

Surpassing Data Limitations with Kyligence On MapR

Case Study

Challenges

- Excessive Memory Resource Demands
- Lacking System Scalability with High Cardinality Dimensions
- Constant Developer Maintenance, Debugging Issues in Production

One Kyligence customer, a leading provider of enterprise-class, cloud-based solutions, helps over 1,500 customers by leveraging a decade of data and insights. This SaaS company offers a sales incentive benchmarking, big data solution that is used to compare sales performance with other companies in the industry, enabling sales teams to boost team success.

The company had adopted NoSQL databases and Apache Spark for the benchmarking solution to work with large, distributed data sets. As the company developed more solutions, more applications accessed the same data sets. Tuning Spark queries and managing memory resources became routine challenges. Issues on production systems often took hours to fix, and maintenance resulted in multi-month delays. The existing system wasn't scalable, and teams couldn't extract insights fast enough. They needed to optimize resources, improve analysis, and adapt quickly.

Kyligence Enterprise

- Petabyte-scale Data Sets
- 100,000+ Concurrent Users
- Sub-second Latency
- Rich API Support
- Visualized Intelligent Modeling
- Monitoring and Optimization
- Seamless BI Integration

Challenges

A new solution was required that could support over 100 TB of data, over 1,000 customers, and 100,000s of concurrent users, with sub-second query latency. It would have to be flexible for querying and making changes to schema and data, have low maintenance costs, and support fast migration and adoption. The company considered Apache Druid, three proprietary solutions, and Kyligence.

Most of the options required too many components or dedicated resources, which posed ongoing challenges for both scalability and optimization. Several had continual maintenance requirements without substantial support, lacked the capability to query raw data or pre-aggregate data, or lacked control over schemas, and models.

Solution

Kyligence was the only on-premise, unified analytics solution that offered quick analysis with granular control, scalability for data and users, optimization with intelligent modelling, simple implementation, and seamless integration with big data technologies, such as MapR. As a MapR user, the company can fully leverage both technologies and

Results

Boosted Query Performance

Queries that previously ran for minutes now finish in sub-seconds.

Reduced Preprocessing Challenges

Less time is spent on data extraction and preparation.

Improved Schema Management

Supports anticipated growth in data volume and changes to business requirements.

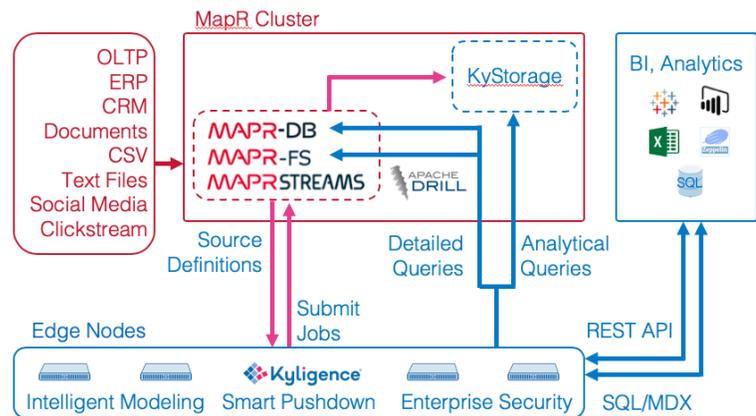
Simplified Updates to Data Models

Data requests can be served in days instead of weeks.

have a total solution for their needs.

Kyligence Enterprise enabled the company to design and implement sub-second queries in a user-friendly interface. The company could pre-join or pre-aggregate data, with MapReduce or Apache Spark, and run many complex queries concurrently without slowing down. The company also didn't need to care where the queries would run because Kyligence intelligently routed queries to the right place, either a pre-calculated result set or any data store that the company had.

Kyligence was fast, big, intelligent, and open enough for the customer's growing business. Through continual improvements and expert support, Kyligence is helping the company drive sales performance success and maintain their high customer satisfaction ratings.



Next Steps

Schedule a short conversation with a member of our team to learn how Kyligence can address your specific challenges. Visit our website and fill out the request form at <https://kyligence.io>, or send an email to info@kyligence.io.

About Kyligence

Kyligence Inc. is the leading data intelligence company that focuses on big data analytics technology innovation, founded by the creators of Apache Kylin. Kyligence is a next-generation big data analytics platform, using machine learning to automate both data preparation and insight discovery at massive scale. With extremely high performance for all data analysts and business users, Kyligence is being used by over 1,000 organizations around the globe, with many leading enterprises among them.

Silicon Valley
99 Almaden Blvd.
Suite 320
San Jose, CA 95113

Washington
2018 156th Ave. Bldg.F
Suite #100
Bellevue, WA 98007

New York
1460 Broadway
New York, NY 10036

Shanghai
4th Floor, Bld Y1 No.112, Liangxiu Road
Shanghai Pudong Software Park
Shanghai, China 201203