

The background of the slide is a blue-tinted image of multiple overlapping data tables. The tables contain various numerical values, some with labels like "D", "N", "SP", "DN", and "NE". The text is semi-transparent and layered, creating a sense of depth and complexity.

The Apache® Ignite™ SQL Grid:

A Hot Blend of Tradition SQL and In-Memory Data Grid

A GridGain Systems In-Memory Computing Application Note

Contents

Apache® Ignite™ Sets the Stage.....	2
SQL Grid: The Way to Replace a Disk-Based RDBMS.....	2
Executing Distributed SQL Queries.....	3
New Tools for Management and Visualization.....	4
Apache Ignite: The Leading Open Source In-Memory Computing Platform.....	5
Bigger, Better, Faster.....	6
Contact GridGain Systems.....	6
About GridGain Systems.....	6

Traditional database management has been turbocharged by the emergence of new technologies and platforms that move the slow work of disk-based data processing to memory. In memory, data processing can happen exponentially faster. Implemented correctly, in-memory computing can be a game changer for database applications across all industries. It is especially effective for industries that have a special need for speed and real-time reporting.

Apache® Ignite™ Sets the Stage

The Apache Ignite in-memory computing platform is the underlying technology that makes this new kind of superfast database management possible. It can run everything in memory to yield the performance and scalability gains that developers require.

Apache Ignite:

- Supports data caching, massive parallel processing, streaming, and much more
- Slides in between the existing application and data layers easily, with minimal changes to code
- Offers ACID-compliant transactions as well as analytics support
- Works with all popular RDBMS, NoSQL, and Hadoop databases and offers a Unified API with support for a wide range of languages
- Can be deployed on premise, in the cloud, or in hybrid environments
- Can be configured for high availability with automatic load balancing and failover
- Remains highly customizable for a multitude of uses as it handles computing and transacting on large-scale data sets in real time

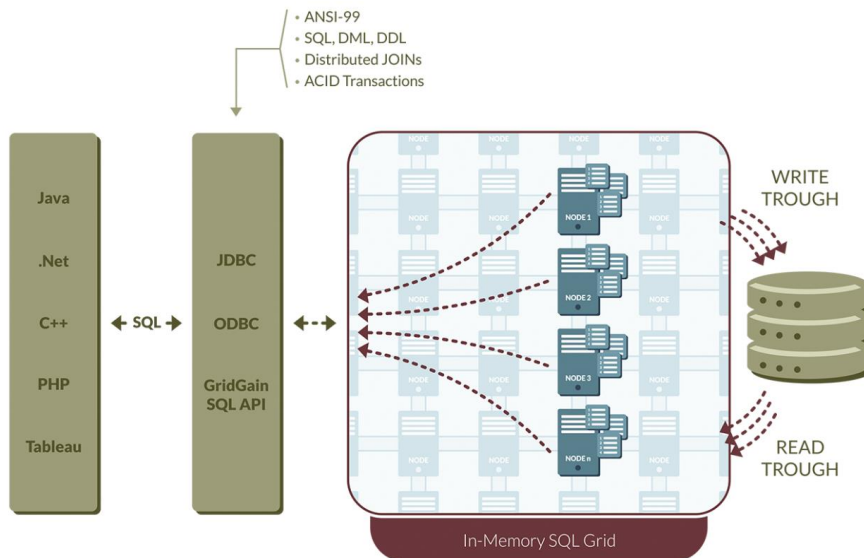
SQL Grid: The Way to Replace a Disk-Based RDBMS

Of all the components of the Apache Ignite in-memory computing platform, one of the most crucial is the **In-Memory SQL Grid**. It is horizontally-scalable, fault-tolerant, and ANSI SQL-99 compliant. It supports all SQL and DML commands and queries such as SELECT, UPDATE, INSERT, MERGE, and DELETE. It is a mature, in-memory solution to supplement a disk-based relational database management system (RDBMS). Geospatial support is built into the product, and all the communication to the SQL grid is done through ODBC and JDBC APIs without custom coding.

SQL Grid is just one part of the overall GridGain in-memory computing platform. It works in concert with the Data Grid to enable advanced clustering across a distributed cluster of machines. Data Grid spreads data across clustered nodes, each of which stores a particular subset of the data. The system can then rebalance the data, making it work in a fault-tolerant and load-balanced way. The inclusion of Service Grid, Compute Grid and Streaming Analytics complete the in-memory computing platform. They boost the data management capabilities of SQL Grid and distinguish it from other memory grid products and solutions on the market today.

Apache Ignite provides many APIs to process, update, and query data. SQL Grid enables SQL-based APIs for Apache Ignite Cluster. Users will find direct APIs for languages including Java, .NET, and C++. That means users of Java, .NET, or C+ libraries can execute Apache Ignite SQL queries directly from their Java,

.NET, or C++ applications without having to make significant changes to the ways they already work.

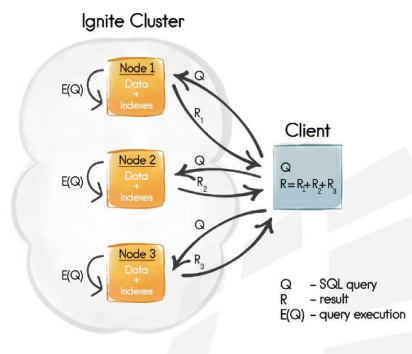


Executing Distributed SQL Queries

Because the SQL Grid is ANSI-99 SQL-compliant, it supports all sorts of aggregations, group by, sorting, JOINS, unions, and more. In its recommended default collocated mode, it can handle any possible fault-tolerant situation. If, for example, a user makes a SQL query while the state of the cluster is simultaneously changing for any reason, the query will yield a consistent result without exception. This is because the query has been sent to every participating node that stores a subset of the data that can be touched as a part of the query execution. The fault tolerance is automatic. At the same time, SQL Grid supports indexing that delivers a powerful performance boost, making redeclarations extremely fast compared to those of disk-based RDBMSs.

Collocated SQL Queries

- Collocated Mode
 - Any kind of JOINS (ANSI-99)
 - Data has to be collocated in advance*
- Recommended mode
 - Avoids data movement
 - Enabled by default



On the rare occasions when it is not feasible to collocate data, a non-collocated mode is also available. In this mode, when a query is sent through all the participating nodes, it might go to other nodes below the data needed to complete the JOIN, potentially triggering some data movement. While that may

result in performance degradation, it guarantees that 100 percent of SQL queries can be successfully completed.

When it comes to indexing, Apache Ignite enables both single-field and group indexes. It supports QuerySqlField annotation that is considered as an index so that SQL Grid will virtualize it as an index field and deliver a very fast response.

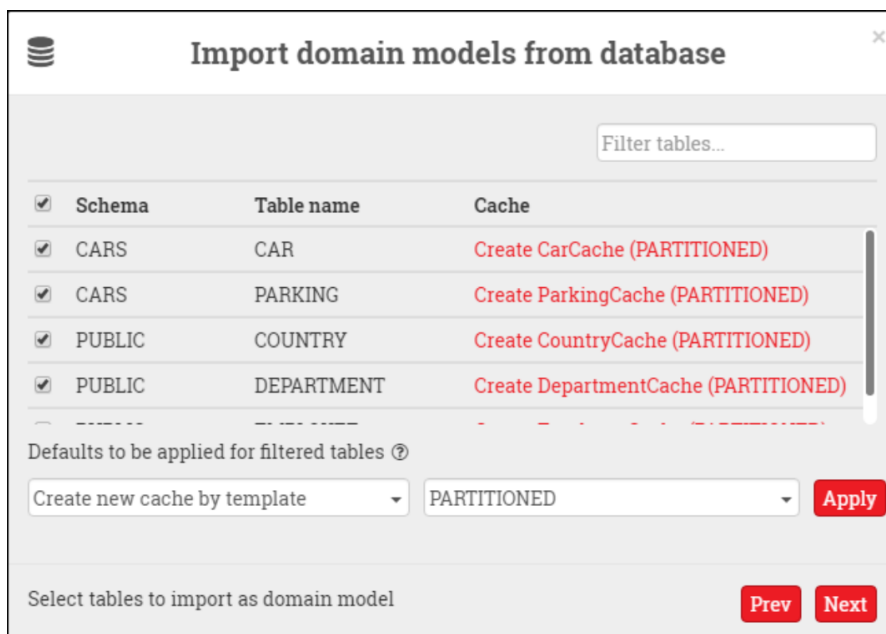
Apache Ignite’s default On-Heap mode, describing where data and indexes reside, implies that all the data will be allocated in the Java Heap. When users use this mode for a particular cache, the indexes will be allocated in the Java Heap as well. Using On-Heap mode covers index maintenance as well. In this mode, Apache Ignite supports only a single implementation, but it should cover most contingencies.

A New Distributed Data Modification Language

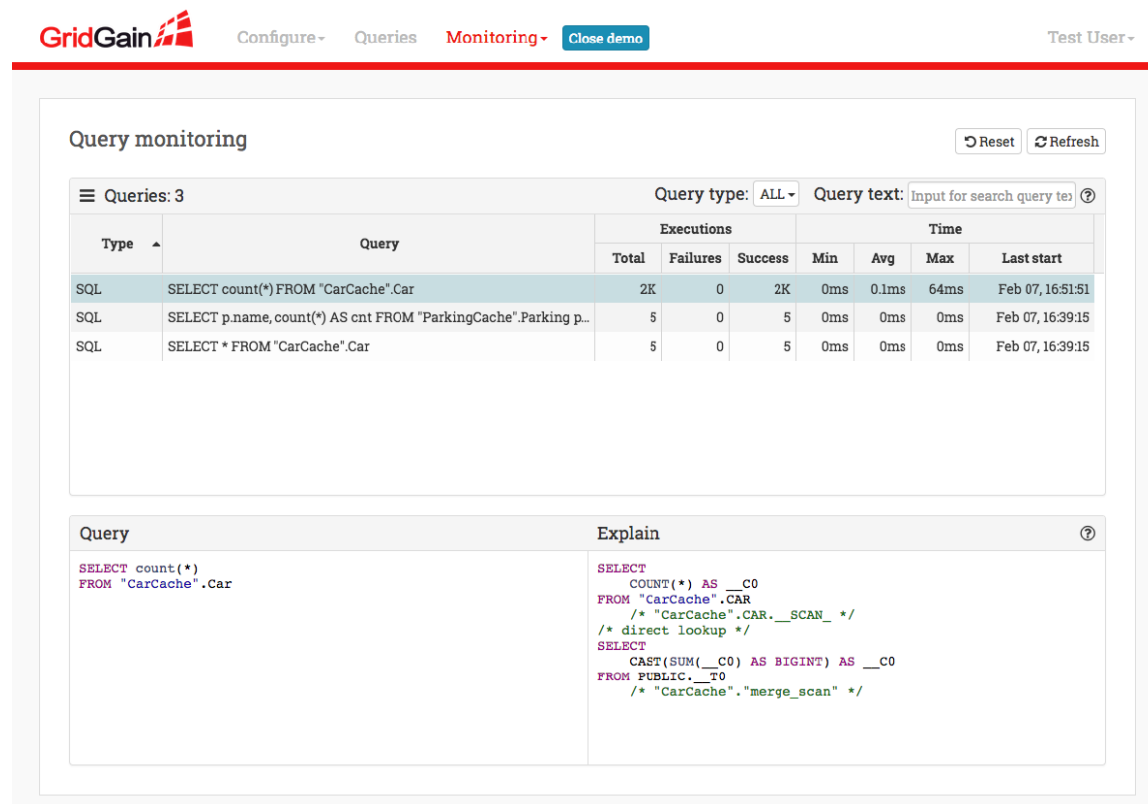
New to SQL Grid is a distributed data modification language (DML), released with Apache Ignite Version 1.9 to support a wide range of queries including MERGE. Apache Ignite 2.0 includes more features. With upcoming data definition language (DDL) support, users will be able to create and drop caches simply by using standard DDL SQL commands.

New Tools for Management and Visualization

SQL Grid connects to existing Apache Ignite APIs and an increasing number of management and visualization tools. These can make life easier for database managers by helping them issue a variety of SQL queries. The Apache Ignite Web Console, for example, is a lightweight Web-based configuration wizard and management and monitoring tool. It helps users prepare cluster configurations, connect to existing databases, and import schemas from them. It lets database managers observe the state of the database, keep an eye on memory virtualization, monitor queries being executed, and more.



In addition, the Web Console can report on query execution in tabular form, in streaming charts, or in graphs. By defining different refresh rates, users can see how the output of a query has changed over time. The tool is also useful for debugging, since users can use issue queries directly from the console and watch their performance.



Query monitoring

Queries: 3 Query type: ALL Query text: Input for search query text

Type	Query	Executions			Time			
		Total	Failures	Success	Min	Avg	Max	Last start
SQL	SELECT count(*) FROM "CarCache".Car	2K	0	2K	0ms	0.1ms	64ms	Feb 07, 16:51:51
SQL	SELECT p.name, count(*) AS cnt FROM "ParkingCache".Parking p...	5	0	5	0ms	0ms	0ms	Feb 07, 16:39:15
SQL	SELECT * FROM "CarCache".Car	5	0	5	0ms	0ms	0ms	Feb 07, 16:39:15

Query	Explain
<pre>SELECT count(*) FROM "CarCache".Car</pre>	<pre>SELECT COUNT(*) AS __C0 FROM "CarCache".CAR /* "CarCache".CAR.__SCAN_ */ /* direct lookup */ SELECT CAST(SUM(__C0) AS BIGINT) AS __C0 FROM PUBLIC.__TO /* "CarCache"."merge_scan" */</pre>

Any other management tools that connect to a database using a JDBC or ODBC driver can be used within Apache Ignite by replacing that driver with the Apache Ignite driver. For instance, Apache Zeppelin, a Web-based notebook that enables interactive data analytics and creates interactive, data-driven documents can be easily integrated with Apache Ignite. In addition, Tableau, the popular data visualization software, can be easily integrated with Apache Ignite.

Apache Ignite: The Leading Open Source In-Memory Computing Platform

Apache Ignite is the leading open source in-memory computing platform. It is a high-performance, integrated and distributed in-memory platform for computing and transacting on large-scale data sets in real-time. It performs orders of magnitude faster than is possible with traditional disk-based or flash technologies. As an in-memory data management software layer, it sits between applications and various data sources, and does not require the rip-and-replacement of existing databases.

Apache Ignite comprises, in one well-integrated framework, a set of key in-memory capabilities, including:

- An in-memory data grid
- An in-memory SQL grid
- An in-memory compute grid
- An in-memory service grid
- In-memory streaming processing
- In-memory acceleration for Hadoop

Despite the breadth of its feature set, Apache Ignite is very easy to use and deploy. There are no custom installers. The code base comes as a single .zip file with only one mandatory dependency: ignite-core.jar. All other dependencies, such as integration with Spring for configuration, can be added to the process a la carte. The project is fully Mavenized and is composed of over a dozen Maven artifacts that can be imported and used in any combination. Apache Ignite is based on standard Java APIs. For distributed caches and data grid functionality, Apache Ignite implements the JCache (JSR107) standard.

The Apache Ignite large scale, distributed in-memory framework offers transactional and analytical applications performance gains of 10 to 1,000 times faster throughput and/or lower latencies. Apache Ignite is an important open source foundation that holds the key to the world of Fast Data across high-volume transactions, real-time analytics and the emerging class of hybrid transactional/analytical workloads (HTAP).

Bigger, Better, Faster

SQL has been a lingua franca for relational database management for decades, and that is unlikely to change. What is changing is the way it can be managed within an in-memory computing platform. It can empower much faster performance, linear scalability, and support for real-time analytics and high-speed transactions that today's database developers demand.

Contact GridGain Systems

To learn more about how the GridGain in-memory computing platform can help your business, please email our sales team at sales@gridgain.com, call us at +1 (650) 241-2281 (US) or +44 (0) 7775 835 770 (Europe), or complete our [contact form](#) to have us contact you.

About GridGain Systems

GridGain Systems is revolutionizing real-time data access and processing by offering enterprise-grade in-memory computing solutions built on Apache Ignite. GridGain solutions are used by global enterprises in financial, software, eCommerce, retail, online business services, healthcare, telecom and other major sectors. GridGain solutions connect data stores (SQL, NoSQL, and Apache™ Hadoop®) with cloud-scale applications and enable massive data throughput and ultra-low latencies across a scalable, distributed cluster of commodity servers. GridGain is the most comprehensive, enterprise-grade in-memory

computing platform for high volume ACID transactions, real-time analytics, and hybrid transactional/analytical processing. For more information, visit gridgain.com

COPYRIGHT AND TRADEMARK INFORMATION

© 2017 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. Windows, .NET and C# are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. JEE and Java are either registered trademarks or trademarks of SUN Microsystems and/or Oracle Corporation in the United States and/or other countries. Apache, Apache Ignite, Ignite, the Apache Ignite logo, Apache Spark, Spark, Apache Cassandra, and Cassandra 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.