Datasheet

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

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 for unlimited scalability
- Elastic resource allocation to reduce cloud resource costs
- Distributed AI engine built for the cloud
- Optimized for cloud storage access

Kyligence Cloud

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

Use Cases

- Migration of OLAP analytics to Cloud
- Consolidation of SSAS, Cognos, cubes and workloads
- Unify data semantics across BI tools
- Adding interactive analytics dashboards to SaaS offerings
- Providing Data-as-a-Service to Machine Learning and AI apps
- Accelerating analytics for cloud

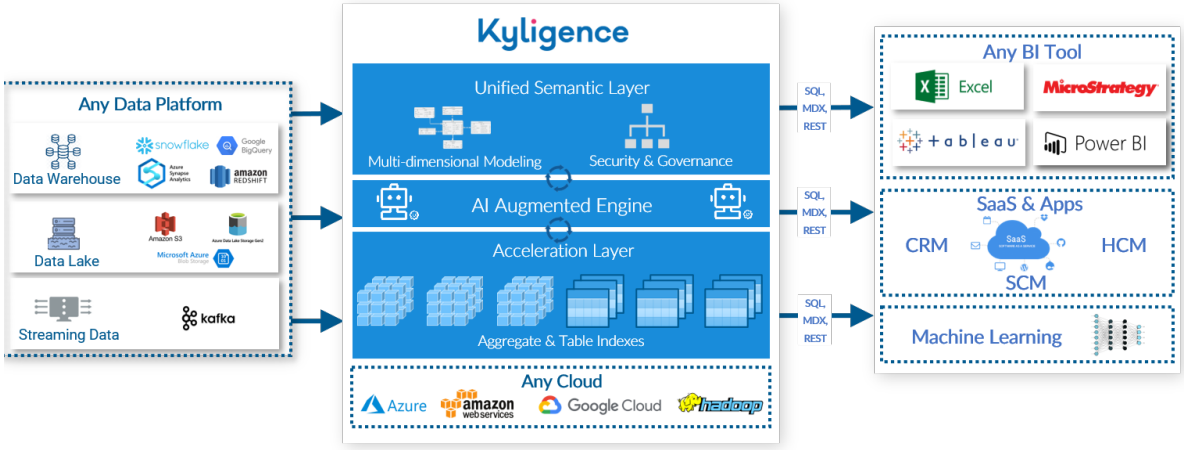

Unified Semantic Layer

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^{tm}

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.

Kyligence feature intelligent query routing that maximizes query performance by instantly choosing the best strategy for delivering the correct results. Kyligence Smart Pushdown^{tm} is designed to service all queries that do not have precomputed results in Kyligence indexes. These

Benefits of Kyligence Cloud

Some of the many benefits of using Kyligence are:

- Unprecedented acceleration of queries for cloud analytics
- Dramatic offload of premium data platforms (Snowflake, Synapse, etc)
- Interactive analytics for hundreds/thousands of concurrent users
- Improved governance, security, efficiency with unified semantics
- System continues learning and improving as it operates
- Extends and enhances investment in legacy OLAP platforms like SSAS, Cognos, and Essbase.

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

Kyligence®