

# GRIDGAIN IN-MEMORY COMPUTING PLATFORM FEATURE COMPARISON: GigaSpaces®

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

## Major Advantages of GridGain vs. GigaSpaces XAP

- Comprehensive in-memory computing platform
- Native ANSI-99 SQL support
- Native 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 (IMDB) with better scalability, availability and reliability
- Support for Apache® Spark™ DataFrames, RDDs, HDFS and SparkSQL acceleration
- Built-in Continuous Learning Framework with support for machine learning and deep learning
- Built on Apache® Ignite™, a top 5 Apache Software Foundation project based on commits and list activity

## GIGASPACEs AND GRIDGAIN IN-MEMORY COMPUTING PLATFORM COMPARED

The GigaSpaces XAP in-memory data grid (IMDG) is used by companies for fast data processing to scale applications. GigaSpaces InsightEdge® bundles XAP with Apache® Spark™ for fast data analytics. Both are available as open source under an Apache 2.0 license (announced 2016) and as commercially supported software.

GridGain, built on the [Apache Ignite™ open source project](#), is an [in-memory computing platform](#) that includes a distributed 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 database or Hadoop database.

GigaSpaces has many of the core capabilities expected in an IMDG, including the ability to distribute and partition data, and scale out across a cluster. But like some other IMDGs, GigaSpaces hasn't evolved much in the last five years. The company has instead chosen to focus more on InsightEdge and fast data analytics, as well as Cloudify, which it recently spun off.

GridGain is better than GigaSpaces as an 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. 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

One of the key innovations is GridGain's ANSI-99 SQL support. 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. GigaSpaces has no way to provide easy access to data for all the tools and applications that rely on SQL. It also requires extensive custom coding in applications to replace existing SQL.

GridGain is also 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. GigaSpaces is a strong IMDG. But it is just an IMDG.

As a database, GridGain provides better scalability, availability and reliability. GigaSpaces is not intended for use as a standalone 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. GigaSpaces only provides backup servers in the clusters. If the servers fail, the data is lost. If you want to use XAP for other use cases such as streaming or machine learning you need to use GigaSpaces InsightEdge with an older version of XAP. This results in using multiple versions and technologies for in-memory computing. In addition, machine and deep learning is not built in. It is provided by Spark and requires data movement over the network that can lead to significant delays.

## GIGASPACEs AND GRIDGAIN EDITIONS

To compare GridGain XAP to Apache Ignite and GridGain, it is important to understand the different GigaSpaces and GridGain editions. GigaSpaces offers XAP Open Source under the Apache 2.0 license; XAP Premium that adds security, deployment, monitoring, management, high availability and cross-language support; and XAP Enterprise that adds MemoryXtend and multi-data center replication. There is also GigaSpaces InsightEdge, a relatively new in-memory insight platform for fast-data analytics and transactional processing that bundles Apache Spark with an older version of GigaSpaces XAP. InsightEdge is also offered either as open source under the Apache 2.0 license or as the InsightEdge Platform under a commercial license.

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 GigaSpaces Enterprise are included unless otherwise noted.

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GIGASPACEs XAP ENTERPRISE 14.0
Native ANSI-99 SQL Support	●	●	●	●
Distributed ACID Transaction Support	●	●	●	●
Slide in-between SQL-based Applications and RDBMSs with support for SQL	●	●	●	● (Requires code for new apps, adding code and data model for 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 as part of DataInsight)
Comprehensive In-Memory Computing Solution	● (IMDG, streaming, machine and deep learning)	● (IMDG, streaming, machine and deep learning)	● (EE + Multi-datacenter data and disaster recovery management)	● (Supports IMDG. Limited IMDB. Streaming part of DataInsight. Deep learning is provided by Spark)
Apache Spark Support for DataFrames, RDDs, HDFS	●	●	●	● (RDD, DataFrame, DataSet support. No HDFS acceleration)
Built-in Machine Learning	●	●	●	● (Machine/deep learning in Spark, not in-place)
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)	● (Only Apache 2.0 license. Not ASF and no big developer ecosystem)

This document provides a detailed feature comparison between the various levels of GridGain and GigaSpaces XAP, 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. GigaSpaces' strengths are as an IMDG. GigaSpaces is not intended for these other use cases. Streaming, machine and deep learning requires the use of GigaSpaces DataInsight, which bundles Apache Spark with an older version of GigaSpaces XAP. Because of this you can end up with multiple technologies, and multiple versions of XAP deployed to support different projects. This ends up fragmenting data and adding developer and administrative costs.

## 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 for Java, .NET, C++, Python and other languages. GigaSpaces does not support SQL. In order to integrate an application with GigaSpaces, 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 GigaSpaces.

## NATIVE 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 Cassandra® or MongoDB®, as well as Hadoop via Spark and HDFS. GigaSpaces requires writing custom code to integrate with third-party databases. It also does not support Cassandra out of the box.

## 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. GigaSpaces XAP 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 GigaSpaces, and more code to enable GigaSpaces to query and write to the underlying database.

## DISTRIBUTED IN-MEMORY DATABASE (IMDB) WITH BETTER 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 scalability, availability and reliability. 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 0-100% 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. With GigaSpaces, indexes can only be stored on-heap, not

off-heap or on disk, 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. GigaSpaces 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 redeploy an existing cluster to any other location on premise or in the cloud, with a different number of nodes and amount of memory. GigaSpaces lacks the same level of backup and recovery. Its only backup are backup servers in the clusters. If the servers fail, the data is lost.

## SUPPORT FOR APACHE® SPARK™ 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. It 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 distributed joins, and GridGain's machine and deep learning capabilities. GigaSpaces provides strong integration with Spark. However, GigaSpaces 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. GigaSpaces does not support built-in MPP-style machine and deep learning that runs in place. With GigaSpaces you need to move the data by streaming it through Spark.

## BUILT ON APACHE® IGNITE™, A TOP 5 APACHE SOFTWARE FOUNDATION PROJECT BASED ON COMMITS AND LIST ACTIVITY

GridGain is built on Apache Ignite, one of the top five Apache Software Foundation projects in terms of commits and community activity. GigaSpaces XAP is open source with an Apache 2.0 license. There is no broad developer ecosystem and it is not an Apache Software Foundation top level project.

## GRIDGAIN AND GIGASPACE'S DETAILED FEATURE COMPARISON

The following table provides a detailed feature comparison between the GridGain Community, Enterprise, and Ultimate Editions, and GigaSpaces XAP Enterprise. 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	GIGASPACE'S XAP ENTERPRISE 14.0
<b>Use Cases</b>				
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	●	●	●	● (No HDFS, SparkSQL acceleration)
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	GIGASPACE XAP ENTERPRISE 14.0
<b>Third Party Database Support, Persistence</b>				
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)
Automatic integration with 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)	●	●	●	●
<b>Native Persistence</b>				
Native Persistence	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (XAP Enterprise only)
Stores Superset of Data	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (Can cache to disk, use disk for overflow)
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 and limited SQL)
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	GIGASPACE XAP ENTERPRISE 14.0
<b>Distributed SQL</b>				
SQL Queries	● (Full ANSI-99 Support)	● (Full ANSI-99 Support)	● (Full ANSI-99 Support)	●
Collocated Distributed Joins	●	●	●	● (Limited SQL)
Non-Collocated Distributed Joins	●	●	●	●
Single Column Indexes	●	●	●	●
Group Indexes	●	●	●	●
Distributed SQL Joins (select * from Person p, Company c where p.c_ id=c.id)	●	●	●	● (With lightweight query language, not SQL)
Query Consistency	●	●	●	●
Query Fault-Tolerance	●	●	●	●
DML (INSERT, UPDATE, DELETE, MERGE)	●	●	●	●
DDL (CREATE, DROP, ALTER)	●	●	●	●
<b>Distributed Queries</b>				
Continuous Queries	●	●	●	●
Predicate-based Queries	●	●	●	●
<b>SQL Drivers</b>				
JDBC Driver	●	●	●	●
ODBC Driver	●	●	●	●
REST API (SQL)	●	●	●	●



FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GIGASPACE XAP ENTERPRISE 14.0
<b>Memory Architecture</b>				
On-Heap Memory	●	●	●	●
Off-Heap Memory	●	●	●	● (XAP Enterprise only)
Off-Heap Indexes	●	●	●	●
Disk as main storage (disk larger than RAM)	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (XAP Enterprise only)
Tiered Storage - On-Heap, Off-Heap and Disk	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (XAP Enterprise only)
<b>ACID Compliant Transactions and Locks</b>				
Atomic Mode (One Operation at a Time)	●	●	●	●
READ_COMMITTED, REPEATABLE_READ, SERIALIZABLE Isolation Levels	●	●	●	●
Deadlock-Free Transactions	●	●	●	●
XA Integration	●	●	●	●
Fault Tolerance (Including Client/Near/ Primary/Backup Node Failures)	●	●	●	●
Optimistic & Pessimistic Concurrency (Two- Phase- Commit)	●	●	●	●
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.)	●	●	●	● (Through JTA and external Transaction Manager)
Explicit Locking	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GIGASPACE XAP ENTERPRISE 14.0
<b>Distributed Architecture</b>				
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, Custom)
Binary Objects	●	●	●	● (Very hard to use)
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	GIGASPACE XAP ENTERPRISE 14.0
<b>Distributed Data Structures</b>				
Queue	●	●	●	●
Set	●	●	●	●
Atomic Log	●	●	●	●
Atomic Ref	●	●	●	●
Atomic Stamped Ref	●	●	●	●
Atomic Sequence	●	●	●	●
Count Down Latch	●	●	●	●
Reentrant Lock	●	●	●	●
Semaphore	●	●	●	●
<b>Data Snapshots (Backups)</b>				
Full Data Snapshots	●	●	● (+ Multi-datacenter data and disaster recovery management)	●
Incremental Data Snapshots	●	●	● (+ 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)	●
<b>Datacenter (WAN) Replication</b>				
Active-Active	●	●	●	●
Active-Passive	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GIGASPACE XAP ENTERPRISE 14.0
<b>Data Rebalancing</b>				
Sync Data Rebalancing (aka Sync Repartitioning)	●	●	●	●
Async Data Rebalancing (aka Async Repartitioning)	●	●	●	●
Delayed Data Rebalancing (Delay Data Rebalancing until All Nodes Have Started)	●	●	●	●
<b>Grid Management and Monitoring</b>				
Rolling Production Updates	●	●	●	●
Management and Monitoring GUI	●	●	●	●
Command line Management Tool	●	●	●	●
<b>Standards</b>				
JCache (JSR-107)	●	●	●	●
SQL (ANSI-99)	●	●	●	● (Limited)
ODBC	●	●	●	● (Use JDBC-ODBC bridge)
JDBC	●	●	●	● (Read only, limited functionality)
XA/JTA	●	●	●	●
OSGI	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GIGASPACE XAP ENTERPRISE 14.0
<b>Out-of-the-Box Integration</b>				
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	●	●	●	●
<b>Cloud and Virtualization Support</b>				
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	●	●	●	●
<b>In-Memory Streaming</b>				
Data Streamers	●	●	●	●
Complex Event Processing (CEP)	●	●	●	●

FEATURE	GRIDGAIN CE 8.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.7	GRIDGAIN UE 8.7	GIGASPACE XAP ENTERPRISE 14.0
<b>Distributed Messaging and Events</b>				
Topic-based Publish/Subscribe Messaging	● (Ordered, Unordered)	● (Ordered, Unordered)	● (Ordered, Unordered)	● (JMS)
Point-to-Point Messaging	●	●	●	● (JMS)
Grid Event Notifications	●	●	●	●
Automatic Batching of Event Notifications	●	●	●	●
<b>Distributed Computing</b>				
Affinity-Aware Execution	●	●	●	●
Executor Service	●	●	●	●
Managed Services	●	●	●	●
Sub-Grid Messaging / Task Execution	●	●	●	●
Zero Deployment Technology	●	●	●	●
Direct API for MapReduce and ForkJoin	●	●	●	●
Early and Late Load Balancing	●	●	●	●
Fault-Tolerance	●	●	●	●
Computation State Checkpoints	●	●	●	●
Distributed Computation (Task) Sessions	●	●	●	●
Cron-like Task Scheduling	●	●	●	●
<b>Security and Audit</b>				
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	GIGASPACE XAP ENTERPRISE 14.0
<b>Data Visualizations</b>				
Hosted Web Console	●	●	●	●
On-Premises Web Console	●	●	●	●
Apache® Zeppelin™	●	●	●	●
Tableau®	●	●	●	●
<b>Client-Server Protocol</b>				
Memcached Support	●	●	●	●
HTTP REST	●	●	●	●
<b>Supported Platforms</b>				
Java and JVM-based Platforms	●	●	●	●
C++ Client	●	●	●	●
.NET/C# Client	●	●	●	●
Scala DSL	●	●	●	●
Node.JS Client	●	●	●	●
Interoperability between .NET/Java/C++	●	●	●	●
<b>Integration with Spark</b>				
Implementation of Spark RDD and DataFrame	●	●	●	●
Native SQL optimization	●	●	●	●
<b>Deployment</b>				
Apache® Mesos™	●	●	●	●
Hadoop® Yarn	●	●	●	●
Apache® BigTop™	●	●	●	●

## Additional Product Comparisons

You can also learn how GridGain compares to other in-memory solutions, including Redis®, Hazelcast®, Terracotta®, Coherence® and Pivotal GemFire® by visiting [www.gridgain.com](http://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](mailto:sales@gridgain.com) or call us at +1 (650) 241-2281 (US) or +44 (0) 208 610 0666 (Europe).

### 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](http://www.gridgain.com).