GridGain

Built On Apache Ignite

FEATURE COMPARISON

GRIDGAIN IN-MEMORY COMPUTING PLATFORM FEATURE COMPARISON: Oracle® Coherence™

This document presents a summary and detailed feature comparison of the GridGain® in-memory computing platform (GridGain) and Oracle® Coherence (Coherence) for use as in-memory data grids (IMDG).

Major Advantages of GridGain vs. Oracle Coherence

- ANSI-99 SQL Support
- ACID Transaction Support
- Slides In-between SQL-based Applications and RDBMSs, eliminating the need to replace SQL with code
- Cross-Language Support for Massively Parallel Processing (MPP) for Java, .NET and C++
- Native Integration with RDBMSs, NoSQL Databases and Hadoop
- Comprehensive In-Memory Computing Platform with IMDG, in-memory database (IMDB) and streaming
- Support for Apache® Spark™ DataFrames, RDDs and HDFS
- Built-in Machine Learning and Deep Learning
- Built on Apache Ignite, a Leading Open Source Project

ORACLE COHERENCE AND GRIDGAIN IN-MEMORY COMPUTING PLATFORM COMPARED

Oracle Coherence is a proprietary an in-memory data grid that is used by companies to scale applications built on an Oracle database. Oracle acquired Tangasol, the company that built Coherence, in 2007 and renamed it Oracle Coherence. Oracle maintained a separate Coherence product, last released as 3.7.1 December 2016. The more up-to-date version is Oracle Coherence 12cR2, branded with the Oracle database and currently at 12.2.1.3.

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), and a stream processing and analytics engine. It can be used with any RDBMS, NoSQL or Hadoop database. GridGain Systems donated the original code to the Apache Ignite project and is the largest contributor.

Oracle Coherence 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 as an IMDG, the Coherence technology is at the same state as many IMDGs were 5 to 10 years ago. Standalone Oracle Coherence 3.7 was released in 2011 and is currently at Release 3.7.1, released December 2016. Oracle Coherence 12cR2 is more up-to-date. 12.2.1 was released Apri 2016, and 12.2.1.3 was released April 2018. But not much has been added since 2011 while most leading vendors have significantly evolved their products in the interim.

GridGain is better than Coherence in almost every IMDG use case. This is due, in part to the innovation that comes from being built on Apache Ignite, one of the top five Apache Software Foundation open source projects. GridGain's ANSI-99 SQL support allows it to slide in-between existing SQL-based applications and RDBMSs and use SQL. Oracle Coherence requires extensive custom coding to replace existing SQL. GridGain has proven high performance support for ACID transactions. Performing transactions with Oracle Coherence is not recommended. GridGain has stronger out-of-the-box support for RDBMSs, NoSQL databases, Spark and HDFS. GridGain supports collocated processing (also known as "massively parallel processing") for Java, .NET and C++. GridGain is also better for every other non-IMDG in-memory computing use case because it is an in-memory computing platform that includes an IMDG, in-memory database (IMDB), streaming analytics, machine learning and deep learning. This includes persistence that enables immediate availability on cluster restarts before memory is loaded. Oracle Coherence loses in-memory data and requires all memory to be loaded first before han-

dling any requests. Oracle Coherence offers none of these additional capabilities and can only be used as an IMDG. Coherence is proprietary software with a more expensive licensing model and far fewer resources focused on future innovation.

Oracle Coherence also has a more traditional, and more expensive licensing model. There is no open source version of Oracle Coherence or a community that helps drive innovation. GridGain is built on Apache Ignite, which is one of the top five Apache Software Foundation projects in terms of commits and community activity.

ORACLE COHERENCE AND GRIDGAIN EDITIONS

To compare Oracle Coherence to Apache Ignite and GridGain, it is important to understand the different GridGain editions. 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 inmemory 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 Oracle Coherence Grid Edition are included unless noted otherwise.

FEATURE	GRIDGAIN CE 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
Native ANSI-99 SQL Support	●	●	●	●
Distributed ACID Transaction Support	●	●	●	● (No explicit locking, performance issues)
Slide in-between SQL-based Applications and RDBMSs with 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++)	●
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. No Spark or HDFS support)
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)	● (IMDG only. No IMDB, streaming, machine or deep learning)
Apache Spark Support for DataFrames, RDDs, HDFS	●	●	●	●
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)	● (No open source project or external development community)

This document provides a detailed feature comparison between the various levels of GridGain and Oracle Coherence, highlighting differences between the open source and enterprise-ready products where relevant. But here are the major differences between the two products which should be considered when evaluating either one for any in-memory computing need, as summarized in the above table:

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 existing application by providing ODBC and JDBC drivers applications to use in place of their existing drivers for Java, .NET, C++, Python and other languages. Oracle Coherence does not support SQL. It requires developers to change existing applications, and to add code for each table and SQL query to the underlying database.

ACID TRANSACTIONS

GridGain has full support for ACID transactions, including OPTIMISTIC and PESSIMISTIC concurrency modes, as well as READ_COMMITTED, REPEATABLE_READ, and SERIALIZABLE isolation levels. Oracle Coherence provides some transaction support, but it is not widely recommended to use transactions due to performance reasons as well as a lack of key features such as explicit locking.

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 Cassandra® or MongoDB®, as well as Hadoop via Spark and HDFS. Oracle Coherence always requires writing custom code to integrate with third-party databases.

SLIDES IN-BETWEEN SQL-BASED APPLICATIONS AND RDBMSS WITH FULL SUPPORT FOR SQL

GridGain's architecture and out-of-the-box integration enables many use cases to be implemented through configuration, not coding. For example, GridGain can slide in between and accelerate SQL-based applications and third-party databases by replacing existing application ODBC or JDBC drivers with GridGain's native drivers, configuring GridGain to connect to the backend database, and supporting ANSI-99 SQL in the application. Oracle Coherence requires 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 Coherence using key-values, and more code that enables Coherence to query and write to the underlying database.

CROSS-LANGUAGE SUPPORT FOR COLLOCATED PROCESSING (MPP)

Both GridGain and Oracle Coherence support multiple languages for developing clients. But Oracle Coherence has no support for collocated computing. GridGain provides general purpose massively parallel processing (MPP) that is used for distributed SQL and machine and deep learning. It also supports user-defined Java, .NET and C++ code. This makes GridGain the only choice for any data intensive applications that cannot wait for data to travel across the network to a client. GridGain is also better suited for lightweight non-Java based clients because it includes a binary protocol that enables client support without requiring a JVM to be deployed with a non-Java client.

COMPREHENSIVE IN-MEMORY COMPUTING SOLUTION

GridGain supports third-party databases as an IMDG and also includes a distributed SQL and key-value hybrid in-memory database that combines memory-centric storage with native. GridGain's 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 and RAM can hold 0-100% of the data and indexes. If a subset of data or an index is not in RAM, it will be used from non-volatile storage. Data in RAM and nonvolatile 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 GridGain, data is immediately available on a cluster startup. 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. When used as a database, GridGain also includes centralized backup and recovery, full and incremental snapshots, continuous archiving and point-in-time recovery, network backups, and heterogeneous recovery that enables a cluster to be restarted in a different location on-premise or in the cloud with a different size. Oracle Coherence cannot be used as a database. Its persistence is only used to back up the in-memory data grid for hot restarts, and requires the data to be loaded back into the cache before becoming operational.

APACHE SPARK™ SUPPORT FOR DATAFRAMES, RDDS AND HDFS

GridGain provides the broadest integration with Spark over any other in-memory computing vendors. It provides native Apache 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 accelerate 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, machine and deep learning capabilities. Oracle Coherence has no support for Apache Spark.

SUPPORT FOR MACHINE LEARNING AND DEEP LEARNING

GridGain includes the GridGain® Continuous Learning Framework, 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 perceptron and tensorflow integration for for deep learning. Developers can develop and deploy their own algorithms across any cluster as well by using the compute grid. Oracle Coherence has no integrated machine or deep learning.

BUILT ON APACHE IGNITE, A LEADING OPEN SOURCE PROJECT

GridGain is built on Apache Ignite, one of the top five Apache Software Foundation projects in terms of commits and community activity. Oracle Coherence has a more traditional and more expensive licensing model. There is no open source version of Oracle Coherence and no community to help drive innovation.

GRIDGAIN AND ORACLE COHERENCE DETAILED FEATURE COMPARISON

The following table provides a detailed feature comparison between the GridGain Community, Enterprise, and Ultimate Editions, and Coherence. 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 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
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	●	●	●	●
Hadoop acceleration	●	●	●	●
In-Memory File System (Hadoop Compliant)	●	●	●	●

FEATURE	GRIDGAIN CE 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
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)
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)	●	●	●	●
Native Persistence				
Native Persistence	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (Enterprise edition 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)	●
Instantaneous Restart (before memory warmup)	●	●	● (+ Multi-datacenter data and disaster recovery management)	●

FEATURE	GRIDGAIN CE 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
Distributed SQL				
SQL Queries	● (Full ANSI-99 Support)	● (Full ANSI-99 Support)	● (Full ANSI-99 Support)	●
Collocated Distributed Joins	●	●	●	●
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)	●	●	●	●
Distributed Queries				
Continuous Queries	●	●	●	●
Predicate-based Queries	●	●	●	●
SQL Drivers				
JDBC Driver	●	●	●	●
ODBC Driver	●	●	●	●
REST API (SQL)	●	●	●	●

FEATURE	GRIDGAIN CE 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
Memory Architecture				
On-Heap Memory	●	●	●	●
Off-Heap Memory	●	●	●	● (with backing maps)
Off-Heap Indexes	●	●	●	●
Disk as main storage (disk larger than RAM)	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (Only persistence for backup)
Tiered Storage - On-Heap, Off-Heap and Disk	●	●	● (+ Multi-datacenter data and disaster recovery management)	● (Disk is for backup only)
ACID Compliant Transactions and Locks				
Atomic Mode (One Operation at a Time)	●	●	●	●
READ_COMMITTED, REPEATABLE_READ, SERIALIZABLE Isolation Levels	●	●	●	● (Transactions are slow and not recommended)
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	●	●	●	● (Not recommended)
Transactional Entry Processor	●	●	●	●
Eviction / Expiration Policies for Transactional Caches	●	●	●	●
Merge with DB Transac-tions (e.g., Oracle DB, MySQL, etc.)	●	●	●	● (only through external Transaction Manager)
Explicit Locking	●	●	●	●

FEATURE	GRIDGAIN CE 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
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)	●	●	●	●
Deadlock-Free Transactions	●	●	●	●
Synchronous and Asynchronous Backup Update	●	●	●	● (Asynchronous Only)
Synchronous APIs	●	●	●	●
Asynchronous APIs	●	●	●	●
Full Sync Mode (Primary and Backups are Sync)	●	●	●	●
Primary Sync Mode Primary is synch, 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, LFU, Hybrid, 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 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
Distributed Data Structures				
Queue	●	●	●	● (Coherence Incubator)
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)	●
Data Recovery rom 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 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
Data Rebalancing				
Sync Data Rebalancing (aka Sync Repartitioning)	●	●	●	●
Async Data Rebalancing (aka Async Repartition-ing)	●	●	●	●
Delayed Data Rebalancing (Delay Data Rebalancing until All Nodes Have Started)	●	●	●	●
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 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
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 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
Distributed Messaging and Events				
Topic-based Publish/Subscribe Messaging	● (Ordered, Unordered)	● (Ordered, Unordered)	● (Ordered, Unordered)	● (Coherence Incubator – Ordered only)
Point-to-Point Messaging	●	●	●	●
Grid Event Notifications	●	●	●	●
Automatic Batching of Event Notifications	●	●	●	●
Distributed Computing				
Affinity-Aware Execution	●	●	●	● (Via Entry Processor)
Executor Service	●	●	●	● (Via Invocation 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	●	●	●	●
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 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
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++	●	●	●	●
Integration with Spark				
Implementation of Spark RDD and DataFrame	●	●	●	●
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®, Terracotta®, GigaSpaces® and Pivotal GemFire® by visiting the GridGain website.

Contact GridGain Systems

To learn more about the GridGain in-memory computing platform, please email our sales team at 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 gridgain.com.