

# Evolution of Open Source Data Infrastructure

Past, Present, and Future

Fangjin Yang  
Cofounder @  imply

# QCon

2016.10.20~22

上海·宝华万豪酒店

## 全球软件开发大会 2016

### [上海站]



购票热线: 010-64738142

会务咨询: [qcon@cn.infoq.com](mailto:qcon@cn.infoq.com)

赞助咨询: [sponsor@cn.infoq.com](mailto:sponsor@cn.infoq.com)

议题提交: [speakers@cn.infoq.com](mailto:speakers@cn.infoq.com)

在线咨询 (QQ): 1173834688

团 · 购 · 享 · 受 · 更 · 多 · 优 · 惠

# 7折

优惠 (截至06月21日)  
现在报名, 立省2040元/张

# Overview

**Simpler times with small data**

**The rise of open source**

**Current open source landscape**

**Where are we headed?**

# Data Insights

Broadly, we care about two use cases:

- OLTP
- OLAP

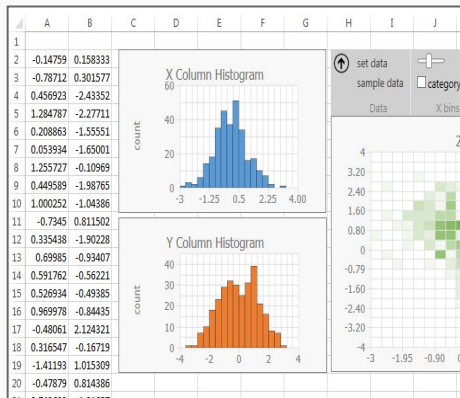
OLTP - business processing - dealing with transactions

OLAP - reporting - business intelligence

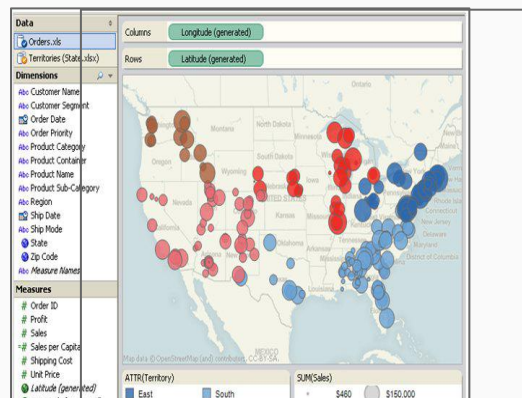
OLAP data - dimensions & measures

# Small Data

# Small Data Analytics



Excel




Tableau



- Solutions are very simple
- Fast and easy to extract insights
- Easy to create different custom visualizations

# Data Growth



MPP databases?

- Oracle, Teradata, IBM, Microsoft, etc.

Proprietary databases are expensive!

# The Rise of Hadoop



# Hadoop

Google GFS paper published in 2003

Google MapReduce paper published in 2004

Nutch project started in 2005 at Yahoo

Nutch became Hadoop and was open sourced in 2006

Community quickly grew

# Early Open Source Stacks



# Hadoop



# Hadoop

When one technology becomes very adopted, its limitations also become more well known

Hadoop is a very flexible solution

Most commonly used for data processing

Not optimized for many things - many inefficiencies!

# Hadoop



# Rise of Open Source Data Infrastructure


Things Hadoop isn't good at:

- Fast queries
- Deliver (streams of) events
- Stream processing
- In-memory computation

These limitations led to new technologies to be created

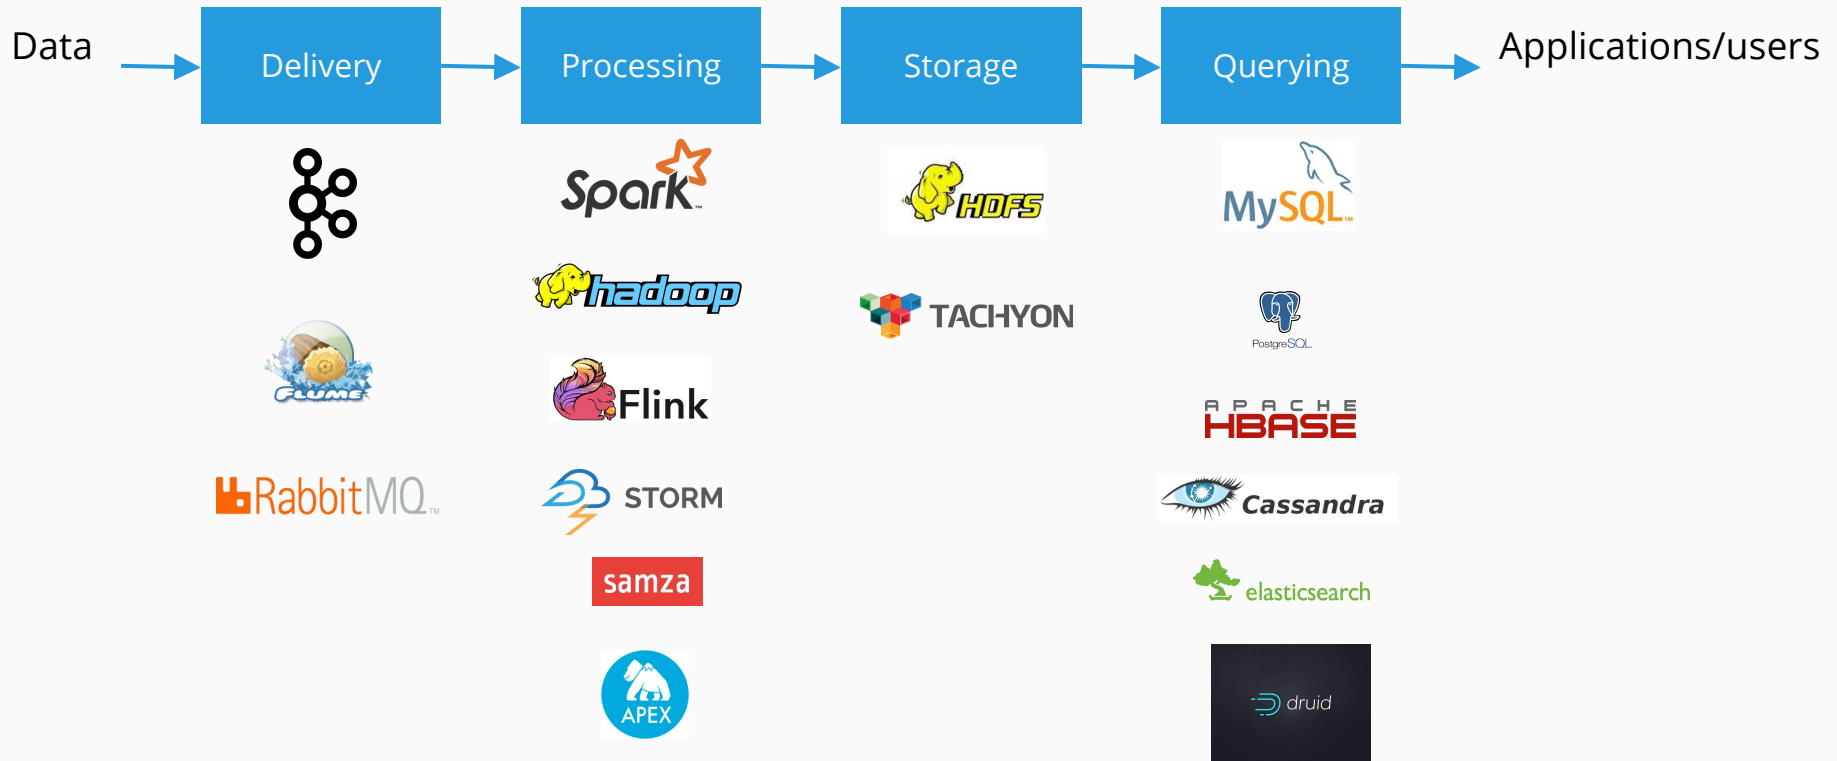
# **Data Infrastructure Space Today**

# Modern Open Source Stacks





# Modern Open Source Stacks



# Data Delivery





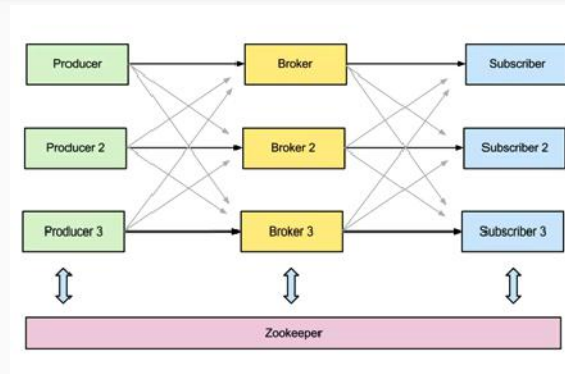
# Data Delivery

Focus is storing data for a limited time and delivering it elsewhere

Three different approaches

- Apache Kafka - publish/subscribe, transaction queues
- RabbitMQ - publish/subscribe, distributed queues
- Apache Flume - push-based event delivery

# Data Delivery



# Storage



# Storage

Distributed file systems

Store data indefinitely

Standard: HDFS

Can overlap with delivery systems (e.g. Kafka)





# Processing

Processing systems are designed to transform data

Has overlap with querying systems

- Query systems: output set smaller than input set
- Processing systems: output set same size as input set
- Having separation is more standard nowadays



# Stream Processing





# Stream Processing

Systems deal with unbounded messages/events


Different approaches

- Spark Streaming
- Storm
- Samza
- Flink
- Kafka Streams
- etc.

# Stream Processing

## Spark Streaming

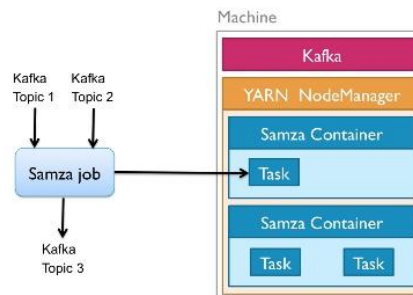
discretized stream processing




records processed in batches with short tasks  
each batch is a RDD (partitioned dataset)

samza

## Samza Architecture



UBER



STORM

# Batch Processing




# Batch Processing

Manipulate (large) static sets of data


Different approaches

- Spark
- Hadoop

# Batch Processing




→ Move data through disk and network (HDFS)



→ User can cache data in memory

# Querying



# Querying

Largest and most complex (broad range of use cases)

Let's focus on the most common use case:

- Business intelligence/data warehousing/OLAP

Significant overlap with storage

- Separation is becoming more common

# SQL-on-Hadoop


Enable ad-hoc queries on different input formats

Examples: Impala, Hive, Spark SQL, Drill, Presto





# SQL-on-Hadoop



# SQL-on-Hadoop

Advantages:

- Flexible /w full SQL support

Disadvantages:

- Slow - serialization/deserialization can have overhead

Many new storage formats

- Apache Parquet, Apache Kudu, Apache Arrow, etc.

# Key/Value Stores

Very fast writes

Very fast lookups

Timeseries databases often have K/V storage engines



OPENTSDDB



# Key/Value Stores

## Pre-computation

- Pre-compute every possible query
- Pre-compute a set of queries
- Exponential scaling costs

ts	gender	age	revenue
I	M	18	\$0.15
I	F	25	\$1.03
I	F	18	\$0.01



Key	Value
I	revenue=\$1.19
I,M	revenue=\$0.15
I,F	revenue=\$1.04
I,18	revenue=\$0.16
I,25	revenue=\$1.03
I,M,18	revenue=\$0.15
I,F,18	revenue=\$0.01
I,F,25	revenue=\$1.03

# Key/Value Stores

## Range scans

- Primary key: dimensions/attributes
- Value: measures/metrics (things to aggregate)
- Still too slow!

ts	gender	age	revenue
I	M	18	\$0.15
I	F	25	\$1.03
I	F	18	\$0.01



Key	Value
I,M,18	0.15
I,F,18	0.01
I,F,25	1.03


# Column stores

Load/scan exactly what you need for a query

Different compression algorithms for different columns

- Encoding for string columns
- Compression for measure columns

Different indexes for different columns



# Druid

Targeted for extremely low latency queries - powering user-facing analytic applications

Custom column format optimized for event data and BI queries

Supports lots of concurrent reads

Streaming data ingestion

# So many choices!

Does the project solve your use case?

Is it stable? Cheap? Fast?

Is there an active and growing community?

10x faster or 10x cheaper -- upgrade!



# **The Next Few Years**

# General Trends

Number of projects reaching saturation point

Streaming computation

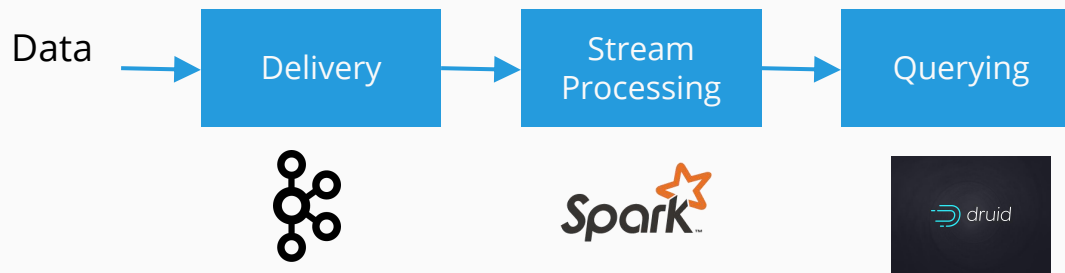
In-memory computation

Standards are slowly emerging


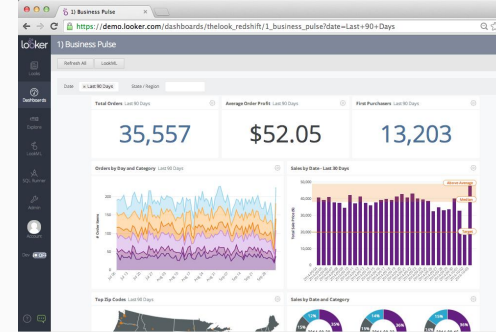
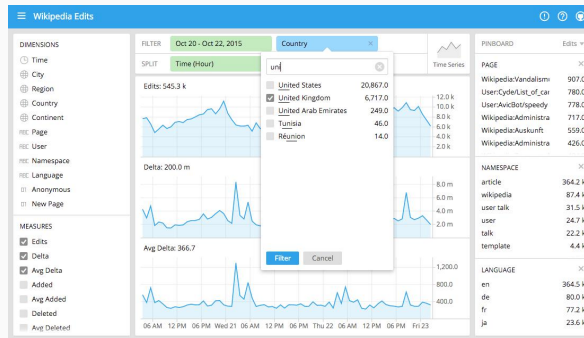
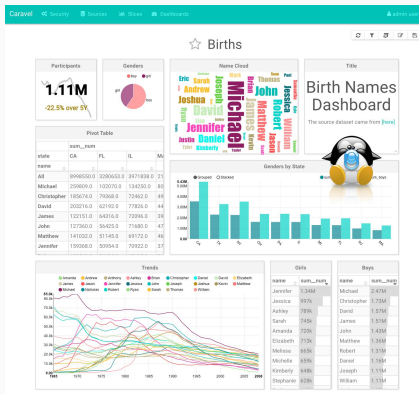
# Future Open Source Stack?



# Future Open Source Stack?



# Applications



DL4J Deep Learning for Java



# Thanks!

[imply.io](http://imply.io)  
[druid.io](http://druid.io)