

GRIDGAIN IN-MEMORY DATA FABRIC

ULTIMATE SPEED AND SCALE FOR ALL APPS AND DATA

The GridGain In-Memory Data Fabric is a proven software solution, which delivers unprecedented speed and unlimited scale to accelerate your business and time to insights. It enables high-performance transactions, real-time streaming and fast analytics in a single, comprehensive data access and processing layer. The In-Memory Data Fabric is designed to easily power both existing and new applications in a distributed, massively parallel architecture on affordable, industry-standard hardware.

The GridGain In-Memory Data Fabric provides a unified API that spans all key types of applications (Java, .NET, C++) and connects them with multiple data stores containing structured, semi-structured and unstructured data (SQL, NoSQL, Hadoop). It offers a secure, highly available and manageable data environment that allows companies to process full ACID transactions and generate valuable insights from real-time, interactive and batch queries.

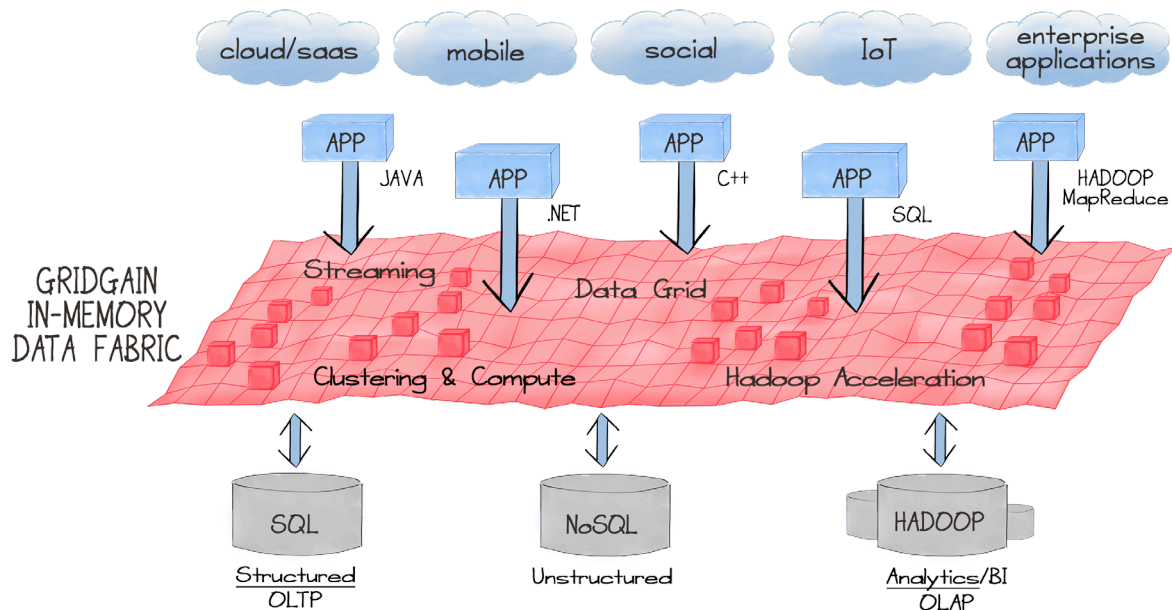
The In-Memory Data Fabric offers a strategic approach to in-memory computing that delivers performance, scale and comprehensive capabilities far above and beyond what traditional in-memory databases, data grids or other in-memory-based point solutions can offer by themselves.

FEATURES

The GridGain In-Memory Data Fabric accesses and processes data from distributed enterprise and cloud-based data stores orders of magnitudes faster, and shares them with today's most demanding transactional, analytical and hybrid applications with varying SLA requirements (real-time, interactive, batch jobs). To that effect, the GridGain In-Memory Data Fabric includes the following four classes of features:

- Cross-platform Data Grid (Java, .NET, C++)
- Clustering and Compute Grid
- Real-time Streaming
- Hadoop Acceleration

The **data grid** feature in the GridGain In-Memory Data Fabric supports local, replicated, and partitioned data sets and allows to freely cross queries between these data sets using standard SQL syntax. It is a fully-featured caching and data grid solution for .NET and C++, besides Java, and includes transactions, client-side caching, etc. It supports standard SQL for querying in-memory data, including support for distributed SQL joins.



The GridGain In-Memory Data Fabric offers an extremely rich set of data grid capabilities, including off-heap memory support, load-balancing, fault tolerance, remote connectivity, support for full ACID transactions and advanced security. Its dynamic protocol provides extensive API-parity between Java, .NET and C++ apps, allowing for easy data exchange and cross-platform dynamic queries (e.g. using SQL).

In-Memory **clustering and compute grids** are characterized by using high-performance, integrated, distributed memory systems to compute and transact on large-scale data sets in real-time, orders of magnitude faster than possible with traditional disk-based or flash technologies.


The **real-time streaming** feature of the GridGain In-Memory Data Fabric uses programmatic coding with rich data indexing support to provide CEP querying capabilities over streaming data. The GridGain In-Memory Data Fabric also provides comprehensive support for customizable event workflow.

Hadoop acceleration included in the GridGain In-Memory Data Fabric features the GridGain in-memory file system (GGFS). It has been designed to work in dual mode as either a standalone primary file system in the Hadoop cluster, or in tandem with HDFS, serving as an intelligent caching layer with HDFS configured as the primary file system.

ULTIMATE SPEED AND SCALE

The GridGain In-Memory Data Fabric accesses and processes data from distributed enterprise and cloud-based data stores orders of magnitudes faster, and shares them with today's most demanding transactional, analytical and hybrid applications. The In-Memory Data Fabric delivers unprecedented throughput and low latency performance in a virtually unlimited, global scale-out architecture for both new and existing applications.

For example, benchmark results that measure atomic distributed cache put operations for randomly generated key-value pairs stored in cache on multiple nodes, show that the GridGain In-Memory Data Fabric performed exceptionally well with approximately 300,000 operations per second.



COMPREHENSIVE AND PROVEN

The GridGain In-Memory Data Fabric is the product of seven years of meticulous research and development, built from the ground up (i.e. no bolt-on's), and is run successfully by hundreds of organizations worldwide. It is a comprehensive in-memory computing solution, which offers a secure and highly available architecture, and also provides a fully-featured, graphical management interface.

The GridGain Data Fabric is used today by Fortune 500 companies, top government agencies as well as innovative mobile and web companies in a broad range of business scenarios, such as automated trading systems, fraud detection, big data analytics, social networks, online gaming, and bioinformatics.

OPEN AND AFFORDABLE

The falling cost of RAM has made In-Memory Computing an economical solution for today's exponentially growing computing and big data demands. The GridGain In-Memory Data Fabric is an effective and affordable way to achieve previously unimaginable performance and scale on inexpensive, industry-standard hardware. GridGain's open source version of the In-Memory Data Fabric offers immediate, unhindered access to the most comprehensive in-memory computing solution on the market, and allows the community to build and deploy hyper-scale, high-performance applications within incredibly short time cycles.

EDITIONS

GridGain Systems is the primary backing organization behind the free, open source GridGain In-Memory Data Fabric, which offers a comprehensive set of data grid, clustering and compute grid, real-time streaming and Hadoop acceleration features. GridGain also continues to develop a fully supported enterprise version of the In-Memory Data Fabric that provides features and services for highly scalable, secure and mission-critical production deployments.

Please visit www.gridgain.com/download to see a feature-by-feature comparison of both editions, and for a free download of the GridGain In-Memory Data Fabric.

ABOUT GRIDGAIN™

GridGain, the leading provider of the open source In-Memory Data Fabric, offers the most comprehensive in-memory computing solution to equip the real-time enterprise with a new level of computing power. Enabling high-performance transactions, real-time streaming and ultra-fast analytics in a single, highly scalable data access and processing layer, GridGain enables customers to predict and innovate ahead of market changes.

GridGain is headquartered in Foster City, California.