

**CLIENT**

A multinational telecommunications conglomerate

**INDUSTRY**

Telecommunications

**BUSINESS NEED**

To establish best practices for the use of NoSQL and open source databases that would result in cost savings.

**SOLUTION**

Pythian delivered a time-series Cassandra solution on Google Cloud virtual machine instances.

**RESULT**

The client now has the confidence that it can migrate a real-world use case off Oracle to reduce its license costs.

**TECHNOLOGIES**

Google Cloud Platform, Apache Cassandra, Apache Spark, Terraform, Puppet

## **PYTHIAN HELPS A TELCO GIANT REDUCE ITS COSTS IN MANAGING TODAY'S EXPLOSION OF UNSTRUCTURED DATA**

As a telecommunications provider with hundreds of millions of customers worldwide, our client has substantial needs in managing its data. Traditionally, the client has been able to rely on Oracle and Microsoft databases. However, today's data is growing exponentially in volume and it is increasingly unstructured. The client recognized that it was time to invest in the performance and cost-effectiveness of non-relational open source databases.

### **SOLUTION**

Pythian worked with the client to establish NoSQL reference architecture patterns, best practices and appropriate databases for common use cases. To save the time and money involved in making existing hardware available, Pythian instead recommended Google Cloud Platform as the host for this project.

The work involved testing the feasibility of replacing Oracle for a specific time-series use case: aggregating billing events to generate invoices at a given cadence (30 to 60 days). To perform the test, Pythian delivered a time-series Cassandra proof of concept hosted on Google Cloud virtual machine instances. Terraform was selected to provision the virtual machines, and Puppet automated the installation of Apache Cassandra. The client's billing events were exported from their Oracle database as a CSV file, which was ingested via Apache Spark.

---

*“ The volume and complexity of today’s data call for open source, non-relational databases that operate in the cloud. Thanks to Pythian and their expertise in data and cloud, we now have a clear strategy for using that technology.”*

*Director of IT Operations*

---

## **RESULT**

By successfully demonstrating queries over a distributed open source NoSQL database, Pythian provided a real-world use case that can be moved away from Oracle to reduce the client’s burden of license costs. Pythian also demonstrated the high availability and elastic scalability capabilities of Apache Cassandra.

## **ABOUT PYTHIAN**

Pythian is a global IT company that helps businesses leverage disruptive data technologies to better compete. Our services and software solutions unleash the power of cloud, data and analytics to drive better business outcomes. Our 20 years in data, commitment to hiring the best talent, and our deep technical and business expertise allow us to meet our client promise of using technology to deliver the best outcomes faster.

## **WORLDWIDE OFFICES**

Ottawa, Canada  
New York City, USA  
London, England  
Sydney, Australia  
San Francisco, USA  
Hyderabad, India