

Migrating from DataSynapse to GridGain and Modern In-Memory Computing

TIBCO acquired DataSynapse in 2009 for \$28 million. DataSynapse GridServer is a specialized compute grid, with strength in executing Java, .NET, C/C++ and R across a cluster of nodes. In 2009 GridServer was at version 5, and there hasn't been much innovation since then. Release 6.2 in November 2014 added TIBCO Enterprise R and TIBCO LogLogic support. Release 7.0 in November 2018 added Kubernetes and Python support.

The distributed computing market has been transformed since then by the development of in-memory computing to solve a host of performance and scalability challenges. GridGain, a leading in-memory computing vendor, was founded in 2010. In 2014, GridGain donated the core code to the Apache Ignite in-memory computing project. Today, thousands of companies use Apache Ignite and GridGain (the commercially-supported version of Ignite) for in-memory computing. Many implementations use GridGain as an in-memory data grid and leverage its compute grid capabilities. GridGain not only covers the capabilities of DataSynapse, it adds a host of other crucial in-memory computing abilities (like co-location) that address today's data- and compute-intensive applications. The combination makes the decision to migrate applications from GridServer to GridGain a simple and straightforward one, and allows you to run most of the latest distributed computing workloads on the same in-memory computing platform.

THE ADOPTION OF IN-MEMORY COMPUTING IN BANKS

The financial crisis of 2008 clearly revealed that banks did not have an adequate measure of "risk." It caused a sea change in both business and regulatory requirements around market risk that created a demand for real-time analytics that overwhelmed existing systems. For example, the pre-crisis goal of intra-day risk visibility became a demand for real-time visibility post-crisis. Banks with a capital market presence now arguably face their biggest regulatory requirements and business challenges in the Fundamental Review of the Trading Book (FRTB) and ("x") valuation adjustments known as XVA. Even compared to Basel 2.5, the performance and scalability challenges of FRTB and XVA are daunting. The computations needed for different stress scenarios and back-testing can increase by up to 50x. Existing high-performance computing and analytics infrastructure such as DataSynapse GridServer can't deliver both the real-time speed and 100x (or greater) scale.

THE GRIDGAIN IN-MEMORY COMPUTING PLATFORM

GridGain has helped leading fintech, investment management, commercial banking, insurance, and other more specialized

financial services companies add speed and scale to applications, APIs, and analytics, and bring a host of innovative new real-time capabilities to market. Examples include the real-time availability of:

- Market and credit risk management
- Portfolio valuation
- Pricing analytics
- Pre-deal limit checking
- Trade settlement
- Cybersecurity and fraud prevention
- Regulatory compliance for Basel I, II, III and FRTB
- Omnichannel banking
- High-performance computing (HPC)

 Omnichannel
Services

 Regulatory
Compliance

 Risk, Fraud
and Security

 High-Performance
Computing

GridGain In-Memory Computing Platform


 In-Memory
Data Grid

 In-Memory
Database

 Streaming
Analytics

 Continuous
Learning Framework

MIGRATING FROM DATASYNAPSE GRIDSERVER TO THE GRIDGAIN IN-MEMORY COMPUTING PLATFORM

GridGain matches DataSynapse's core capabilities, including the ability to distribute and run any Java, .NET or C/C++ code across a cluster along with any needed data. But instead of moving the data each time with the code, GridGain distributes the data exactly once and then collocates the code to run it in-place against the data. This is a critical difference for many newer workloads, as the data is often too big (terabytes or more) to move in near real-time over the network. Most corporate networks can move one gigabyte a second, or 3.6 terabytes an hour. This makes GridGain a better solution for large data sizes (even up to petabytes of data).

GridGain uses cloud-native computing principles to handle any number of requests. It can automatically scale up and down across private or public cloud infrastructures (including Amazon Web Services (AWS), Google Cloud or Microsoft Azure)

as needed to handle different workloads. It can manage data and computing as a single or multiple clusters across multiple private and public data centers as well.

USING GRIDGAIN AS A COMPUTE GRID TO REPLACE DATASYNAPSE

GridGain is a proven in-memory computing platform used for grid computing by investment banks, fintech companies, insurance companies, and others. It matches DataSynapse's capabilities as a general-compute grid, and provides full support for the automatic distribution and deployment of classes and resources required to execute a specific job written in Java, .NET or C/C++. Capabilities include:

- **Topology Resolution**—Provisioning of nodes based on any node characteristic or user-specific configuration, such as type of operating system, time window, CPU load or available heap memory.
- **Collision Resolution**—Control over job execution such as which jobs get executed, which jobs get rejected, how many jobs can be executed in parallel, or the order of overall execution.
- **Load Balancing**—Fine-grained control of load balancing across the grid using policies such as Round Robin, Random, Adaptive rules, or Affinity Load Balancing.
- **Failover**—Configurable automatic or manual failover for ensuring that jobs automatically failover onto other nodes in case of node crashes or other exceptions.
- **Distributed Closures**—Support for broadcast or load-balanced distribution of closures using call/run/apply methods.
- **Distributed Executor Service**—Use of ExecutorService as a distributed, fault-tolerant cluster-enabled implementation.
- **MapReduce and ForkJoin**—Support for ComputeTask, a lightweight in-memory MapReduce/ ForkJoin implementation.
- **Continuous Mapping**—MapReduce/ ForkJoin equivalent for continuous stream processing of jobs.
- **Per-Node Shared State**—Shared concurrent node-local-map to help share the node state across compute jobs and services.

- **Collocated Computing**—Collocate code with data via Affinity Call and Run methods.
- **Fault Tolerance**—Automatic job failover by transferring jobs to other available nodes following predefined conditions.
- **Checkpointing**—Save intermediate job state for faster recovery.
- **Job Scheduling**—Support for random, FIFO, one-at-a-time (ordered), and priority-based execution.
- **Cron-Based Scheduling**—Periodic execution using a scheduler based on cron syntax.
- **Task Deployment**—Multiple deployment implementations including peer-class, Uri and LocalDeployment.
- **Aspect-Oriented Programming (AOP)-Based Enabling**—Support for annotations to change methods using AOP.

By combining the power of a general-purpose compute grid with an in-memory data grid, GridGain enables leading fintech and financial services companies to perform real-time risk, analytics, and compliance processing with higher speed and scale than what is possible with DataSynapse GridServer. It successfully supports billions of present value (PV), "X" value adjustment (XVA) or Monte Carlo calculations per second.

GRIDGAIN AS A GENERAL-PURPOSE IN-MEMORY COMPUTING PLATFORM

GridGain provides a host of other capabilities that add speed and scale beyond grid computing. GridGain's native support for ANSI-99 SQL and ACID transactions enables it to slide between leading databases and applications as an in-memory data grid (IMDG) without having to rip-and-replace trading or risk infrastructure. It can be used as an distributed in-memory database (IMDB) for data ingestion, transactions and risk computations. It also supports streaming analytics, with the broadest support for Apache Spark™, for processing trade or market data. Its Continuous Learning Framework leverages machine and deep learning across trading, fraud or portfolio management.

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 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, Finasträ, 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.