



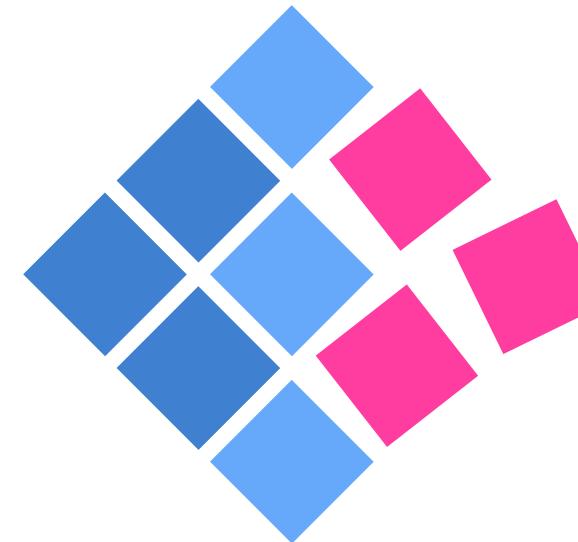
Apache Kylin 架构与案例

HadoopCon 2016 Taipei

韩卿 | Luke Han
Apache Kylin co-creator & PMC chair
lukehan@apache.org

About me...

- 韩卿 | Luke Han
 - Kyligence Inc Co-founder & CEO
 - Apache Kylin Creator, PMC Chair
 - ASF Member, Microsoft MVP
 - lukehan@apache.org
 - @lukehq



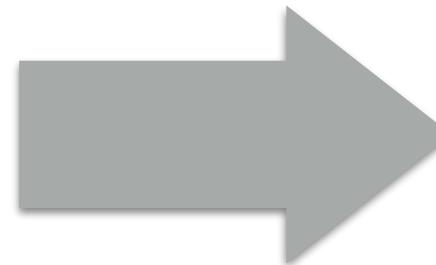
Kyligence

=

Kylin + Intelligence

A leading intelligent big data analytics company

Startup



cloudera®



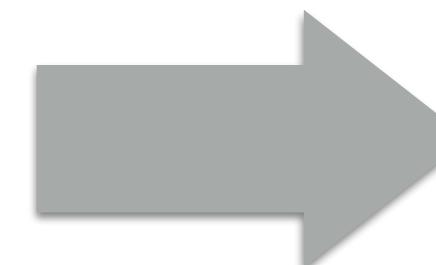
databricks™



Apache Kafka

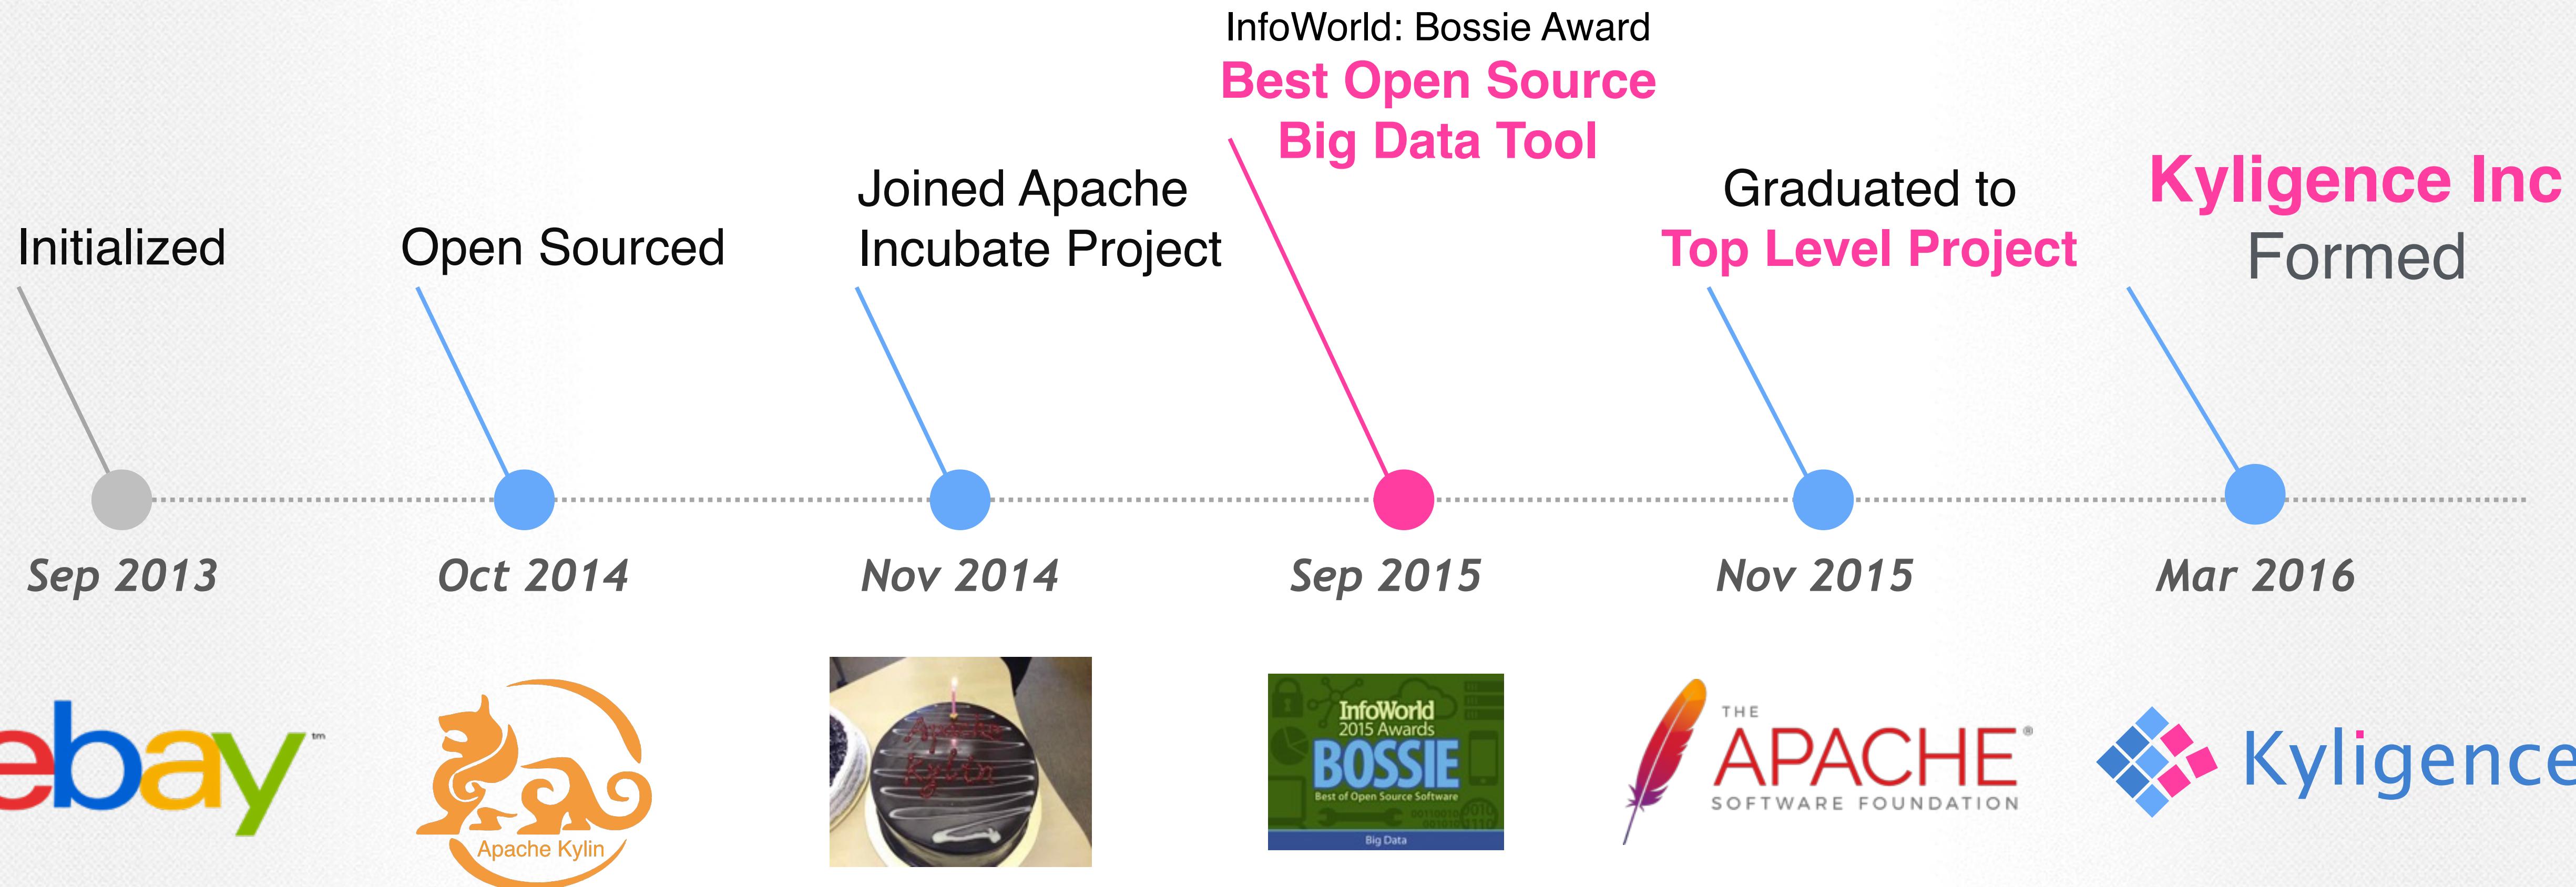


confluent



Kyligence

History



Apache Kylin

- ✓ Leading open source OLAP on Hadoop
- ✓ Fast growing open source community
- ✓ Adopted by 100+ global organizations
- ✓ Only born in China Apache Top Level Project
- ✓ Bossie Award: Best Open Source Big Data Tool

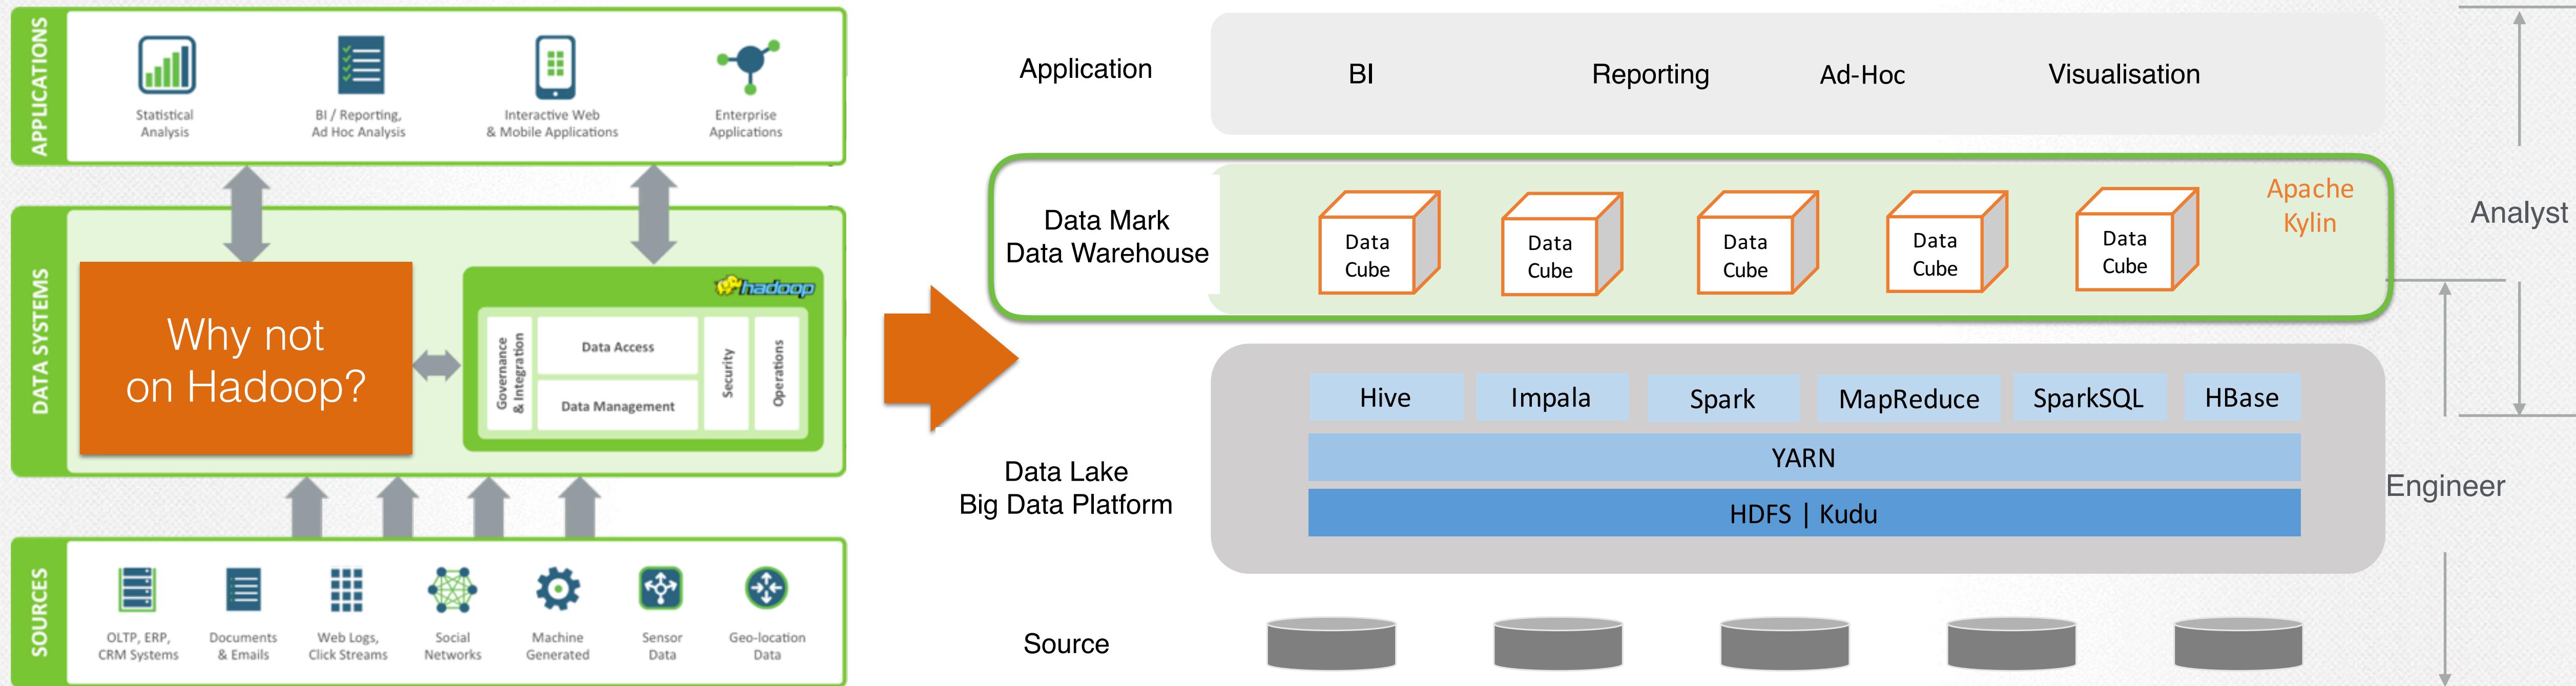


The technical aspects of Kylin are exciting, of course, but just as exciting is the way that Kylin represents a growing involvement of Asian countries like China in the Open Source community.

—Ted Dunning, VP of Apache IPMC



Industry Needs



Global Users





Why Kylin?

Volume

- **Data Volume**
 - eBay: 100+B per cube
 - Baidu Map: 10+B user behaviour data
 - Guandong Mobile: 10+B per day

Speed

- **Performance**
 - eBay: 90% < 1.18s
 - Meituan : 95%<1s, 99%<3s
 - Beijing Mobile: 36+ times fast than SparkSQL solution

Users

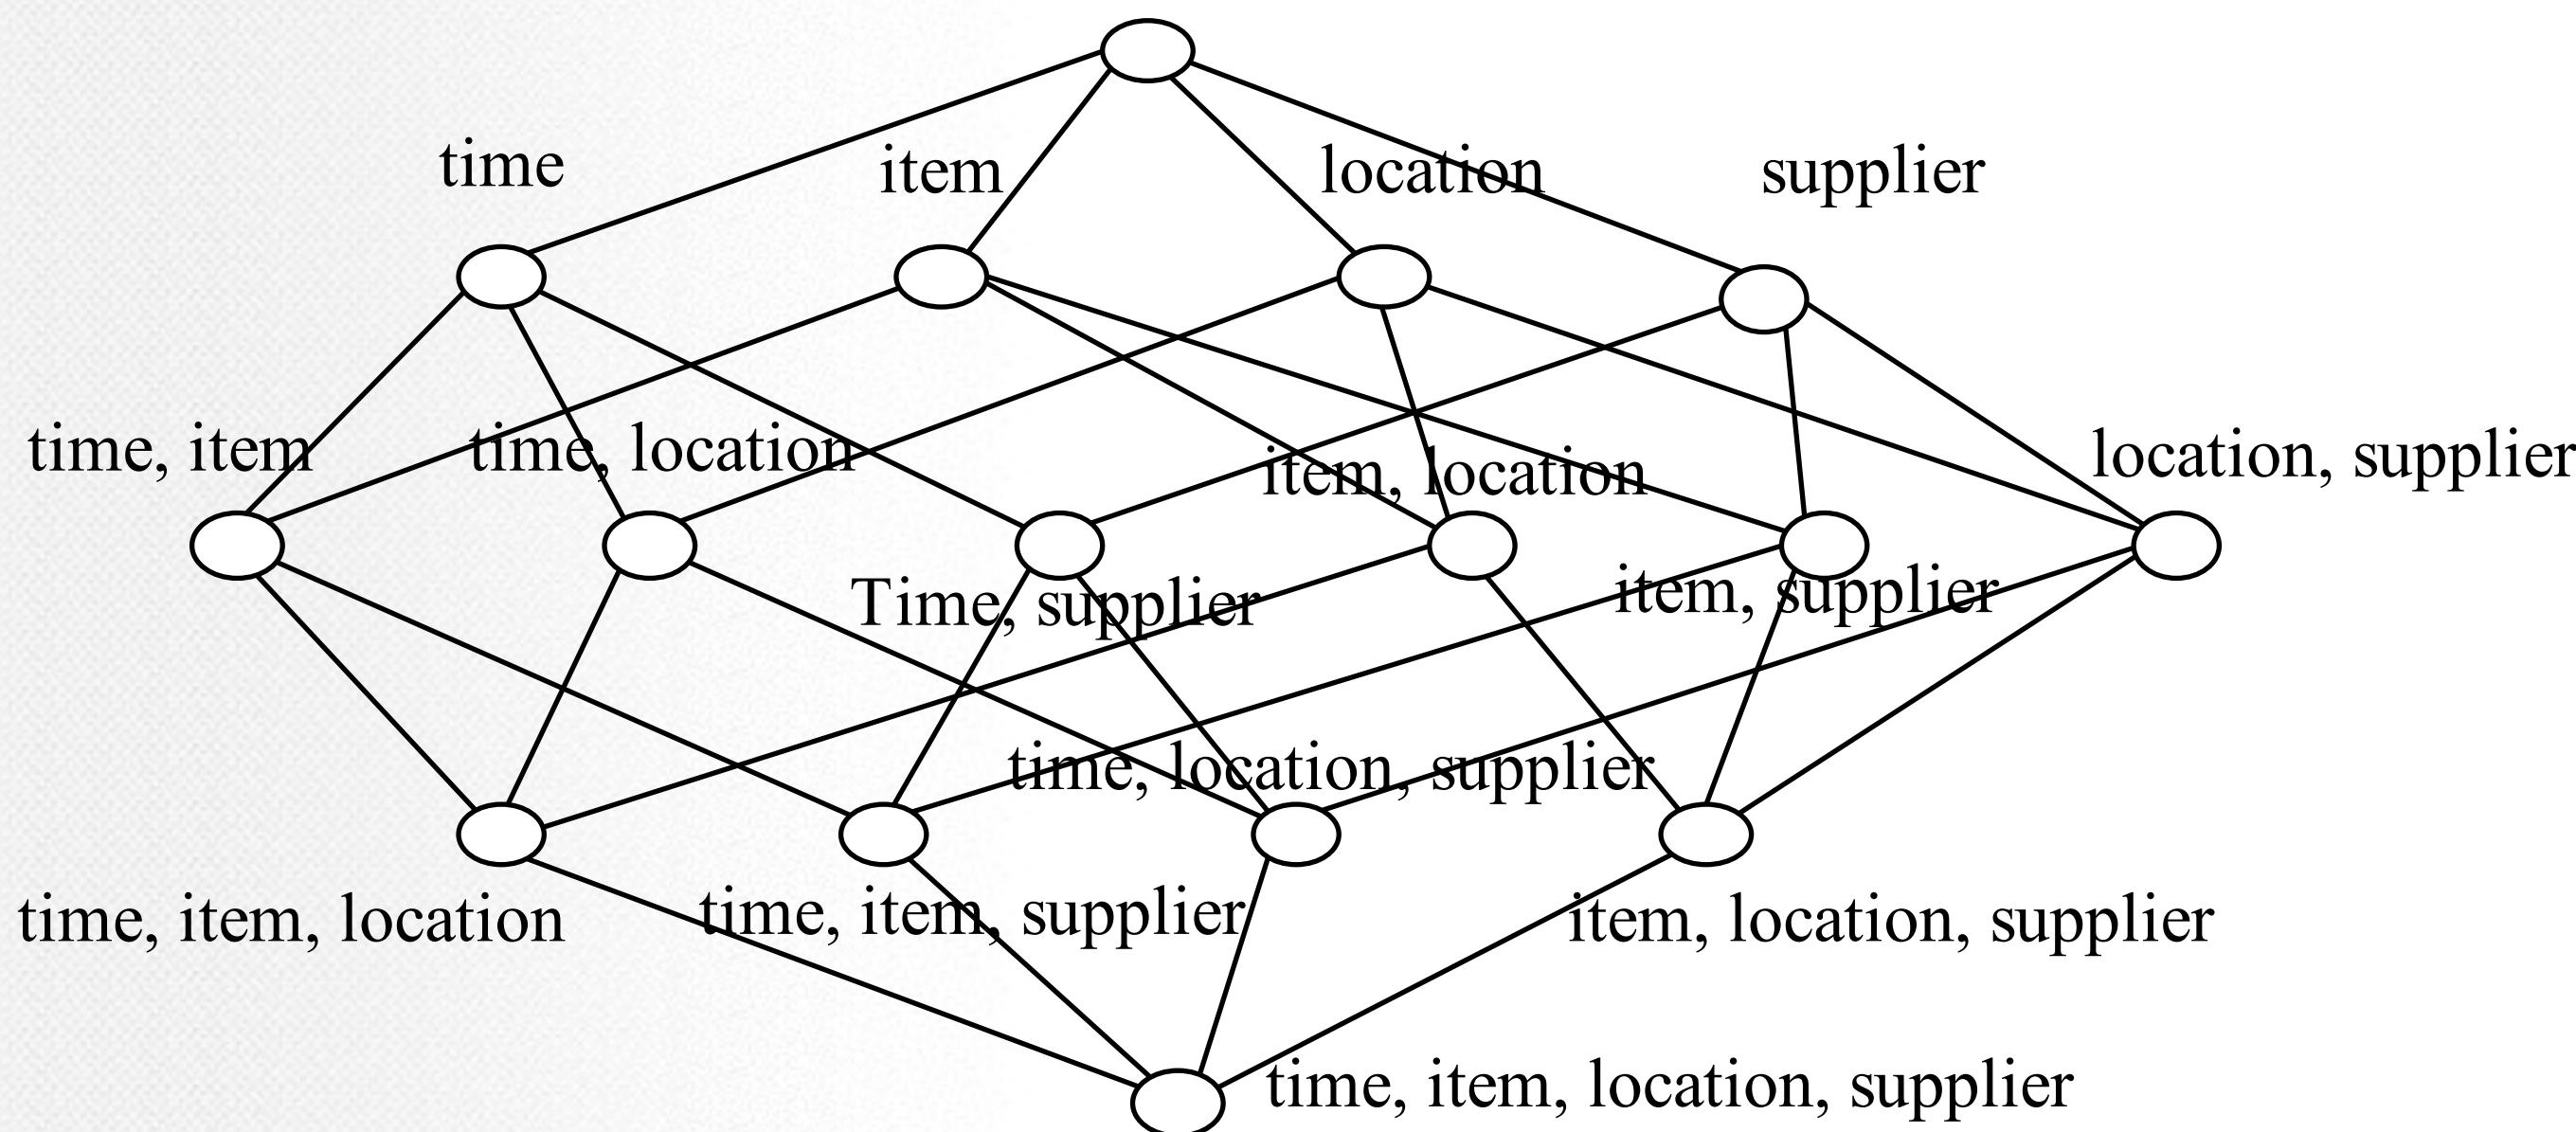
- **Users**
 - 100+ real use cases
 - Only one born in China Apache Top Level Project
 - “Best Open Source Big Data Tool”

Cost

- **Cost**
 - eBay: From Teradata to Kylin/Hadoop
 - VIP.com: Replaced Greenplum with Kylin/Hadoop
 - Guomei: 3x fast implementation

Architecture

Balance between Space & Time



0-D(apex) cuboid

1-D cuboids

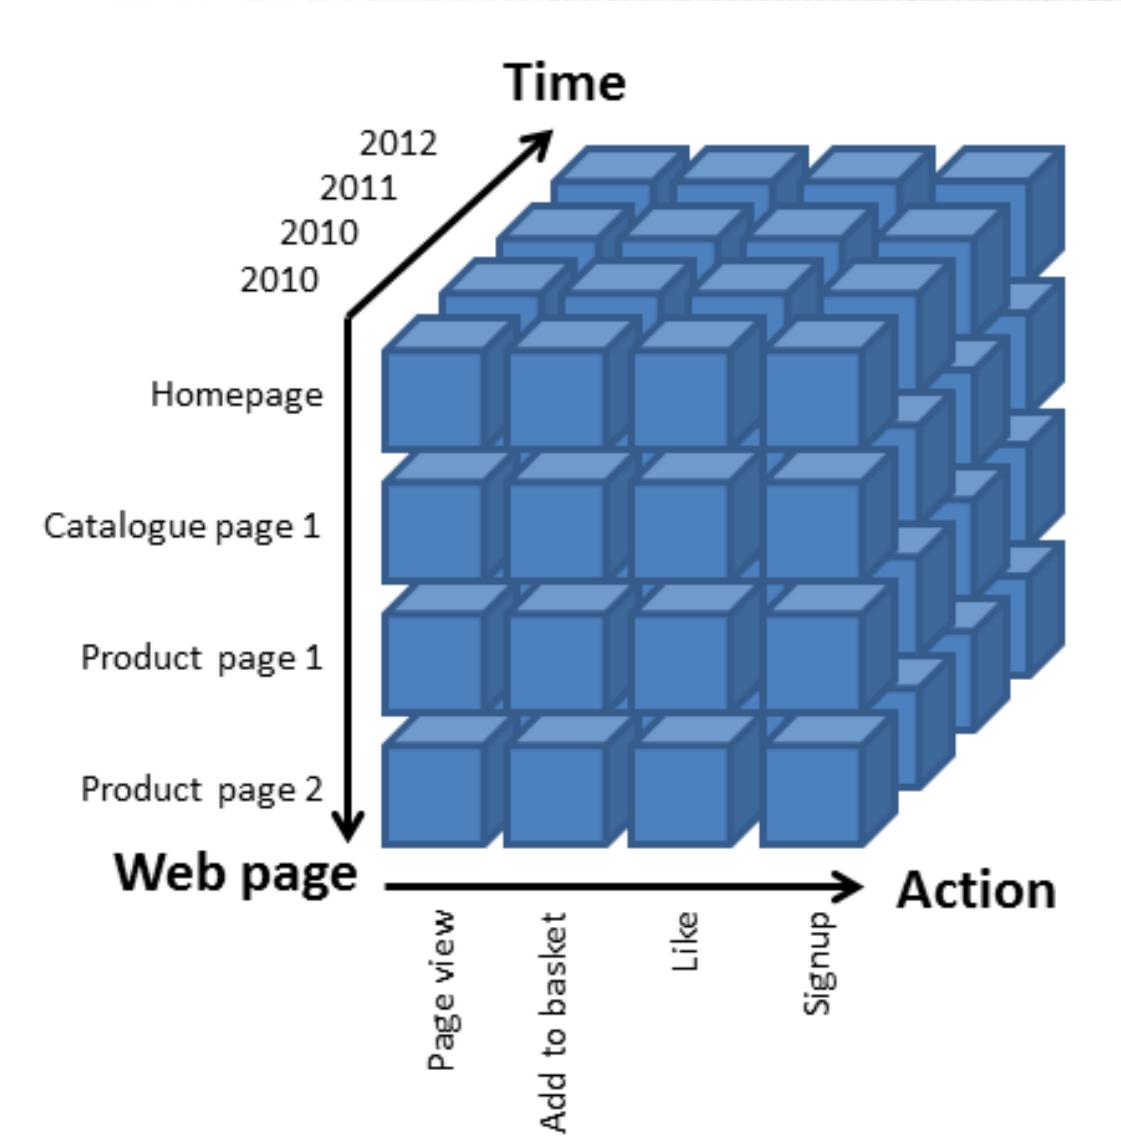
2-D cuboids

3-D cuboids

4-D(base) cuboid

OLAP Cube

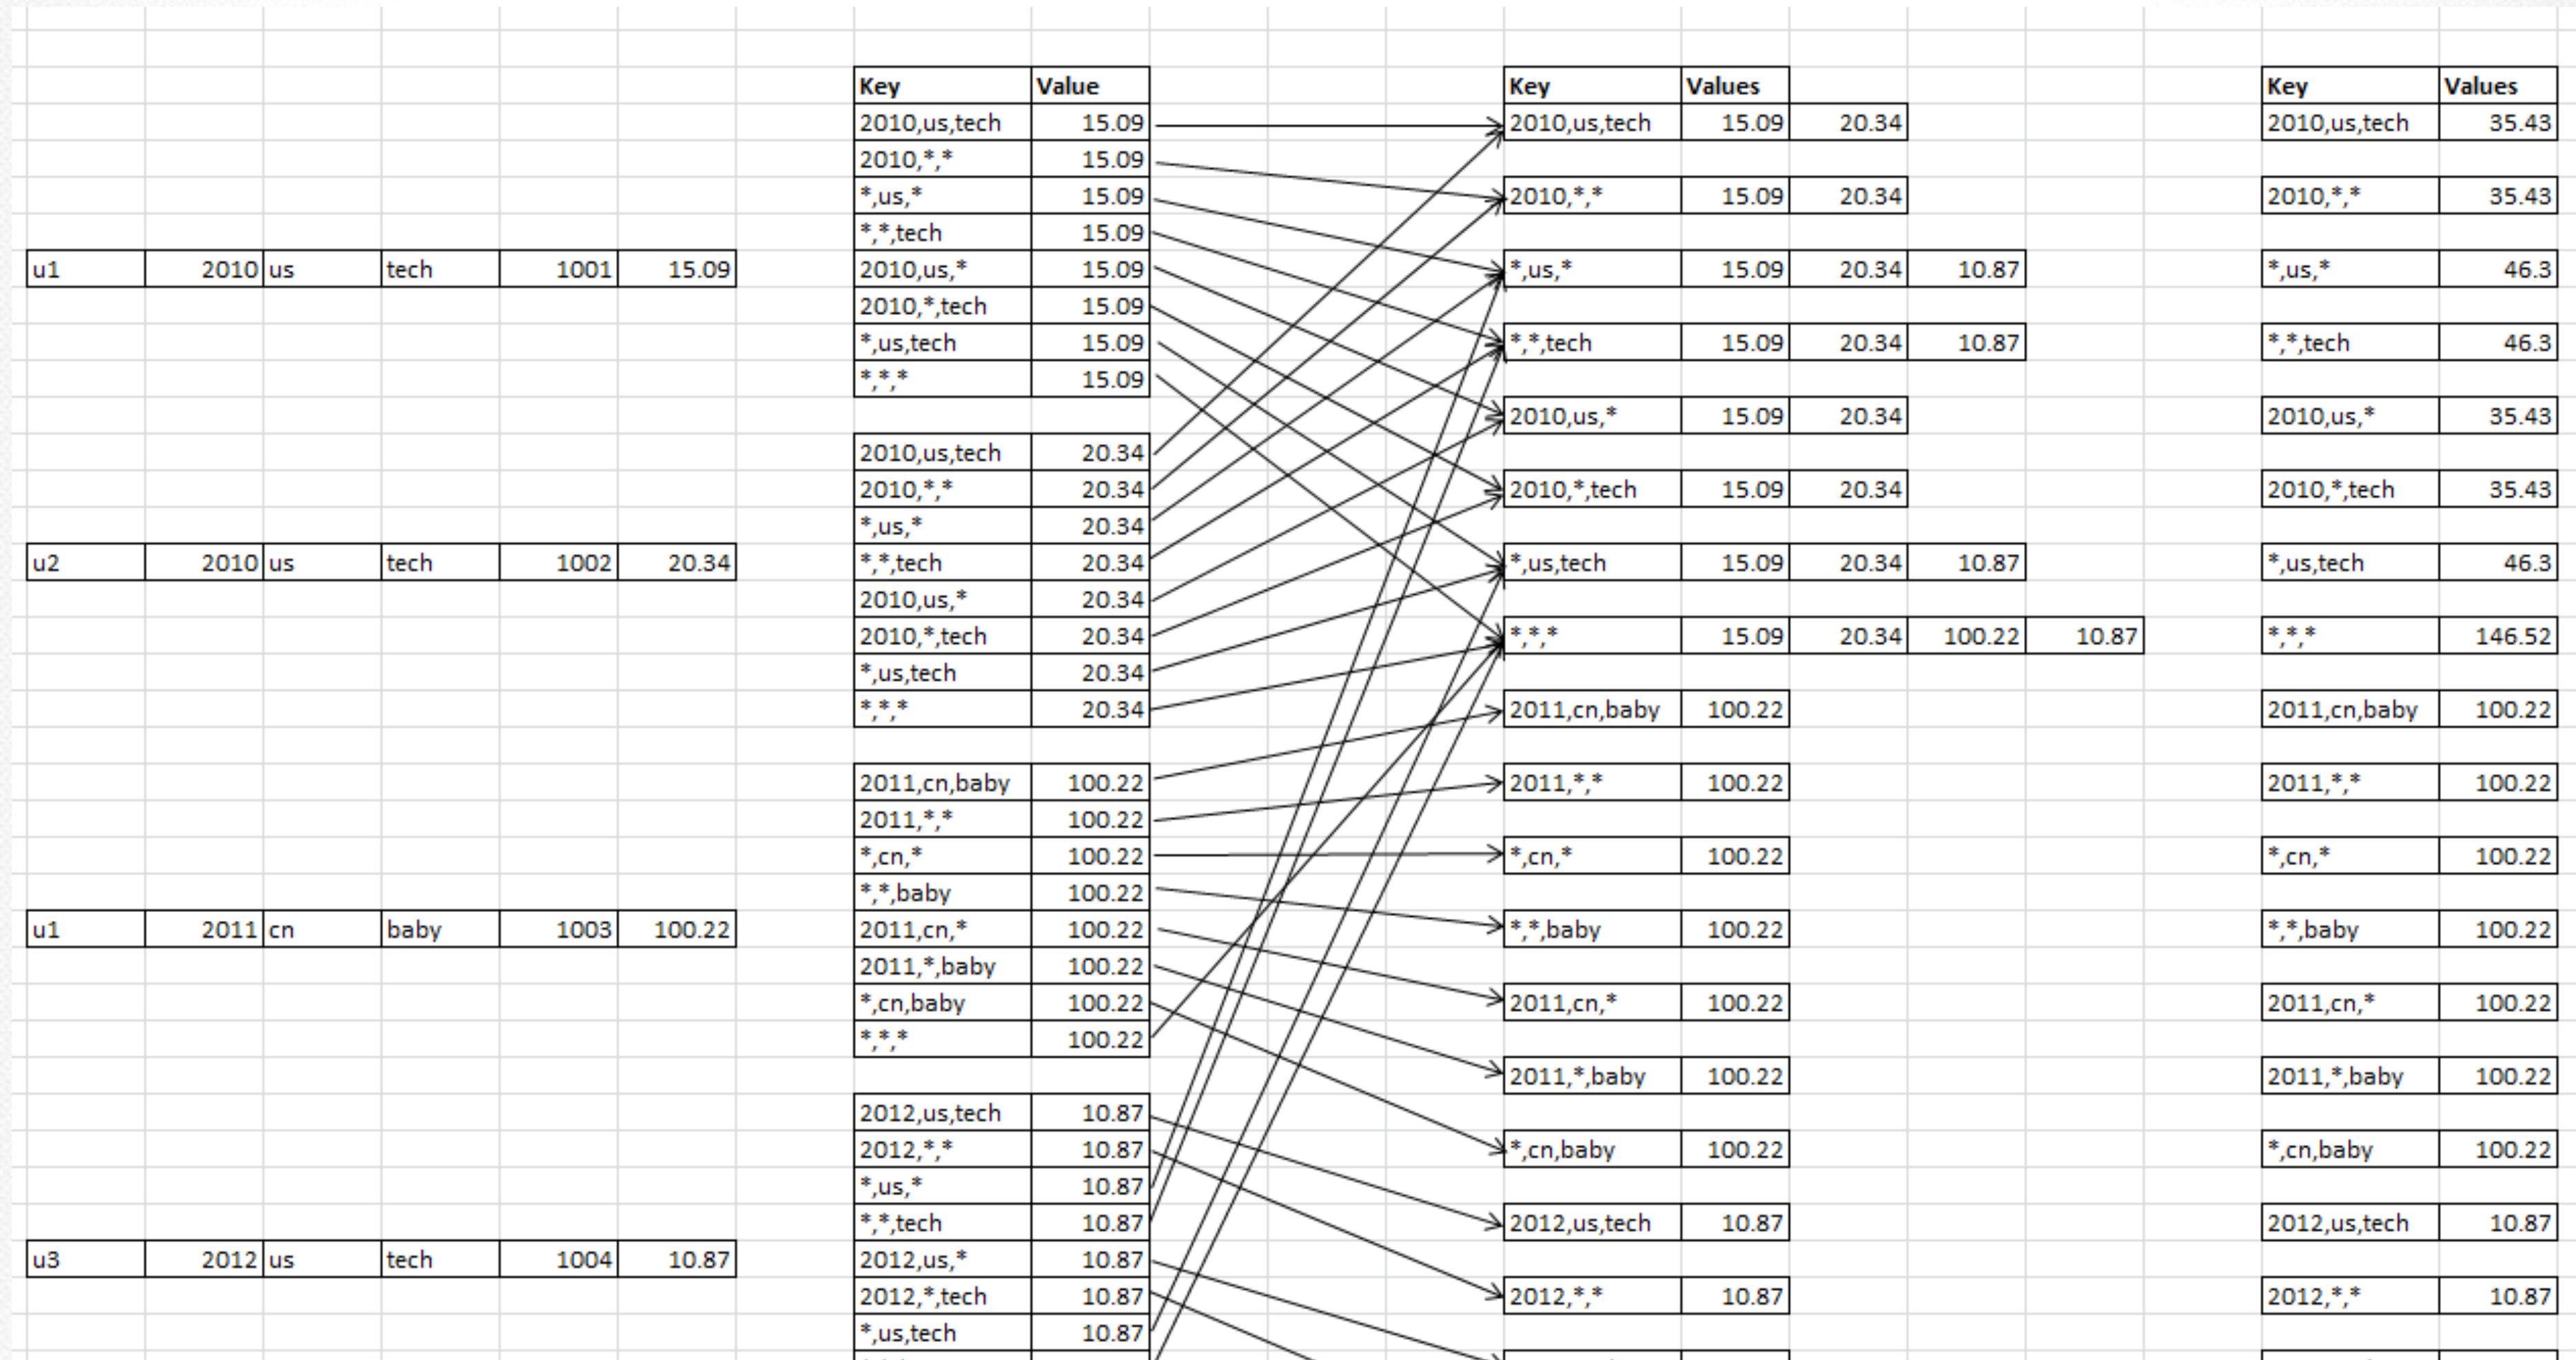
- Cuboid = one combination of dimensions
- Cube = all combination of dimensions (all cuboids)



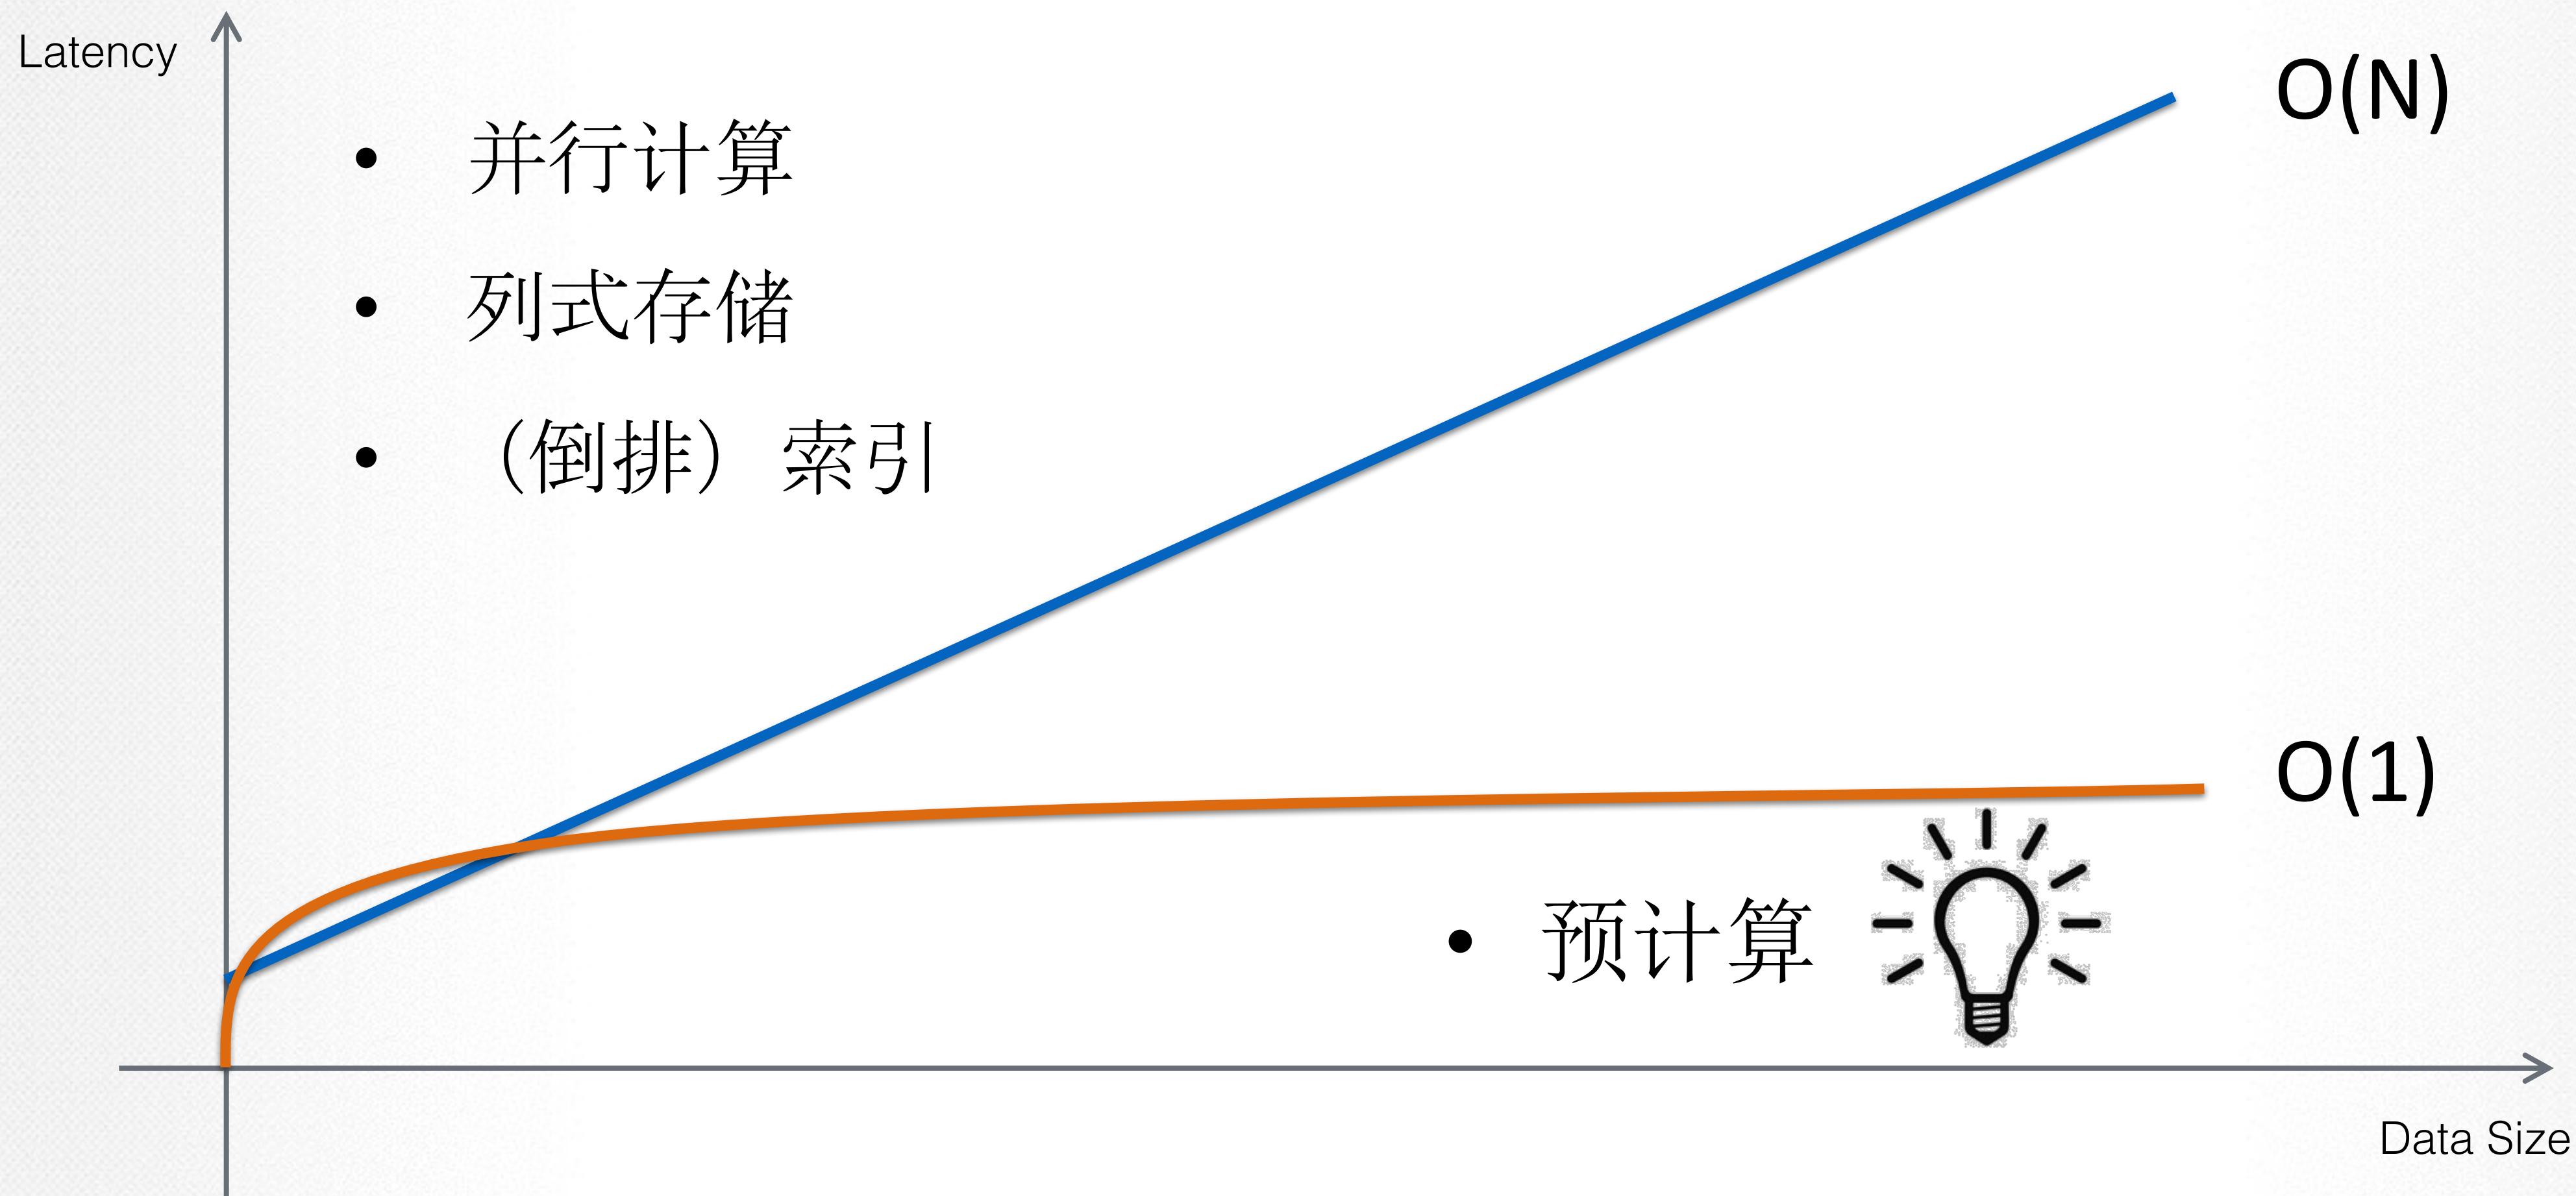
- Base vs. aggregate cells; ancestor vs. descendant cells; parent vs. child cells

1. (9/15, milk, Urbana, Dairy_land) - <time, item, location, supplier>
2. (9/15, milk, Urbana, *) - <time, item, location>
3. (*, milk, Urbana, *) - <item, location>
4. (*, milk, Chicago, *) - <item, location>
5. (*, milk, *, *) - <item>

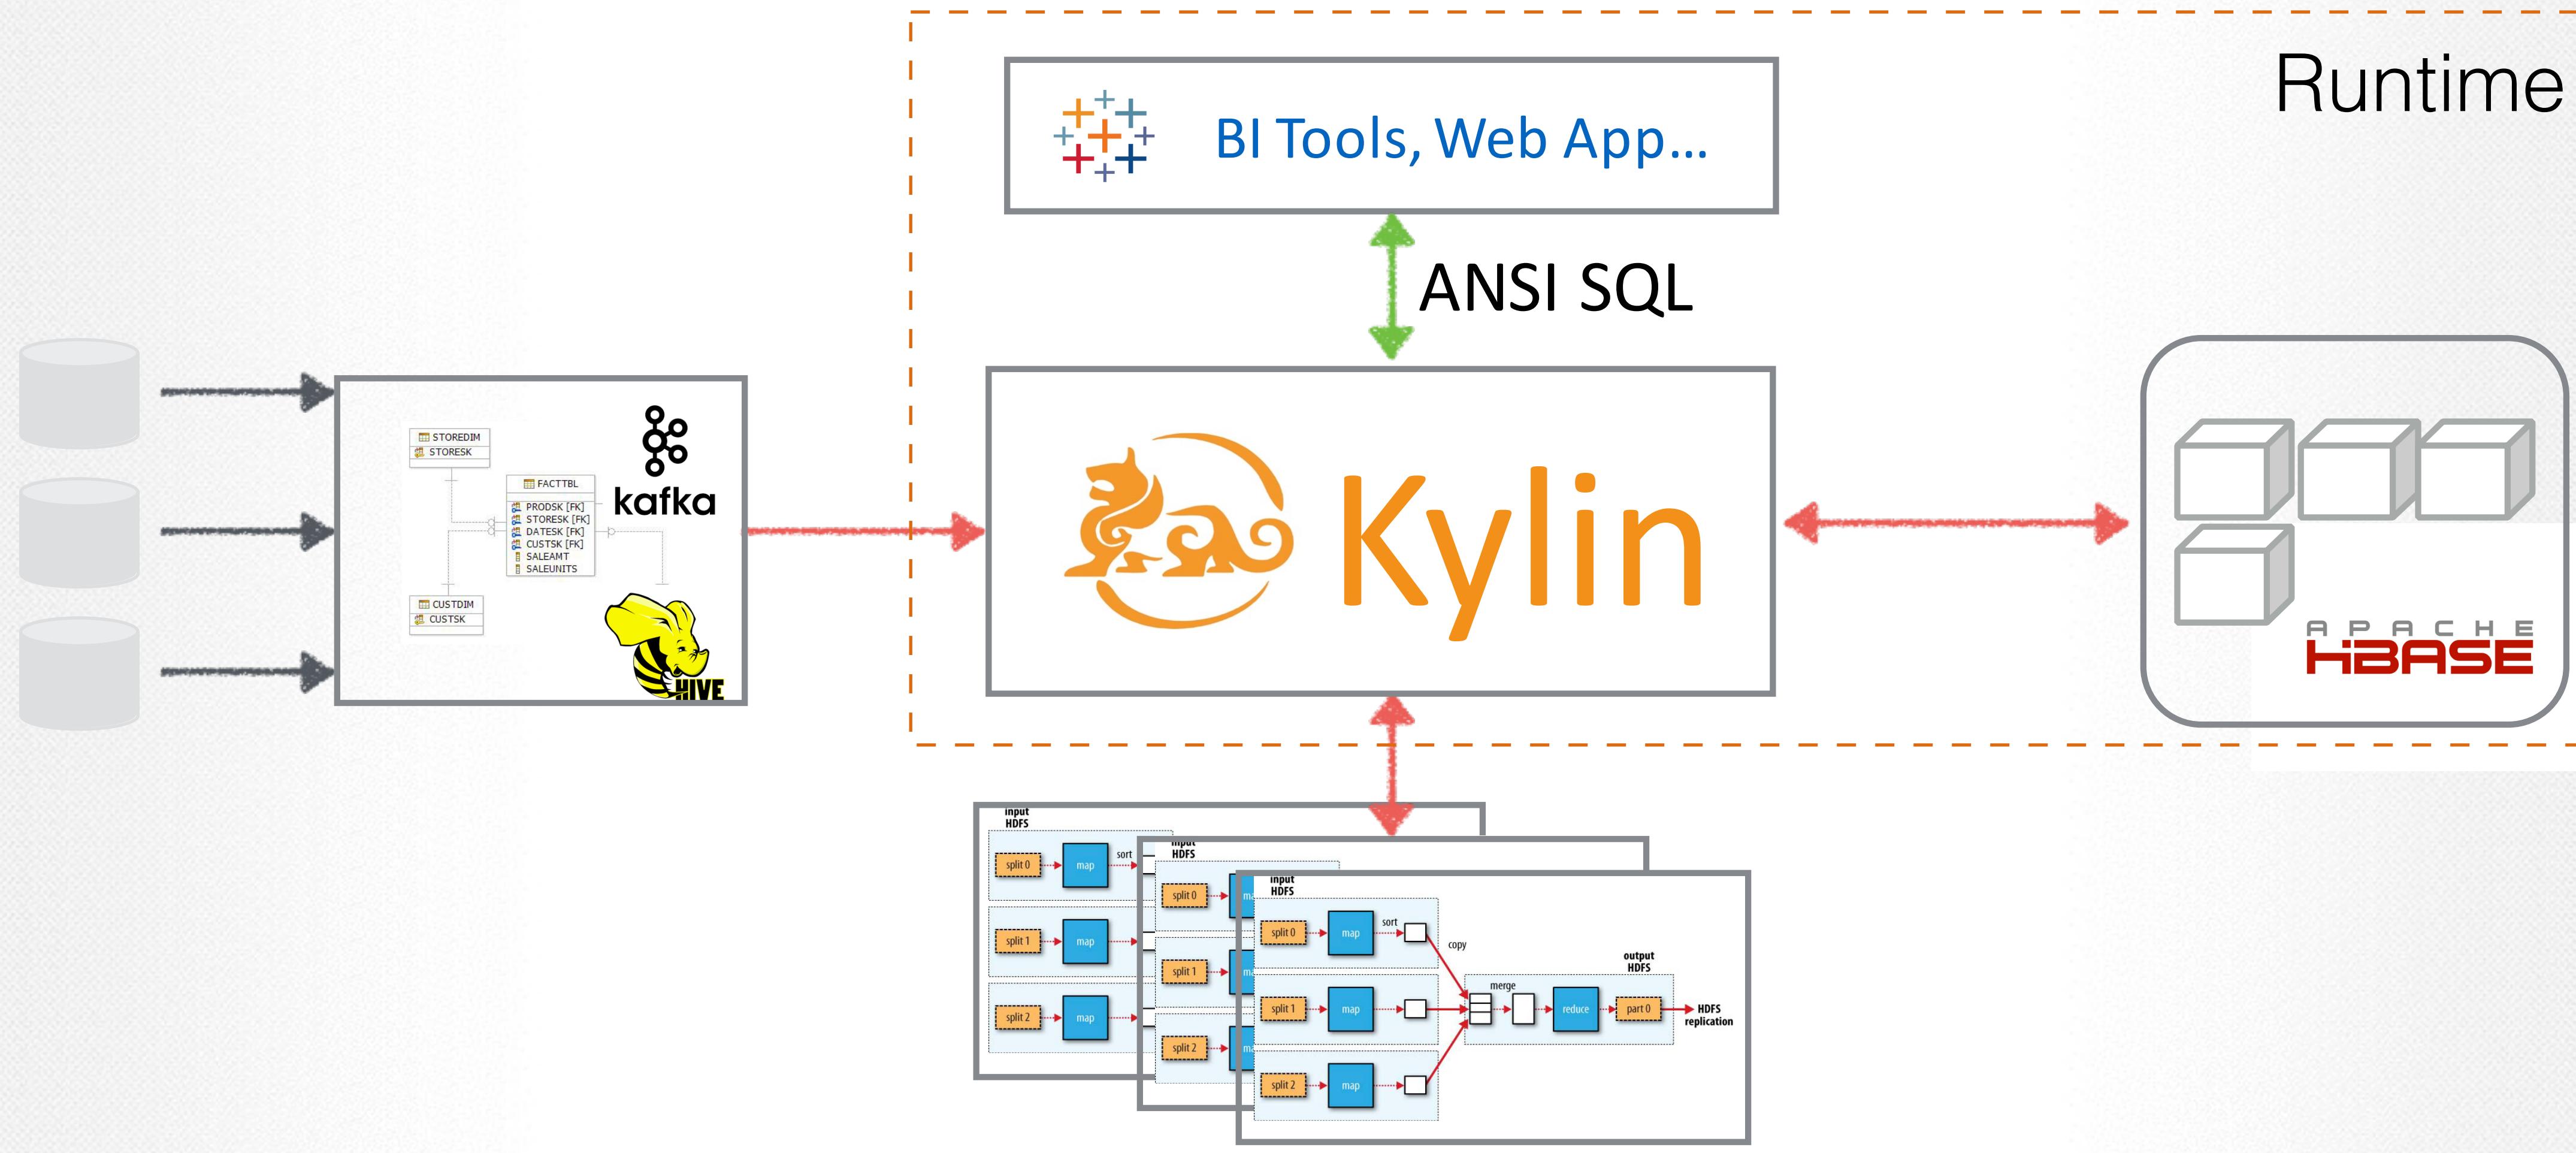
From Relational to KV



$O(n)$ vs $O(1)$

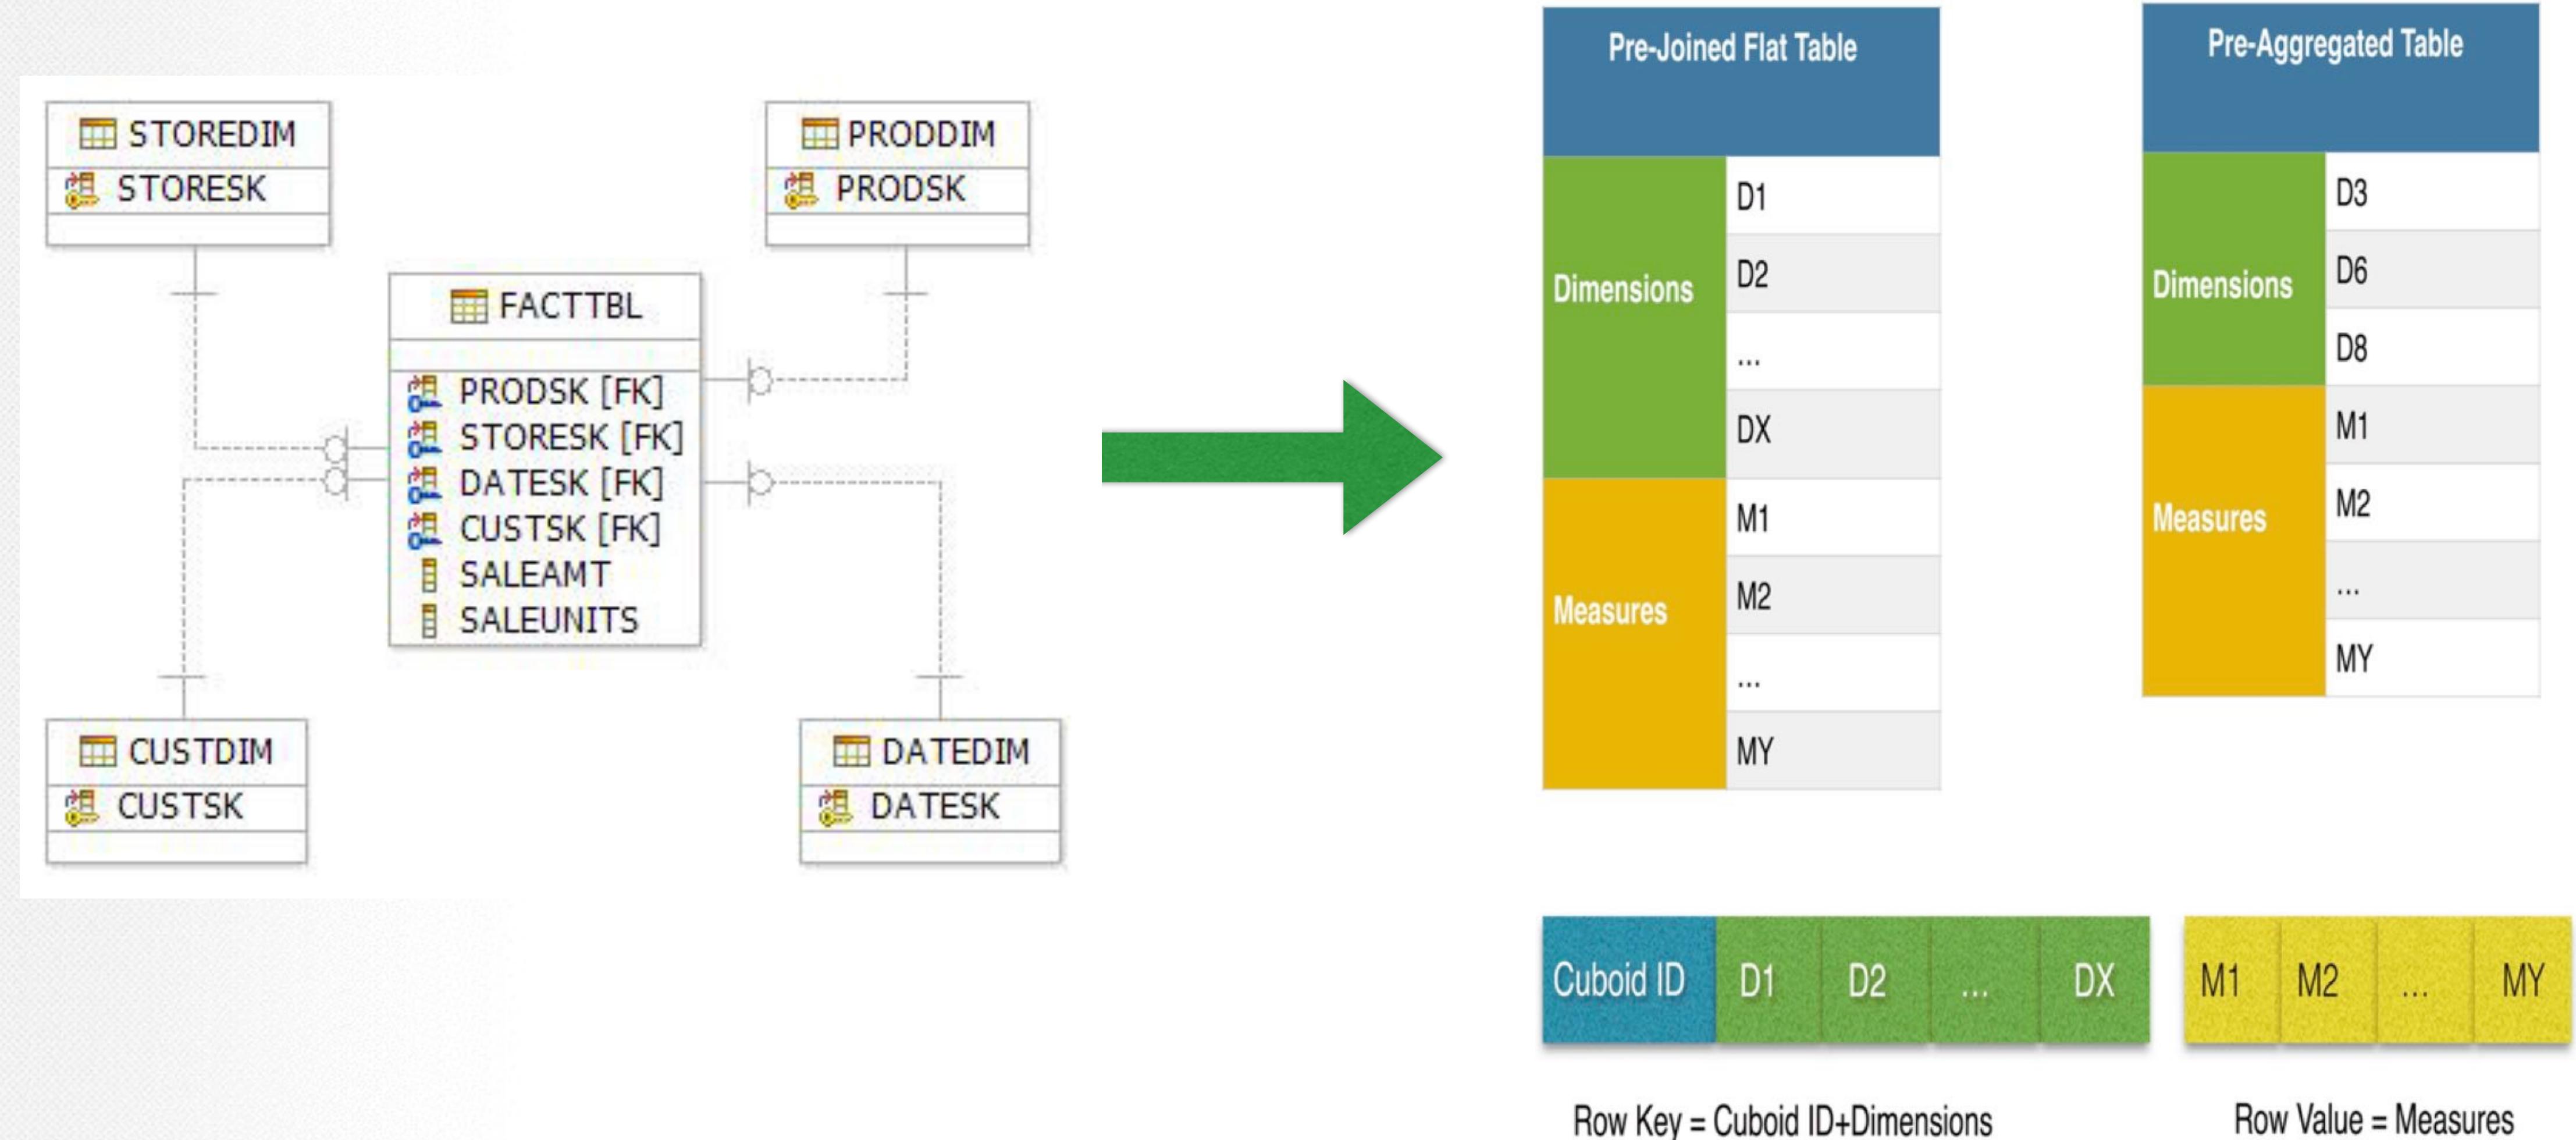


Architecture



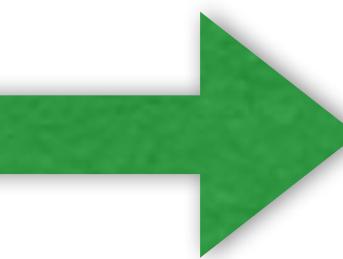
Map Reduce/Spark/Streaming...

Storage



SQL

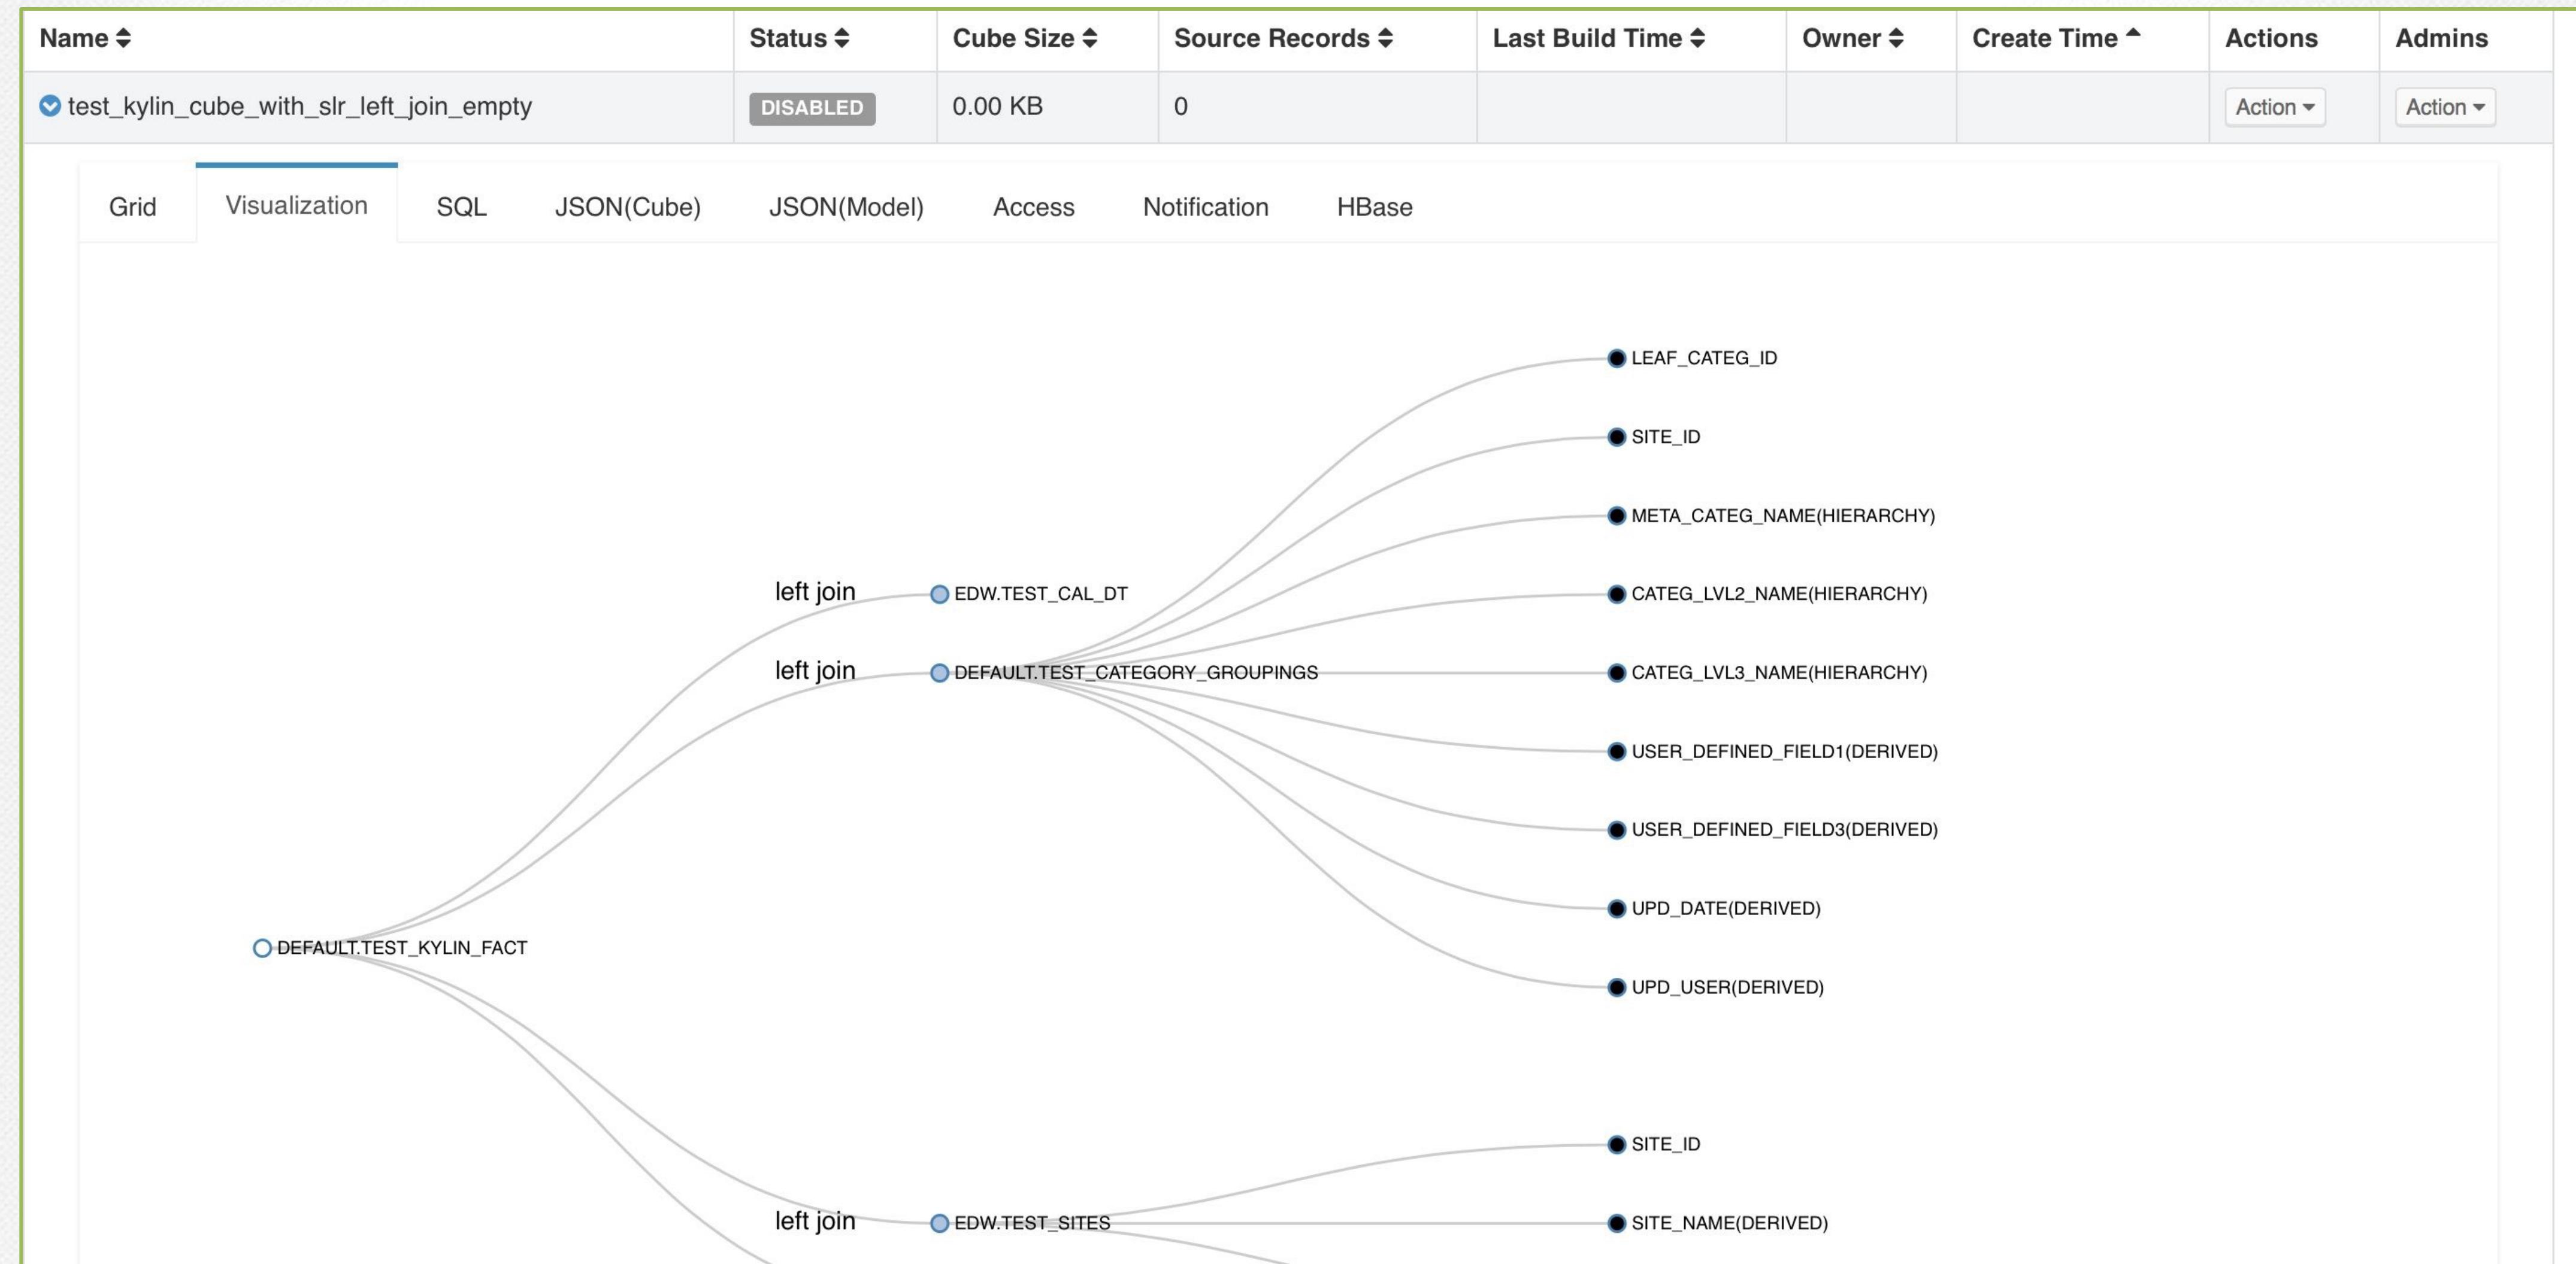
```
SELECT test_cal_dt.week_beg_dt,  
test_category.category_name,test_category.lvl2_name,  
test_category.lvl3_name,  
test_kylin_fact.lstg_format_name, test_sites.site_name,  
SUM(test_kylin_fact.price) AS GMV, COUNT(*) AS  
TRANS_CNT  
FROM test_kylin_fact  
LEFT JOIN test_cal_dt ON test_kylin_fact.cal_dt =  
test_cal_dt.cal_dt  
LEFT JOIN test_category ON test_kylin_fact.leaf_categ_id  
= test_category.leaf_categ_id AND  
test_kylin_fact.lstg_site_id = test_category.site_id  
LEFT JOIN test_sites ON test_kylin_fact.lstg_site_id =  
test_sites.site_id  
WHERE test_kylin_fact.seller_id = 123456 OR  
test_kylin_fact.lstg_format_name = 'New'  
GROUP BY test_cal_dt.week_beg_dt,  
test_category.category_name,test_category.lvl2_name,  
test_category.lvl3_name,  
test_kylin_fact.lstg_format_name,test_sites.site_name
```



OLAPToEnumerableConverter

OLAPProjectRel(WEEK_BEG_DT=[\\$0], category_name=[\\$1],
CATEG_LVL2_NAME=[\\$2], CATEG_LVL3_NAME=[\\$3], LSTG_FORMAT_NAME=[\\$4],
SITE_NAME=[\\$5], GMV=[CASE=(\\$7, 0, null, \\$6)], TRANS_CNT=[\\$8])
OLAPAggregateRel(group={[0, 1, 2, 3, 4, 5]}, agg#0=[\\$SUM0(\\$6)],
agg#1=[COUNT(\\$6)], TRANS_CNT=[COUNT()])
OLAPProjectRel(WEEK_BEG_DT=[\\$13], category_name=[\\$21],
CATEG_LVL2_NAME=[\\$15], CATEG_LVL3_NAME=[\\$14],
LSTG_FORMAT_NAME=[\\$5], SITE_NAME=[\\$23], PRICE=[\\$0])
OLAPFilterRel(condition=[OR(=(\\$3, 123456),=(\\$5, 'New'))])
OLAPJoinRel(condition=[=(\\$2, \\$25)], joinType=[left])
OLAPJoinRel(condition=[AND(=(\\$6, \\$22),=(\\$2, \\$17))], joinType=[left])
OLAPJoinRel(condition=[=(\\$4, \\$12)], joinType=[left])
OLAPTableScan(table=[[DEFAULT, TEST_KYLIN_FACT]], fields=[[0, 1, 2, 3,
4, 5, 6, 7, 8, 9, 10, 11]])
OLAPTableScan(table=[[DEFAULT, TEST_CAL_DT]], fields=[[0, 1]])
OLAPTableScan(table=[[DEFAULT, test_category]], fields=[[0, 1, 2, 3, 4, 5,
6, 7, 8]])
OLAPTableScan(table=[[DEFAULT, TEST_SITES]], fields=[[0, 1, 2]])

Design Data Model



Manage Job

Jobs						
<input type="checkbox"/> NEW	<input type="checkbox"/> PENDING	<input type="checkbox"/> RUNNING	<input type="checkbox"/> FINISHED	<input type="checkbox"/> ERROR	<input type="checkbox"/> DISCARDED	
Job Name	Cube	Progress	Last Modified Time	Duration	Actions	
PC_SESSION_COPY - 20150325000000_20150326000000 - BUILD - PDT 2015-04-12 16:56:38	PC_SESSION_COPY	<div style="width: 100%;">100%</div>	2015-04-12 18:20:41 PST	144.00 mins	N/A	
PC_SESSION_COPY - 20150324000000_20150325000000 - BUILD - PDT 2015-04-12 06:49:57	PC_SESSION_COPY	<div style="width: 100%;">100%</div>	2015-04-12 09:07:41 PST	197.00 mins	N/A	
PC_SESSION_COPY - 20150323000000_20150324000000 - BUILD - PDT 2015-04-09 22:34:08	PC_SESSION_COPY	<div style="width: 100%;">100%</div>	2015-04-09 23:57:09 PST	143.00 mins	N/A	

Detail Information

Job Name	PC_SESSION_COPY - 20150325000000_20150326000000 - BUILD - PDT 2015-04-12 16:56:38
Job ID	bae5c8b9-78e0-4849-a579-042a3f05cab1
Status	FINISHED
Duration	144.00 mins
MapReduce Waiting	3.62 mins

Output

```
2015-04-12 18:17:52.434 - State of Hadoop job: job_1427705526386_166781:FINISHED - SUCCEEDED
Counters: 52
File System Counters
FILE: Number of bytes read=27406033130
FILE: Number of bytes written=56621871289
FILE: Number of read operations=0
FILE: Number of large read operations=0
FILE: Number of write operations=0
HDFS: Number of bytes read=11690233914
HDFS: Number of bytes written=1143439
HDFS: Number of read operations=25315
HDFS: Number of large read operations=0
HDFS: Number of write operations=17
Job Counters
Killed map tasks=73
Launched map tasks=6401
Launched reduce tasks=1
Other local map tasks=109
```

MapReduce Job job_1415168056392_0001

Logged in as: dr.who

Kylin 127.0.0.1 MapReduce Job job_1415168056392_0001

Job Overview

Job Name: INSERT OVERWRITE TABL...E_DIM.SELLER_TYPE_CD(Stage-10)
User Name: root
Queue: default
State: SUCCEEDED
Uberized: false
Submitted: Wed Nov 05 01:31:34 PST 2014
Started: Wed Nov 05 01:32:19 PST 2014
Finished: Wed Nov 05 01:34:16 PST 2014
Elapsed: 1mins, 57sec
Diagnostics:
Average Map Time 1mins, 51sec

ApplicationMaster

Attempt Number	Start Time	Node	Logs
1	Wed Nov 05 01:32:01 PST 2014	sandbox.hortonworks.com:8042	logs

Task Type

Map	Reduce	Total
1	0	1

Attempt Type

Maps	Reduces	Failed	Killed	Successful
0	0	0	1	0

#7 Step Name: Update Cube Info
Duration: 0.00 mins

End 2015-04-12 18:20:40 PST

//kyligence.io

Explore Data

Kylin Query Cubes Jobs Tables Admin Help Welcome, ADMIN

Project: onlyinner

EDW

- TEST_CAL_DT
CAL_DT (date)
WEEK_BEG_DT (date)
- TEST_CATEGORY_GROUPINGS
CATEG_LVL3_NAME (varchar)
CATEG_LVL2_NAME (varchar)
UPD_USER (varchar)
SITE_ID (integer)
UPD_DATE (varchar)
USER_DEFINED_FIELD
USER_DEFINED_FIELD
META_CATEG_NAME (varchar)
LEAF_CATEG_ID (integer)
- TEST_KYLIN_FACT
PRICE (decimal(38, 16))
SLR_SEGMENT_CD (smallint)
LSTG_SITE_ID (integer)
SELLER_ID (bigint)
CAL_DT (date)
LSTG_FORMAT_NAME
LEAF_CATEG_ID (integer)
- TEST_SELLER_TYPE_DIM
- TEST_SITES

New Query Saved Queries Query History

```
1 select lstg_format_name, week_beg_dt, meta_categ_name, sum(price) as price
2 from test_kylin_fact
3 inner join test_category_groupings on test_category_groupings.leaf_categ_id = test_kylin_fact.leaf_categ_id
4 and test_category_groupings.site_id = test_kylin_fact.lstg_site_id
5 inner join test_cal_dt on test_cal_dt.cal_dt = test_kylin_fact.cal_dt
6 where week_beg_dt >= date'2013-01-01' and week_beg_dt <= date'2013-05-31'
7 group by week_beg_dt, test_category_groupings.meta_categ_name, lstg_format_name
```

Results (1215)

Drag a column header here and drop it to group by that column.

LSTG_FORMAT_NAME	WEEK_BEG_DT	META_CATEG_NAME	PRICE
FP-GTC	2013-01-01		
Auction	2013-01-01		
Others	2013-01-01		
Auction	2013-01-01		
Auction	2013-01-01		
Others	2013-01-01		
FP-GTC	2013-01-01		
Auction	2013-01-01		

Results

Status: All Start Time: 2014-11-12 17:33:48 Duration: 3.10s Rerun Save

Query String: `select lstg_format_name, week_beg_dt, meta_categ_name, sum(price) as price from test_kylin_fact inner join test_category_groupings on test_category_groupings.leaf_categ_id = test_kylin_fact.leaf_categ_id and test_category_groupings.site_id = test_kylin_fact.lstg_site_id inner join test_cal_dt on test_cal_dt.cal_dt = test_kylin_fact.cal_dt where week_beg_dt >= date'2013-01-01' and week_beg_dt <= date'2013-05-31' group by week_beg_dt, test_category_groupings.meta_categ_name, lstg_format_name`

Project: onlyinner Cubes: test_kylin_cube_with_slr_empty

Graph Type: Pie Chart

Dimensions: LSTG_FORMAT_NAME

Metrics: PRICE

Results

Status: All Start Time: 2014-11-12 17:33:48 Duration: 3.10s Rerun Save

Query String: `select lstg_format_name, week_beg_dt, meta_categ_name, sum(price) as price from test_kylin_fact inner join test_category_groupings on test_category_groupings.leaf_categ_id = test_kylin_fact.leaf_categ_id and test_category_groupings.site_id = test_kylin_fact.lstg_site_id inner join test_cal_dt on test_cal_dt.cal_dt = test_kylin_fact.cal_dt where week_beg_dt >= date'2013-01-01' and week_beg_dt <= date'2013-05-31' group by week_beg_dt, test_category_groupings.meta_categ_name, lstg_format_name`

Project: onlyinner Cubes: test_kylin_cube_with_slr_empty

Graph Type: Line Chart

Dimensions: WEEK_BEG_DT

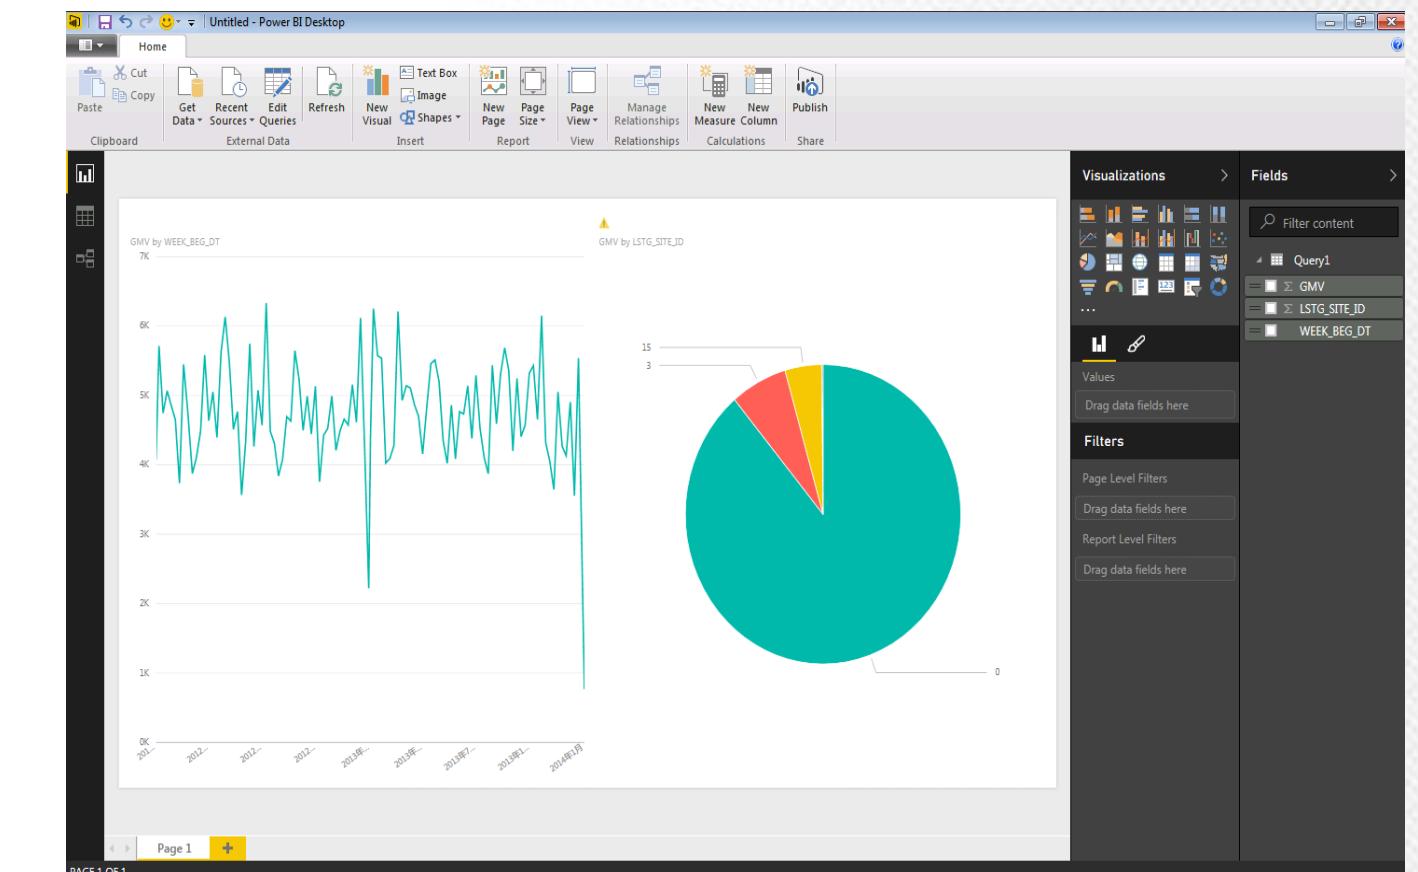
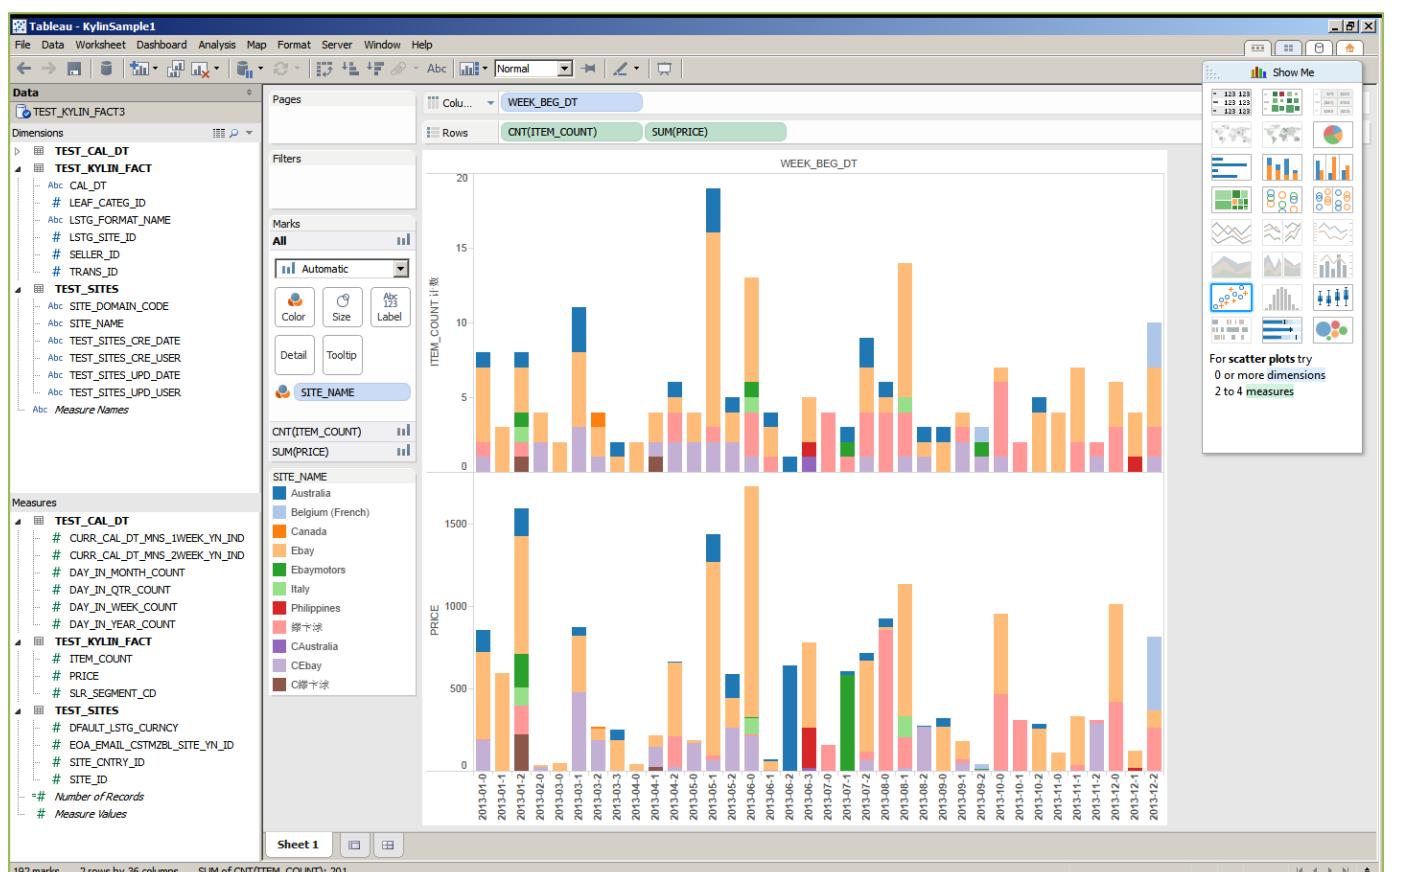
Metrics: PRICE

Home Page | Google Group

Home Page | Google Group

BI Tools

- Tableau
- Excel/PowerBI
- Zeppelin
- Caravel



This screenshot shows an Excel spreadsheet with a large data grid. A 'Query Editor' window is open, showing a table named 'DEFAULT_TEST' with 4,550 rows loaded. The table contains columns such as CAL_DT, LEAF_CATEG_ID, LSTG_SITE_ID, LSTG_FORMAT_NAME, SLR_SEGMENT_CD, PRICE, ITEM_COUNT, SELLER_ID, and TRANS_ID.

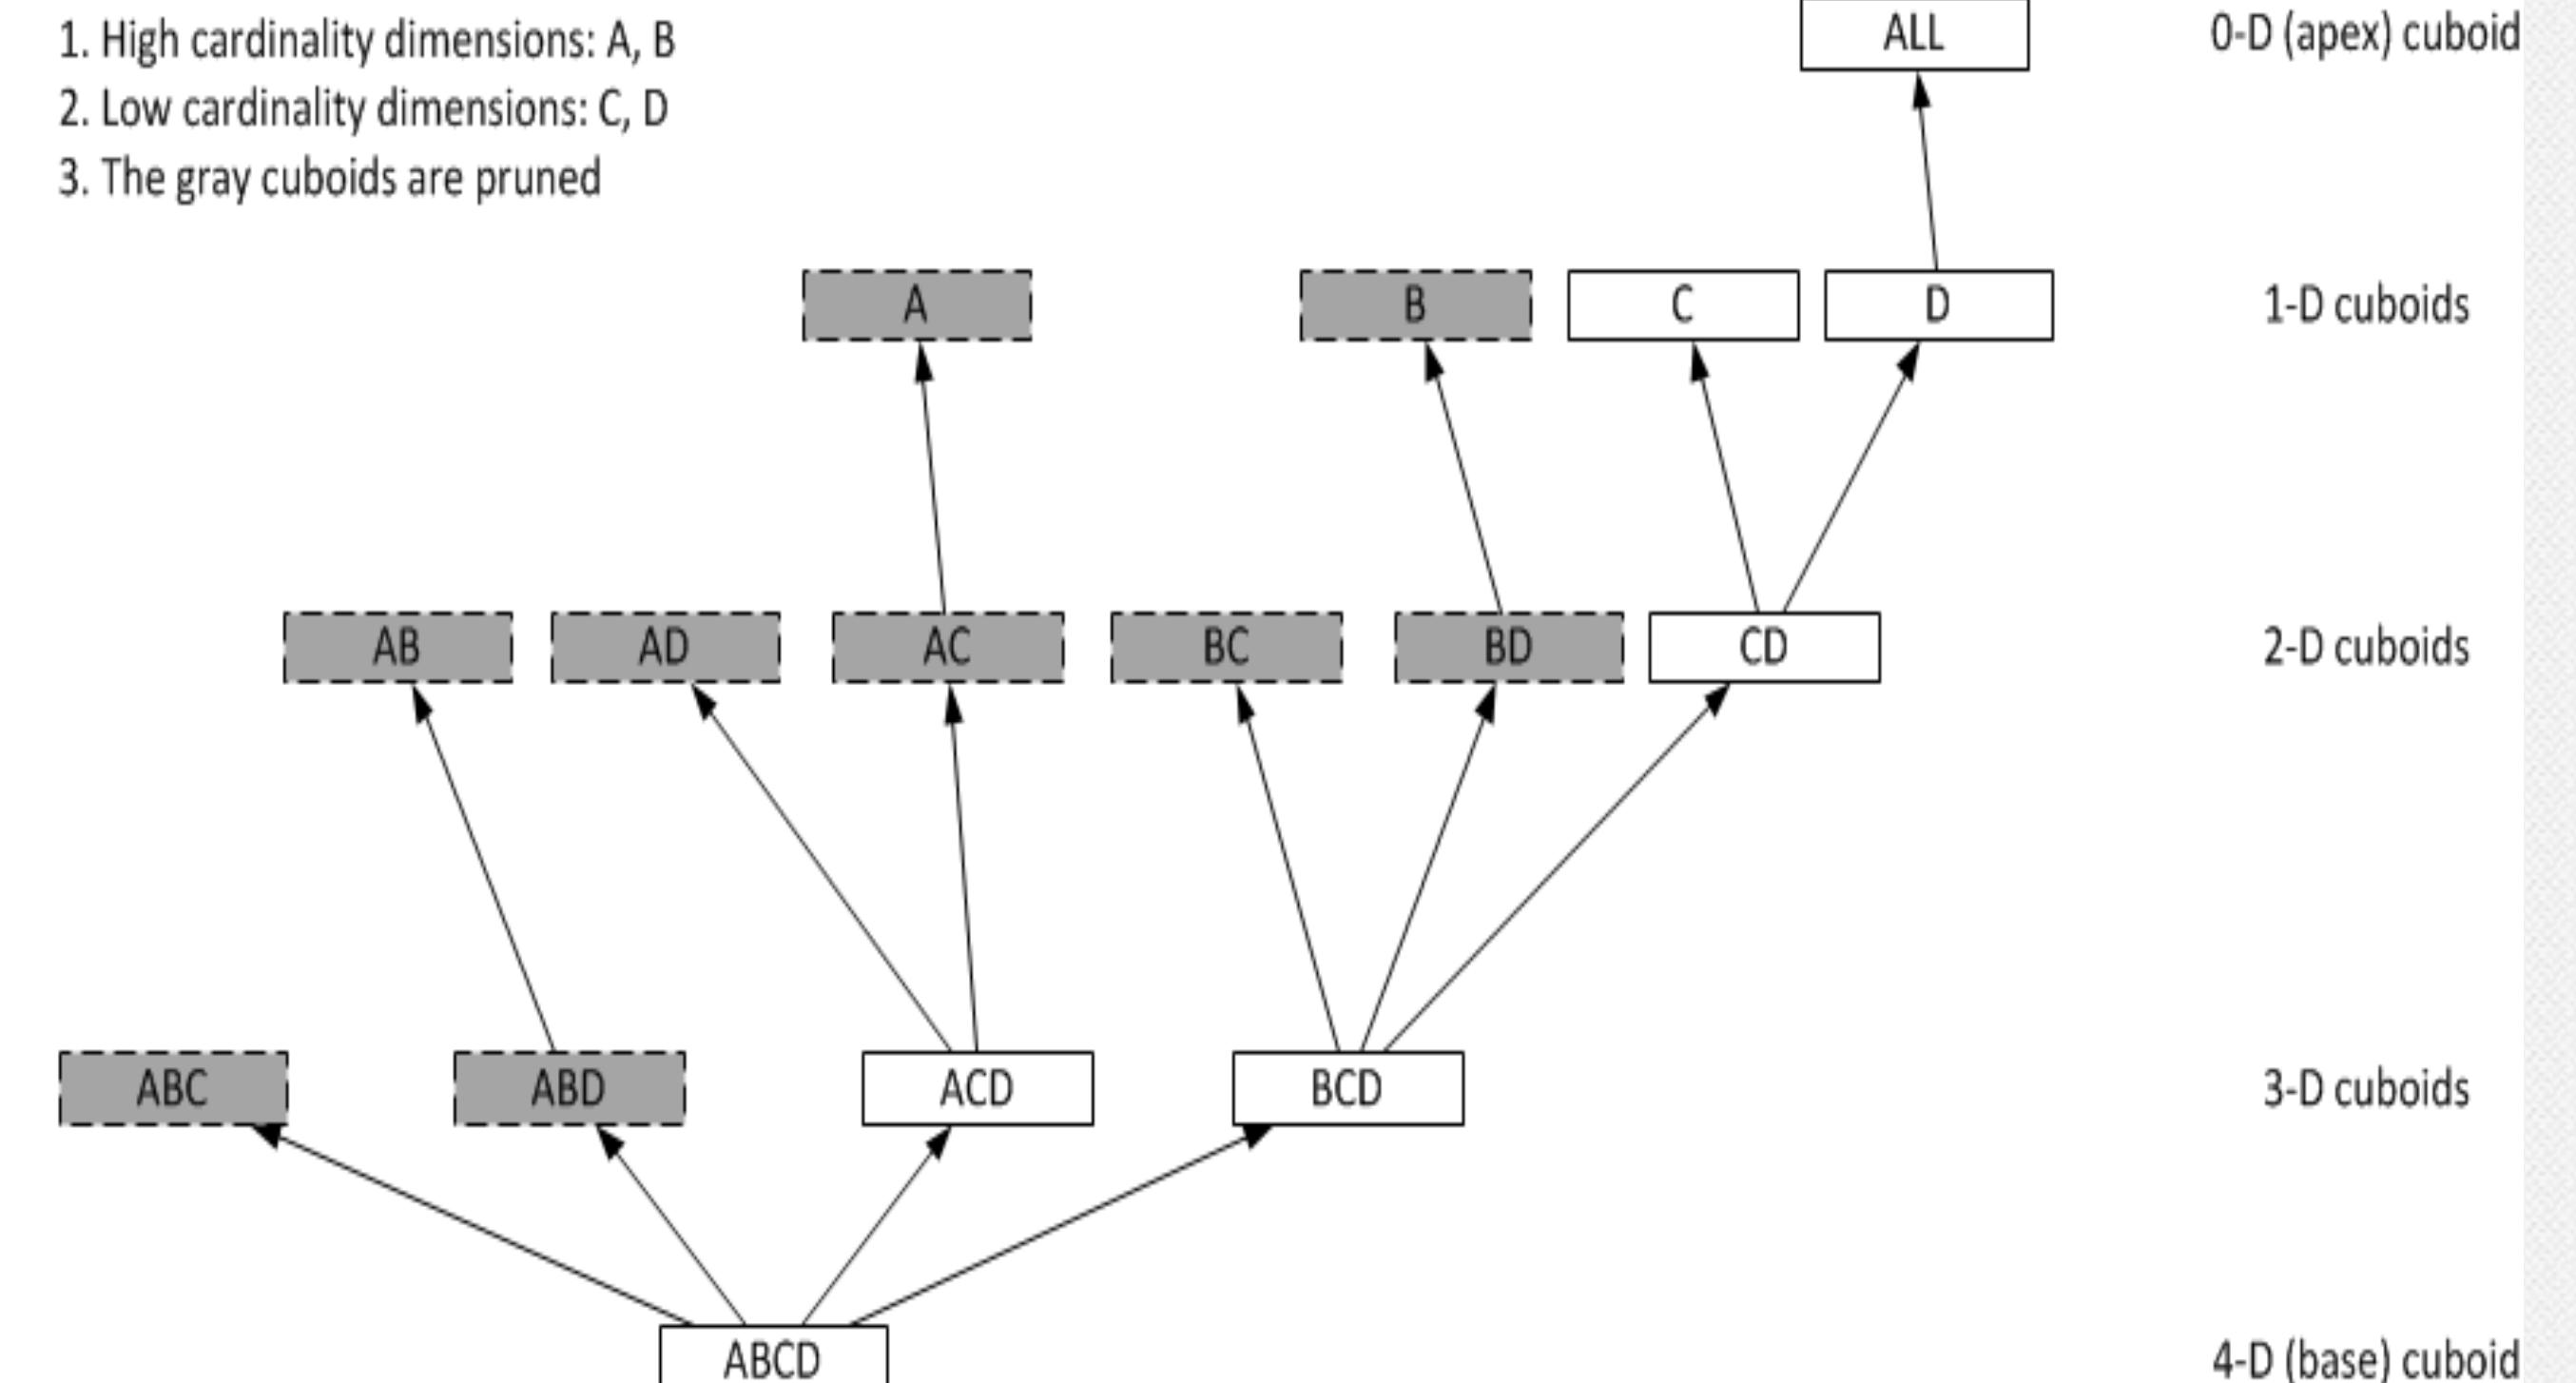


How to deal with data expansion

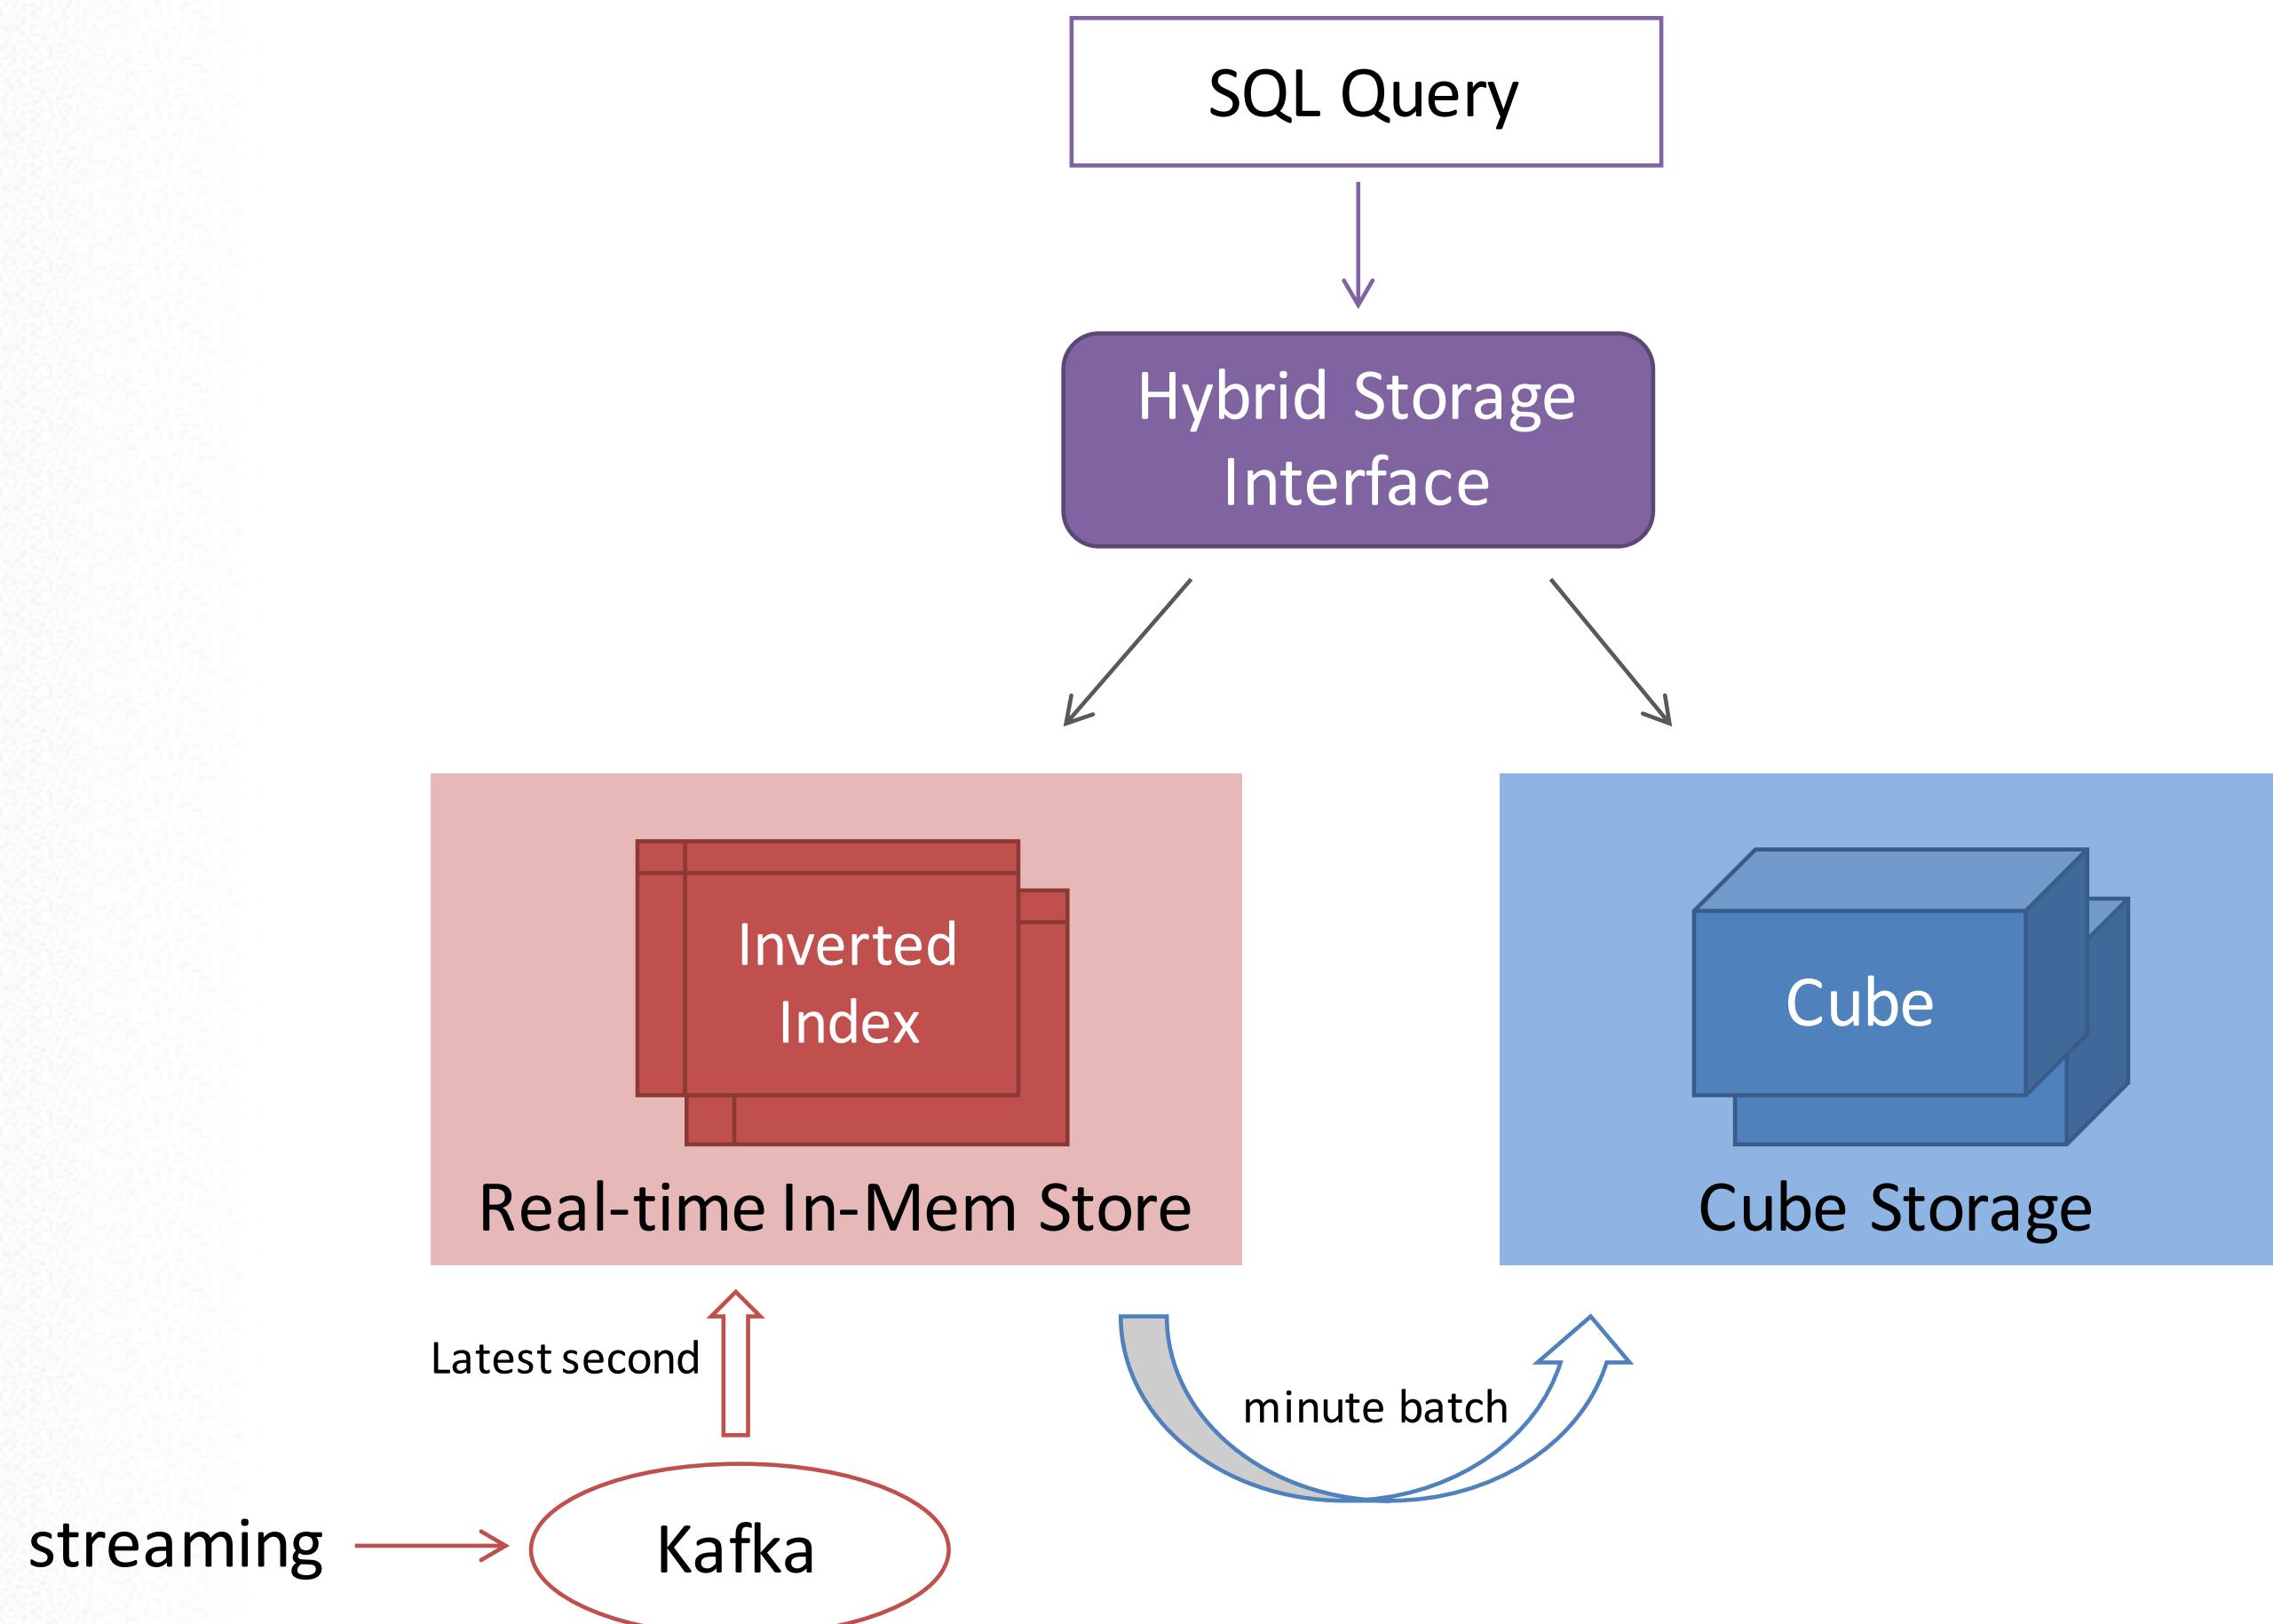
- Reduce Dimensions
 - Hierarchy
 - Derived
 - Extended
- Aggregation Group

$$2^{30} \rightarrow 2^{10} + 2^{10} + 2^{10}$$

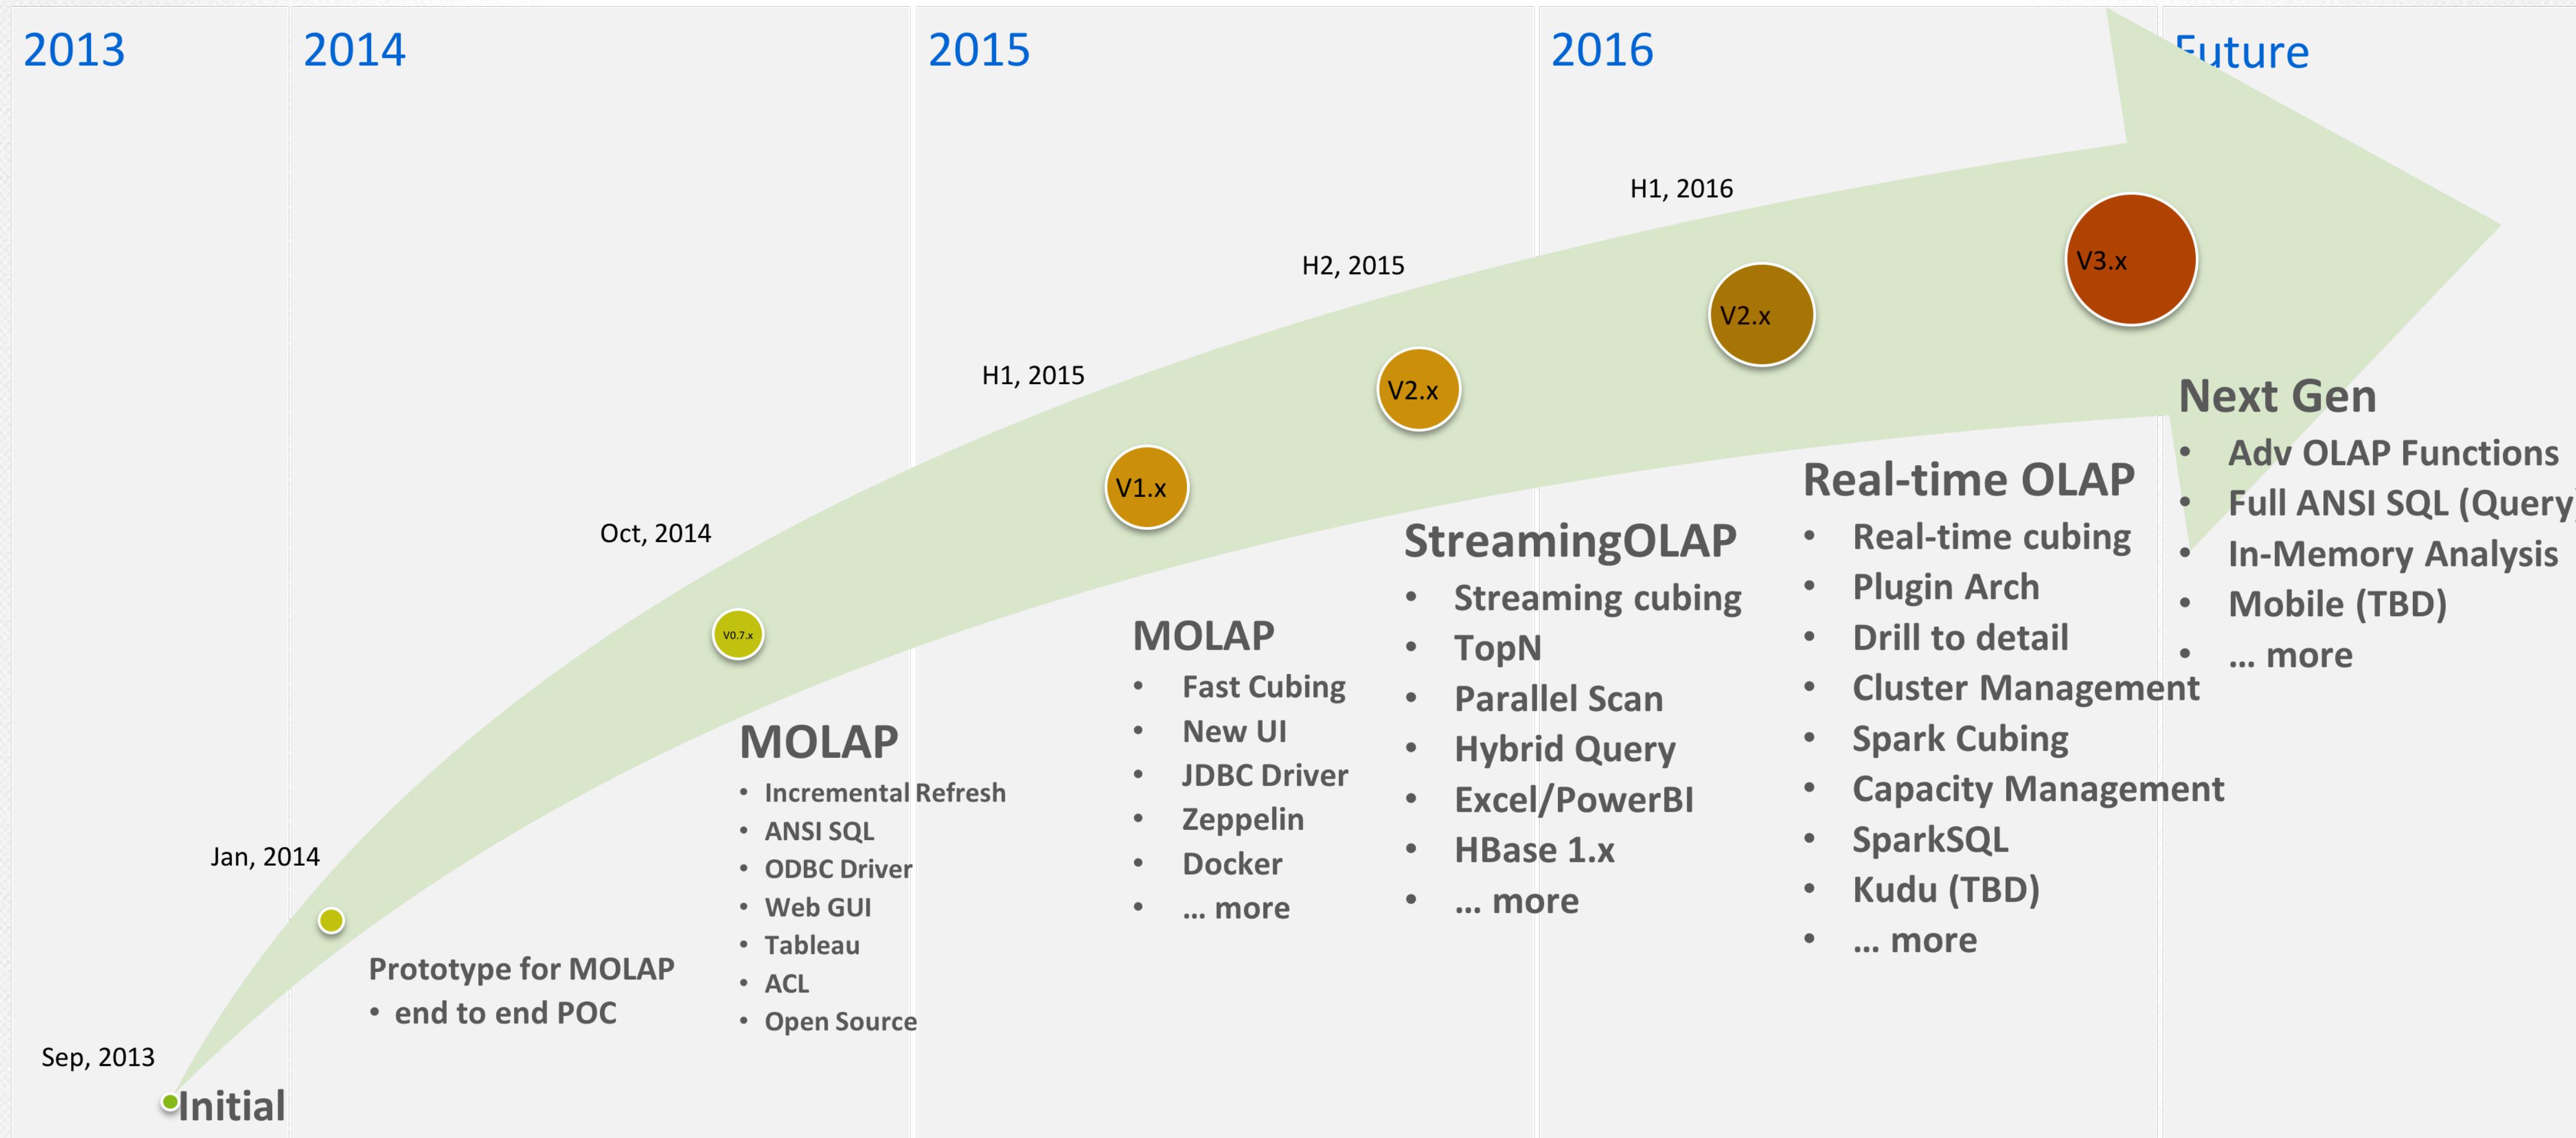
- Partial Cube



Lambda Architecture



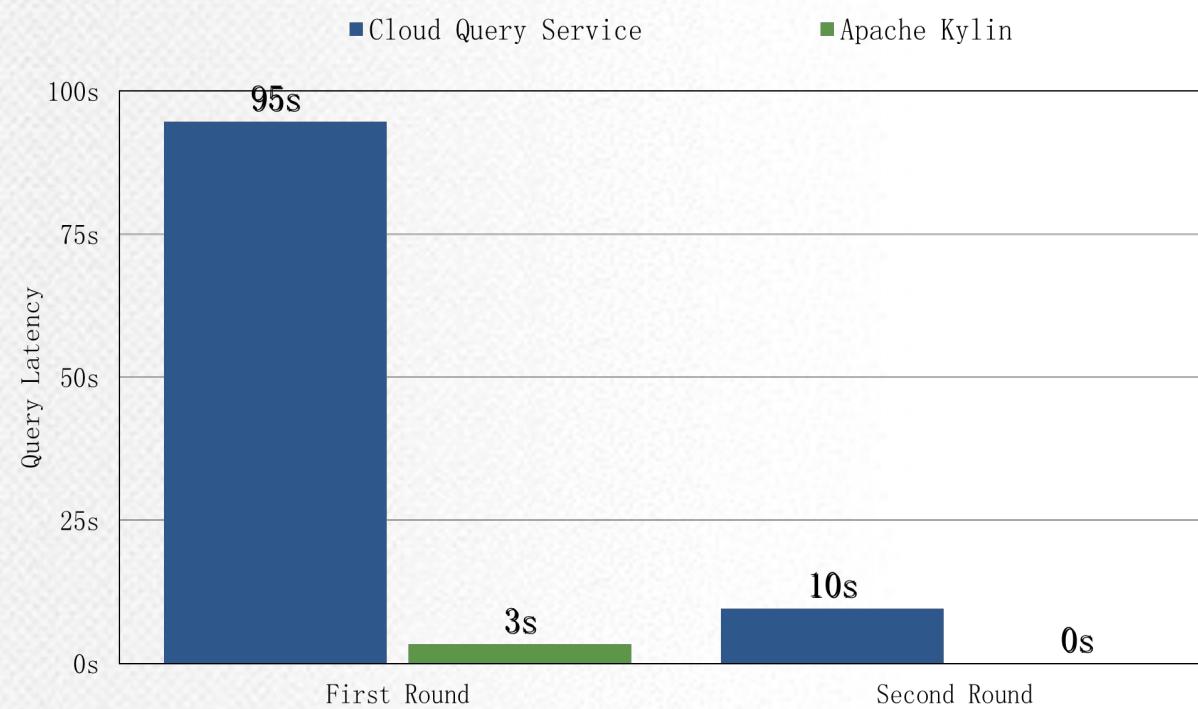
Roadmap



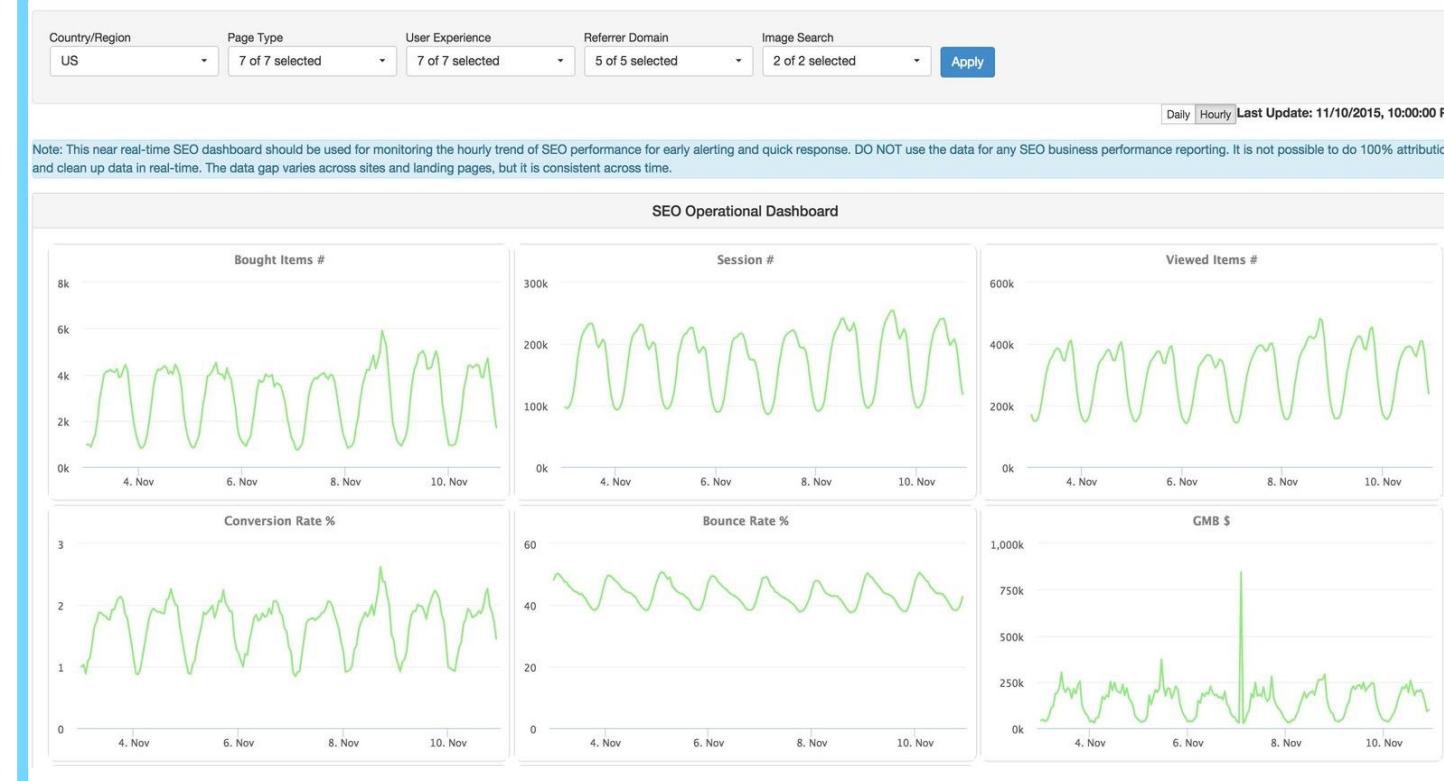
Use Cases

eBay

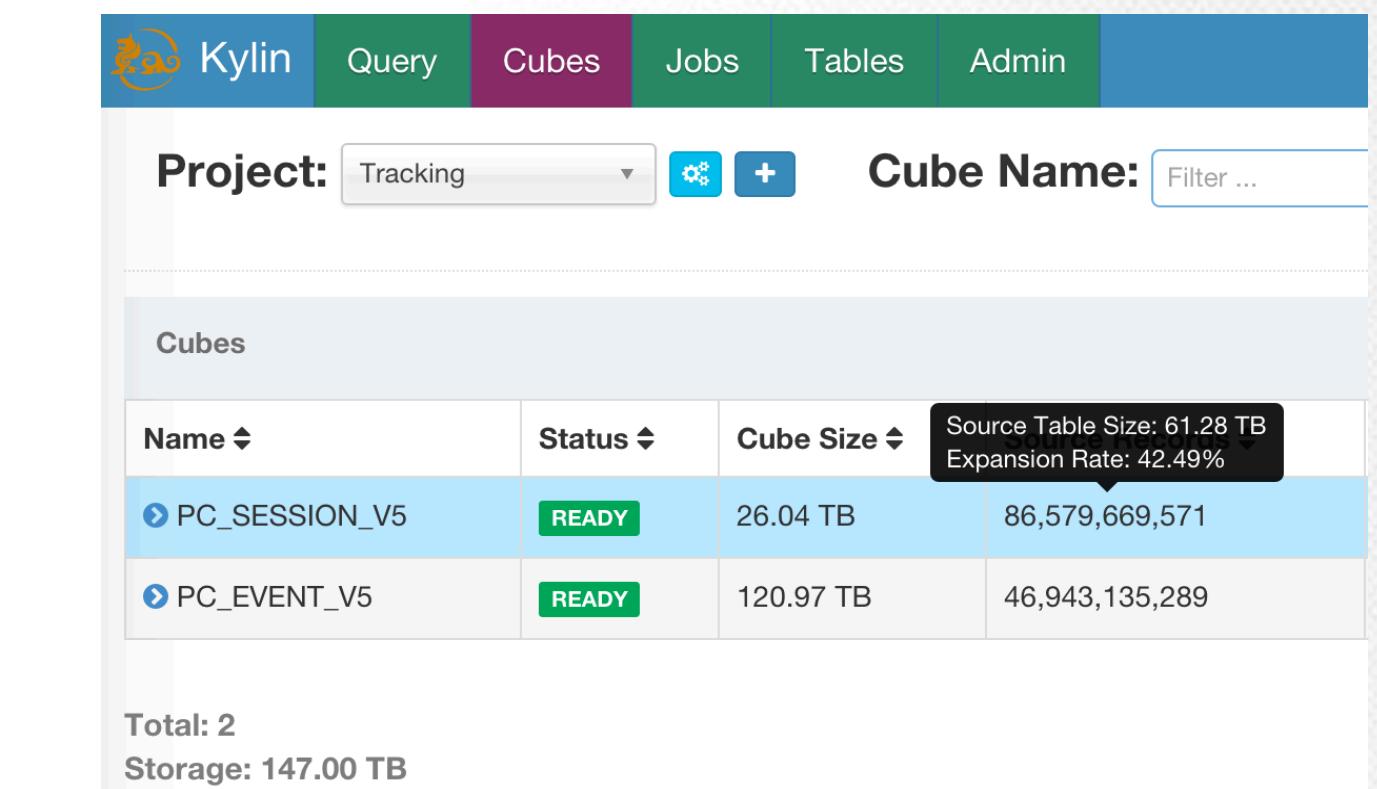
- eCG
 - GA Deep Dive
 - Pilot Case
 - Extreme challenge for storage and building time
 - Enabled Tableau



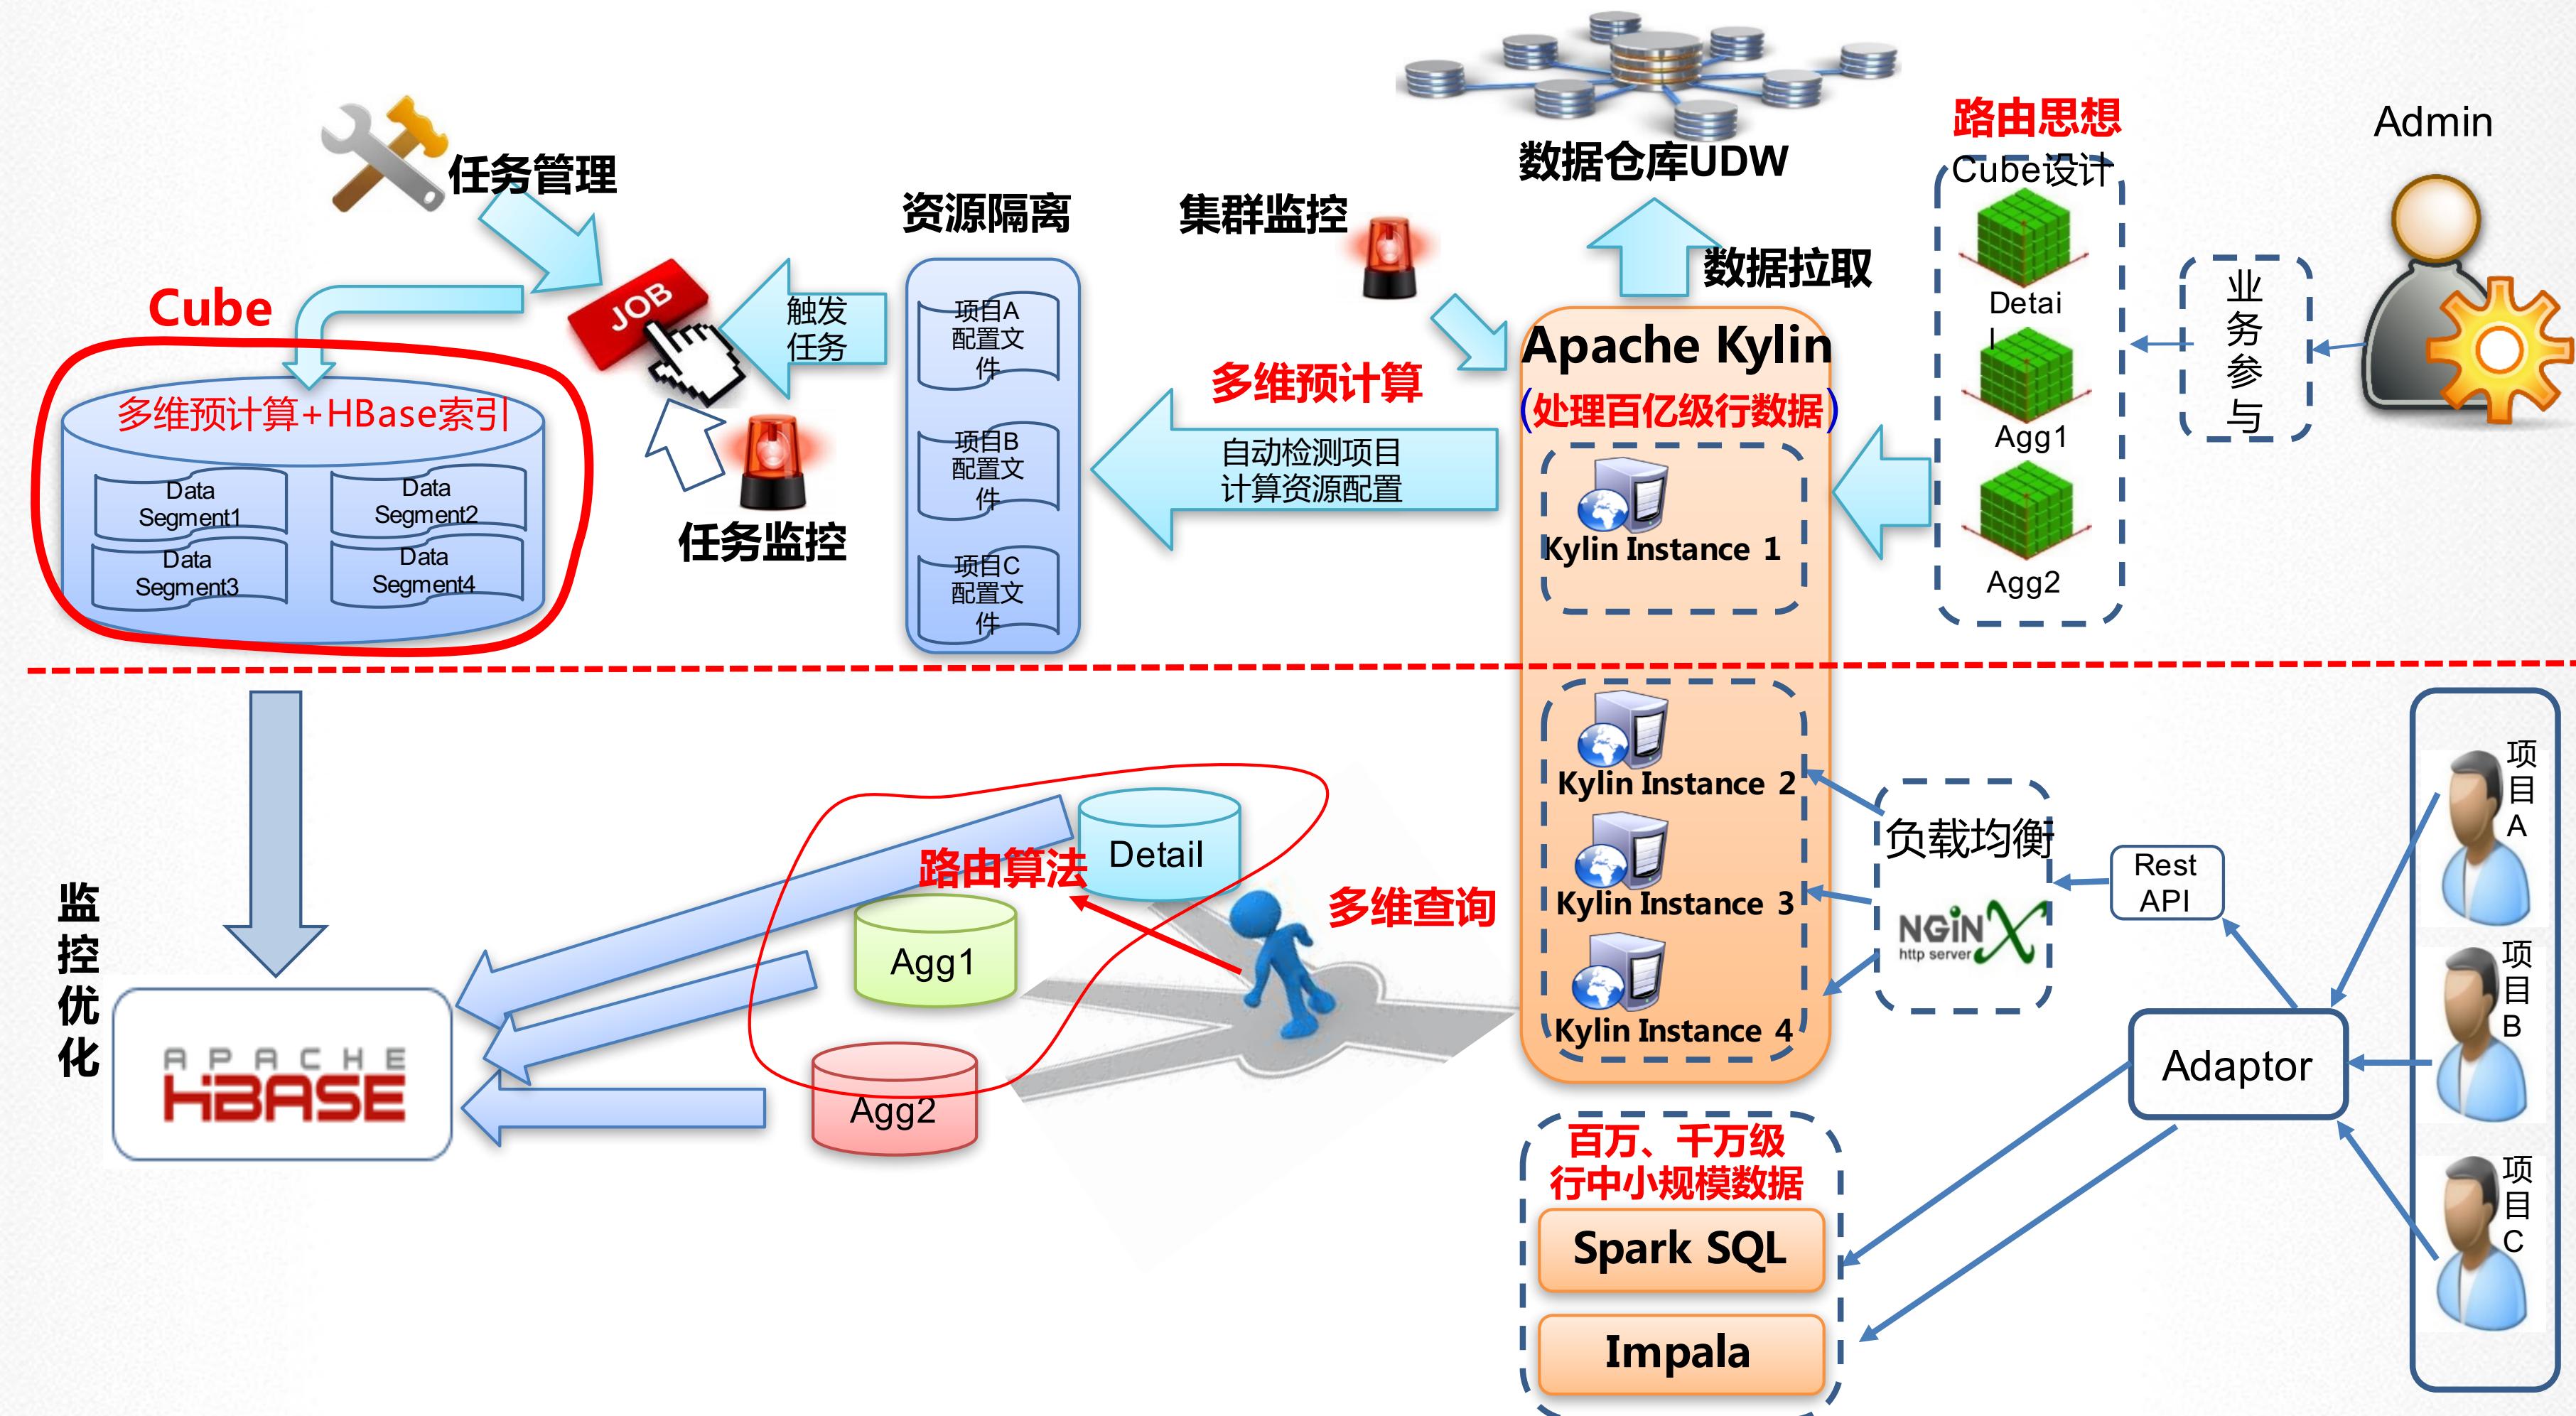
- DSS Nous
 - NRT SEO Dashboard
 - Near Real-time and Historical together



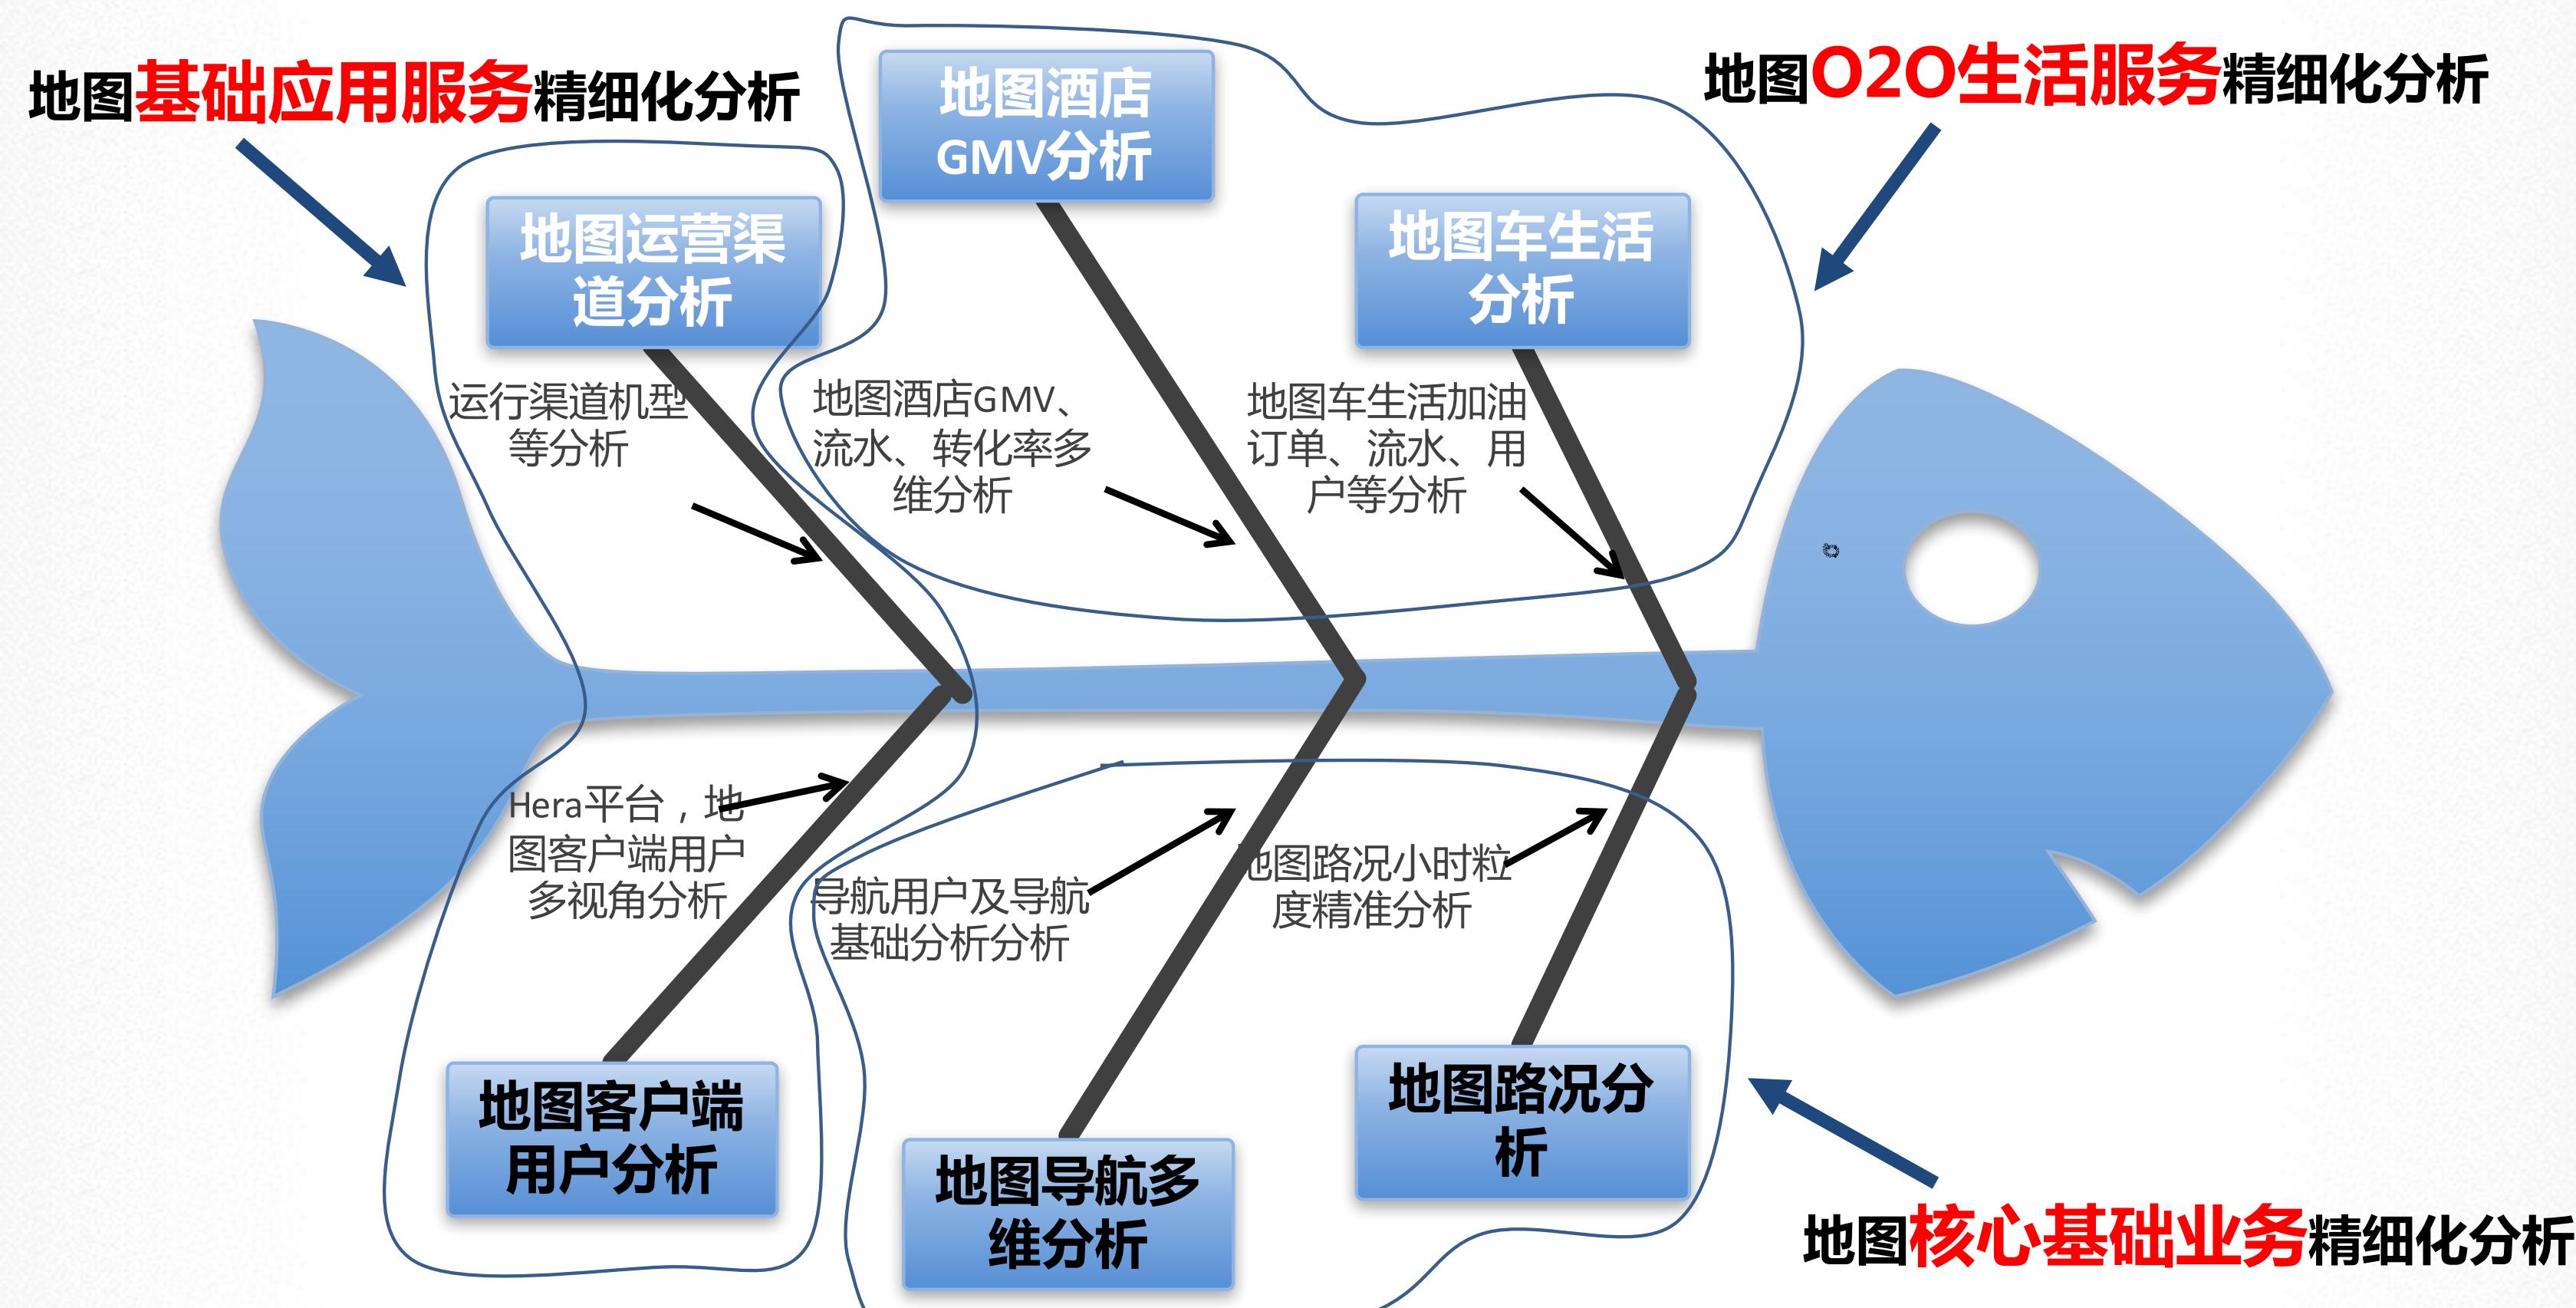
- eCS
 - Tracking Report
 - Hugest Cube: 100+B raw records in one cube



百度地图



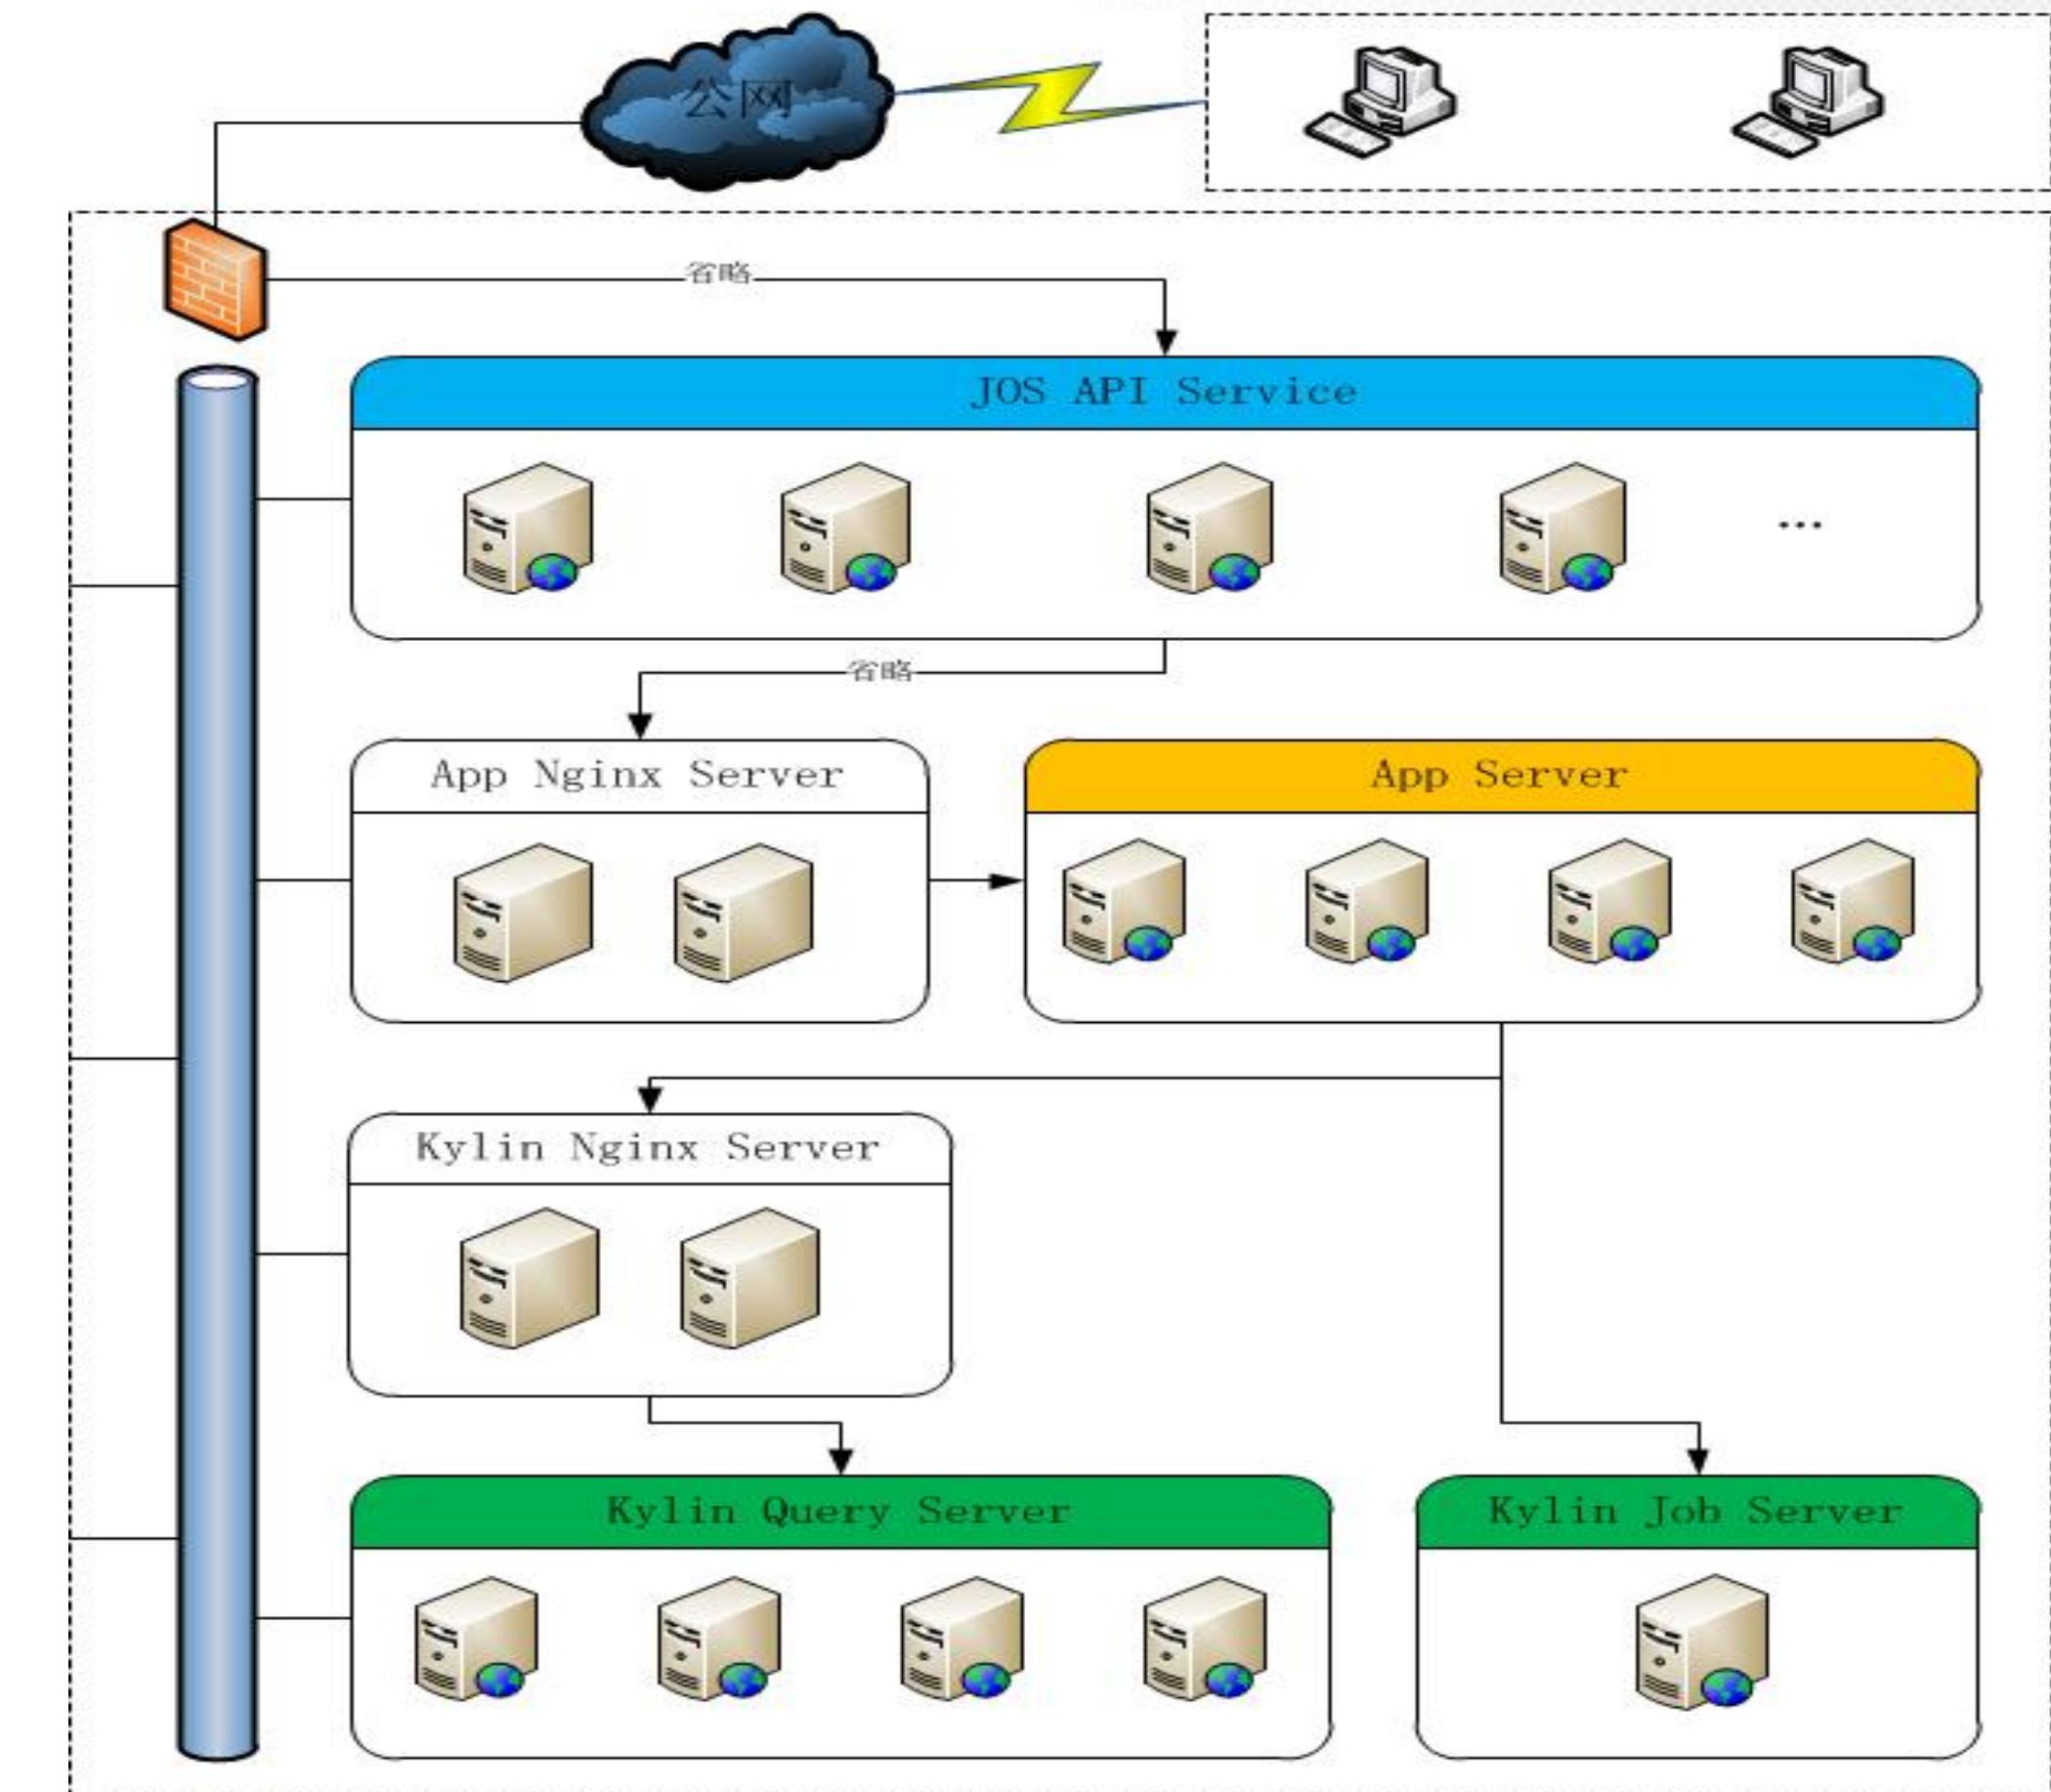
百度地图



共计**50+亿**行源数据，最大**单表20+亿**行记录

京东

- Powered by Apache Kylin
 - JCloud JOS - API Statistics
 - JCloud 云海 - Low latency query engine
 - JCloud Data Cloud - OLAP platform for online analytics tool



北京移动

Data Overview



20+M Users Data Volume 30+B/day ETL 3TB/day Cluster 400TB Jobs 800+/day

I. Fast and Easy deployment

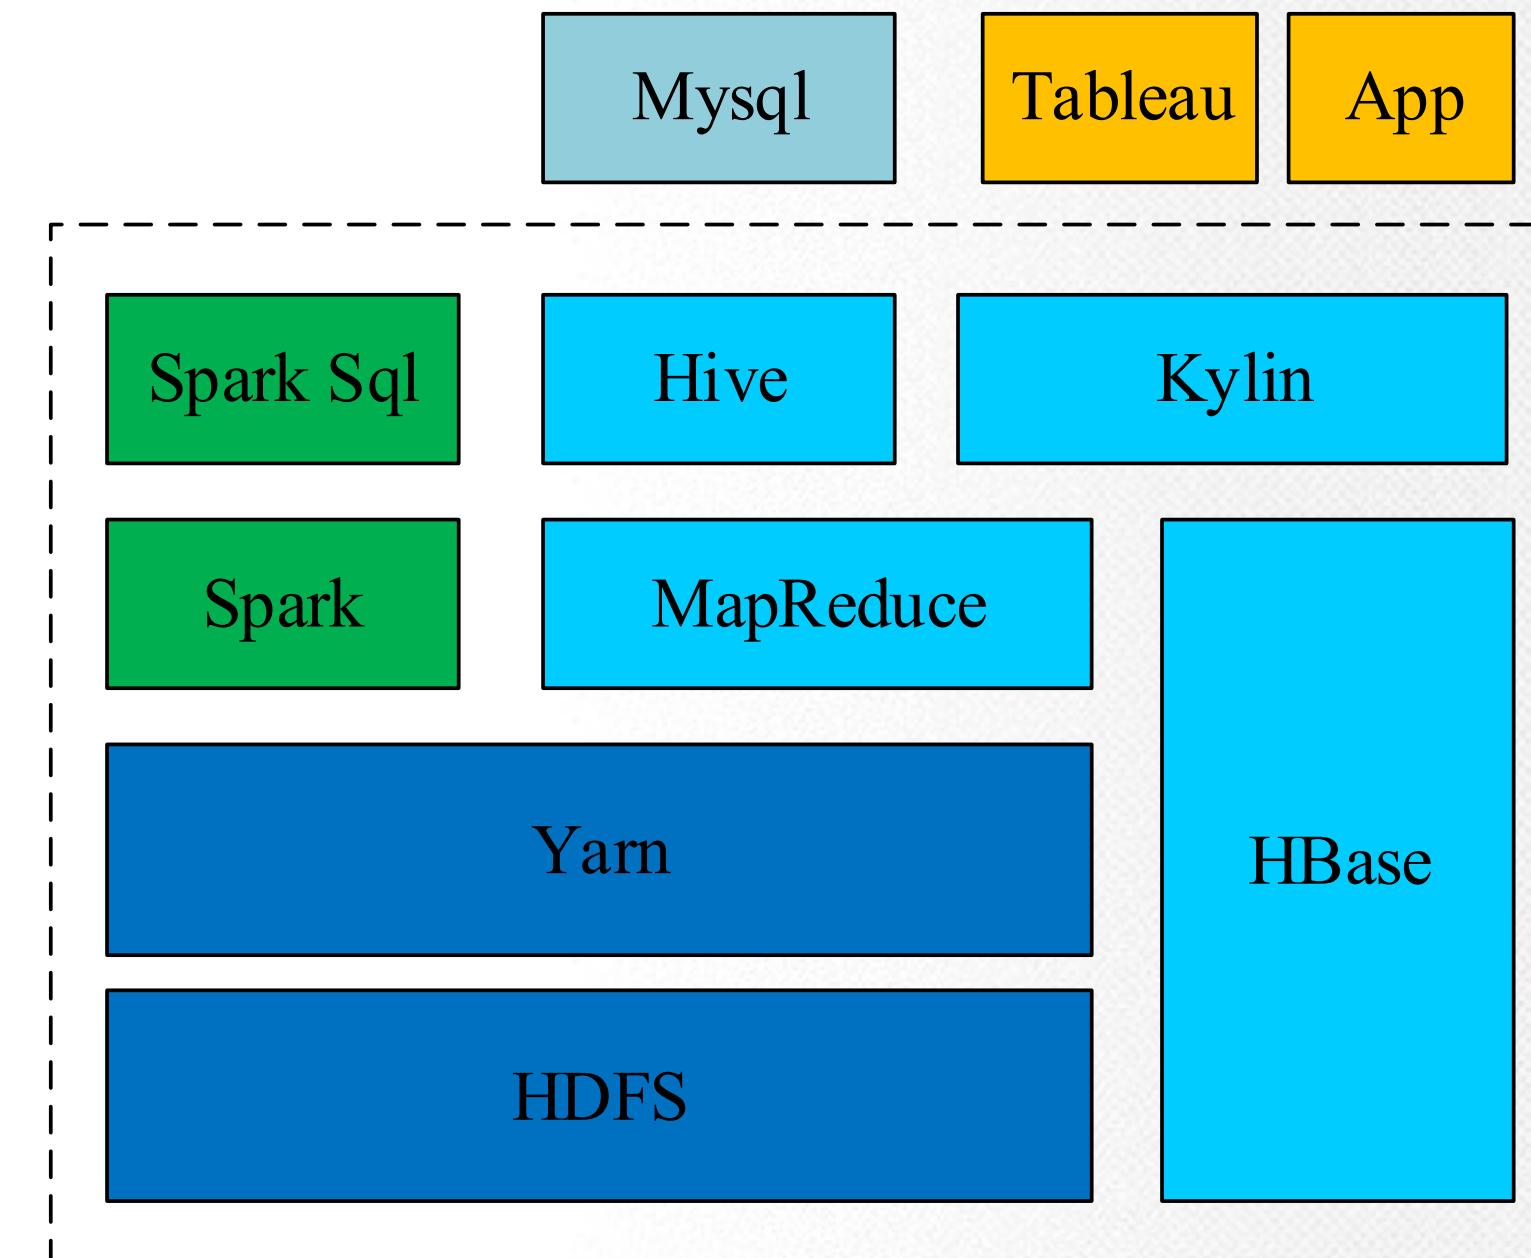


II. Low query latency

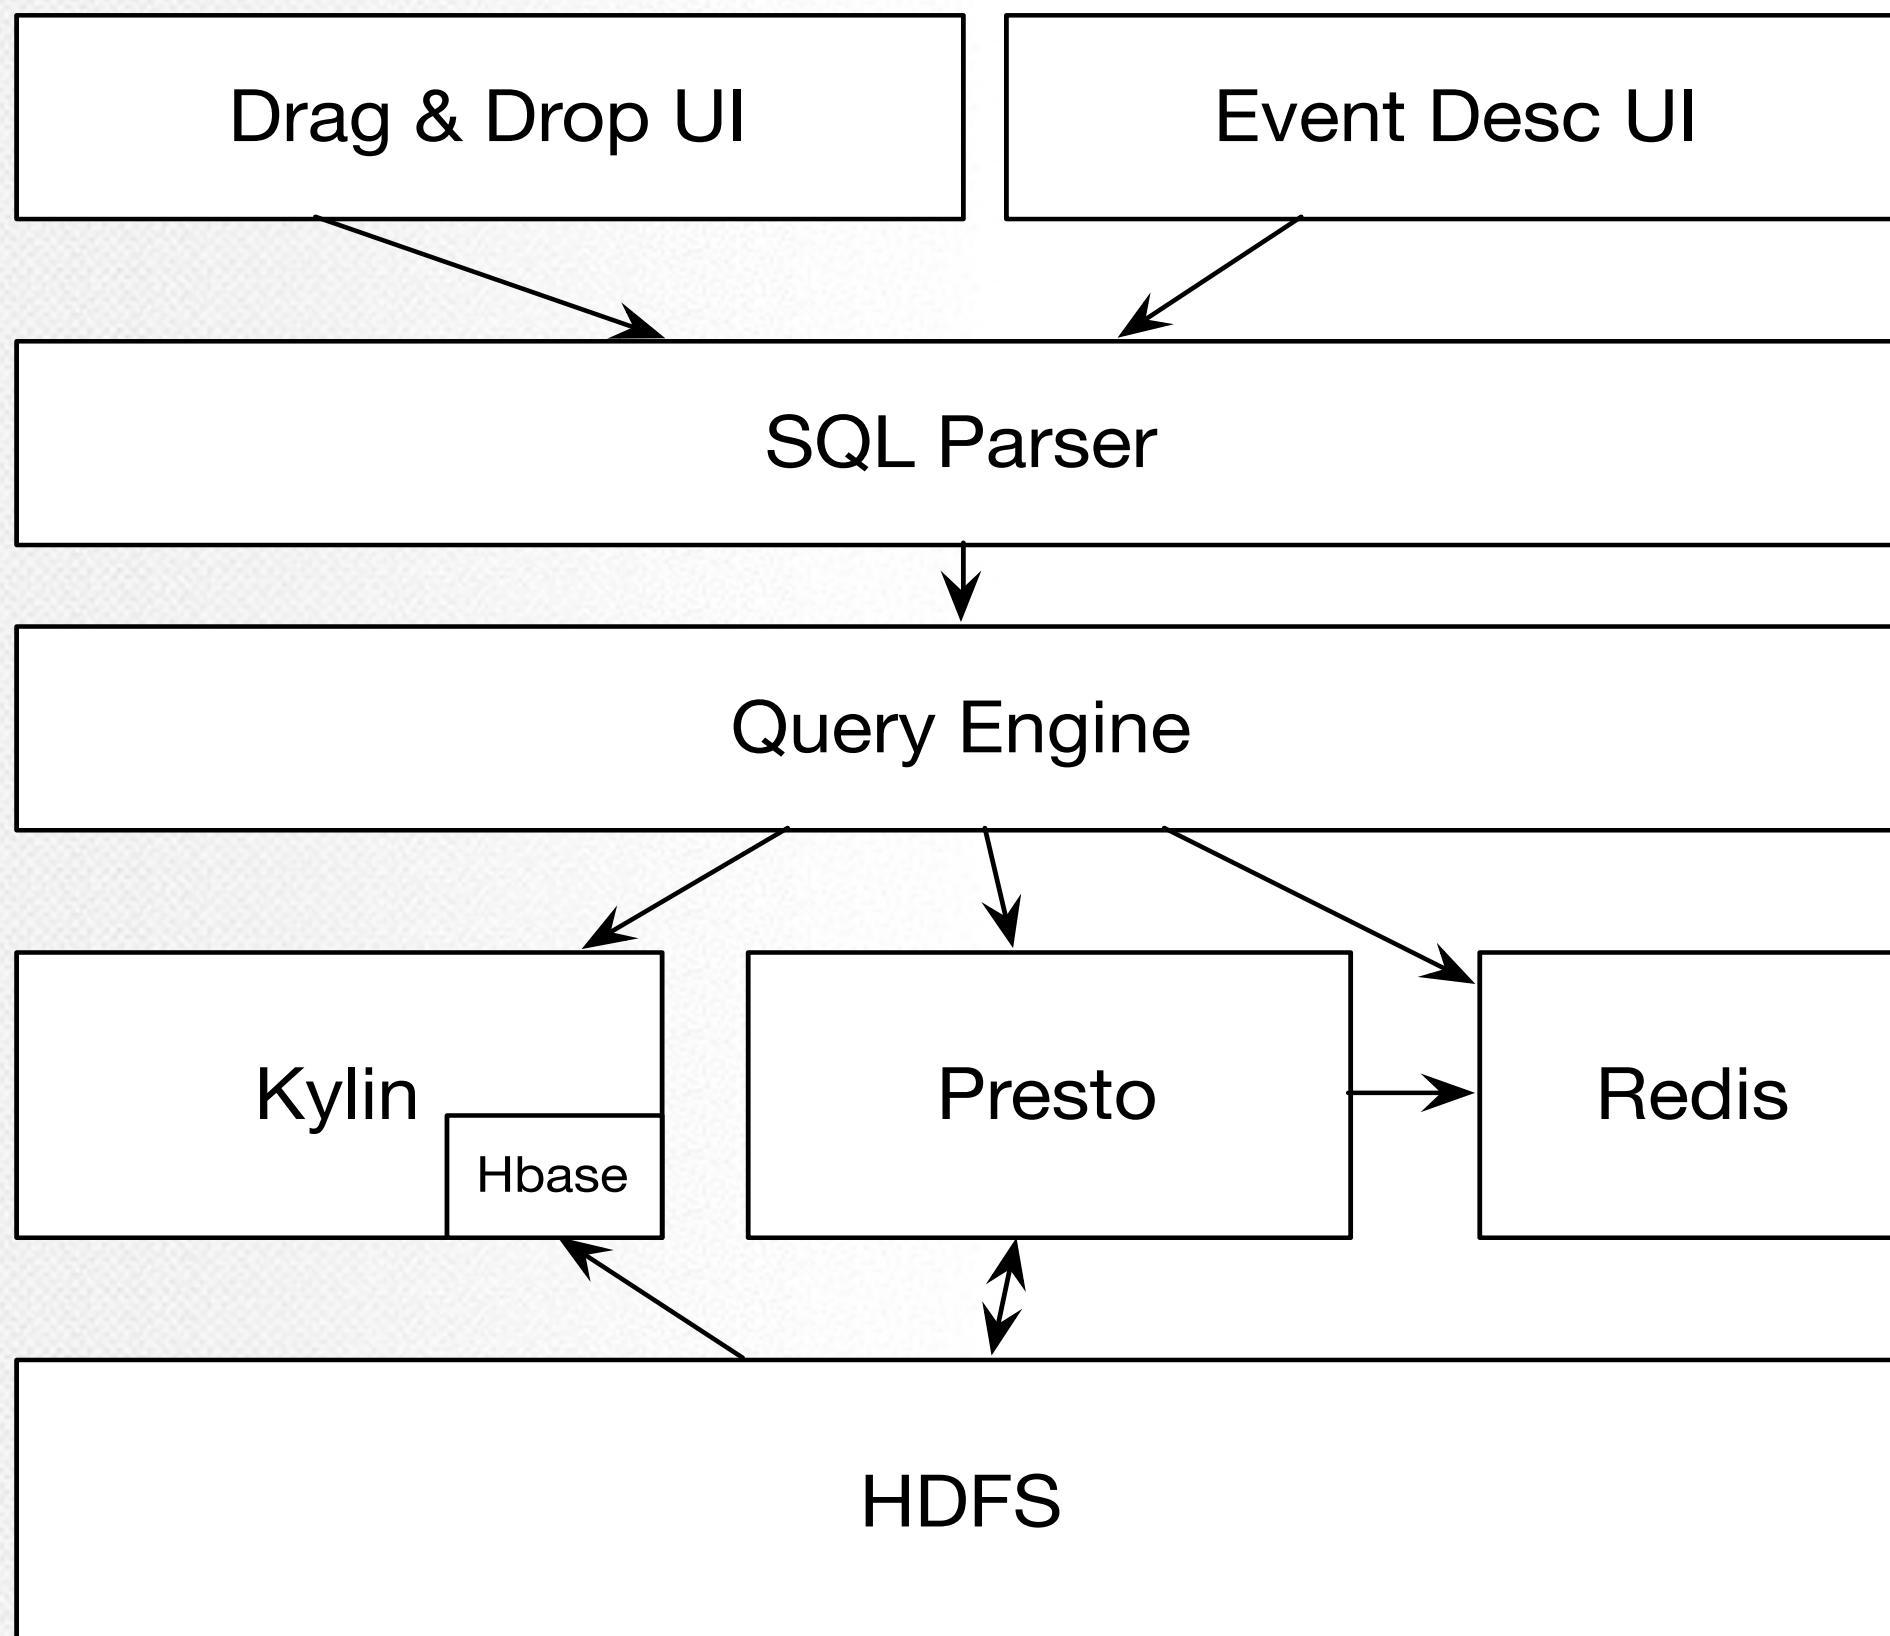
	Reousrces	Latency	comments
hive	86vcores+380GBMEM	1522s	orc+zlib
spark sql	131vcores+912GBMEM	125s	orc+zlib
kylin	Hbase5 nodes	3.43s	

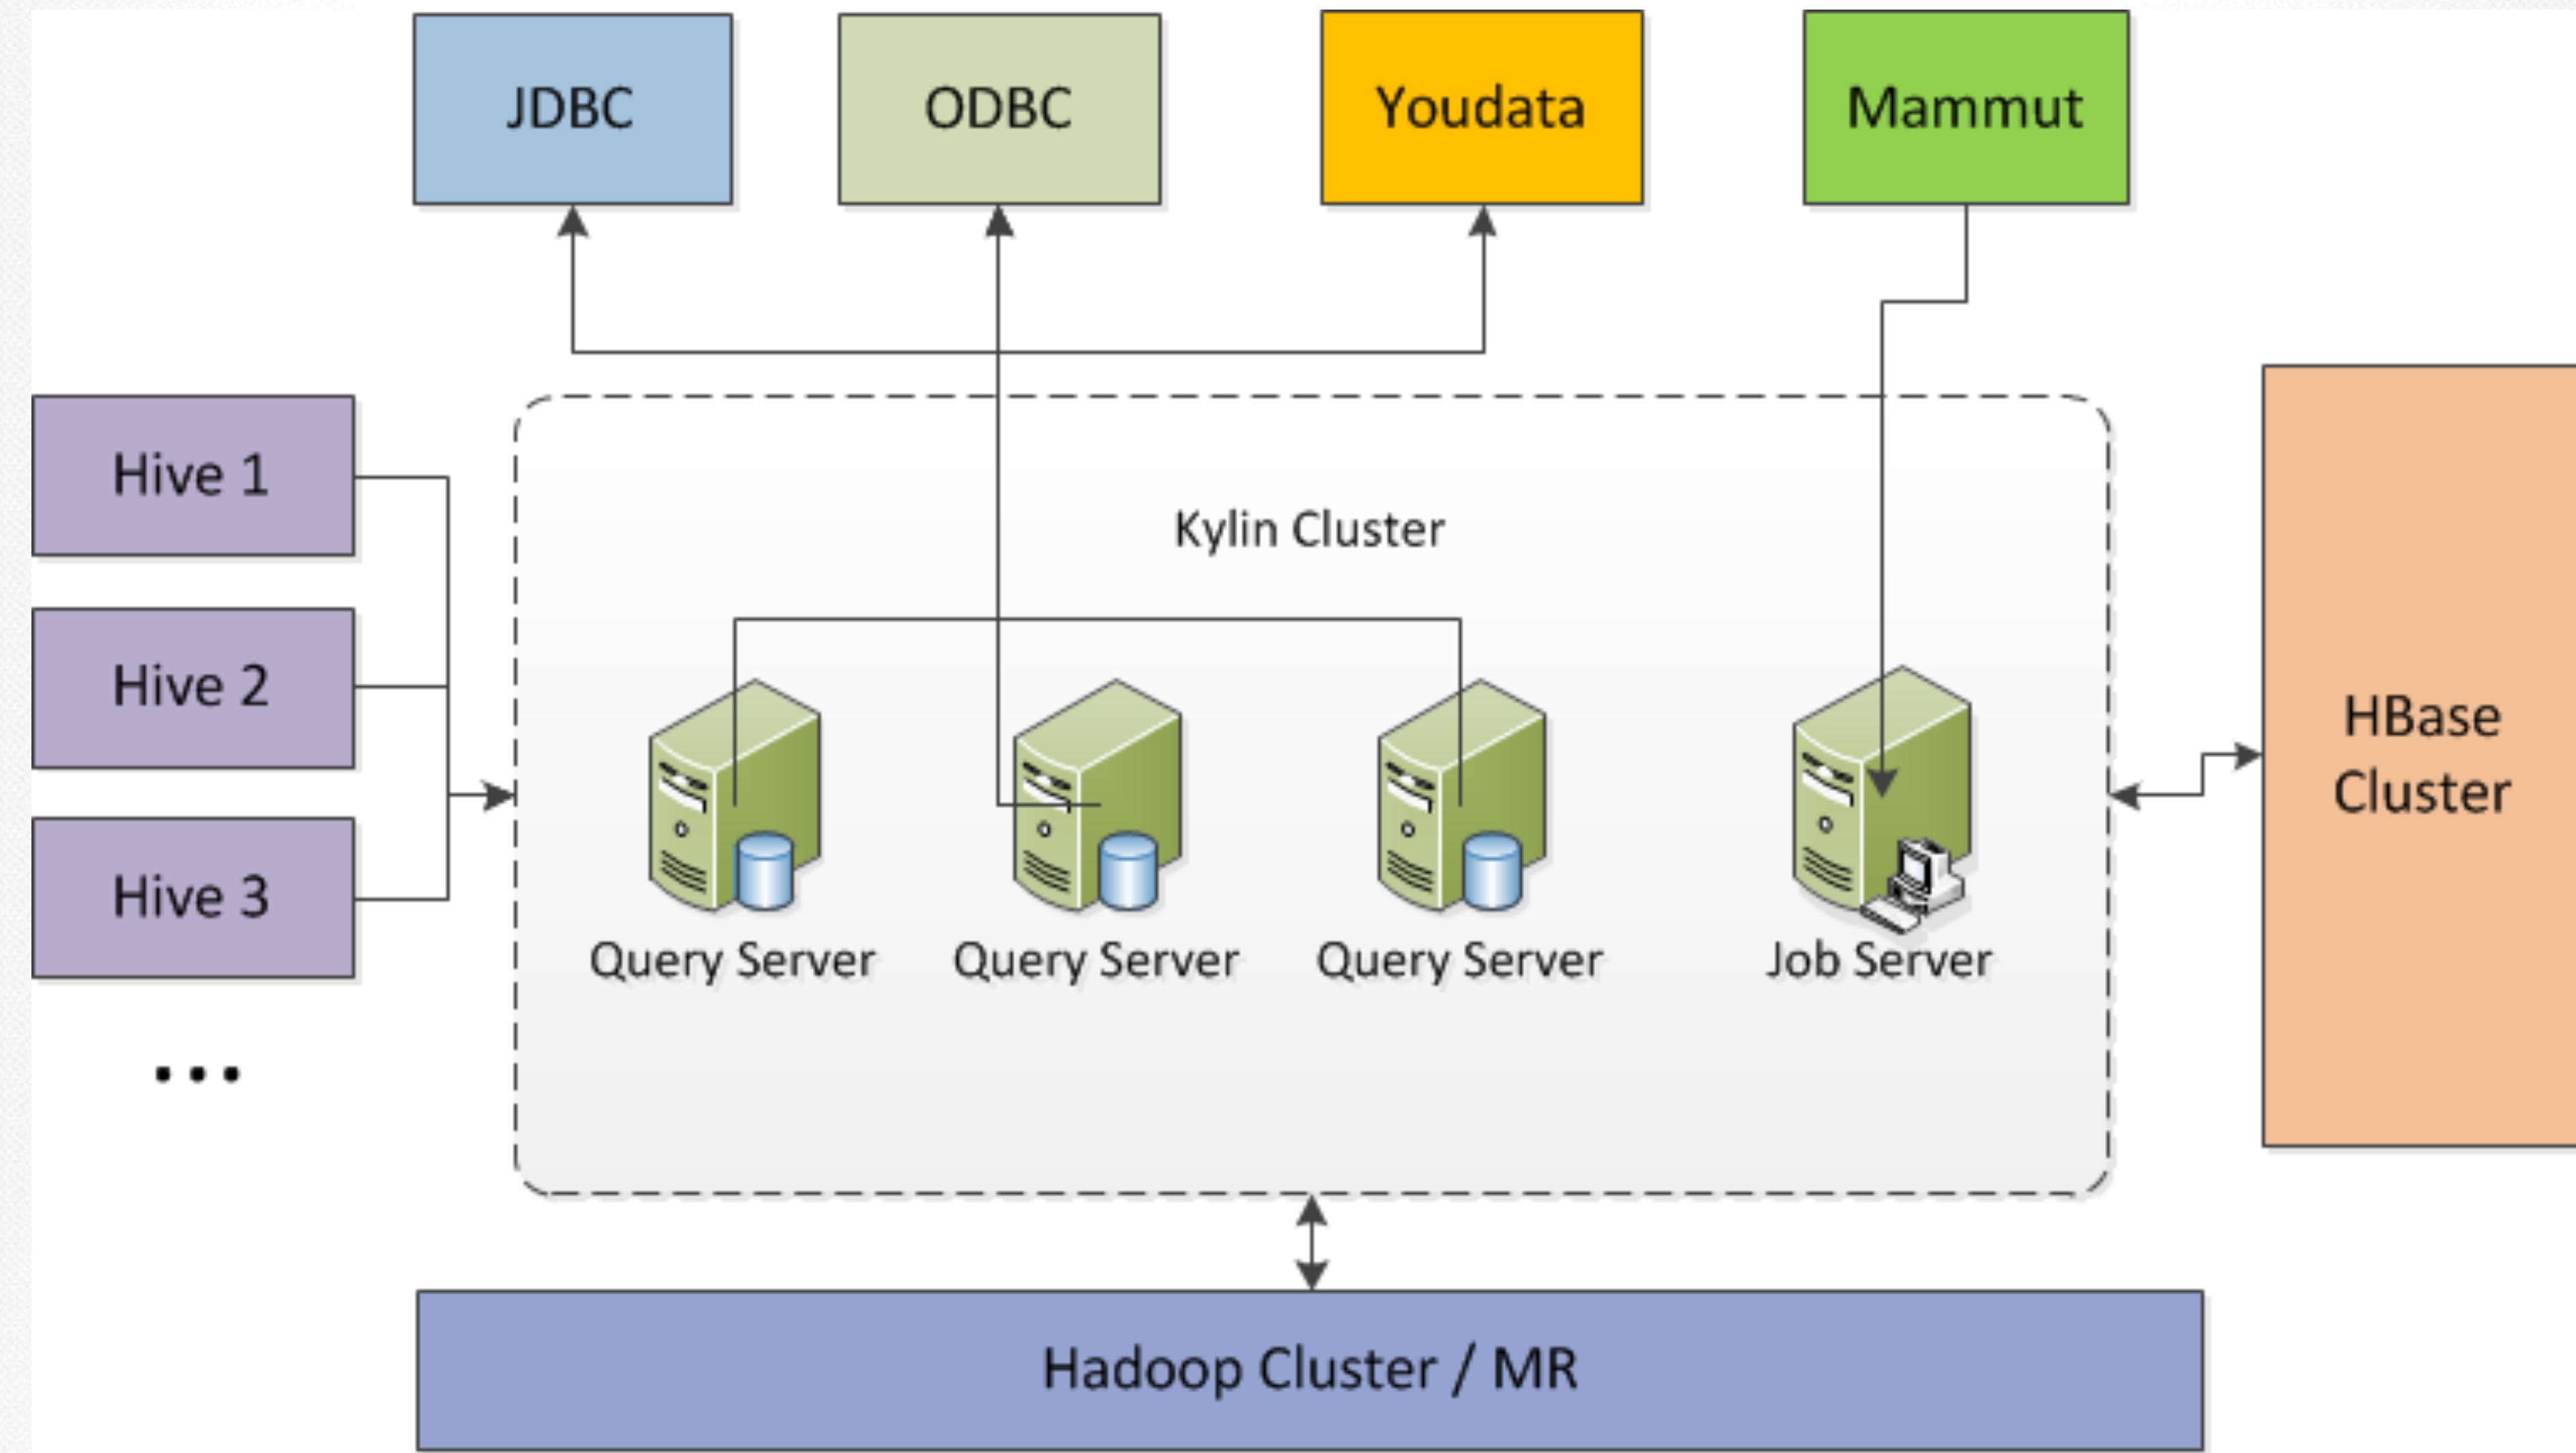
*Test SQL : select rat,count(distinct msisdn) from phone_usertmp where reportdate='20160225' group by rat;

*Data size: 103GB , records: 1.1+B

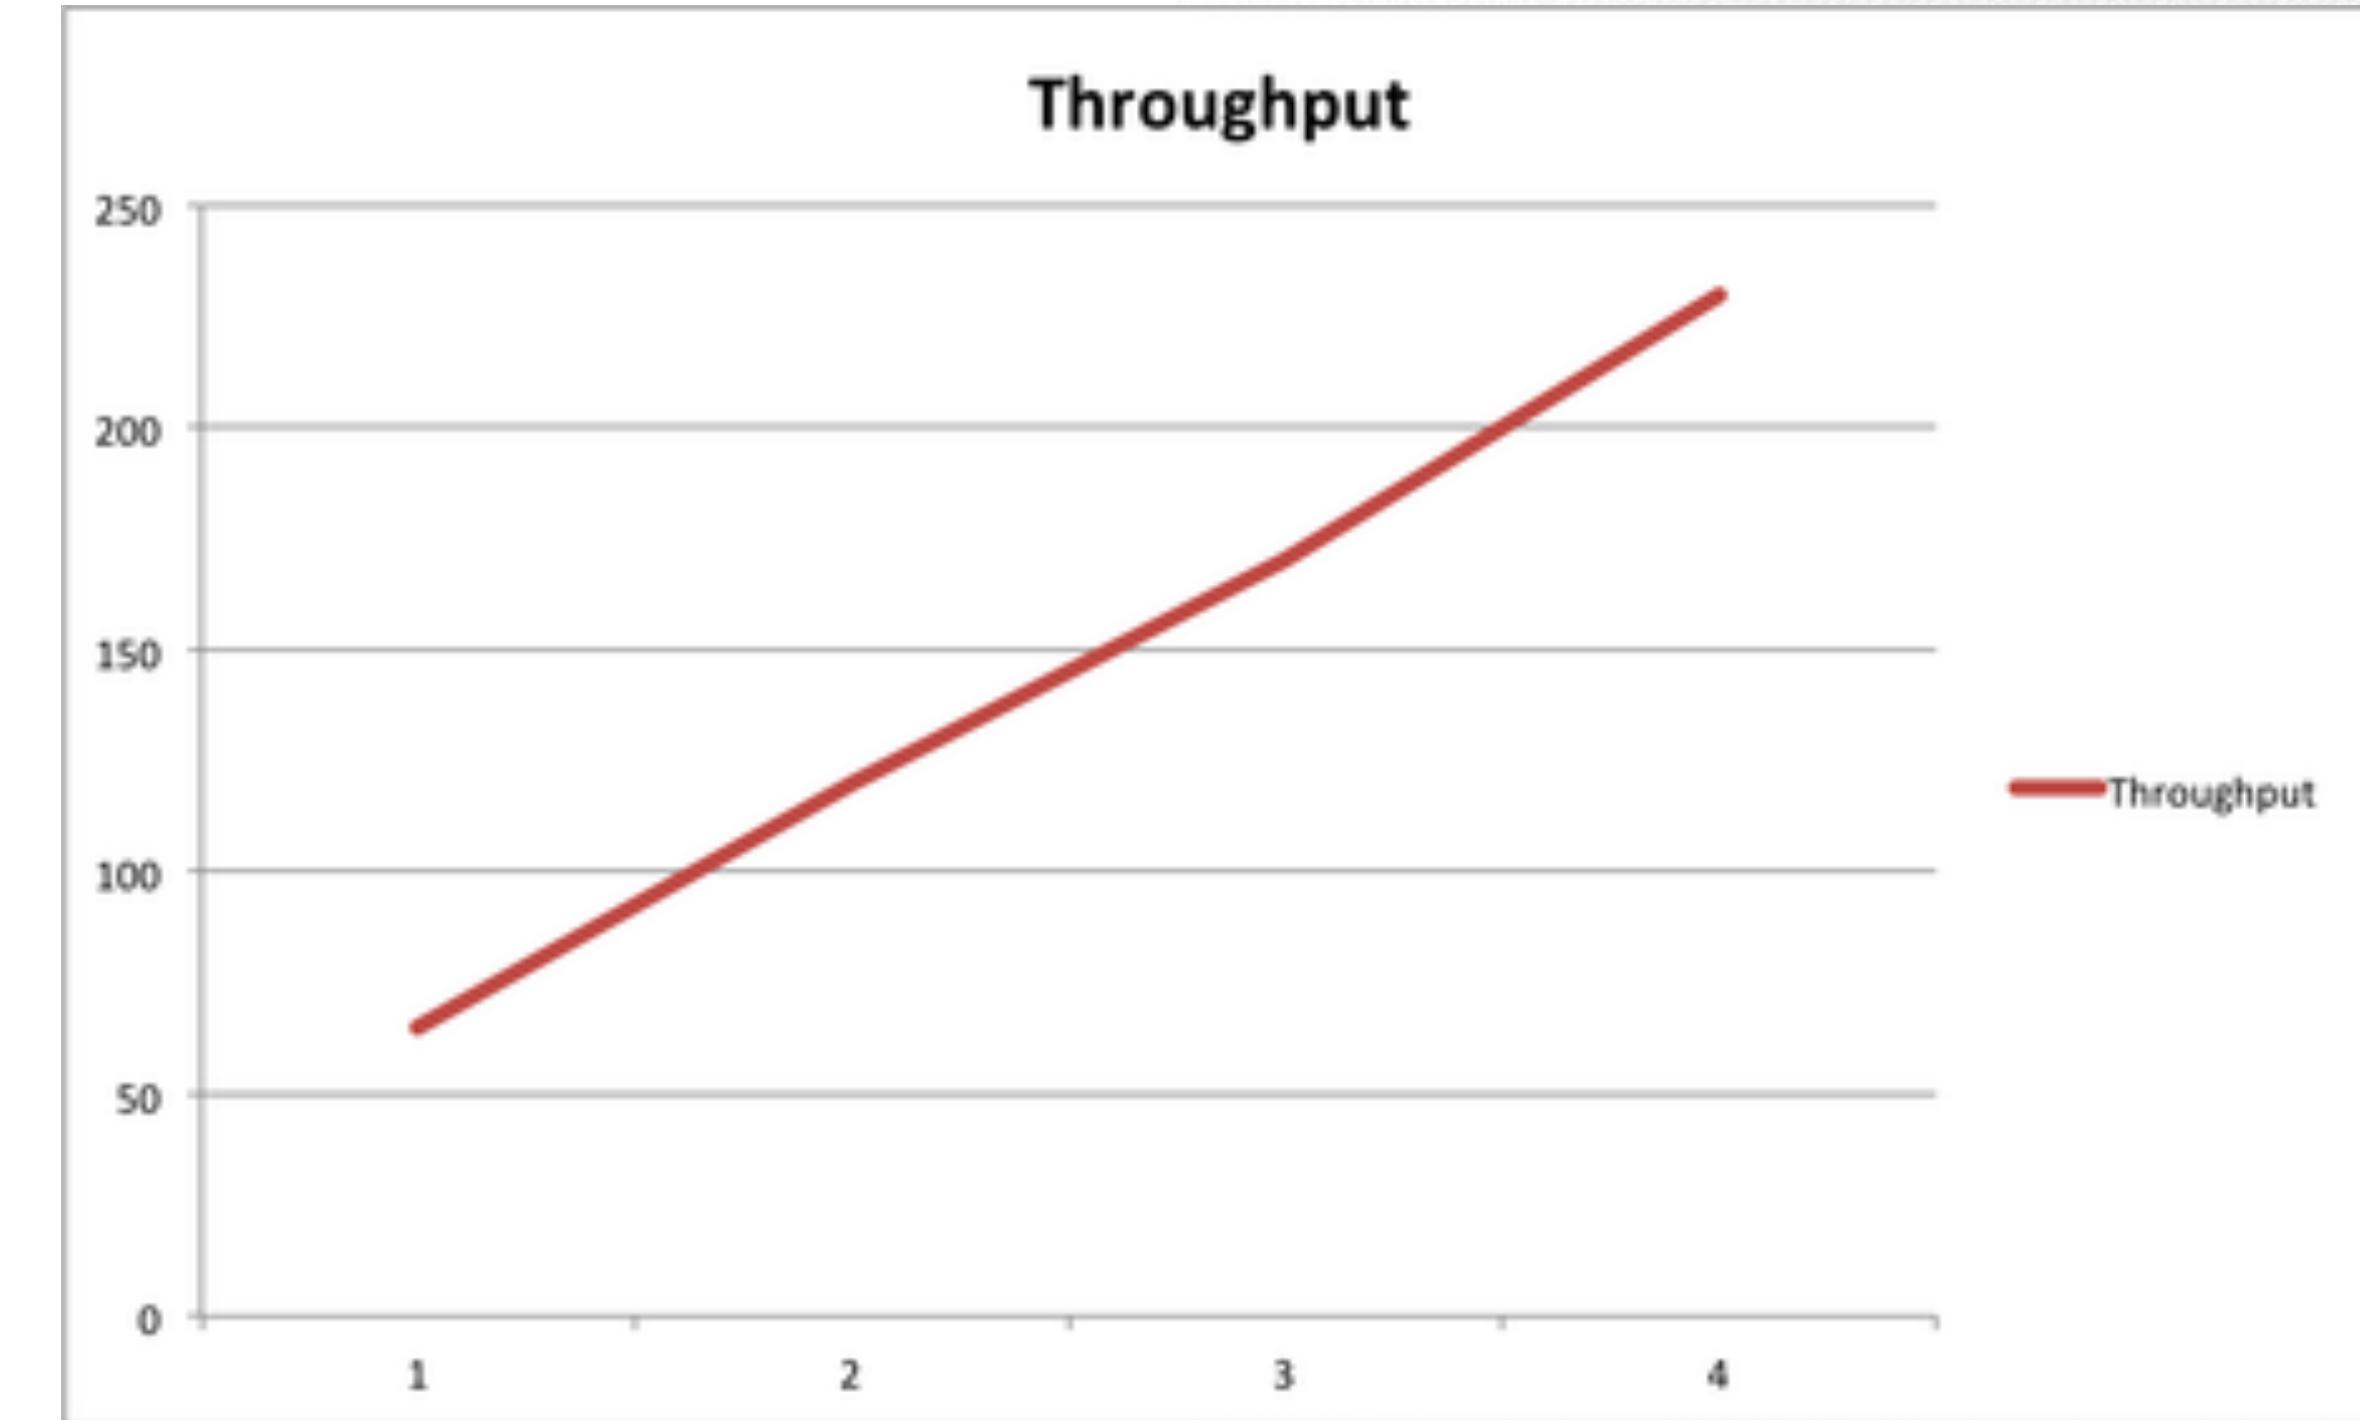
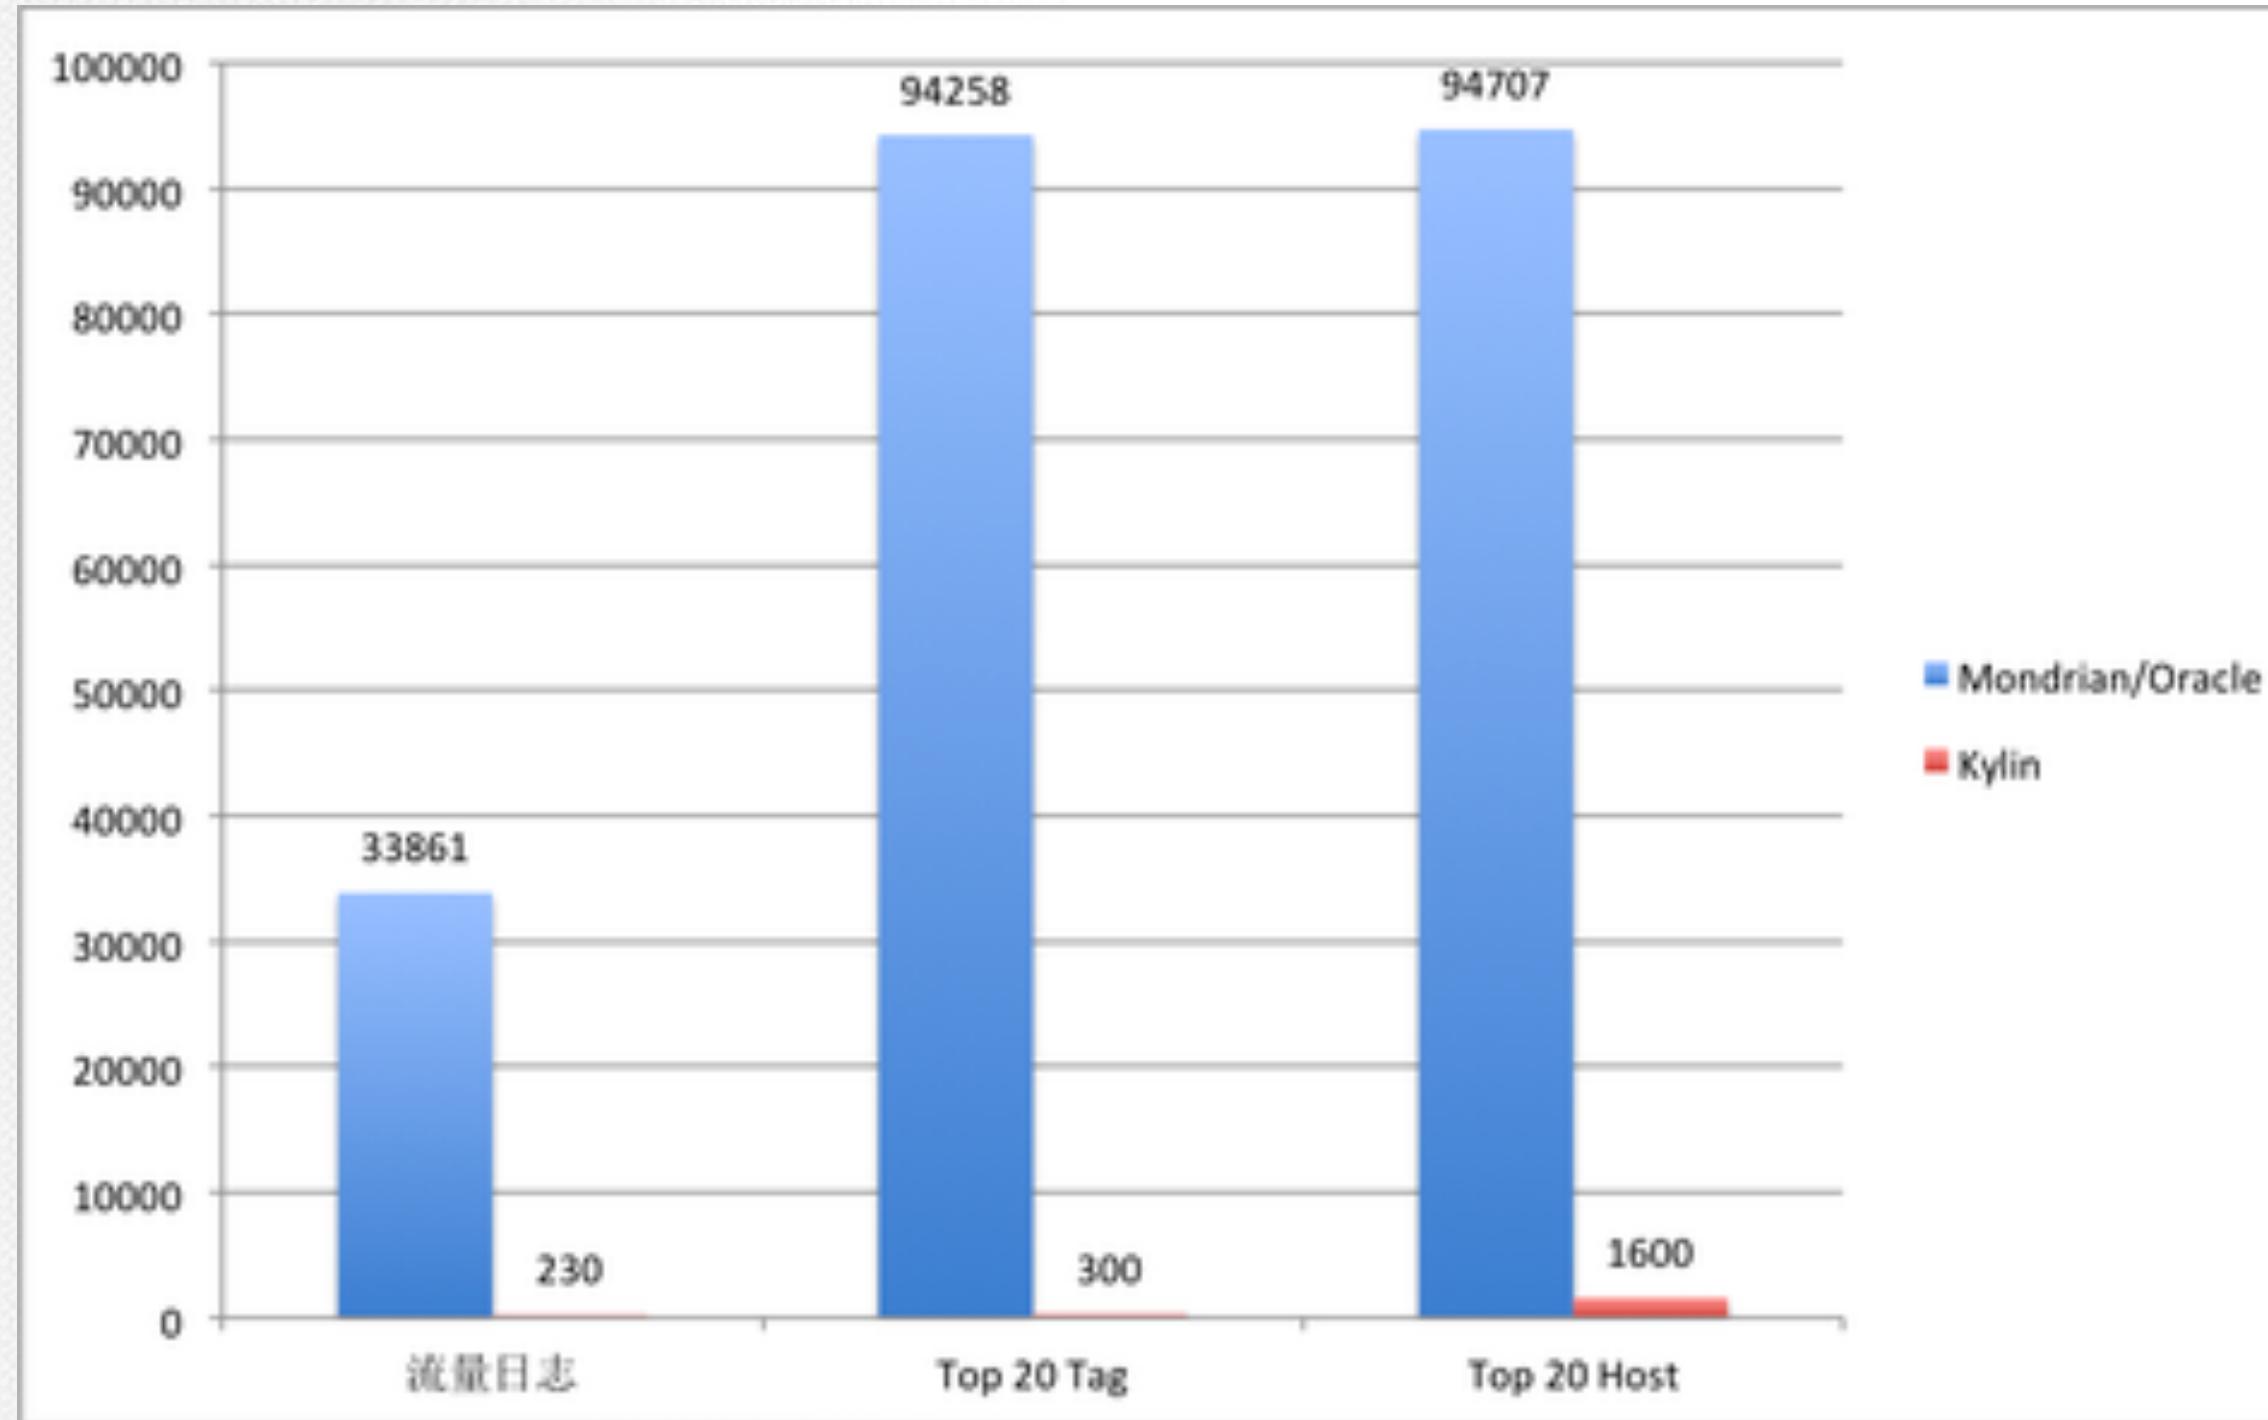


唯品会





Performance Benchmark



By NetEase:
<http://www.bitstech.net/2016/01/04/kylin-olap/>

Apache Kylin

- 千亿数据，亚秒级查询延迟
- 标准SQL，交互式分析
- 无缝集成，整合主流BI工具
- 非侵入式部署，可扩展架构
- 水平扩展，从容应对高并发
- 快速实施，无需编码



Kyligence Inc

- ✓ Formed by creator of Apache Kylin
- ✓ Kylin PMC members: 6
- ✓ Contributed over 90% Kylin code
- ✓ Success stories from variety industries
- ✓ Global partnership
- ✓ Certified training
- ✓ 5x8 or 7x24 expert support



Ecosystem

App



Data Mart
Data Warehouse



Data Lake
Big Data Platform



Cloud



Solutions

SI

Q & A

- **For open source Kylin**
 - dev@kylin.apache.org
 - Twitter: @ApacheKylin
 - <http://kylin.apache.org>
- **For Kyligence Inc**
 - info@kyligence.io
 - Twitter: @Kyligence
 - <http://kyligence.io>



WeChat: Kyligence



WeChat: ApacheKylin