

GridGain In-Memory Computing Platform Feature Comparison: Pivotal GemFire

This document presents a summary and detailed feature comparison of the [GridGain® in-memory computing platform](#) (GridGain) and Pivotal GemFire® (GemFire) for use as in-memory data grids (IMDG) and for in-memory computing use cases.

Major Advantages of GridGain vs. GemFire

- Comprehensive in-memory computing platform
- Native ANSI-99 SQL support
- Distributed ACID transaction support
- Automatic integration with RDBMSs, NoSQL databases and Hadoop
- Sits in-between SQL-based applications and RDBMSs, eliminating the need to replace SQL with code
- Distributed in-memory database with better transaction support, scalability, availability and reliability
- Apache® Spark™ Support for DataFrames, RDDs, HDFS and SparkSQL acceleration
- Built-in Continuous Learning Framework for machine learning and deep learning
- Built on Apache Ignite, a top 5 Apache Software Foundation

PIVOTAL GEMFIRE AND GRIDGAIN IN-MEMORY COMPUTING PLATFORM COMPARED

Pivotal GemFire is an in-memory data grid (IMDG) powered by the Apache Geode™ open source project. It is used by companies to scale data services on demand to support high-performance, real-time applications.

GridGain, built on the [Apache Ignite open source project](#), is an [in-memory computing platform](#) that includes a distributed in-memory data grid (IMDG), a hybrid SQL and key-value in-memory database (IMDB), a stream processing and analytics engine, and a continuous learning framework that supports real-time machine and deep learning. It can be used with any RDBMS, NoSQL or Hadoop database.

GemFire has many core IMDG capabilities, including the ability to distribute and partition data, and scale out across a cluster. But GemFire hasn't evolved much in the last few years. Gemstone, the original company, was acquired in 2010 by VMWare and later spun out as part of Pivotal in 2012. In 2015 Pivotal donated GemFire code to the Apache Software Foundation for the Apache Geode project. Geode became a top-level Apache project the end of 2016.

GridGain is a superior IMDG for the majority of existing applications. This is due in part to being built on the innovative Apache Ignite project. GridGain Systems donated the original Apache Ignite code to the Apache Software Foundation (ASF) in 2014 and remains the most active contributor. Ignite became a top level ASF project in 2015. Ignite is now one

of the top five Apache Software Foundation open source projects in commits and list activity, with over twice the commits of Apache Geode. In the last three years, Apache Ignite and GridGain have added major new capabilities, including:

- Distributed SQL and ACID transaction support
- All the capabilities of a SQL and key-value IMDB
- The broadest Spark support
- A Continuous Learning Framework for machine and deep learning

Both the GemFire and Apache Geode releases have been incremental, maintenance releases for the last three years. Since Pivotal contributed the initial GemFire code that became Apache Geode, the only significant additions have been some integration between Geode/GemFire and other Pivotal technologies which Pivotal sells with GemFire.

One key area of innovation is GridGain's support for ANSI-99 SQL and ACID transactions. These capabilities support GridGain usage as an IMDG on top of existing databases and as a distributed in-memory database. GridGain is the only IMDG with support for SQL. As a result, it is the only IMDG that allows applications, analytics and reporting tools that rely on SQL to easily access data. GridGain is also the only IMDG that can sit in-between existing SQL-based applications and RDBMSs and use ANSI-99 SQL instead of code for the integration. GemFire has no way to provide easy access to data for tools and applications that rely on SQL. It also requires extensive custom coding in applications to replace existing SQL.

GridGain also has proven, high performance support for distributed ACID transactions across nodes and partitions. GemFire can perform multi-partition transactions but only within a single node of a cluster, which is too limiting for large-scale distributed applications.

GridGain is better for supporting IMDG and many other non-IMDG in-memory computing use cases on a single common deployment. It is a broader in-memory computing platform that includes an IMDG, in-memory database (IMDB), streaming analytics, machine learning and deep learning. GemFire is a strong IMDG with support for streaming, but it is just an IMDG.

As a database, GridGain provides better transactional support, scalability, availability and reliability. GemFire is not in-

tended for use as a database. It doesn't support SQL or fully distributed transactions. It doesn't scale as well across memory and disk. It doesn't provide immediate availability. And it doesn't completely ensure consistency or recoverability.

GEMFIRE AND GRIDGAIN EDITIONS

To compare products, it is important to understand the differences between Geode, GemFire, Apache Ignite, and the GridGain editions. This document will compare GemFire to the various GridGain editions. This can also be used to compare Apache Ignite to Geode since, like GridGain, GemFire contains everything in Geode that is considered production-ready.

The [GridGain Community Edition](#) (CE) includes the current version of Apache Ignite with LGPL dependencies, as well as bug fixes that have not yet been released in Ignite. The [GridGain Enterprise Edition](#) (EE) adds enterprise-grade security, deployment, and management capabilities needed for most mission critical in-memory data grid applications. The [GridGain Ultimate Edition](#) (UE) includes the Enterprise Edition features plus advanced data management and disaster recovery features for using GridGain as an in-memory database. For this comparison, all the features of Terracotta BigMemory Max are used unless otherwise noted.

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Native ANSI-99 SQL Support	●	●	●	●
Distributed ACID Transaction Support	●	●	●	● (No support for distributed transactions across members/nodes)
Slide in-between SQL-based Applications and RDBMSs with no Custom Coding	●	●	●	● (No SQL support. Requires code to replace SQL in existing apps)
Cross-Language Support for Collected Processing (MPP)	● (Supports Multiple Languages, MPP for Java, .NET, C++)	● (Supports Multiple Languages, MPP for Java, .NET, C++)	● (Supports Multiple Languages, MPP for Java, .NET, C++)	● (Supports Java, .NET, C++)
Integration with RDBMSs, NoSQL Databases and Hadoop	● (Out-of-the box support for RDBMSs, NoSQL Databases, HDFS, Spark)	● (Out-of-the box support for RDBMSs, NoSQL Databases, HDFS, Spark)	● (Out-of-the box support for RDBMSs, NoSQL Databases, HDFS, Spark)	● (Requires coding. Does support Spark)
Comprehensive In-Memory Computing Solution	● (EE + Multi-datacenter data and disaster recovery management)	● (EE + Multi-datacenter data and disaster recovery management)	● (EE + Multi-datacenter data and disaster recovery management)	● (Supports IMDG only)
Apache Spark Support for DataFrames, RDDs, HDFS	●	●	●	● (Supports RDDs, returning OQL as DataFrames)
Built-in Machine Learning	●	●	●	●
Built on a Leading Open Source Project	● (Built on Apache Ignite, a top 5 Apache Software Foundation open source project)	● (Built on Apache Ignite, a top 5 Apache Software Foundation open source project)	● (Built on Apache Ignite, a top 5 Apache Software Foundation open source project)	● (Built on Apache Geode, which is not a top 10 ASF project, little innovation)

This document provides a detailed feature comparison between the various versions of GridGain and GemFire, highlighting differences between the open source and enterprise-ready products where relevant.

COMPREHENSIVE IN-MEMORY COMPUTING PLATFORM

GridGain can support multiple third-party databases as an IMDG, act as an IMDB, support streaming analytics, and perform machine and deep learning for different applications, all on the same deployment. GemFire's strengths are as an IMDG and its support for streaming use cases. GemFire is not meant to be used as an IMDB or for machine learning.

NATIVE ANSI-99 SQL SUPPORT

GridGain supports ANSI-99 compliant SQL, including distributed SQL JOINS for querying, updating, indexing and defining data. GridGain works with the SQL from an application by providing ODBC and JDBC drivers to use for Java, .NET, C++, Python and other languages. GemFire does not support SQL. In order to integrate an application with GemFire, developers have to change existing applications. They have to add code that replaces access to each table and each SQL query or update to the underlying database. A lack of SQL support also means various SQL-based analytics and reporting tools cannot access the data in GemFire.

DISTRIBUTED ACID TRANSACTION SUPPORT

GridGain has full support for distributed ACID transactions, including OPTIMISTIC and PESSIMISTIC concurrency modes, as well as READ_COMMITTED, REPEATABLE_READ, and SERIALIZABLE isolation levels both as an IMDG and IMDB. GemFire has extensive support for transactions but not distributed transactions beyond a single node in a cluster.

AUTOMATIC INTEGRATION WITH RDBMSs, NOSQL DATABASES AND HADOOP

GridGain can automatically integrate with all leading RDBMSs including IBM DB2®, Microsoft SQL Server®, MySQL®, Oracle® Database and Postgres®. It also automatically integrates with leading NoSQL databases, such as Apache Cas-

sandra® or MongoDB®, as well as Hadoop via Spark and HDFS. GemFire requires writing custom code to integrate with third-party RDBMSs and does not support Cassandra, MongoDB or HDFS.

SITS IN-BETWEEN SQL-BASED APPLICATIONS AND RDBMSs, ELIMINATING THE NEED TO REPLACE SQL WITH CODE

GridGain's architecture and out-of-the-box integration enables many use cases to be implemented by working with SQL with no custom coding. GridGain can sit in-between and accelerate SQL-based applications and third-party databases by allowing the SQL applications to use SQL to query and insert/update data in GridGain. The applications use GridGain's native JDBC/ODBC drivers in place of existing drivers. You can then configure GridGain to connect to the backend database. GridGain can use the schema of the underlying RDBMS during configuration to create its own matching SQL schema. GridGain supports any ANSI-99 SQL requests from the application through the drivers, which means developers can focus on SQL and not have to write custom code that replaces existing SQL. GemFire may require significant code changes to an application because it does not support SQL. It requires developers to add code that replaces SQL calls to the underlying database with calls to GemFire, and more code to enable GemFire to query and write to the underlying database.

DISTRIBUTED IN-MEMORY DATABASE (IMDB) WITH BETTER TRANSACTION SUPPORT, SCALABILITY AVAILABILITY AND RELIABILITY

GridGain includes a distributed SQL and key-value hybrid in-memory database that combines memory-centric storage with built-in persistence. As an IMDB, GridGain provides better transactional support, scalability, availability and reliability. It supports distributed transactions across all data, partitions and nodes. GemFire does not support distributed transactions across nodes. GridGain's built-in persistence is a distributed ACID and SQL-compliant disk store for storing data and indexes on SSD, Flash, 3D XPoint, and other types of non-volatile storages. With persistence enabled, nonvolatile storage

houses the full data set while RAM holds 0% to 100% of the data and indexes. GridGain scales better by allowing indexes, not just data, to be stored outside of RAM. If a subset of data or an index is not in RAM, it will be used from non-volatile storage. Data in RAM and non-volatile storage is stored and treated exactly the same way. Any changes are written to a write-ahead log and then to non-volatile storage to ensure low latency. GemFire can't store indexes outside of RAM, which limits its performance and scalability. GridGain also has immediate availability on cluster restarts for SQL and key-value operations before any data or indexes are loaded into RAM. It becomes fully operational once all the cluster nodes are interconnected with each other. There is no need to warm up the memory by preloading data from the disk. GemFire does not support SQL and, as a key-value store, requires all keys to be loaded first before handling requests. GridGain also includes centralized backup and recovery to ensure data is never lost. Capabilities include full and incremental snapshots, continuous archiving and point-in-time recovery down to the transaction. It also includes network backups and heterogeneous recovery of a cluster where you can re-deploy an existing cluster to any other location on premise or in the cloud, with a different number of nodes and amount of memory. GemFire supports full and incremental backups but cannot ensure consistency unless all writes are stopped during the backup process. It cannot backup or restore down to the level of an individual transaction and it cannot easily do heterogeneous recovery.

APACHE® SPARK™ SUPPORT FOR DATAFRAMES, RDDS, HDFS AND SPARKSQL ACCELERATION

GridGain provides the broadest integration with Spark compared to other in-memory computing vendors. It provides native Apache Spark DataFrame, RDD and HDFS support. This integration simplifies the access, writing and saving of data. It also enables state to be shared across Spark jobs. The integration can also dramatically improve SQL performance compared to standalone Apache Spark because, unlike Spark, GridGain supports primary and secondary indexes. GridGain also improves overall analytics and machine learning performance by providing access to GridGain's MPP capabilities, which include built-in dis-

tributed joins, and GridGain’s machine and deep learning capabilities. GemFire provides support for RDDs and enables OQL to be returned as a DataFrame. However, GemFire is not a DataFrame provider that allows reads or writes through the DataFrame interface, and it does not accelerate SparkSQL or HDFS.

BUILT-IN CONTINUOUS LEARNING FRAMEWORK WITH SUPPORT FOR MACHINE LEARNING AND DEEP LEARNING

GridGain includes the GridGain Continuous Learning Framework, which provides built-in machine learning and deep learning with real-time performance on petabytes of data. GridGain provides several out-of-the-box machine learning algorithms optimized for MPP-style

collocated processing including linear and multi-linear regression, k-means clustering, decision trees, k-NN classification and regression. GridGain also includes a multilayer perception and TensorFlow integration for deep learning. Developers can create and deploy their own algorithms across any cluster as well by using the compute grid. GemFire provides no support for machine or deep learning.

BUILT ON APACHE IGNITE, A TOP 5 APACHE SOFTWARE FOUNDATION

GridGain is built on Apache Ignite, one of the top five Apache Software Foundation projects in terms of commits and community activity. GemFire contributed the code for Apache Geode, which is

the core of GemFire, but it is not a top 10 ASF project. GemFire has less than half the commits of Ignite and has not added any major functionality in the 3 years since it was donated.

GRIDGAIN AND GEMFIRE DETAILED FEATURE COMPARISON

The following table provides a detailed feature comparison between the GridGain Professional, Enterprise, and Ultimate Editions, and GemFire. This comparison is based on our best knowledge of the features available at the time this document was created for the product versions indicated.

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Use Cases				
In-Memory Data Grid	●	●	●	●
Third party Database Caching and Persistence (Inline)	●	●	●	●
SQL Database	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
In-Memory Database	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Web Session Clustering	●	●	●	●
Apache Spark Acceleration	●	●	●	● (Spark RDD support provided through Spark Connector. Also allows OQL to return DataFrame. SnappyData also combines GemFire and Spark.)
Hadoop acceleration	●	●	●	●
In-Memory File System (Hadoop Compliant)	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Third Party Database Support, Persistence				
Automatic support for Leading RDBMSs (Oracle, IBM DB2, Microsoft SQL Server, MySQL, Postgres ...)	●	●	●	● (It works well, but requires coding to implement. It's not out of the box)
Inline Support for Apache Cassandra	●	●	●	●
Inline support for MongoDB	●	●	●	●
Write-Through and Read-Through Caching	●	●	●	●
Write-Behind Caching	●	●	●	●
Auto-Loading of SQL Schema/Data	●	●	●	●
Store Loader (Optimized BulkDB Load)	●	●	●	●
Native Persistence				
Native Persistence	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Stores Superset of Data	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Store Indexes on Disk	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
SQL or Key-Value over Disk	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (key-value, OQL only)
Instantaneous Restart (before memory warm-up)	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (Faster than others but requires keys preloaded)

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Distributed SQL				
SQL Queries	● (Full ANSI-99 Support)	● (Full ANSI-99 Support)	● (Full ANSI-99 Support)	● (OQL)
Collocated Distributed Joins	●	●	●	● (OQL)
Non-Collocated Distributed Joins	●	●	●	●
Single Column Indexes	●	●	●	●
Group Indexes	●	●	●	●
Distributed SQL Joins (select * from Person p, Company c where p.c_ id=c.id)	●	●	●	●
Query Consistency	●	●	●	●
Query Fault-Tolerance	●	●	●	●
DML (INSERT, UPDATE, DELETE, MERGE)	●	●	●	●
DDL (CREATE, DROP, ALTER)	●	●	●	●
Distributed Queries				
Continuous Queries	●	●	●	●
Predicate-based Queries	●	●	●	●
SQL Drivers				
JDBC Driver	●	●	●	●
ODBC Driver	●	●	●	●
REST API (SQL)	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Memory Architecture				
On-Heap Memory	●	●	●	●
Off-Heap Memory	●	●	●	●
Off-Heap Indexes	●	●	●	●
Disk as main storage (disk larger than RAM)	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Tiered Storage - On-Heap, Off-Heap and Disk	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
ACID Compliant Transactions and Locks				
Atomic Mode (One Operation at a Time)	●	●	●	●
READ_COMMITTED, REPEATABLE_READ, SERIALIZABLE Isolation Levels	●	●	●	● (READ_COMMITTED ONLY)
Deadlock-Free Transactions	●	●	●	●
XA Integration	●	●	●	●
Fault Tolerance (Including Client/Near/ Primary/Backup Node Failures)	●	●	●	●
Optimistic & Pessimistic Concurrency (Two- Phase- Commit)	●	●	●	● (Optimistic only)
One-Phase-Commit Optimization	●	●	●	●
Near Cache Transactions (i.e., Client Cache Transactions)	●	●	●	●
Cross-Partition Transactions	●	●	●	●
Transactional Entry Processor	●	●	●	●
Eviction / Expiration Policies for Transactional Caches	●	●	●	●
Merge with DB Trans- actions (e.g., Oracle DB, MySQL, etc.)	●	●	●	●
Explicit Locking	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Distributed Architecture				
Key-Value Store	●	●	●	●
Partitioning and Replication	●	●	●	●
Elasticity (add/remove nodes on demand)	●	●	●	●
Client-side (Near / inline) Cache	●	●	●	● (no transactional capabilities)
Dynamic Cache Creation	●	●	●	●
EntryProcessor, aka Delta (Partial) Updates	●	●	●	●
Data Redundancy (Key Backups)	●	●	●	●
Synchronous and Asynchronous Backup Update	●	●	●	● (Synchronous only)
Synchronous APIs	●	●	●	●
Asynchronous APIs	●	●	●	●
Full Sync Mode (Primary and Backups are Sync)	●	●	●	●
Primary Sync Mode (Primary is sync, Backups are Async)	●	●	●	●
Full Async Mode (Primary and Backups are Async)	●	●	●	●
Network Segmentation (Split Brain)	●	●	●	●
Data Conflict Resolution	●	●	●	●
Data Affinity and Collocation	● (Rich Support)	● (Rich Support)	● (Rich Support)	●
Custom affinity (partitioning) function	●	●	●	●
Data Eviction and Expiration	● (LRU, FIFO, Random, Sorted, Custom)	● (LRU, FIFO, Random, Sorted, Custom)	● (LRU, FIFO, Random, Sorted, Custom)	● (LRU, async only)
Binary Objects	●	●	●	●
Pluggable Interfaces (SPIs) to Customize Grid Subsystems	●	●	●	●
Dynamic Object Version Change (allowing dynamic change to an object's structure)	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Distributed Data Structures				
Queue	●	●	●	●
Set	●	●	●	●
Atomic Log	●	●	●	●
Atomic Ref	●	●	●	●
Atomic Stamped Ref	●	●	●	●
Atomic Sequence	●	●	●	●
Count Down Latch	●	●	●	●
Reentrant Lock	●	●	●	●
Semaphore	●	●	●	●
Data Snapshots (Backups)				
Full Data Snapshots	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Incremental Data Snapshots	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Continuous Archiving (WAL)	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Data Recovery from Snapshots	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Snapshots Scheduling	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Tools for Snapshotting	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Datacenter (WAN) Replication				
Active-Active	●	●	●	●
Active-Passive	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Data Rebalancing				
Sync Data Rebalancing (aka Sync Repartitioning)	●	●	●	●
Async Data Rebalancing (aka Async Repartitioning)	●	●	●	●
Delayed Data Rebalancing (Delay Data Rebalancing until All Nodes Have Start- ed)	●	●	●	●
Grid Management and Monitoring				
Rolling Production Updates	●	●	●	●
Management and Monitoring GUI	●	●	●	●
Command line Management Tool	●	●	●	●
Standards				
JCache (JSR-107)	●	●	●	●
SQL (ANSI-99)	●	●	●	●
ODBC	●	●	●	●
JDBC	●	●	●	●
XA/JTA	●	●	●	●
OSGI	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Out-of-the-Box Integration				
Automatic RDBMS integration	●	●	●	●
Spring Framework	●	●	●	●
Apache® Maven™	●	●	●	●
Web Session Clustering	●	●	●	●
Hibernate L2 Cache	●	●	●	●
MyBatis L2 Cache	●	●	●	●
Vert.x	●	●	●	●
JMS	●	●	●	●
Apache® Flume™	●	●	●	●
MQTT	●	●	●	●
Twitter	●	●	●	●
Apache® Kafka™	●	●	●	●
Apache® Camel™	●	●	●	●
Apache® Storm™	●	●	●	●
Spring Caching	●	●	●	●
Oracle® Golden Gate	●	●	●	●
Cloud and Virtualization Support				
TCP/IP Cluster Protocol	●	●	●	●
Pluggable Discovery	●	●	●	●
Amazon® Web Services	● (S3-Based IP Finder)	● (S3-Based IP Finder)	● (S3-Based IP Finder)	●
Google® Compute	●	●	●	●
Microsoft Azure	●	●	●	●
Apache® JClouds™	●	●	●	●
Docker Container	●	●	●	●
Kubernetes	●	●	●	●
In-Memory Streaming				
Data Streamers	●	●	●	●
Complex Event Processing (CEP)	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Distributed Messaging and Events				
Topic-based Publish/Subscribe Messaging	● (Ordered, Unordered)	● (Ordered, Unordered)	● (Ordered, Unordered)	●
Point-to-Point Messaging	●	●	●	●
Grid Event Notifications	●	●	●	●
Automatic Batching of Event Notifications	●	●	●	●
Distributed Computing				
Affinity-Aware Execution	●	●	●	●
Executor Service	●	●	●	●
Managed Services	●	●	●	●
Sub-Grid Messaging / Task Execution	●	●	●	●
Zero Deployment Technology	●	●	●	●
Direct API for MapReduce and ForkJoin	●	●	●	● (MapReduce only)
Early and Late Load Balancing	●	●	●	●
Fault-Tolerance	●	●	●	●
Computation State Checkpoints	●	●	●	●
Distributed Computation (Task) Sessions	●	●	●	●
Cron-like Task Scheduling	●	●	●	●
Security and Audit				
SSL Support	●	●	●	●
Client Authentication	●	●	●	●
Cluster Member Authentication	●	●	●	●
ACL-Based Passcode Authentication	●	●	●	●
JAAS Authentication	●	●	●	●
Authorization and Permit	●	●	●	●
Audit (Trace Events)	●	●	●	●
Multi-Tenancy	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GEMFIRE 9.7
Data Visualizations				
Hosted Web Console	●	●	●	●
On-Premises Web Console	●	●	●	●
Apache® Zeppelin™	●	●	●	●
Tableau®	●	●	●	●
Client-Server Protocol				
Memcached Support	●	●	●	●
HTTP REST	●	●	●	●
Supported Platforms				
Java & JVM-based Platforms	●	●	●	●
C++ Client	●	●	●	●
.NET/C# Client	●	●	●	●
Scala DSL	●	●	●	●
Node.JS Client	●	●	●	●
Interoperability between .NET/Java/C++	●	●	●	● (Limited support)
Integration with Spark				
Implementation of Spark RDD and DataFrame	●	●	●	● (Can execute OQL and return a DataFrame. But does not make GemFire a DataFrame provider)
Native SQL optimization	●	●	●	●
Deployment				
Apache® Mesos™	●	●	●	●
Hadoop® Yarn	●	●	●	●
Apache® BigTop™	●	●	●	●

Additional Product Comparisons

You can also learn how GridGain compares to other in-memory solutions, including Redis®, Hazelcast®, Pivotal GemFire®, Oracle Coherence® and GigaSpaces® by visiting us at www.gridgain.com.

Contact GridGain Systems

To learn more about the GridGain in-memory computing platform, 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 fill out our contact page at 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.