

Comparison Guide

Apache Kylin vs Kyligence

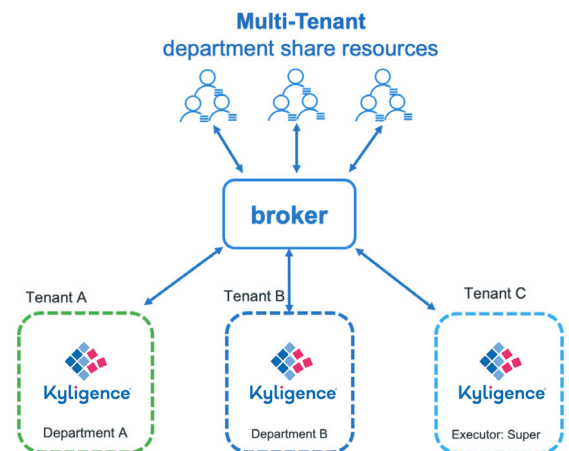
Apache Kylin is the leading open source OLAP (Online Analytical Processing) engine for Big Data capable offering sub-second query latency on trillions of records. Today, it's used by thousands of organizations worldwide as their critical analytics application for Big Data.

In order to provide a solution capable of meeting the needs of demanding, enterprise-scale applications, the team behind Apache Kylin created Kyligence. Compared to Apache Kylin, Kyligence offers significant innovations and comprehensive improvements in performance, stability and security, while supporting DevOps and integrating seamlessly with customers' existing IT platforms and systems.

Deployment for High-Availability and Multi-Tenancy

Both Apache Kylin and Kyligence Enterprise are designed based on a scale-out architecture that enables them to easily serve demanding, enterprise-scale analytical applications. However, only Kyligence Enterprise supports multitenancy for sharing resources across multiple groups without negatively impacting service-levels and security. This helps share development and maintenance costs which can accelerate adoption.

To achieve an extreme, highly available, and stable query service, Kyligence also supports physical separation of query and data loading resources called "Read/Write Separation Deployment". This ensures continuous 24/7 stability and high performance for queries even during ETL time.



Smart Query Engine for Hybrid Workloads

Besides queries on aggregates, Kyligence builds up a smart engine for more workloads, providing a unified analytics platform for all groups of users.

Query Engine / Workloads	Apache Kylin	Kyligence
Aggregates	✓	✓
Drill-down to Raw Data	✗	✓
Smart Pushdown	✗	✓
Massive Result Set Export	✗	✓

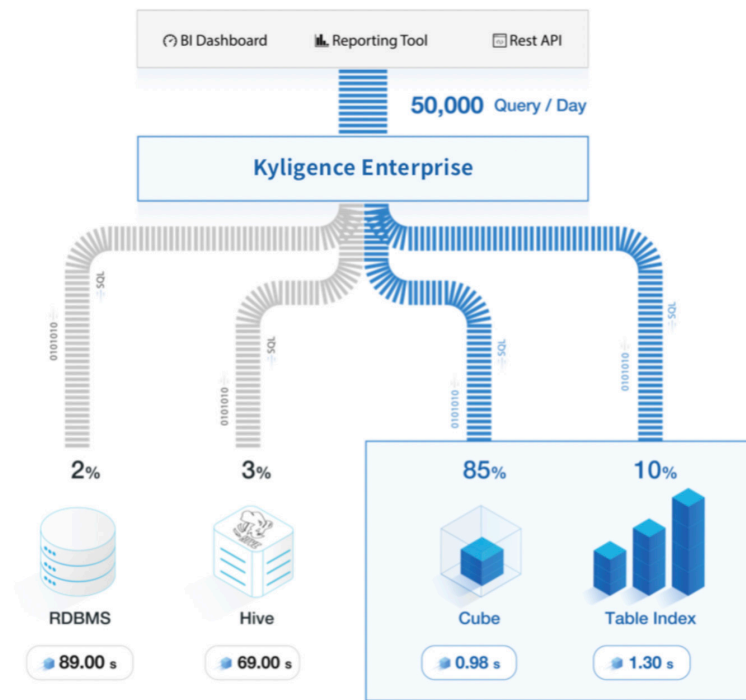


Figure 2 Kyligence Smart Query Engine

BI Compatibility

Kyligence Enterprise seamlessly integrates with a variety of mainstream tools, such as PowerBI and Tableau, and empowers business users by enabling them to use their preferred tools to access big data.

24/7 Support

Kyligence provides dedicated support to Fortune 100 companies and leading enterprises across every industry including finance, telecom, manufacturing, and retail. All Kyligence products receive professional 24/7 support to help customers navigate their digital transformation process quickly.

Enterprise-grade Security

Kyligence Enterprise provides granular access controls, and enables the assignment of different data access permissions to specific users down to the cell level. All while operating on a single, unified data source and semantic layer.

Kyligence Enterprise also supports integrations with existing Identity and Access Management (IAM) systems like LDAP and Kerberos.

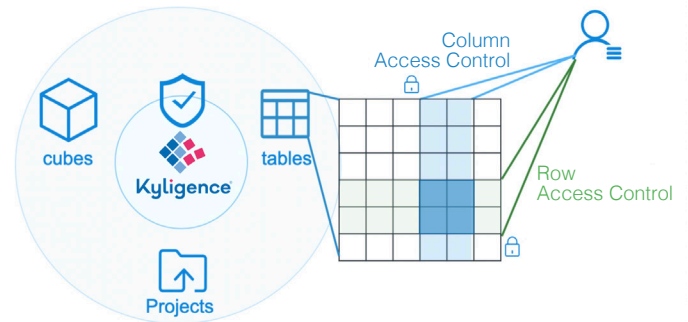


Figure 3 Kyligence Cell-level Data Security

Connect to all Data Sources

Apache Kylin was designed as a Hadoop big data platform and supports Hive as its main data source. But for enterprises looking for a way to migrate their traditional analytical solutions to a new platform, Kyligence provides various connectors for commercial databases and data warehouses like Oracle, Greenplum, and SQL on Hadoop solutions like Impala.

Data Sources	Aph	Kyligence
Hive	✓	✓
Streaming	✓	✓
SQL on Hadoop	✗	✓
Commercial RDBMS & Data Warehouse	✗	✓

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

Appendix: Apache Kylin vs. Kyligence Feature Matrix

		Apache Kylin	Kyligence
Positioning		Open Source Technology	Enterprise-Level Application
Deployment	Read / Write Separation	✗	✓
	Multi-Tenancy	✗	✓
Query	Aggregates	✓	✓
	Drill-Down to Raw Data	✗	✓
	Smart Pushdown	✗	✓
	Massive Result Set Export	✗	✓
Modeling	Manual Modeling	✓	✓
	Intelligent Auto Modeling	✗	✓
	Multi-Partition	✗	✓
Data Source	Hive	✓	✓
	Streaming	✓	✓
	SQL on Hadoop	✗	✓
	Commercial RDBMS and Data Warehouse	✗	✓
Storage	Columnar Storage Engine	✗	✓
	Multi-Level Index	✗	✓
	Metadata RDBMS and Data Warehouse	✗	✓
Security	User / Role-Based Authorization	✓	✓
	Cell-Level Data Access Control	✗	✓
	Enterprise Security Integration (LDAP, Kerberos)	✗	✓
BI Compatability	Mainstream BI Vendor Certification	✗	✓
	Commercial ODBC Driver	✗	✓
DevOps	Kyligence Manager	✗	✓
Support	24 / 7 Technical Support	✗	✓