


Kyligence Cloud

Kyligence Cloud is a cloud-native, distributed analytics platform that provides low-latency query performance for a wide range of enterprise use cases. Based on the popular open source OLAP project Apache Kylin, Kyligence delivers unprecedented analytics performance to hundreds or even thousands of concurrent analysts against petabyte data sets.

Kyligence uses intelligent precomputation of result sets to deliver sub-second query response times at incredible scale. Kyligence constructs *aggregate indexes* - similar to traditional OLAP cubes - that can be distributed across a cluster of servers to achieve this level of scale and concurrency. Kyligence is an AI-augmented analytics platform that employs machine learning to create data models automatically, and then build, tune and maintain cubes quickly and with minimal effort.



Kyligence Cloud 4 Architecture

Cloud-Native Architecture

Kyligence Cloud has been engineered for use in cloud environments. The key characteristics of this cloud-native architecture are:

- Separation of compute and storage to achieve unlimited scalability
- Elastic resource allocation to reduce cloud resource cost
- Distributed AI engine built for the cloud
- Optimized for cloud storage access
- Accelerate queries on cloud data lakes or data warehouses

Kyligence Cloud

- High-performance cloud analytics
- Cloud-native architecture
- Unified Semantic Layer
- Smart Pushdown™
- AI-assisted modeling
- Intelligent precomputation
- Enterprise security and access control

Use Cases

- Accelerate BI reporting, dashboards, and analytics
- Migration of OLAP analytics to cloud
- Add cost effective scaling and concurrency to Snowflake, Synapse, and Redshift
- Unify data semantics across BI tools
- Consolidation of SSAS, Cognos cubes and workloads
- Adding interactive analytics dashboards to SaaS offerings
- Providing Data-as-a-Service (DaaS) to machine learning and AI implementations
- Deploy interactive analytics for cloud warehouses and data lakes
- Extend the life of legacy analytics platforms - Teradata, Netezza, Greenplum, etc.

Unified Semantic Layer

Kyligence is designed to seamlessly integrate with cloud data sources such as data warehouses and data lakes as well as all popular BI tools including Excel, Tableau, MicroStrategy, Power BI, and more.

Kyligence also provides a powerful Unified Semantic Layer that uniformly maintains business logic, hierarchies, and calculations that frees business users from concerns about the technical complexity and implementation of the underlying data source.

This addresses significant data governance issues without needless copying or moving datasets and without having to reinterpret data models and BI semantics across numerous BI systems.

Intelligent Precomputation

Kyligence employs intelligent precomputation of query results to deliver the sub-second response times needed for truly interactive analytics/dashboards on very large datasets. But simply caching all possible results is not practical as the resulting cubes would be huge and inefficient. Kyligence AI-Augmented Engine provides the intelligence by continually analyzing system usage and determining which result sets should and shouldn't be made available in the cube.

Smart Pushdown™

The intelligence of Kyligence Cloud extends beyond cube building and querying. Traditional OLAP systems had no notion of query pushdown to the source systems. While the goal of configuring Kyligence Cloud is to maximize performance by serving data from cubes, there is an entire class of complex or ad hoc queries that can not be serviced by cubes. These types of queries require a “pushdown” to the source system (such as a data warehouse or RDBMS).

Kyligence features intelligent query routing that maximizes query performance by instantly choosing the best strategy for delivering the correct results. Kyligence Smart Pushdown™ is designed to service all queries that do not have precomputed results in Kyligence cubes and indexes. These more complex queries may not achieve the same level of performance as precomputed results, but they will deliver accurate results as fast as the source systems can deliver.

Supported Data Sources

- Cloud Object Storage including
 - AWS S3
 - Azure ADLS gen2
Azure Blob Storage
 - Google Cloud Storage
 - Hadoop/HDFS
- Cloud Data Warehouses including
 - Snowflake
 - Azure Synapse
 - AWS Redshift

Open Source and Standards

- Apache Kylin for distributed OLAP
- Apache Spark for aggregate index/cube building
- Apache Parquet for columnar storage of data
- MDX protocol
- ANSI SQL
- REST APIs

Requirements

Kyligence Cloud Server

- Azure: Standard_D3, Standard_D5_V2*
- AWS: m4.xlarge, m5.4xlarge*

Cluster: Master/Worker/Edge

- Azure: Standard_D3/Standard_D4/
Standard_D3, Standard_D3_v2/
Standard_D5_v2*
- AWS: m4.xlarge/m4.2xlarge/
m4.xlarge, m5.xlarge/m5.4xlarge/
m5.4xlarge*

* recommended