



# GRIDGAIN IN-MEMORY DATA FABRIC

Feature Comparison: Oracle® Coherence™

July, 2016

This document presents a feature comparison of the [GridGain In-Memory Data Fabric](#) (GridGain) and Oracle® Coherence™ (Coherence) for in-memory computing applications.

## About the GridGain In-Memory Data Fabric

GridGain, built on [Apache® Ignite™](#), is an [in-memory computing platform](#) that enables you to dramatically accelerate and scale out your existing data-intensive applications without ripping and replacing your existing databases. GridGain can reduce query times by 1,000x versus disk-based systems. You can scale out by adding new nodes to your GridGain cluster, which can handle hundreds of terabytes of data from multiple databases.

You can modernize your existing data-intensive architecture by inserting GridGain between your existing application and database layers. GridGain integrates seamlessly with RDBMS, NoSQL and Apache® Hadoop™ databases. It features a unified API which supports SQL, C++, .NET, PHP, MapReduce, JAVA/Scala/Groovy, and Node.js protocols for the application layer. Your GridGain cluster, applications, and databases can run on premise, in a hybrid environment, or on a cloud platform such as AWS® or Microsoft Azure.

GridGain is available as a Professional Edition and an Enterprise Edition. The [GridGain Professional Edition](#) includes the current version of Apache Ignite plus bug fixes that are not yet released in Ignite. The [GridGain Enterprise Edition](#) includes enterprise-grade security, deployment, and management features which have been added to the core features of Apache Ignite. A subscription to the GridGain Professional Edition is part of a GridGain Standard Support package and a subscription to the GridGain Enterprise Edition is included in GridGain Enterprise Support packages.

## Key Differences Between GridGain and Coherence

The GridGain In-Memory Data Fabric includes an in-memory data grid feature. Coherence is an in-memory data grid. The data grid capabilities of both products include functionality which partitions and caches data in memory. Both of the data grid solutions can be scaled out across distributed clusters. However, there are many differences in the way caching, transactions, and data querying are supported. The GridGain in-memory computing platform also includes many additional features not included in Coherence that are often highly valuable for companies that are moving to in-memory computing.

The following are major differences between the two products which should be considered when choosing an in-memory solution:

- **Open Source vs. Proprietary/Closed Source** – The GridGain In-Memory Data Fabric is built on Apache Ignite, an open source Apache Software Foundation project originally contributed by GridGain. Its open source feature set is available to a robust, global community and is continuously enhanced and updated through feedback and contributions. Oracle Coherence is a proprietary solution.
- **Memory Formats** – GridGain supports storing data in on-heap or off-heap memory, depending on the configuration. It supports storing query indexes off-heap as well. Oracle Coherence provides a limited option to store the data in off-heap memory and does not support off-heap indexes.

- **SQL Queries** – GridGain supports complete SQL (ANSI-99) syntax, including distributed SQL JOINS ([collocated](#) and [non-collocated](#)) for querying in-memory data. Oracle Coherence does not support SQL and users have to perform JOINS manually by combining multiple query results.
- **ACID Transactions** – GridGain can process hundreds of thousands of transactions per second per server. It is not recommended using transactions in Coherence due to poor performance.
- **Data Streaming** – GridGain provides support for in-memory streaming, including maintaining and querying sliding windows of streaming data. Coherence does not offer support for streaming.

## GridGain and Coherence Detailed Feature Comparison

The following table provides a detailed feature comparison between the GridGain Professional Edition, GridGain Enterprise Edition, and Oracle 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 PE 1.7	GridGain EE 7.6	Oracle Coherence 12.2.1
<b>Distributed Caching</b>			
Key-Value Store	●	●	●
Partitioning and Replication	●	●	●
Client-side (Near) Cache	●	●	● (No Transactional Capabilities)
Dynamic Cache Creation	●	●	●
EntryProcessor, aka Delta (Partial) Updates	●	●	●
Data Redundancy (Key Backups)	●	●	●
Synchronous and Asynchronous Backup Update	●	●	●
Synchronous APIs	●	●	●
Asynchronous APIs	●	●	●
Full Async Mode (Primary and Backups are Async)	●	●	●
Data Affinity and Collocation	● (Rich Support)	● (Rich Support)	●
Data Eviction and Expiration	● (LRU, FIFO, Random, Sorted, Custom)	● (LRU, FIFO, Random, Sorted, Custom)	● (LRU, LFU, Hybrid, Custom)
Binary Objects	●	●	● (Not Used to Use)
Pluggable Interfaces (SPIs) to Customize Grid Subsystems	●	●	●
<b>Memory Formats</b>			
On-Heap Memory	●	●	●

Feature	GridGain PE 1.7	GridGain EE 7.6	Oracle Coherence 12.2.1
Off-Heap Memory	●	●	● (Only with Backing Maps)
Off-Heap Indexes for Off-Heap data	●	●	●
Disk Overflow	●	●	●
Tiered Storage - On-Heap to Off-Heap to Disk	●	●	●
<b>ACID Compliant Transactions and Locks</b>			
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	●	●	●
Custom Affinity (Partitioning) Function	●	●	●
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 Transactions (e.g., Oracle DB, MySQL, etc.)	●	●	● (Only through External Transaction Manager)
Explicit Locking	●	●	●
<b>Distributed Data Structures</b>			
Queue	●	●	● (Coherence Incubator)
Set	●	●	●
Atomic Long	●	●	●
Atomic Ref	●	●	●
Atomic Stamped Ref	●	●	●
Atomic Sequence	●	●	●
Count Down Latch	●	●	●

Feature	GridGain PE 1.7	GridGain EE 7.6	Oracle Coherence 12.2.1
ReentrantLock	●	●	●
Semaphore	●	●	●
<b>Distributed Queries (Searches)</b>			
SQL Queries	● (Full ANSI-99 Support)	● (Full ANSI-99 Support)	●
Continuous Queries	●	●	●
Predicate-based Queries	●	●	●
Single Column Indexes	●	●	●
Group Indexes	●	●	●
Affinity Collocation Based SQL Joins	●	●	●
<a href="#">Non-Collocated Distributed SQL Joins</a>	●	● (Will be available in 7.7)	●
Query Consistency	●	●	●
Query Fault-Tolerance	●	●	●
Custom Query API	●	●	● (Custom Coherence Query Language)
<b>SQL Drivers</b>			
JDBC Driver	●	●	●
ODBC Driver	●	●	●
<b>Persistence and Data Loading</b>			
Write-Through and Read-Through	●	●	●
Write-Behind Caching	●	●	●
Auto-Loading of SQL Schema/Data	●	●	●
Local Recoverable Store	●	●	●
Data Streamer (Optimized Bulk Put or Load Operations)	●	●	●
Store Loader (Optimized Bulk DB Load)	●	●	●
<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>Standards</b>			
JCache (JSR-107)	●	●	●

Feature	GridGain PE 1.7	GridGain EE 7.6	Oracle Coherence 12.2.1
SQL (ANSI-99)	●	●	●
ODBC	●	●	●
JDBC	●	●	●
XA/JTA	●	●	●
OSGI	●	●	●
<b>Integrations</b>			
Automatic RDBMS integration	●	●	●
Spring Framework	●	●	● (Very Complicated API)
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)	●
Google® Compute	● (Implemented Using Google Compute Engine Storage Based IP Finder)	● (Implemented Using Google Compute Engine Storage Based IP Finder)	●
Apache® JClouds™	●	●	●



Feature	GridGain PE 1.7	GridGain EE 7.6	Oracle Coherence 12.2.1
Docker Container	●	●	●
Apache® Mesos™	●	●	●
Hadoop Yarn	●	●	●
<b>In-Memory Streaming</b>			
Data Streamers	●	●	●
Complex Event Processing (CEP)	●	●	●
Sliding Data Windows	●	●	●
Continuous Queries over Data Windows	●	●	●
<b>Distributed Messaging and Events</b>			
Topic-based Publish/Subscribe Messaging	● (Ordered and Unordered)	● (Ordered and Unordered)	● (Coherence Incubator-Ordered Only)
Point-to-Point Messaging	●	●	●
Grid Event Notifications	●	●	●
Automatic Batching of Event Notifications	●	●	●
<b>Distributed Compute</b>			
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	●	●	●
<b>Configuration &amp; Grid Management</b>			
Spring XML Configuration	●	●	● (Has Proprietary XML Syntax)
Programmatic Configuration	●	●	●
Elasticity (Ability to Add/Remove Grid Nodes on Demand)	●	●	●

Feature	GridGain PE 1.7	GridGain EE 7.6	Oracle Coherence 12.2.1
Dynamic Schema Changes (Allowing Dynamic Change to an Object's Structure)	● (Portable Objects)	● (Portable Objects)	●
Datacenter (WAN) Replication (Active-Active, Active-Passive)	●	●	●
Rolling Production Updates	●	●	●
Network Segmentation (Split Brain)	●	● (Transactional)	● (Non-Transactional)
Data Conflicts Resolution	●	●	●
Management and Monitoring GUI	●	●	●
Command-Line Management Tool	●	●	●
<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	●	●	●
<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 & JVM-based Platforms	●	●	●
C++	●	●	●
.NET/C#	●	●	●
Scala DSL	●	●	●
Interoperability between .NET/Java/C++	●	●	●
<b>Integration with Spark</b>			
Implementation of Spark RDD	●	●	●



Feature	GridGain PE 1.7	GridGain EE 7.6	Oracle Coherence 12.2.1
SQL Queries	●	●	●
<b>Hadoop Deployment</b>			
Apache® Mesos™	●	●	●
Hadoop® Yarn	●	●	●
Apache® BigTop™	●	●	●
<b>Hadoop Acceleration</b>			
Hadoop Accelerator	●	●	●
In-Memory File System (Hadoop Compliant)	●	●	●

## Additional Product Comparisons

You can also learn how GridGain compares to other in-memory solutions, including Hazelcast®, Pivotal GemFire®, Terracotta®, Redis®, and GigaSpaces® by visiting the [Product Comparisons page](#) on our website.

## Contact GridGain

To learn more about the GridGain In-Memory Data Fabric, 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)7775 835 770 (Europe).

### ABOUT GRIDGAIN

GridGain is revolutionizing real-time data access and processing by offering the enterprise-grade GridGain In-Memory Data Fabric built on Apache Ignite™. The solution is used by global enterprises in financial, tech, retail, healthcare and other major sectors. GridGain solutions connect traditional and emerging data stores (SQL, NoSQL, and Hadoop) with cloud-scale applications and enable massive data throughput and ultra-low latencies across a scalable cluster of commodity servers. A converged data platform, the GridGain In-Memory Data Fabric offers the most comprehensive, enterprise-grade in-memory computing solution for high-volume transactions, real-time analytics and hybrid data processing. The company is funded by Almaz Capital, MoneyTime Ventures and RTP Ventures. For more information, visit [gridgain.com](http://gridgain.com).

## **COPYRIGHT AND TRADEMARK INFORMATION**

© 2016 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. Coherence, JEE and Java are either registered trademarks or trademarks of SUN Microsystems and/or Oracle Corporation in the United States and/or other countries. Windows, Azure, .NET and C# are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Apache, Apache Spark, Spark, Apache Hadoop, Hadoop, 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.