



Multiple Graph Database Challenges, One Solution

[Massiv.io](#), is the company behind **Integral**, a platform delivering real-time architecture and visualization of enterprise applications to architects, developers and operations professionals. Integral is built on the idea that “*Code is the Source of Truth*” and aims to put an end to architecture problems from the day it’s implemented.

When initially developing Integral, Massiv.io knew the platform needed to store and process information, or source code, that is not only complex and interrelated but must also be queried extremely quickly, **typically less than 100 milliseconds**. Furthermore, it needed to be capable of clustering data in a *containerized* cloud environment, where instances are regulated according to varying loads.

Addressing Key Issues in a Graph Database Market

After an extensive investigation, Massiv.io came to the conclusion that a graph database was the best fit for their requirements and began researching nearly every graph database on the market. **One of the key issues faced was that Integral required highly interconnected data to be traversed incrementally** (e.g. return edges and vertices based on computed values, traversal depth may vary according to these computations). While easier to do directly in Java, this proved difficult to do with a query language. Throughout their investigative journey, they concluded that most graph databases appeared to be architecturally divided into two camps which presented issues:

1. A programmatic API that is executed remotely over a clustered data store. In these cases, the clustering technology was solid, but Massiv.io quickly ran into performance problems, having to issue many API calls which were remotely executed. **Query time was tracked in seconds, which was too slow for their needs.**
2. The ability to run the database In-memory. This eliminated performance issues. However, by testing a graph database that implemented this architecture, they concluded it had limited support for clustering which would not meet Massiv.io's requirements. This wasn't a complaint that

Industry

Software

Challenge

Quickly and efficiently process & query complex, interrelated source code while clustering data in a cloud environment.

Approach

Investigate every available Graph database vendor to find the right solution to incrementally traverse data with a clustering solution that can hold up in demanding production environments.

Solution

Use OrientDB's graph database, a straightforward setup with clustering and embedding capabilities to complement Integral.

Result

Fast setup resulting in maximum performance with successful clustering and queries with times in the range of sub-milliseconds.



CASE STUDY

the clustering features are not open source (they would have happily paid for those), but that upon investigation they believed this specific vendor had not delivered a clustering solution which was equipped for the demands of Integral's deployment environment.

Multiple Challenges, One Solution

“OrientDB’s DESIGN PROVED TO BE THE PERFECT COMPLIMENT TO OUR PLATFORM. OrientDB SAVED US COUNTLESS TIME IN PROVIDING CRITICAL FEATURES FOR OUR PLATFORM THAT MORE TRADITIONAL DATABASE TECHNOLOGIES DO NOT DELIVER. IT DID SO IN A WELL-DESIGNED, EASY-TO-INTEGRATE APPROACH.”

- **JIM MARINO, ARCHITECT & TECHNOLOGY STRATEGIST, MASSIV.IO**

Amongst the several graph databases researched and tested by Massiv.io was OrientDB. Their list of technical features as well as their approach seemed to match all of Integral's requirements. *“OrientDB’s design proved to be the perfect compliment to our platform.”* said Jim Marino, Architect & Technology Strategist for Massiv.io. *“OrientDB saved us countless time in providing critical features for our platform that more traditional database technologies do not deliver. It did so in a well-designed, easy-to-integrate approach.”*

Furthermore, OrientDB addressed the architectural issues presented by other vendors and proved easy to install. Embedding OrientDB for maximum performance presented no challenges and its clustering approach was reliable and straightforward to set up. **As a result, Integral’s queries are as low as sub-milliseconds!** *“It took us a couple of days to remove the previous graph database and replace it with OrientDB. The process was straightforward.”*

Surpassing Set Goals

“INTEGRAL ALSO WOUND UP WITH A MUCH SIMPLER DEPLOYMENT TOPOLOGY. OrientDB SAVED US COUNTLESS DEVELOPMENT TIME AND EFFORT BY DELIVERING A CRITICAL PIECE OF OUR INFRASTRUCTURE.”

- **JIM MARINO, ARCHITECT & TECHNOLOGY STRATEGIST, MASSIV.IO**

Through this collaboration, Massiv.io's end goals were not only achieved but surpassed **with performance going from seconds (sometimes up to 30) down to milliseconds.** *“Integral also wound up with a much simpler deployment topology. OrientDB saved us countless development time and effort by delivering a critical piece of our infrastructure.”* mentioned Mr. Marino. Integral now successfully analyzes source code, configuration and deployment artifacts to produce an always accurate, continuously updated view of your enterprise architecture.

About OrientDB

The native multi-model database combines the connectedness of graphs, the agility of documents and a familiar SQL dialect. Fortune 500 companies, government entities and startups all use the technology to build large-scale innovative applications. Some of their clients include Comcast, Ericsson, the United Nations, Verisign, Pitney Bowes, Sky, Diaku, CenturyLink and Sonatype.

OrientDB won the prestigious 2015 Infoworld Bossie award and has been covered by multiple media outlets.

www.orientdb.com

Contact Info

✉ enquiries@orientdb.com

About Massiv.io

Massiv IO delivers real-time architecture and visualization of enterprise applications to architects, developers and operations professionals through Integral with no run-time agents and no code changes.

<http://www.massiv.io/>