

Shaking it Up: How CGG GeoSoftware is Improving Performance on a Seismic Scale

Introduction

CGG, a leader in fully integrated geoscience products and services, delivers geological, geophysical and reservoir capabilities to a broad base of customers primarily in the global oil and gas industry. Since its creation in 1931, CGG has been pushing the limits of technology to record, manage and analyze “large” volumes of data. In many ways it entered the “Big Data” market in 1971, when it acquired the industry’s first 3D marine seismic survey with three parallel recording streamers behind the vessel.

Today, the company provides seismic equipment, acquisition services and advanced technologies to acquire, process, and analyze massive amounts of data. As an example, their largest of around 40 globally networked, processing and imaging centers works with an active data set of over 100PB/day.

This case study peers into the challenges the CGG GeoSoftware Business is facing today to meet its demanding business requirements and how its evaluation of a variety of technology solutions led to selecting Objectivity’s ThingSpan™.

The Business Challenge

IT systems are constantly and rapidly evolving, and the scale and pace of change are magnified in the Big Data infrastructure domain. Although CGG and its CGG GeoSoftware business had been tackling Big Data with proven, mature technology decades before the term became mainstream, it anticipated scalability challenges in the near future to ensure that its seismic and reservoir characterization algorithms and applications could keep pace with exponentially growing volumes of data.

From a purely technical standpoint, this was a monumental step on the road to full data center-readiness and Lambda architecture support, but along with the technical challenges, it also required a path that enabled current solutions to maintain their commercial viability along the way. The ideal approach for CGG GeoSoftware was to be as agnostic as possible to the ever growing array of deployment environments, all while continuously commercializing and supporting their current technology and developing next-generation solutions to provide a superior level of technical and economic performance.

“

Over the past 10 years, Objectivity has been an excellent partner in proactively supporting CGG’s ability to fully leverage various IT advances in our commercial oil and gas software solutions. Its ThingSpan™ platform continues that trend by providing a low-risk avenue to harness the Hadoop ecosystem, thereby increasing the deployment options available to our clients.

”

- Kamal Al-Yahya
Senior VP, CGG GeoSoftware



Use Case: Parallel Seismic Data Analysis

For the oil and gas industry, much of seismic analysis involves data formats and algorithms that do not lend themselves to modern parallel architectures. The generally accepted practice today is to decimate or otherwise subset the data – a time-consuming process that can prevent analytics from fully capitalizing on the nuances of the full dataset.

By adopting enabling technology to support a parallel seismic format and a parallel access pattern, a Lambda-compliant framework could provide the ability to leverage all available data and perform more analysis in less time, thereby achieving more accurate scientific results.

The Technical Evaluation

When determining the criteria for selecting an ideal solution, CGG GeoSoftware's primary priority has been to maintain the ability to abstract its technology away from a growing selection of hardware platforms. Having already begun the migration to distributed compute environments, it was important that the company's software remain agnostic and capable of running on all potential deployments.

In the technical evaluation process, CGG considered major relational data stores, all of which it determined as currently cost-prohibitive in relation to the scalability required. The organization also evaluated several popular NoSQL databases available on Hadoop; although those alternatives provided better price/performance compared to relational solutions, they were also eliminated due to the lack of ACID compliance and product support.

Once it was clear that CGG GeoSoftware required a well-supported platform that could agnostically operationalize its applications in a scalable, data center ready Lambda architecture environment while lowering the total cost of ownership, Objectivity's ThingSpan became a leading contender as the right solution for its needs.

The ThingSpan™ Solution

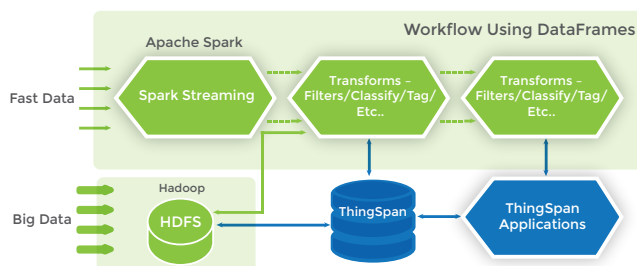
Mitigating adoption risk is critical for CGG GeoSoftware. As a long-time Objectivity customer, CGG was already familiar with Objectivity's best-of-class support with turnarounds within 24 hours. CGG GeoSoftware's existing deployments on Objectivity's legacy product Objectivity/DB also continued to propel the migration to ThingSpan.

ThingSpan is designed to accelerate time-to-production of Industrial Internet of Things applications. It runs natively on top of the Hadoop Distributed File System (HDFS) as a YARN application while using Apache Spark for workflow and data transformation.

"Despite the complexity of our current system, ThingSpan only required minor changes to enable our applications to run on Hadoop, which could not be said of other solutions. The conversion was, in many ways, the software equivalent of flipping a switch," said Marco Ippolito, Software Architect, Data Model, at CGG GeoSoftware.

Because geoscience involves modelling physical objects and locations, a platform that enables object data modelling like ThingSpan is a natural fit. ThingSpan's support for Lambda architecture and its portability across topologies and hardware solutions means that researchers and analysts across all scientific disciplines can easily collaborate on the same Data platform.

ARCHITECTURE DIAGRAM



Future Plans

CGG GeoSoftware continues to innovate on ThingSpan as it develops its path forward for Big Data technologies. Its focus is on introducing more advanced interpretation and analysis on a wider variety of geoscience data while improving the quality of results and overall user experience in reservoir interpretation and characterization. The company is pleased to work with its clients and partners on Objectivity's ThingSpan as it develops ways to cost-effectively extract significantly more value from data.



CGG (www.cgg.com) is a fully integrated Geoscience company providing leading geological, geophysical and reservoir capabilities to its broad base of customers primarily from the global oil and gas industry. Through its three complementary business Segments of Equipment, Acquisition and Geology, Geophysics & Reservoir (GGR), CGG brings value across all aspects of natural resource exploration and exploitation. CGG employs over 7,000 people around the world, all with a Passion for Geoscience and working together to deliver the best solutions to its customers.



Objectivity, Inc. is a pioneer in high-performance distributed database platforms that power mission-critical applications for the most demanding and complex data sources in the enterprise. Objectivity enables organizations to accelerate time-to-value of their data assets at scale by enriching Big Data with Fast Data. With a rich history serving Global 1000 customer and partners, Objectivity holds deep domain expertise in fusing vital information from massive data volumes to capture new revenue opportunities, drive competitive advantages and deliver better business value. Objectivity is privately held with headquarters in San Jose, California. Visit <http://www.objectivity.com> to learn more.