



# Kyligence Analytics Platform

## 2017 Datasheet





# Apache Kylin: A Leading Open Source OLAP Engine for Big Data

Apache Kylin is a leading open source OLAP on Hadoop that provides ANSI-SQL interface and multi-dimensional OLAP for massive datasets. With our powerful technology of pre-calculation, Apache Kylin enables sub-second query latency on petabyte-scale dataset. Apache Kylin consumes data through batch and streaming, and provides query interfaces through ODBC, JDBC and REST API.

Apache Kylin: A Top-Level Project in Apache Software Foundation

Originated from eBay China, Apache Kylin joined Apache Software Foundation (ASF) and graduated as a top-level project in November 2015.

Apache Kylin: Global Users



As a part of our big data platform LAMBDA, Kylin powers inventory, campaign, behavior and demand analysis for advertising. Kylin has replaced our old systems for its great capabilities in handling data volume, velocity, and variety. With the first version in production, we are looking forward to pushing more data into Kylin and achieve near real-time updates.

[Exponential](#)

Glispa is using Apache Kylin as an OLAP component within its data management platform. It enabled us to meet latency requirements for real time analytics over big amounts of data. That helps our data scientists and analysts make sense of the DMP's incomprehensible data in a quick ad-hoc manner.

[Glispa](#)

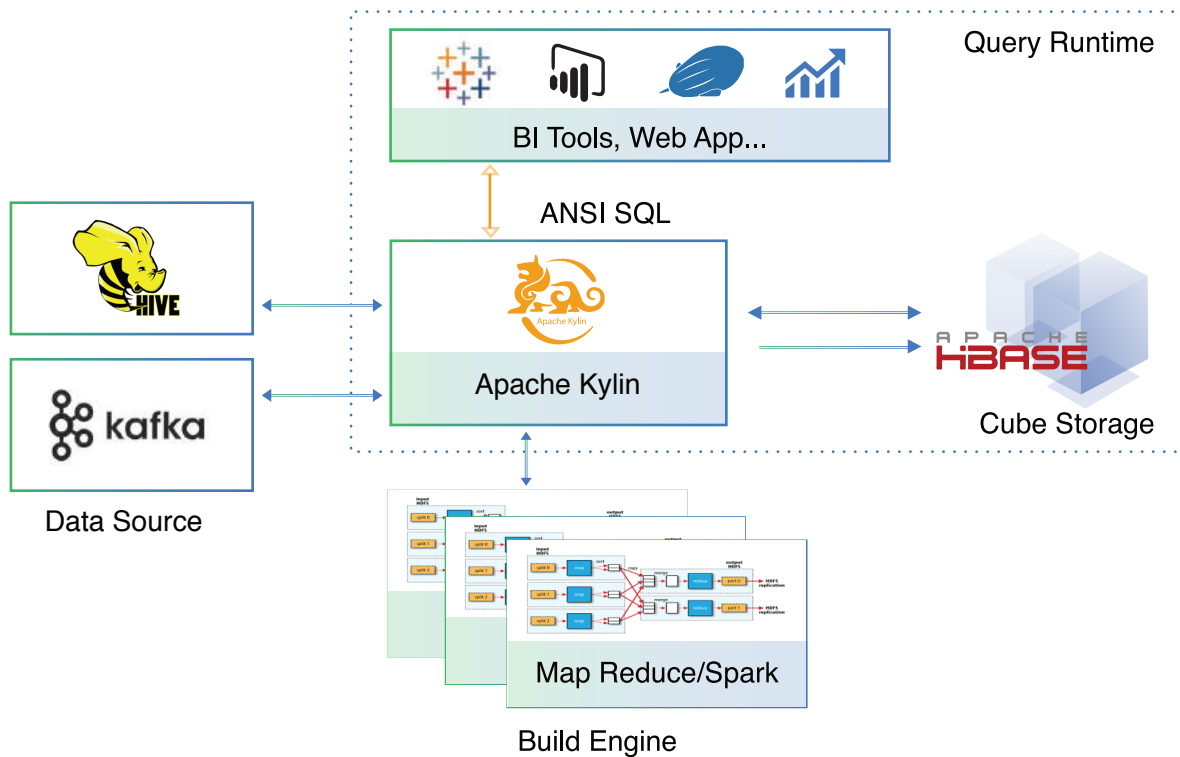
Apache Kylin has helped us overcome our challenges with massive datasets, supported analytics with sub-second latency, and satisfied thousands of analysts' need.

[Toutiao](#)

Traditional data warehouse and BI applications are now migrated to enterprise platform powered by Apache Kylin. Our big data business scenario is fully supported by Apache Kylin.

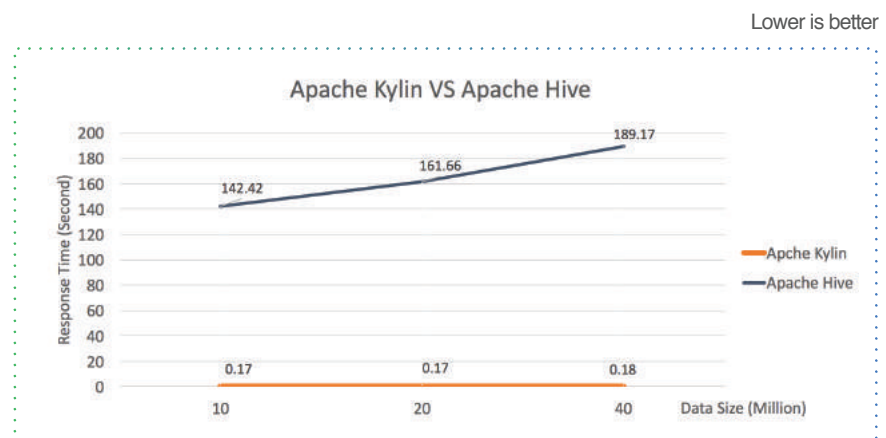
[CPIC](#)

## Architecture




## Performance

- 100+ times faster than Hive
- Sub-second SQL latency at PB-scale
- Web-scale concurrency



\*Based on star schema benchmark

## Advantage



Sub-second Query Latency



ANSI SQL



Native on Hadoop



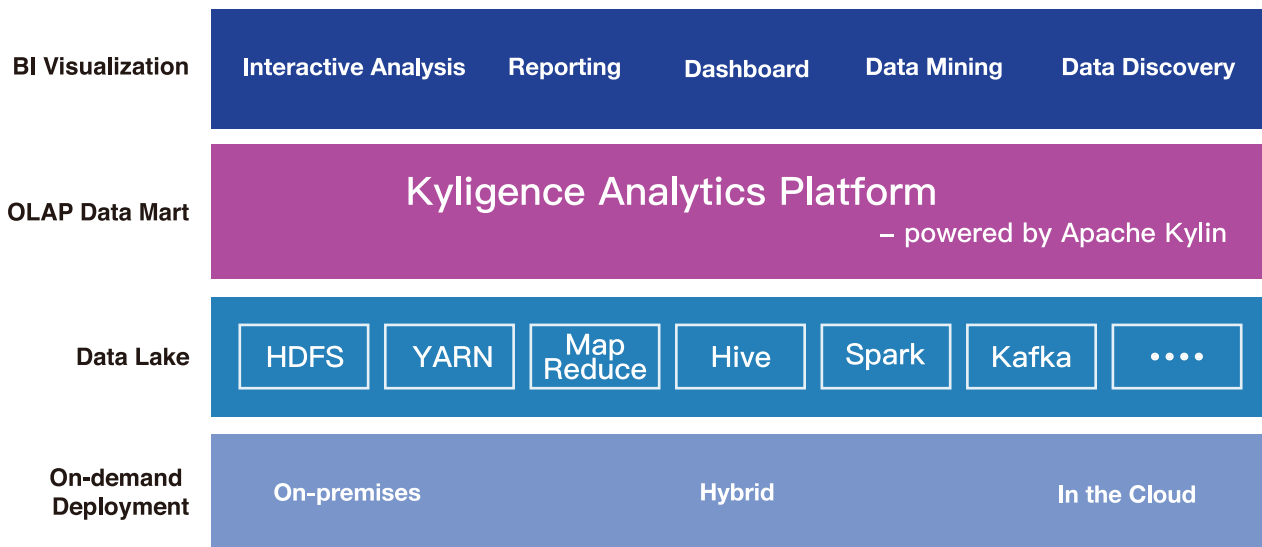
Batch and Streaming Data



# Kyligence Analytics Platform

Kyligence Analytics Platform (KAP) is an enterprise OLAP on Hadoop. Powered by Apache Kylin, KAP enables sub-second SQL query latency on petabyte-scale dataset, provides high concurrency at internet scale, and empowers analysts to architect BI on Hadoop with industry-standard data warehouse and business intelligence methodology.

KAP is a unified analytics platform that simplifies Big Data Analytics for business users, analysts, and engineers. Self-service, seamless integration with BI tools, and no requirements for programming skill. KAP is a native OLAP solution on Hadoop and it interacts with cluster only via standard APIs. KAP supports main Hadoop distributions on-premises and in the cloud.



## KAP: Enterprise Apache Kylin



Enterprise  
OLAP on Hadoop



Intelligent  
Data Modeling



Seamless  
BI Integration



Enterprise Features

## Enterprise OLAP on Hadoop

Hybrid OLAP architecture enables both mission critical analytics in sub-second latency at PB-scale data, and flexible discovery capability over various SQL data sources.

### Query Pushdown

With smart query router, SQL can be pushed down to underlying engines such as Spark SQL, Impala and Hive from KAP automatically

### Patented Columnar Storage Engine

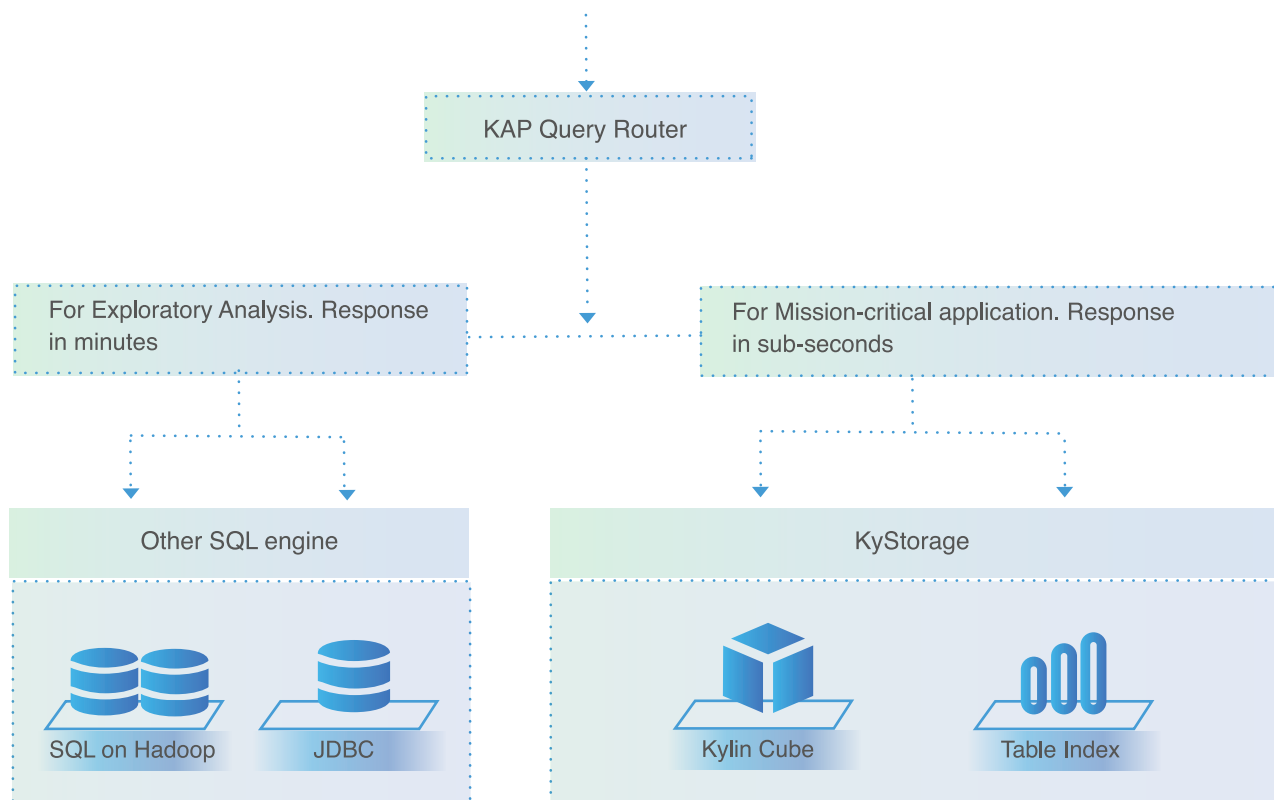
KyStorage is a HDFS based columnar storage engine designed for KAP. With KyStorage, KAP can speed up performance by dozens of times and reduce storage space by over 50% than open source Apache Kylin

### Secondary Index

KyStorage provides secondary index to optimize extreme use cases like ultra high cardinality dimensions, complex filtering condition, and huge table joins.

### Detail Query

KAP fully supports query against detail data and aggregated data.



## Intelligent Data Modeling

KAP's Intelligent Assistant provides auto suggestion for data model optimization and performance tuning.

### Intelligent Cubing

The Intelligent engine supports auto generation of dimensions and measures, provides one-click model optimization. It helps to reduce cluster utilization and storage consumption.

### Data Model Designer

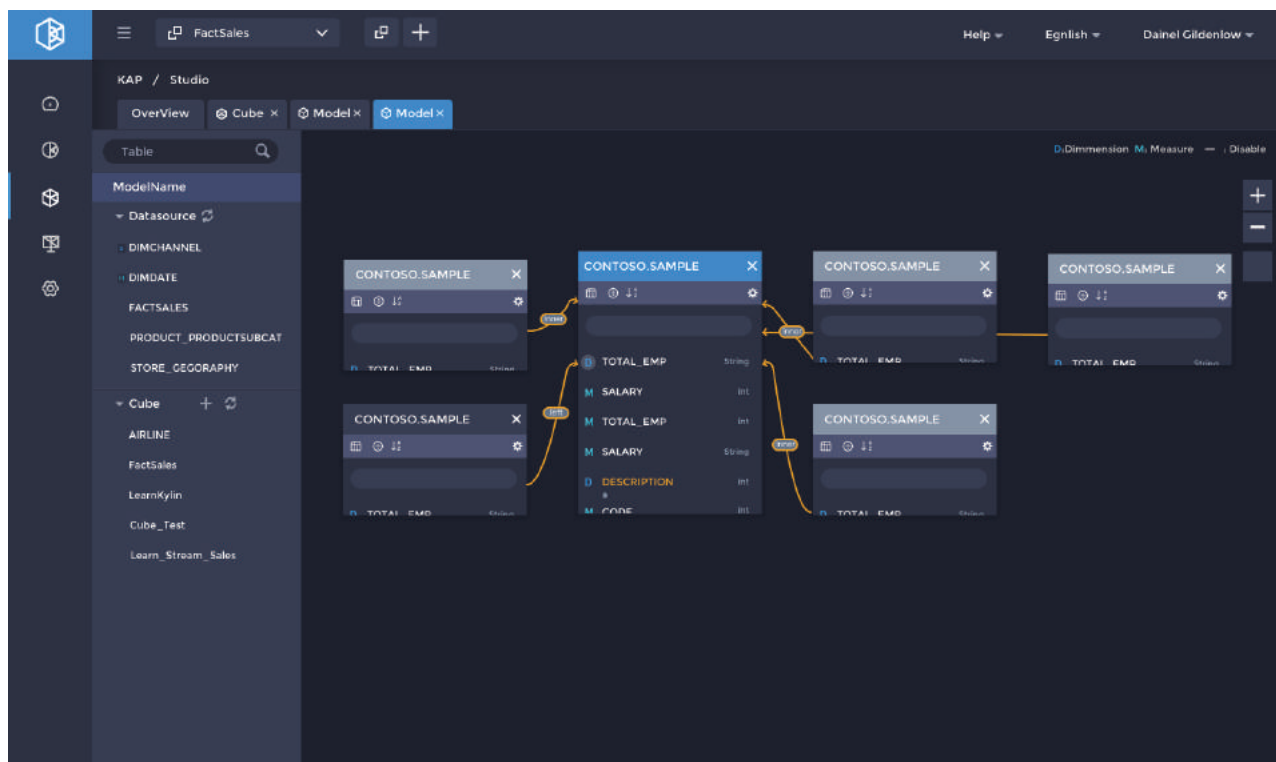
KyStudio, KAP's data model designer, is an intuitive modeling component enables the analysts to well design data model with drag-and-drop experience and wizard optimization

### Semantic Layer

Enriched semantic layer, the Computed Column feature empowers business user to transform data using embedded function or Hive User Defined Function (UDF) by themselves.

### Advanced Functions

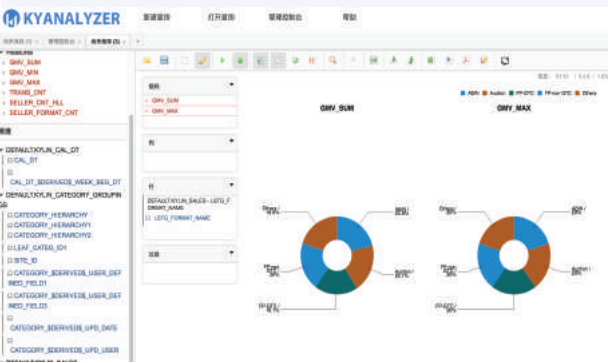
KAP supports advanced OLAP analytics functions including Top-N, Precision Distinct Count at PB dataset, Percentile and other custom-functions.



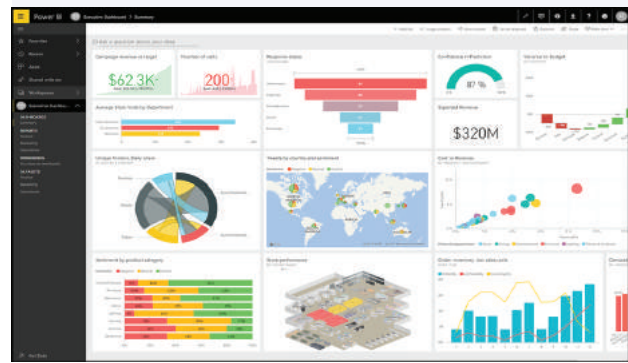
## Seamless BI Integration

KAP is fully integrated with BI tools from open source projects to enterprise products.

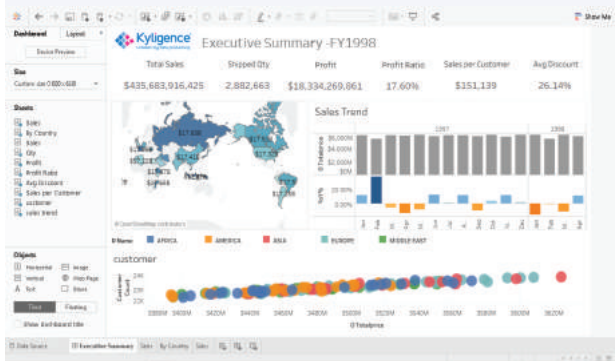
### • KyAnalyzer



### • Power BI



### • Tableau



### • More



in alphabetical order

## Seamless Integration with BI

KAP supports most of the leading BI softwares including Tableau, Cognos, MicroStrategy, Excel, Power BI/Excel, Zeppelin, and Superset. KAP continues your BI investment to extend for Big Data Analytics without pulling data to your BI servers.

## Out-of-box Agile BI

KyAnalyzer is an out-of-box agile BI tool that KAP provides to enable users gain insight directly from KAP.

## ODBC/JDBC Driver

KAP interacts BI with standard interface including ODBC/JDBC driver and REST API.

## Enterprise Features

KAP is designed for enterprise users. With its reliability and easy DevOps, KAP can smoothly blend into enterprise IT systems.

### Certification

KAP is fully compatible with Apache Hadoop and certified on all major Hadoop commercial distributions including Cloudera, Hortonworks, MapR, Huawei FusionInsight, Azure HDInsight and AWS EMR.

### Cell Level ACL

Control your big data access on different level of granularity at project level, table level, row level or cell level.

### Enhanced Security

Supports Kerberos, SSO and LDAP for high security requirement environment and use cases.

### Easy DevOps

Dev Ops is easy with simplified installation, upgrade and enhanced metadata management. KAP also supports high-availability and load-balance deployment.

	Apache Kylin	KAP	KAP Plus
Position	Open Source OLAP on Hadoop	Enterprise OLAP on Hadoop	Enterprise OLAP on Hadoop
Architecture	MOLAP	HOLAP	HOLAP
Query Performance	Sub-second latency	Sub-second latency	Sub-second latency
Parallel Computing	HBase Coprocessor	HBase Coprocessor	Spark
Storage Engine	HBase	HBase	KyStorage*
Push Down	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Detailed Data Query	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Visualized Modeling	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Assisted Modeling	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
LDAP/ Kerberos	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Cell Level Security	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Built-in BI Tool	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
KyBot self-service	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Technical Support	Community without SLA	24 X 7 SLA	24 X 7 SLA

\* KyStorage is a columnar storage engine based on HDFS developed by Kyligence.





# Unleash Big Data Productivity



Kyligence Inc.  
info@kyligence.io  
+1 408-524-2964  
Santa Clara, CA