

# Accelerate Big Data Exploration with Kyligence and Qlik

## Technology Solution Overview

### 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

### Strengths

- Multi-dimensional Analysis
- Detailed Pushdown Queries
- Hundreds of Dimensions
- Hundreds of Indicators
- Self-service Deployments
- No Coding Required

In today's big data world, business analysts or data scientists need to explore petabytes of data to get the most accurate results. Many enterprises have built data lakes on the Hadoop platform. To explore data on Hadoop, Qlik typically connects to the ODBC driver for Hive and loads the data into memory through the special on-demand apps.

The challenges of this architecture are:

- Hive is meant to be a data warehouse solution, not an interactive ad-hoc query tool. Even though the performance of Hive has improved tremendously in the past couple of years, it still is not suitable for interactive queries when users expect answers within a couple of seconds.
- The aggregation operations scan and calculate across billions of rows of data, adding more delays to the query.
- The amount of aggregation loaded in memory is limited to the size of memory on the server, which cannot provide the speed and agility that Qlik users demand for large volumes of data.

### Solution

Instead of calculating aggregations on the fly, or storing them in memory, Kyligence pre-calculates the aggregations on the Hadoop cluster, saving the results in the cube data-structure. The cubes are stored on the cluster, as well, so the size of the cubes is no longer a concern.

Qlik applications will connect to Kyligence instead of Hive store. If the query is an aggregated query, Kyligence will fetch the result from the cube, which usually takes less than a half second. If the customer is looking for a specific record, Kyligence will either execute a look up in Table Index or push the query down to other query engines for an answer.

Kyligence is the new generation of OLAP technology for the Big Data era.

## Why Kyligence?

### Get more value out of data

Hundreds of thousands of analysts or data scientists can explore petabyte-levels of data, taking full advantage of the Associative Data Engine.

### See the full picture

Customers can work on the full data set, instead of just the partial data that fits in memory.

### Insights in real time

Kyligence can bring stream data from sources like Kafka into the cube in real time, giving customers the most up-to-date insights.

### Faster time-to-market

There is no coding or complex ETL involved. The AI-driven intelligent modeling by Kyligence significantly shortens project delivery times.

## Architecture



- Qlik gets 100x performance boost on Hadoop.
- Live-connect to data on Hadoop or run the in-memory analysis with on-demand apps.
- GUI-based modeling, cubing, and management; no coding required.
- Supports iterative and real-time cube-building.

## 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](mailto:info@kyligence.io).

## About Kyligence

Kyligence Inc. is the company accelerating productivity with managed data services – on-premise to multi-cloud. Kyligence provides an AI-augmented data platform for analysts and data engineers, automating data management, discovery, interaction, and insight. Kyligence was founded by early contributors to Apache Kylin, and the Kyligence platform 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