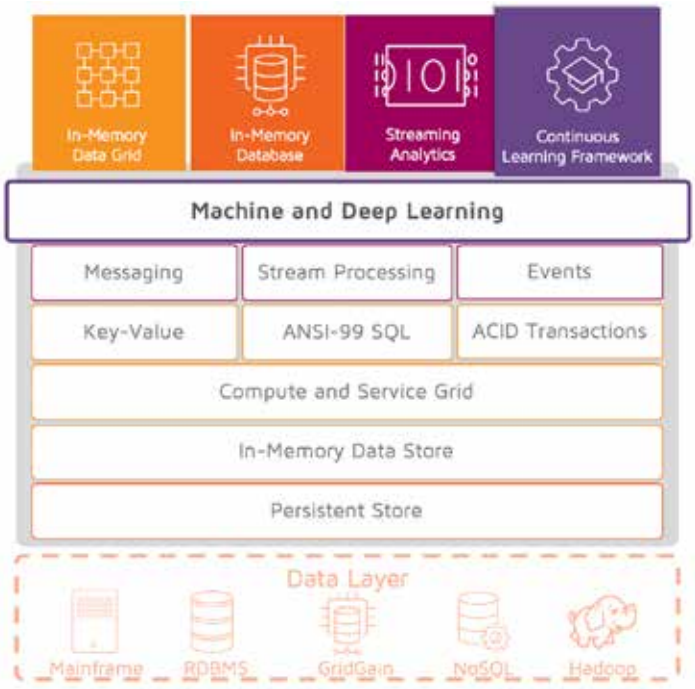

Figure 4. GridGain for Machine Learning and Deep Learning

ENTERPRISE-LEVEL FEATURES

Some features that enterprise-class companies seek are not available in the open source version. GridGain builds on the capabilities of Apache Ignite with features such as:

- **Management and Monitoring**. GridGain Web Console, the GUI-based Management and Monitoring tool, provides a unified operations, management and monitoring system for GridGain deployments. GridGain Web Console provides management and monitoring views into all aspects of GridGain operations. This includes HPC, Data Grid, Streaming, and Hadoop acceleration via standard dashboards, advanced charting of performance metrics, and grid health (telemetry) views, among many other features.
- **Data Center Replication**. GridGain reliably replicates data on a per-cache basis across two or more regions connected by wide area networks. This allows geographically remote data centers to maintain consistent views of data. With GridGain reliability and predictability, Data Center Replication ensures business continuity and can be used as part of a disaster recovery plan. Data Center Replication integrates with your application so that caches marked for replication are automatically synchronized across the WAN link. Predictable performance includes:
  - Asynchronous data replication to maximize performance
  - Fault tolerance
  - Conflict resolution when needed
  - Scaling (at the local cluster, and in the number of WAN-linked clusters)
  - Multiple types of topologies supported
  - Brokerless (no JMS broker needed for replication)
- **Oracle GoldenGate Integration**. The Oracle GoldenGate integration in the GridGain Enterprise and Ultimate Editions provides real-time data integration and replication into a GridGain cluster from different environments. When configured, the GridGain in-memory computing platform will automatically receive updates from the connected source database, converting the data from a database relational model to cache objects.
- **Centralized Backup and Recovery Management**. The GridGain Ultimate Edition provides centralized backup and recovery using either the GridGain Web Console or Snapshot Command Line Tool. You can perform, schedule and manage backups, and then recover to any point in time on any cluster using a combination of full and incremental snapshots with continuous archiving. This includes the ability to backup remotely using network backups, and then (re)deploy to a different cluster of any size anywhere on premise or in the cloud. You can also deploy backups to support testing in development, quality assurance (QA) and staging environments.
- **Full, Incremental and Continuous Backups**. Within the GridGain Web Console or Snapshot Command Line Tool you can centrally perform or schedule full and incremental snapshots across a distributed cluster. You can then use them as backup and restore points for later recovery. You can also use continuous archives of write-ahead log (WAL) files to backup down to each transaction. The combination of full and incremental snapshots with continuous archiving helps ensure data is never lost. The GridGain Ultimate Edi-

tion minimizes the administrative and hardware costs of backup and recovery by providing the ability to schedule full and incremental snapshot creation and deletion, as well as the movement of snapshots to another (remote) lower cost location.

- **Network Backups**. With the GridGain Ultimate Edition, snapshots do not need to be stored locally on the same cluster machines used to handle the operational load. They can also be managed and stored remotely on premise or in the cloud. When combined with the remote storage of continuous archives, this capability helps ensure a cluster can be quickly recovered even if an entire data center disappears.
- **Point-in-Time Recovery**. You can quickly restore a GridGain Ultimate Edition cluster to any point in time through the combination of full and incremental backups with continuous archiving. Point-in-time recovery can be used to restore a system up to any change without having to manually resubmit or replay existing transactions that occurred following a full or incremental snapshot. Continuous archiving helps spread network loads to minimize peak network traffic. It also allows recovery to be more granular and up-to-date, which helps reduce overall downtime needed to restore a cluster to a current, valid state.
- **Heterogeneous Recovery**. The GridGain Ultimate Edition also allows you to restore an existing cluster to another location, on premise or in the cloud, with a different size and topology. GridGain already allows you to dynamically add nodes to a cluster for scalability, and create a hybrid cluster across any collection of nodes or data centers on premise or in the cloud. Heterogeneous Recovery enables you to rapidly bring up a different size cluster the moment an existing cluster goes down or bring a new cluster up so that you can take an existing cluster down. This helps reduce downtime and increase availability.

MEETING THE CHALLENGES OF HIGH-FREQUENCY TRADING

As competition intensifies in the field of high-frequency trading, financial services firms need a new level of transactional speed and analytic power to beat their competition. In-memory computing can provide that speed and power. And GridGain in-memory computing platform offers a scalable, comprehensive, and affordable solution—an elegant and efficient way to give traders the high-performance edge they need.

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 the form at www.gridgain.com/contact to have us 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.