

A photograph of a person in a dark suit and striped tie, holding a purple umbrella. In front of them, on a dark surface, are white paper cutouts of a family (a car, two adults, and two children) and a house. The scene is dimly lit, with the umbrella providing a focal point of light.

Empowering Insurance Companies with In-Memory Computing

A GridGain Systems In-Memory Computing White Paper

November 2017

Contents

Challenges Facing Insurance Providers	2
Gaining an Edge Through New Technology	3
The Crucial Element: State-of-the Art Data Technology.....	5
Why Now is the Time for In-Memory Computing.....	6
GridGain Systems: A Leader in In-Memory Computing	8
In-Memory Computing: The Best Policy for Success	9
Contact GridGain Systems	10
About GridGain Systems	10

In a time when innovative technologies are upending many industries, the insurance industry faces great opportunities – and significant challenges. Growth continues to be strong: in 2014, the insurance industry was responsible for about 2.6% of U.S. GDP, and net premiums in the U.S. are continuing to rise beyond the 2015 total of \$1.2 trillion. Online and mobile apps are enabling insurance companies to reach a wider range of potential customers, and capabilities such as smart contracts, Robotic Process Automation (RPA), machine learning, and telematics (tracking risk parameters by means of data-transmitting IoT devices) are revolutionizing internal processes and risk assessment.

At the same time, providers are facing increased competition as well as increased customer expectations for immediate response on all of their digital channels. The current climate of economic, political, regulatory, and climate-related uncertainty steps up the pressure to operate efficiently and strategically.

As insurance providers seek to stand out in an increasingly crowded field, technology is a crucial part of maintaining a competitive edge. With consumers expecting instant turnaround on quotes, underwriting, and claims processing, successful providers must process transactions in real-time. They must also be able to scale their capacity quickly to meet usage growth. Plus, they need robust streaming capabilities and sophisticated analytics to help them quickly monitor and analyze the many data streams now relevant to their business.

To achieve this level of performance, scalability, and analytical sophistication, many payments providers are turning to in-memory computing solutions. This white paper will discuss the challenges facing today's insurance industry, the opportunities new technologies can offer, and the crucial edge that providers can gain with solutions such as the GridGain in-memory computing platform.

Challenges Facing Insurance Providers

With the insurance industry changing in significant ways, providers face a host of new challenges, including:

- Increased customer expectations for instant, digital interactions
- Heightened competition from both traditional and nontraditional providers
- Cultural changes affecting insurance models, such as the sharing economy
- Vulnerability to hacking, fraud, and other forms of cybercrime in the digital marketplace
- A climate of uncertainty with respect to the regulations, the economy, and weather

Let's look at each of these challenges in more detail.

Increased customer expectations. As new methods of digital interaction become the norm, more and more customers expect fast, easy, mobile communication with their insurance providers. Whether they are looking for quotes, policy underwriting, or claims processing and payments, they want the process

to be available 24/7 from their tablets and smartphones, as well as from their desktop and laptop computers – and they want the interactions to occur in real time. Plus, they want to interact with their insurance company through a range of portals, including social media.

Heightened competition. Legacy insurance providers seeking to satisfy these increased customer expectations must now compete with agile, tech-savvy startups such as Lemonade, Trov, Fabric, and Ladder. Lemonade uses machine learning and bots to deliver home and renters insurance through a quick-response app. Trov provides on-demand insurance for assets such as computers and cameras through a smartphone app, while Fabric and Ladder sell life insurance through a fast, app-based process.

While some of this competition is limited by geography (Trov, for instance, was available only in the U.K. and Australia at the time this was written), these limits are quickly loosening. In general, the availability of insurance online is providing customers with access to a wider range of providers, not just local ones. Plus, comparison-shopping online is getting easier. For example, in the U.K., customers can compare rates and coverage among providers with the help of an A.I.-based insurance-advisory app called Brolly. This ease of comparison puts increased pressure on insurance providers to keep rates competitive, as well as to improve their technology, customer service, and customer retention.

Cultural changes. In addition to coping with increased competition and customer expectations, insurance providers are facing some cultural changes that require new insurance models – particularly with respect to the sharing economy. Homeowners renting out rooms through AirBnB and drivers using their vehicles as Uber or Lyft taxis require special considerations with respect to risk assessment. Many providers are looking at usage-based insurance for such situations. For example, the startup Slice, which sells on-demand, short-term insurance for homeshare hosts, who can turn the coverage on and off through an app as needed.

Cyber risk. With the digital revolution and the migration of many business activities online, businesses are now much more vulnerable to hacking, online fraud, and other forms of cyber-crime. On one hand, this change presents a new opportunity to insurers, with businesses seeking cyber insurance against losses related to this type of online activity. However, insurance providers must at the same time figure out how to protect themselves – and the sensitive data of their customers – against cybercrime.

A climate of uncertainty. As an industry based on predictive modeling, the insurance industry is affected particularly strongly by uncertainty. Currently, some significant areas of uncertainty include: the slowly growing economy, with its low interest rates; unusual global weather patterns, with catastrophic events; and political uncertainty in the U.S., with laws regarding health insurance and financial regulations in flux. Insurance providers are facing additional uncertainty around new or strengthened regulatory requirements relating to investments, consumer protection, fraud prevention, and other issues.

Fortunately, even as these challenges have grown for insurance providers, the insurance technology (or “InsurTech”) to help meet these challenges has also evolved quickly, creating new opportunities and innovative solutions.

Gaining an Edge Through New Technology

Some key ways in which insurance providers are gaining success through new InsurTech include the following:

- Using software robots, artificial intelligence (A.I.), and machine learning to automate and speed up processes
- Using blockchain-based digital ledgers and smart contracts to reduce process steps and increase efficiency
- Using data from sensors, drones, and other smart devices within the Internet of Things (IoT) to provide data-driven coverage that improves risk management and safety
- Implementing high-performance data solutions to enable all of the above, and more

Let's take a closer look at these InsurTech-enabled trends.

Automating processes with robots and A.I. To help address customer expectations of instant, digital interactions, many insurance providers are turning to techniques such as Robotic Process Automation (RPA), artificial intelligence, and machine learning. RPA, which uses software to replicate easily definable process steps that humans perform on computers, is well suited to many rule-based steps within claims processing. This technique becomes even more powerful when combined with A.I. and machine-learning capabilities that can evaluate large amounts of process data and infer subjective decision-making capabilities. The speed at which A.I.-based bots can respond to humans makes possible the type of results advertised by Lemonade, the app-based home and renter insurance provider mentioned earlier, which promises “90 seconds to get insured” and “3 minutes to get paid.”

Increasing efficiency with blockchains. The blockchain technology behind the digital currency Bitcoin is being adapted to many applications that can benefit from the efficiency of a decentralized, distributed, peer-to-peer public ledger – including in the insurance industry. With blockchains, many process steps can happen simultaneously, with details centrally posted and verified, rather than going through a series of “middlemen.” (For an in-depth discussion of how blockchains streamline processing even for non-bitcoin transactions, see [“Enable Bitcoin and Blockchain Technology with In-Memory Computing.”](#)) This technology is being used to create peer-to-peer *smart contracts* – capable of automating many administrative processes – in the Ethereum blockchain.

One example of how this technology is benefitting the insurance industry is the decentralized crop-insurance application developed by Etherisc. This application uses a smart contract to check for weather-related crop damage and promises to automatically pay claims to insured parties.

Improving risk management and safety with IoT data. The growth of the Internet of Things has revolutionized the risk-management aspect of the insurance industry, enabling companies to measure risk-related parameters in cars, people, and homes in real time using sensor-enabled, data-transmitting devices. The term *telematics*, which combines “telecommunications” and “informatics,” is sometimes

used to describe this phenomenon, as is *usage-based insurance* (UBI), since the parameters being measured can paint a picture of how an insured asset is being used.

Auto insurance is one of the first areas in which this practice has been embraced. In fact, the National Association of Insurance Commissioners has reported that approximately 70% of all auto insurance carriers are expected to use telematics UBI by 2020. That is, auto insurance companies will use sensors in cars to monitor risk-related data, such as miles driven, times of day driven, locations visited, airbag deployment, and acceleration and braking patterns. Collecting such data allows insurers to improve their risk assessment and better calculate premiums for individual customers. This process can also improve safety and reduce claims, as drivers adapt to lower-risk driving practices to reduce their premiums.

As wearable IoT devices and “smart homes” become more common, telematics is also starting to make an impact in life and health (L & H) insurance, as well as property and casualty (P & C). John Hancock Life Insurance has a program that involves giving Fitbit bands to policy holders – and encouraging wellness – in exchange for discounted life-insurance coverage. In the area of smart home devices, insurers are looking at how to use such devices to detect – and even prevent or mitigate – damage and loss from occurrences such as fire, theft, or floods. For example, a smart thermostat could detect a heater failure during cold weather and notify absent homeowners and insurers to prevent frozen pipes and eventual flooding. Anticipating significant benefits from data-transmitting IoT devices, 60% of insurers in a 2016 survey cited in ia magazine reported they had already launched or piloted programs for incorporating connected devices.

Implementing high-performance data solutions. As insurance providers strive to provide real-time, 24/7 services and to quickly process and analyze an avalanche of new data streams, many are finding that their legacy data systems are overwhelmed. To deal with the “performance gap” between their legacy systems and the demands of today’s customers and processes, providers are looking toward high-performance, state-of-the-art data technology, as described in the next section.

The Crucial Element: State-of-the Art Data Technology

To implement the new technologies needed to succeed in today’s insurance market, one ingredient above all is crucial: a technology that can process and analyze enormous amounts of data in real-time.

All of the new technologies being embraced by insurers rely on fast processing of massive amounts of data. Robotic interactions that enable online quotes, policies, and claim-handling must deliver the real-time response that customers expect – even while analyzing large amounts of risk-related and process-related data. The data volume can be especially high when machine learning is involved, since this technique requires learning by example, inferring patterns in large amounts of historical data.

Digital ledgers and smart contracts, which can streamline insurance processing, also involve very large amounts of data, since the underlying blockchain technology is peer-to-peer. Subscribers to the blockchain must continually process arriving data blocks, since there is no central repository.

Finally, the IoT data necessary for telematics and usage-based insurance is likely to increase dramatically as smart cars, wearables, and smart home devices become more and more mainstream. Insurers will need to process tremendous amounts of IoT data in real time in order to use the data effectively – especially where risks and safety threats are involved.

When you combine these data needs with other data-processing and analysis needs that are part of the insurance business, such as those related to regulatory compliance, it's clear that the underlying data processing technology must be state-of-the-art.

In other words, it should have the following characteristics:

- **High-performance and low latency:** Speed is essential for providing the instant insurance transactions that today's customers expect, as well as for providing real-time risk analysis and processing high volumes of streaming data from IoT devices, as discussed above
- **Scalable:** The ability to easily scale a system up is essential for accommodating fast growth and maintaining high performance even as the volume of customers, policies, and claims increases
- **Flexible and interoperable:** Flexibility (e.g., being open source and cloud-capable) and interoperability are important for offering insurance services from a range of devices, coordinating services such as payments via hubs, and enabling access to diverse data sources
- **Fault-tolerant and highly available:** Data supporting payment of premiums and claims must be maintained with the high level of reliability, availability, and recoverability required for all financial data
- **Secure:** Fully implemented security features are vital for keeping private customer data safe, particularly as IoT devices increase the amount of data collected per customer
- **Capable of sophisticated big data analytics:** The ability to effectively use predictive modeling, complex event processing (CEP), machine learning, and statistical and A.I. techniques is extremely important for supporting real-time risk analysis, fraud prevention, and data-driven business intelligence

Fortunately, there is a data processing technology that is available with these features and is well suited to this type of high speed, big data use case: in-memory computing, as implemented in the GridGain in-memory computing platform.

Why Now is the Time for In-Memory Computing

In-memory computing has become the technology of choice for applications such as insurance that require high-speed data processing and analysis. Because in-memory computing involves keeping data in RAM for extremely fast access, with no slowdowns resulting from moving data between RAM and disk, it is faster than systems built on disk-based databases.

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

1000x faster. The move from disk-based databases to systems which hold all data in RAM is a key factor in improving performance. However, simply moving to memory is not sufficient to guarantee extremely high processing speeds. Maximum speed requires distributed computing, with multiple machines performing analytics in parallel across a distributed data set. 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 GridGain have found that they can process transactions up to 1,000 times faster.

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. As a result, it is now feasible to build systems based on hybrid transactional/analytical processing (HTAP) architectures. HTAP can eliminate the need for a separate OLAP system and allow companies to standardize on a single architecture. This can result in dramatic decreases in hardware, software, development and maintenance costs, which reduces TOC and improves ROI.

Risk-analysis case study: Finastra Financial Software. One example of a GridGain client excelling in an area relevant to insurance is Finastra, maker of a financial software product focusing on risk management and risk analysis. This software is used by 48 of the world's 50 largest banks. While Finastra is oriented toward banking, its situation is similar to what insurance providers face: large amounts of data arriving and requiring analysis in real time – analysis that is used, in this case, to create reports for customers.

With GridGain, Finastra was able to build a cluster and serve many customers in a multi-tenant setup that integrates trading systems with cloud-based components. Customers can use this common platform to build their models and analyze their positions across a variety of instruments in real time, thanks to GridGain technology. Felix Grevy, the director of product management for FusionFabric.cloud at Finastra, described the GridGain advantage this way: "With GridGain, we have achieved real-time processing of massive amounts of trade and transaction data, eliminating bottlenecks and enabling us to offer next-generation financial services to our customers."

Financial-services case study: Sberbank. Another noteworthy GridGain customer is Sberbank, the largest bank in Russia and a top 25 bank in Europe. Sberbank was faced with a similar problem to the one currently facing insurance companies, transitioning from person-to-person branch office interactions to 24/7 digital, online environments. The bank was expanding from strictly a traditional, brick-and-mortar approach – 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. Their new blended approach included in-branch transactions as well as 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 solutions from more than ten potential 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 their next-generation platform with the performance and scalability necessary for a 100x increase in transactions per second.

The GridGain in-memory computing platform provided several other important capabilities that Sberbank's next-generation platform would require such as real-time analytics, flexible pricing, ease of deployment, hardware independence of cluster components, and ACID compliant transactions. 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 the GridGain 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 grappling with the challenges resulting from increasing fast data and big data workloads, demand for the GridGain in-memory computing platform is growing dramatically. This comprehensive platform contains a complete feature set that surpasses the capabilities of in-memory database or in-memory data grid point solutions, and it is well suited for OLAP, OLTP and HTAP use cases.

GridGain allows users to consolidate on a single high performance and highly scalable solution for transactions and analytics, resulting in a lower TCO versus running separate transactional and analytical architectures. 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 businesses to build less expensive systems that perform thousands of times better.

Users enjoy the following:

- **A unified high-performance architecture.** The GridGain in-memory computing platform consists of multiple features connected by a clustered, in-memory file system. The In-memory Data Grid, In-Memory Database, Streaming Analytics, and Service Grid are interconnected. Computations occur as close as possible to the data. 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 platform is massively scalable, allowing companies to add nodes and memory to the cluster in real-time with automatic data rebalancing. Because it is hardware agnostic, users can choose their own preferred hardware for scaling out or up.
- **Full SQL support.** GridGain is ANSI SQL-99 compliant and supports DML and DDL so users can leverage their existing SQL code using the GridGain JDBC and ODBC APIs. Users with existing code bases that are not based on SQL can leverage their existing code through supported APIs for Java, .NET, C++, and more.
- **A Distributed ACID and ANSI-99 SQL-Compliant Disk Store.** The optional GridGain [Persistent Store](#) is a distributed ACID and ANSI-99 SQL-compliant disk store which addresses the issues of memory overflow and fast restart for the GridGain in-memory data grid. The Persistent Store may incorporate spinning disks, solid state drives (SSDs), Flash, 3D XPoint or other similar storage technologies. If used, the optional Persistent Store keeps the superset of data and all the SQL indexes on disk, which allows GridGain to be fully operational from disk. The combination of this new feature and the platform's advanced SQL capabilities allows GridGain to serve as a distributed transactional SQL database, spanning both memory and disk, while continuing to support all of the traditional in-memory only use cases. Persistent Store allows organizations to maximize their return on investment by defining their own optimal tradeoff between infrastructure costs and application performance by adjusting the amount of data that is kept in-memory.
- **High availability.** GridGain provides essential high-availability features such as data-center replication, automatic failover, fault tolerance, and quick recovery on an enterprise scale.
- **Transaction processing.** The GridGain platform supports ACID-compliant transactions in a number of user-configurable modes.
- **Security features.** GridGain supports authentication, authorization, multiple encryption levels, tracing, and auditing.
- **Open Source framework.** GridGain is based on Apache® Ignite™, a popular Apache Software Foundation open source project. GridGain Systems created the code which was contributed to the Apache Software Foundation that became Apache Ignite and fully supports the open source project. The GridGain Enterprise Edition extends the features in Apache Ignite to provide enterprise-level capabilities and services, such as enhanced security, data center replication, management and monitoring tools, rolling production upgrades, and network segmentation protection. The GridGain Ultimate Edition includes all the features of the GridGain Enterprise Edition plus a [Cluster Snapshots](#) feature for automated backups of the GridGain Persistent Store.
- **Production Support.** [GridGain Support](#), available to GridGain [Professional](#), [Enterprise](#), and [Ultimate](#) Edition users, includes faster availability of software patches and 24/7 enterprise-grade support.

In-Memory Computing: The Best Policy for Success

The rapid changes in the insurance industry in recent years have brought a host of new challenges to providers, along with new opportunities. To successfully navigate the competitive new terrain, providers need to provide real-time quotes, underwriting, and claims-processing from a wide variety of digital and mobile portals, with the help of techniques such as robotic process automation (RPA) and machine learning. They also need to make strategic use of blockchain-based options such as distributed public

ledgers and smart contracts that can help them process and pay claims efficiently. Plus, as IoT devices become an accepted part of risk management through telematics, providers will need to monitor and analyze increasingly massive amounts of streaming data. These changes demand a new level of transactional speed and analytic power.

Fortunately, in-memory computing solutions can now provide the level of performance these companies need. Combining an open source framework with enterprise-level features, the GridGain in-memory computing platform offers a scalable, comprehensive, secure, and affordable solution – an elegant and efficient way to give insurance providers the high-performance edge they need.

Contact GridGain Systems

To learn more about how the GridGain in-memory computing platform can help your business, please email our sales team at sales@gridgain.com, call us at +1 (650) 241-2281 (US) or +44 (0) 7775 835 770 (Europe), or complete our [contact form](#) and we will contact you.

About GridGain Systems

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

COPYRIGHT AND TRADEMARK INFORMATION

© 2017 GridGain Systems. All rights reserved. This document is provided “as is”. Information and views expressed in this document, including URL and other web site references, may change without notice. This document does not provide you with any legal rights to any intellectual property in any GridGain product. You may copy and use this document for your internal reference purposes. GridGain is a trademark or registered trademark of GridGain Systems, Inc. Windows, .NET and C# are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. JEE and Java are either registered trademarks or trademarks of SUN Microsystems and/or Oracle Corporation in the United States and/or other countries. Apache, Apache Ignite, Ignite, the Apache Ignite logo, Apache Spark, Spark, Apache Cassandra, and Cassandra are either registered trademarks or trademarks of the Apache Software Foundation in the United States and/or other countries. All other trademarks and trade names are the property of their respective owners and used here for identification purposes only.