



Powering Financial Spread Betting with In-Memory Computing

A GridGain Systems In-Memory Computing White Paper



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 white paper 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 smart phones 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.
- **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.

The next sections 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 benefit from heavy analytics and real-time processing of big data, 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.

It is not surprising that many financial institutions are turning to the GridGain in-memory computing platform for big data applications such as spread betting because of these substantial improvements in speed and ROI.

FINANCIAL APPLICATIONS OF IN-MEMORY GRID COMPUTING

Many customers are using the GridGain in-memory computing platform for financial applications that have similar characteristics to spread betting. This technology powers the actual trading platforms where large numbers of transactions needed to be processed in real time. At the same time, GridGain also provides a highly capable framework for mathematical and statistical modeling, for limiting risk, and for collecting information from multiple markets. So, customers are also using our technology for financial analytics, risk management, big data analytics and the types of compliance and monitoring that have recently been introduced into spread betting.

GRIDGAIN SYSTEMS: A LEADER IN IN-MEMORY COMPUTING

With companies grappling with the challenges resulting from increasing volumes of data needing real-time processing, demand for the GridGain in-memory computing platform is growing dramatically. This comprehensive platform contains a complete feature set that surpasses the capabilities of

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. 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.

in-memory database point solutions, making GridGain well suited to financial use cases involving mathematical modeling, machine learning, risk management, complex event processing, and other capabilities that can provide a crucial edge in spread betting.

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 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.

GAINING THE EDGE IN SPREAD BETTING

With spread betting emerging as a rapidly growing trend, participants looking for an edge are turning toward technology that can support 24/7 transactions, real-time responses, and the intensive analytics needed to hone predictive models and risk-management strategies. Fortunately, in-memory computing solutions can provide the level of performance and scale needed by brokerages, asset managers, online gambling firms, and other players seeking to lead in this market. The GridGain in-memory computing platform offers a scalable, comprehensive, and affordable solution—a sure bet for those who want to beat the spread.

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 visit www.gridgain.com/contact.

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, Finastrå, 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.