



Boost Software and SaaS with In-Memory Computing

A GridGain Systems In-Memory Computing White Paper



Every business relies on software and the data it collects and analyzes. As data volumes continue to increase dramatically, and as businesses demand software solutions that are faster, more robust, and easier to implement and maintain—developers are constantly in search of new techniques and tools that can address all their needs.

Software-as-a-service, commonly known as SaaS, typically delivers on the promise of ease of implementation and maintenance. When it is implemented in an in-memory computing environment, the speed and scale concerns can be taken care of as well. Businesses have a long wish list for their software solutions. They want stability, reliability, security, scalability, and speed. They can get there today with serverless architectures that rely heavily on virtualization and containerization, distributed systems, and microservice-based architectures.

Unfortunately, many software and SaaS development projects have been hardware constrained. In particular, memory used to be a gating factor. The high-cost component made achieving optimal application speed, especially in real-time analytics environments, extremely expensive. Now, however, with memory prices lower than ever, it is possible for businesses to use memory to house data and perform super-fast processing and analytics. This gives businesses real-time window into a single source of truth that helps them make better and more timely decisions.

IN-MEMORY COMPUTING: THE RIGHT CHOICE NOW

In-memory technology has been around for decades, ever since developers realized that working with data stored in memory was faster than working with data stored on disk. Today, that speed means everything. Digital transformation has pushed companies closer to customers who demand real-time interactions, and the amounts of Internet traffic, connected devices, and data—all of which require instant analysis—continue to grow. It's estimated that total global data will grow from eight zettabytes in 2015 to 35 or more zettabytes in 2020.

Luckily, the cost of memory continues to tumble—down 30 percent annually over time—so moving to in-memory platforms has gotten more cost-effective with each passing year. According to Gartner, the in-memory technology market will grow to \$10 billion by the end of 2019, representing 22 percent compound annual growth over time.

The combination of these two trends means that the time is right to leverage the improved performance and scale provided by in-memory computing to make software and SaaS implementations run optimally.

APACHE® IGNITE™ TACKLES SAAS ISSUES

The growth of the in-memory computing market includes in-memory databases, in-memory analytics, in-memory data grids, in-memory app servers, in-memory messaging, event stream processing, and analytic servers. The GridGain and Apache Ignite in-memory computing platforms provide complete support for all of these components.

Apache Ignite:

- Supports data caching, massive parallel processing, in-memory SQL, streaming, and much more.
- Slides in between the existing application and data layers easily, with minimal changes to code.
- Offers ACID-compliant transactions as well as analytics support.
- Works with all popular RDBMS, NoSQL, and Hadoop databases and offers a Unified API with support for a wide range of languages.
- Can be deployed on premise, in the cloud, or in hybrid environments.
- Can be configured for high availability with automatic load balancing and failover.
- Remains highly customizable for a multitude of uses as it handles computing and transacting on large-scale data sets in real time.

The GridGain and Apache Ignite in-memory computing platforms run everything in memory to yield the performance and scalability gains that software and SaaS developers need to be successful. The GridGain Enterprise Edition adds features such as better security, better network segmentation protection, data center replication, and rolling updates that allow updating of multiple nodes on a cluster without taking the whole cluster down.

SAAS CAN TAKE ADVANTAGE OF IN-MEMORY BENEFITS

GridGain has numerous customers that are actively engaged in developing and selling software and SaaS solutions. In a 2016 survey, GridGain asked them what they were considering using in-memory computing for. The top three answers, each at or above 60 percent, were database caching, application scaling, and high-speed transactions. Also

on the list: Spark acceleration, real-time streaming analysis, and database scaling. Clients said they most frequently use these features in private clouds, on premise, in AWS®, or in Microsoft Azure®.

Java was the clear leader among the protocols they use to access their data, with SQL and MapReduce® following behind. The data stores they are most likely to use with GridGain or Apache Ignite are MongoDB®, HDFS, and Oracle®.

It is interesting to understand how GridGain's SaaS customers choose to deploy the in-memory computing platform. A powerful starting point is the In-Memory Data Grid, inserted between the application and database layers to cache the disk-based data from RDBMS, NoSQL, or Hadoop databases in RAM, delivering a performance increase up to 1,000 times. This in-memory key value store can replicate and partition data caches across multiple nodes and deliver elastic on-demand scalability to add new nodes. Distributed in-memory transactions can also be ACID-compliant. The data grid offers support for all popular data stores, with read-through, write-through and write-behind. Setup remains completely flexible to address any use case.

The **In-Memory SQL Grid** is the horizontally-scalable, fault-tolerant, ANSI SQL-99 compliant portion of GridGain and Ignite, with support for all SQL and DML commands such as SELECT, UPDATE, INSERT, MERGE, and DELETE queries. Geospatial support is built into the product, and all the communication with the SQL grid is done through ODBC and JDBC APIs without custom coding.

The **In-Memory Compute Grid** enables distributed parallel processing of resource-intensive compute tasks. It offers adaptive load balancing, automatic fault tolerance, linear scalability, and custom scheduling. Built around a pluggable SPI design, it offers a direct API for Fork-Join and MapReduce processing.

The **In-Memory Service Grid** provides control over services deployed on each cluster node and guarantees continuous availability of all deployed services in case of node failures. It can automatically deploy services on node startup, deploy multiple instances of a service, and terminate any deployed service. It is a load-balanced and fault-tolerant way of running and managing services across the grid.

In-memory streaming and continuous event processing establish windows for processing and run either one-time or continuous queries against these windows. The event workflow is totally customizable, and is often used for real-time analytics. Data can be indexed as it is being streamed to make it possible to run extremely fast distributed SQL queries against the streaming data.

In-memory Hadoop acceleration provides easy-to-use extensions to disk-based HDFS and traditional MapReduce, delivering up to 10 times faster performance. GridGain and/or Ignite can be layered on top of an existing disk-based HDFS and used as a caching layer offering read-through and write-through while the Compute Grid can run in-memory MapReduce.

With all these features and tools enabled and running optimally, SaaS users can vastly improve their applications by:

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

DEPLOYING IN-MEMORY COMPUTING PLATFORMS TO IMPROVE ECOMMERCE

Many enterprise software and web-scale SaaS users across all kinds of industries —everything from financial services and mobile to biotech genomic modeling— have adopted GridGain's in-memory computing solution to improve performance and functionality. Clients including Workday, Microsoft, IBM, and Siemens have all taken advantage of GridGain's in-memory computing speed.

High-speed transactions: In the case of Misys, a financial services software provider with 2,000 clients, including 48 of the world's 50 largest banks, the issue was how to eliminate data processing bottlenecks. Its clients deal with huge amounts of trading and accounting data, and they need to manage high-speed transactions and conduct real-time reporting, something that Misys's software architecture could not provide. There was a true need for speed. Nighttime batch processing was no longer acceptable.

The solution: To move to a Java-based IT stack that taps into a data lake as opposed to a traditional data warehouse, and to do that while creating an environment that can address the individual regulations of up to 130 countries where their clients conduct business. The new Misys commodity servers each contain up to 256GB RAM to store and process transaction and market data in-memory and manage parallel processing across clusters.

Misys has also used in-memory computing tactics to develop and launch its new FusionFabric Connect product, a cloud-based SaaS collection of modules that integrates many trading systems.

Faster search: Another GridGain client markets more than 20,000 information products worldwide. It needed to implement much faster search across five million different information assets, a number that may soon grow to 50 million. With a potential \$275 million of annual revenue in the balance, the client turned to GridGain to deliver the search speed that only in-memory computing can provide.

IoT management: GridGain worked with Silver Spring Networks, a company specializing in Internet-connected devices, to address a large real-time data processing challenge. The company's connected energy meters deliver data from millions of homes, and it forecasted that its existing database infrastructure would soon no longer be able to keep up. Not only did it need to ingest massive amounts of data constantly, it also had to overlay current electricity rates and perform compute-heavy calculations on the fly in order to produce regular analytical reports combining current state and recent history.

Silver Spring Networks was committed to keeping as much of its database infrastructure as it could, so the solution was to add GridGain's In-Memory Data Fabric on top of it, creating a fast space in which to do the required real-time data processing and analytics and providing as much scalability as the company needs—taking their cluster from 100 to 1,000 nodes if necessary—as its business grows. Today the company can store a week of its current data entirely in RAM, vastly improving its ability to deal with the never-ending rush of data it must process moment by moment.

GRIDGAIN SYSTEMS: A LEADER IN IN-MEMORY COMPUTING

The GridGain in-memory computing platform helps SaaS 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. This, coupled with the rapidly decreasing cost of memory, boosts ROI for in-memory computing initiatives, enabling online retailers to build less expensive systems that perform thousands of times better. Any eCommerce company wrestling with the need for high-performance transactions and/or analytics across large data volumes can benefit from the GridGain in-memory computing platform.

A Unified High-Performance Architecture

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

Scalability

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

Full SQL Support

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

High Availability

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

Transaction Processing

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

Security Features

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

Open Source Framework

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

Production Support

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

FLEXIBILITY, SCALABILITY AND SPEED HAVE COME TOGETHER

Developers require flexibility to build the perfect custom software and SaaS solution at the right time.

An in-memory computing platform delivers a highly scalable and high-performance experience featuring high-speed transactions and real-time analytics across a variety of data stores and applications. When those two toolkits come together, the possibilities for developers to point their businesses on a path to growth and success are limitless.

Contact GridGain Systems

To learn more about how GridGain can help your business, please email our sales team at sales@gridgain.com, call us at +1 (650) 241-2281 (US) or +44 (0)208 610 0666 (Europe), or complete our [contact form at www.gridgain.com/contact](http://www.gridgain.com/contact) and we will contact you.

About GridGain Systems

GridGain Systems is revolutionizing real-time data access and processing with the GridGain in-memory computing platform built on Apache® Ignite™. GridGain and Apache Ignite are used by tens of thousands of global enterprises in financial services, fintech, software, e-commerce, retail, online business services, healthcare, telecom and other major sectors, with a client list that includes ING, Raymond James, American Express, Societe Generale, Finastra, IHS Markit, ServiceNow, Marketo, RingCentral, American Airlines, Agilent, and UnitedHealthcare. GridGain delivers unprecedented speed and massive scalability to both legacy and greenfield applications. Deployed on a distributed cluster of commodity servers, GridGain software can reside between the application and data layers (RDBMS, NoSQL and Apache® Hadoop®), requiring no rip-and-replace of the existing databases, or it can be deployed as an in-memory transactional SQL database. GridGain is the most comprehensive in-memory computing platform for high-volume ACID transactions, real-time analytics, web-scale applications, continuous learning and hybrid transactional/analytical processing (HTAP). For more information on GridGain products and services, visit www.gridgain.com.

© 2019 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. Java, JMS and other Java-related products and specifications are either registered trademarks or trademarks of Oracle Corporation and its affiliates in the United States and/or other countries. Apache, Apache Ignite, Ignite, the Apache Ignite logo, Apache Spark, Spark, Apache Hadoop, Hadoop, Apache Camel, Apache Cassandra, Cassandra, Apache Flink, Apache Flume, Apache Kafka, Kafka, Apache Rocket MQ, Apache Storm 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.