

Delivering a Great OmniChannel Banking Experience with In-Memory Computing

Today, the majority of banking customers use digital channels. But being a digital bank by itself is clearly not enough. According to one study, customers who only use digital channels are among the least satisfied customers, alongside customers who only use branches. But customers who use multiple channels acquire 80% more products and generate more than 2x the revenues compared to customers who use a single channel. This should not be surprising. We all want a fast, simple omnichannel banking experience that allows us to use our channel of choice whenever we want and get the most out of our banking experience.

Omnichannel banking requires more than a consistent API strategy across channels. It requires a single, real-time view of the customer that can be seamlessly shared across channels; infrastructure that can handle 10x or greater loads created by the increased interactions on digital channels; and the ability to proactively personalize, promote and improve each customer's experience. Some companies have successfully consumed new sources of big data and inserted real-time analytics and automation into transactions and interactions to help improve the experience. Gartner calls this approach hybrid transactional/analytical processing (HTAP). The foundation that helped these banks deliver greater speed, scale and in-process HTAP for omnichannel banking is in-memory computing.

THE GRIDGAIN IN-MEMORY COMPUTING PLATFORM FOR OMNICHANNEL BANKING

GridGain is the leading in-memory computing platform for real-time business. It is built on Apache® Ignite™, one of the top five Apache open source projects. GridGain Systems contributed the code that became Ignite to the Apache Software Foundation and continues to be a leading contributor to the project.

Leading banks rely on GridGain to help them deliver an integrated omnichannel banking experience. By using GridGain, companies have not only added speed and scale to digital channels. They've opened up previously siloed data for seamless sharing across channels, and implemented in-process HTAP using real-time streaming analytics, machine and deep learning to help proactively monitor and improve the end-to-end banking experience.

CREATING A DIFFERENTIATED BANKING EXPERIENCE WITH IN MEMORY COMPUTING

A leading global bank was struggling to increase their pace of innovation, both to service changing customer needs and to move beyond traditional banking with new services and business models. Part of the challenge was that core banking services resided on a mainframe that was hard to change. They were also unable to continue to scale their infrastructure and handle increasing loads. Mobile traffic was growing 25% per year, and each mobile user interacted 7x more than a Web user.

Using GridGain, the bank added an in-memory computing layer in front of their mainframe and end-to-end middleware that helped open up and share previously siloed data and functionality across channels. GridGain not only easily scaled to offload from the mainframe and handle the growth. It helped cut end-to-end latency to less than 100ms. It also helped the bank rapidly create new APIs for consumers and 3rd parties, and be first to market with new services for PSD2, SEPA and instant payments.

ADD IN-MEMORY SPEED AND UNLIMITED HORIZONTAL SCALABILITY TO EXISTING APPLICATIONS AND APIS

GridGain is the only in-memory computing platform that slides in-between existing applications and databases as an in-memory data grid (IMDG) with no rip-and-replace of the existing systems. GridGain is able to accomplish this without major architectural changes because it supports ANSI-99 SQL and distributed ACID transactions. GridGain can sit on top of leading RDBMSs including IBM DB2®, Microsoft SQL Server®, MySQL®, Oracle® and Postgres® as well as NoSQL databases such as Apache Cassandra™ and MongoDB®. GridGain loads the underlying database schema, then the data, and then handles all SQL requests via JDBC/ODBC. GridGain also includes a Service Grid that natively hosts microservices and APIs. For all writes, GridGain commits transactions to the underlying database and then GridGain to ensures data consistency. This enables GridGain to offload all reads from the database and run distributed SQL in memory on its shared nothing, scale-out architecture. The result is up to 1000x lower latency and unlimited horizontal scalability.

BUILD COMMON OMNICHANNEL DATA SERVICES FOR SEAMLESS SHARING ACROSS CHANNELS

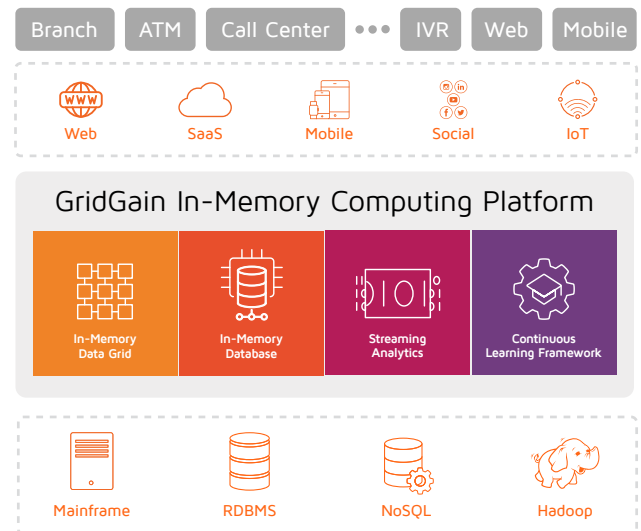
Several banks use GridGain to seamlessly share data across different channels in real-time. Once GridGain is used as an IMDG for an application or database, any of the data accessed by GridGain is immediately available for any other use. This allows data across applications to be used together using SQL or any custom code as though they were in the same database. Applications can use GridGain's built-in features or general-purpose in-memory computing to implement and run any new real-time HTAP computations during a transaction or interaction to help share and enrich information across channels. GridGain's unlimited horizontal scalability enables new computations to be added without impacting existing application performance.

IMPROVE THE END-TO-END BANKING EXPERIENCE WITH HTAP

It is no longer enough to run faster, add new channels or provide more self-service. Customers expect a better end-to-end experience. This requires banks to adopt new approaches such as HTAP that help identify issues and take action during each interaction to help improve the customer experience and customer satisfaction. GridGain provides the broadest in-process HTAP support that is used by several banks to improve the end-to-end customer experience. Companies can ingest just about any data from a host of streaming technologies including Apache Camel™, Kafka™, Spark™, JMS or IBM® MQ, and run just about any real-time processing mid-stream or mid-transaction. You can leverage GridGain's built-in streaming and messaging capabilities, GridGain's broad support for Apache Spark™, the

GridGain Continuous Learning Framework for machine and deep learning, or the GridGain Compute Grid to run any SQL, Java, .NET, or C++. GridGain sends code across the cluster to the data and executes it locally using massively parallel processing (MPP) to achieve in-memory speed and unlimited linear horizontal scale.

By using GridGain's in-process HTAP support, banks have been able to deliver new applications that help identify important events and automate responses to improve the end-to-end experience. Examples include real-time analytics, personalization, pricing, processing monitoring and alerting. One company relies on GridGain to deliver a personalized experience by performing nearly 500,000 pricing calculations per second during peak load times to promote the best products based on each customer's preferences.



Learn More

[Introducing the GridGain In-Memory Computing Platform](#)

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, 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.

© 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. Apache, Apache Ignite, Ignite and the Apache Ignite logo 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.