

Cloud Computing and Advanced Relationship Analytics

Using Objectivity/DB to Discover the Relationships in your Data

By Brian Clark
Vice President, Product Management
Objectivity, Inc.
408-992-7136
brian.clark@objectivity.com
www.objectivity.com

April 2010

Cloud Computing and Advanced Relationship Analytics

Using Objectivity/DB to Discover the Relationships in your Data

There is a wealth of information, connections and relationships within the terabytes and petabytes of data being collected by organizations on distributed cloud platforms. Utilizing these complex, multi-dimensional relationships will be the key to developing systems to perform advanced relationship analysis. From predictive analytics to the next generation of business intelligence, “walking” the social and professional graphs will be critical to the success of these endeavors.

Most applications and data layers today are only capable of simple analytics, finding out who-did-what or who-bought-what. Advanced analytics yields far more knowledge at a much deeper level, understanding who, where, what, how and why – right now.

Google, Yahoo and other leaders in online and personalized services are looking to use complex, multi-dimensional graph data to improve everything they do -- from supporting vast indexes and catalogs of content, to providing the best value and return to their advertising users. Graph technology offers a superior solution for relationship analytics requirements, enabling organizations to traverse any number and complexity of relationships in virtually any amount of distributed data, from any number of sources and types, in near real-time, and on the same commodity hardware obtained through cloud computing platform providers.

Solve the problem.

The No SQL (or “not only SQL”) movement is defined by a simple premise: **Use the solution that best suits the problem and objectives.** If the data structure is more appropriately accessed through key-value pairs, then the best solution is likely a dedicated key-value pair database. If the objective is to quickly find connections within data containing objects and relationships, then the best solution is a graph database that can get results without any need for translation (O/R mapping). Today’s availability of numerous technologies that finally support this simple premise are helping to simplify the application environment and enable solutions that actually exceed the requirements, while also supporting performance and scalability objectives far into the future.

No SQL Technologies


Key-Value Pair	Column Databases
Graph Databases	Document Database

Cloud computing has adopted a broad variety of No SQL technologies to support these leading-edge requirements. By using solutions designed to support specific tasks and requirements, organizations can more easily achieve often significant reductions in complexity and costs associated with their systems. And by using more targeted and capable components, these systems are also able to achieve even greater levels of performance and scalability.

Graph databases may be the most important part of the No SQL movement.

Graph databases typically solve problems related to complexity of data, while key-value and column-store solutions seek to address common issues encountered as data volumes grow in size. Our technology addresses both the complexity and scalability requirements to give you the best of both worlds, managing complex and big data.

Graph data is represented by nodes or vertices and edges, where any node could be connected to any number of other nodes via the edges between them.



Graph database technologies can support rapid traversal of these edges to get results in a matter of seconds (or less). Because the data is persisted where relationships are first class citizens, performance is no longer an issue. But when this type of work is done in a relational database or key-value environment, there are very expensive constraints and limitations to performance. And, of course, if the graph database architecture is distributed, then scalability limits are also addressed very nicely.

As social media, personalized web and advertising services, business intelligence and organizations in other spaces understand the importance of the deeper relationships within their data, their success in utilizing this information will depend on the technology. Trying to scale out graph data and relationship analysis using relational technology is simply not the answer. In addition, the complex custom code, high-end server hardware, map reduce layers, and administration overhead required to support these architectures can significantly increase costs and overhead.

The days of compromising requirements to support centralized database server architectures or relational systems that simply aren't designed to solve some problems, are over. Today, there are other options and solutions.

Objectivity/DB is the original graph database.

Today many people are looking for graph technologies that can support for their applications and online, distributed systems. While several solutions focus on the graph interface or API, the underlying database is even more critical. The graph database must support your requirements for performance and scalability, ingest and fusion.

Objectivity/DB is the leading, enterprise-proven, distributed data management solution and persistence layer that supports the most demanding requirements, and graph computing needs. Benefits include:

- Objectivity/DB persists complex object models where the object and relationships can be any degree of complexity. Objects are persisted natively in the database, and there is no mapping layer. References between objects are stored as direct persistent pointers (object identifiers) making navigation very efficient and high performance.
- Complex graphs (networks) of objects and references can be persisted and navigated (queried) efficiently. The objects and references in the graph can be distributed across tens of thousands of databases in each federation. The navigational queries can be broken down into sub-tasks and executed in parallel.
- Objectivity/DB has a fully distributed (peer-to-peer) architecture that allows virtually unlimited scaling of both data and processing. This federated architecture (a collection of distributed databases) supports a single logical view across all objects in the federation, no matter where the data is located, allows elastic scalability of both processing and data in the cloud.
- Objectivity/DB has been deployed in a variety of configurations, including embedded devices and sensors, telecommunications network element management, enterprise client/server and large distributed systems, and on grid and cloud environments.
- Objectivity/DB supports several object languages, including Java, C#, C++, Python and Smalltalk. SQL access is also supported via an ODBC connection. Objects and application components created in any language can be accessed from any language.
- Objectivity/DB runs on a number of hardware platforms and operating systems, including Linux, Windows, and major Unix operating systems. Databases created on any platform can be accessed heterogeneously from any platform.

For More Information

Please visit us online at www.objectivity.com or call (U.S.) 408 992-7100.

Objectivity, Inc.

640 West California Avenue, Suite 210

Sunnyvale, CA 94086-3624 U.S.

Phone: U.S. (408) 992-7100

FAX: U.S. (408) 992-7171

www.objectivity.com