

OW2 Open Source Corporate Network Meeting

Play with Big Data on the Shoulders of Open Source

Liu Jie

Technology Center of Software Engineering

Institute of Software, Chinese Academy of Sciences

2012-10-19

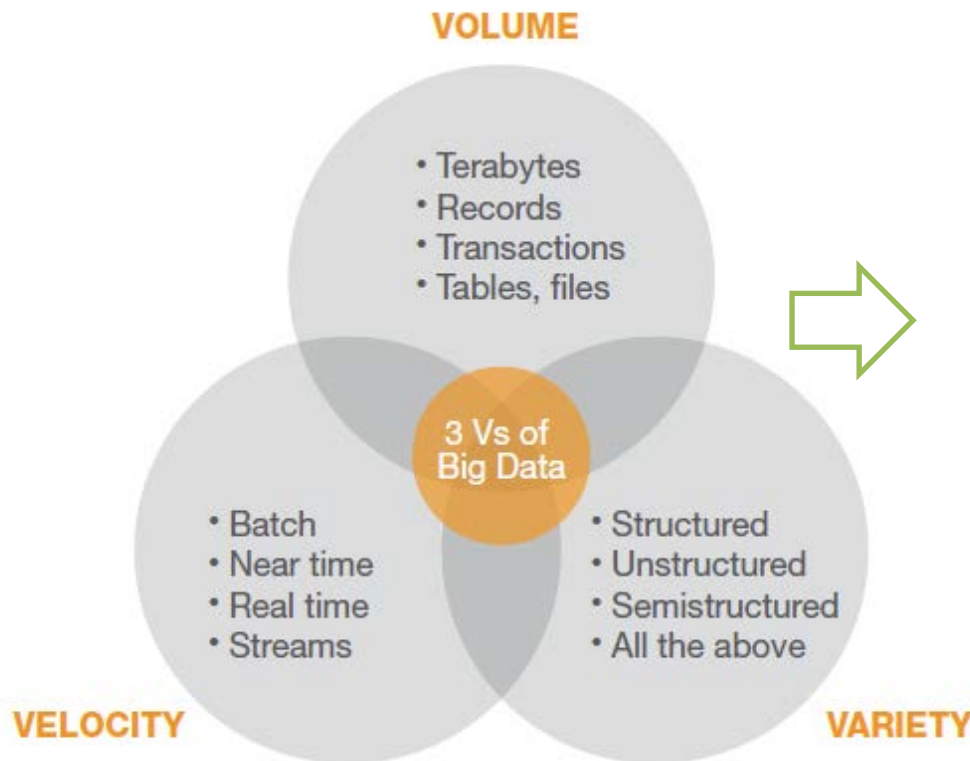

Big data changing the world

Expanding Data Sources

- Science and research
 - Gene sequences
 - LHC accelerator
 - Earth and space exploration
- Enterprise applications
 - Email, documents, files
 - Applications log
 - Transaction records
- Web 2.0 data
 - Search log / click stream
 - Twitter/ Blog / SNS
 - Wiki
- Other unstructured data
 - Video/Movie
 - Graphics
 - Digital widgets



Big data bigger challenges



Challenges

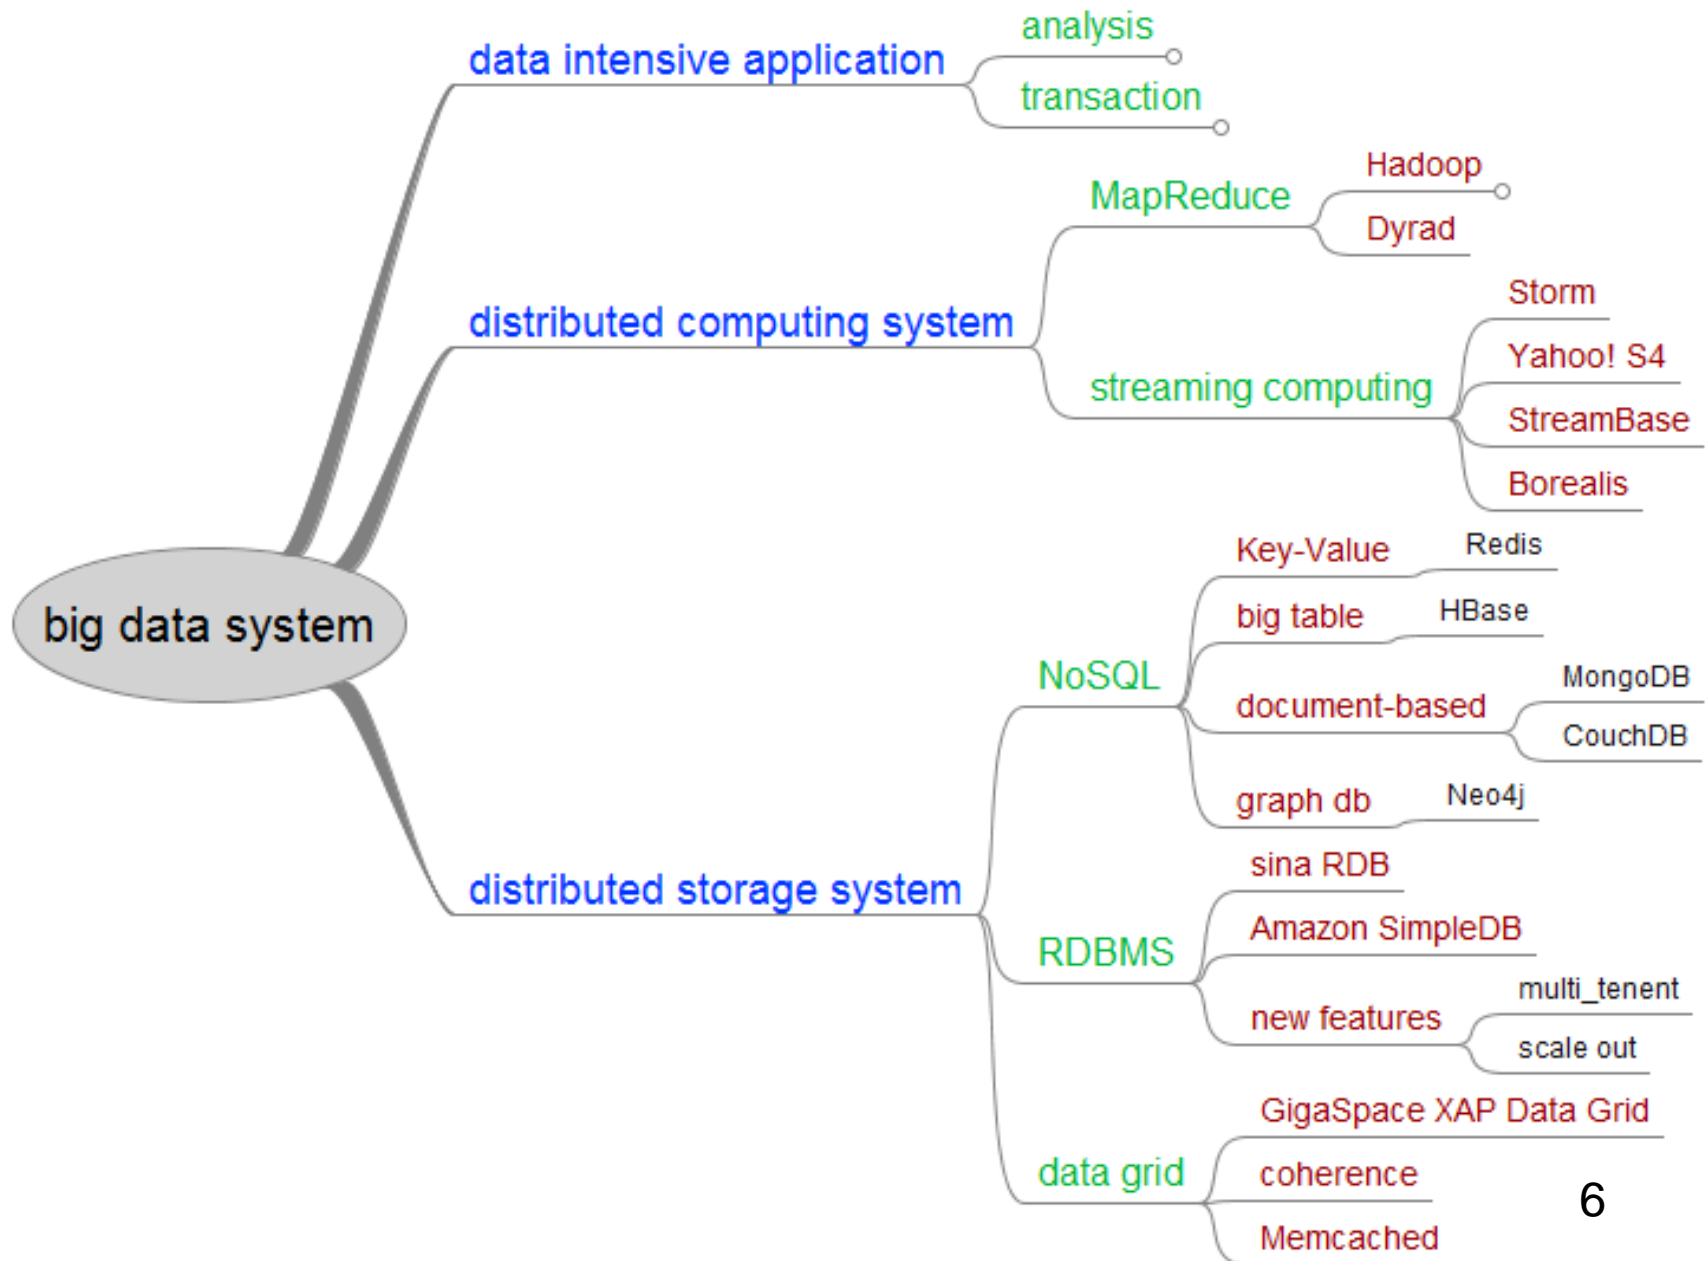
- Scale out automatically
 - Vs. scale up manually
- More capacity and bigger pool
 - E.g., 10 PB in a single file system
- New process capability
 - Loading, Analyzing, Moving data
 - Intelligence
- Better performance
 - Linear vs. exponent
 - Faster
- Autonomous
 - Fewer human interference
 - Lower cost

Outline

- Open source for big data
- MDD-MR
- RR-HBase
- Conclusion

Open Source for Big data

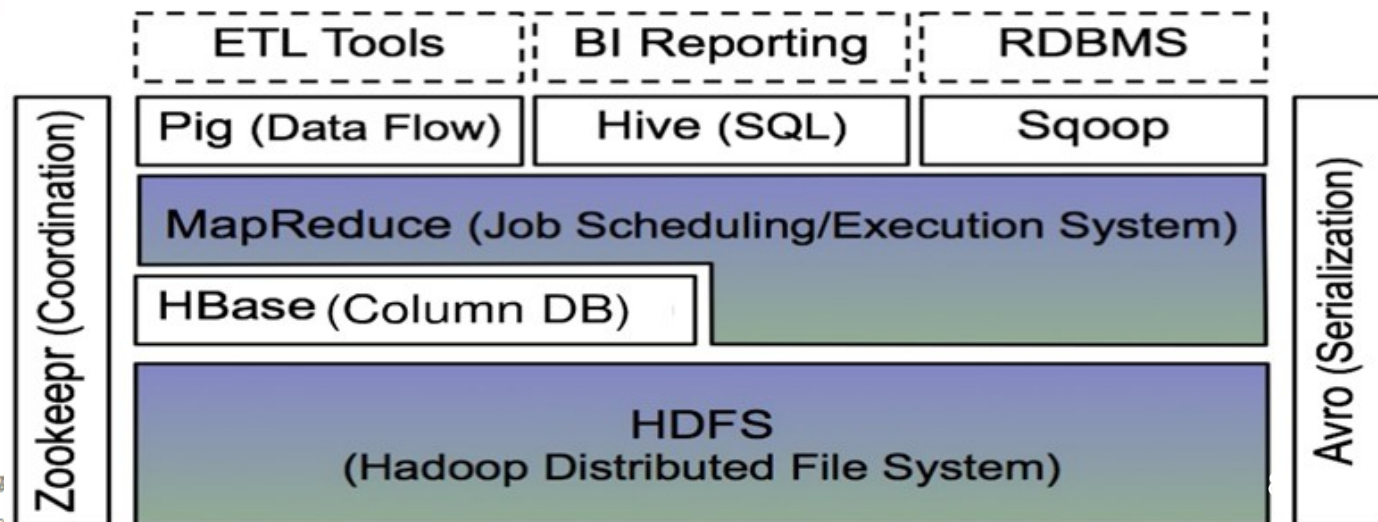
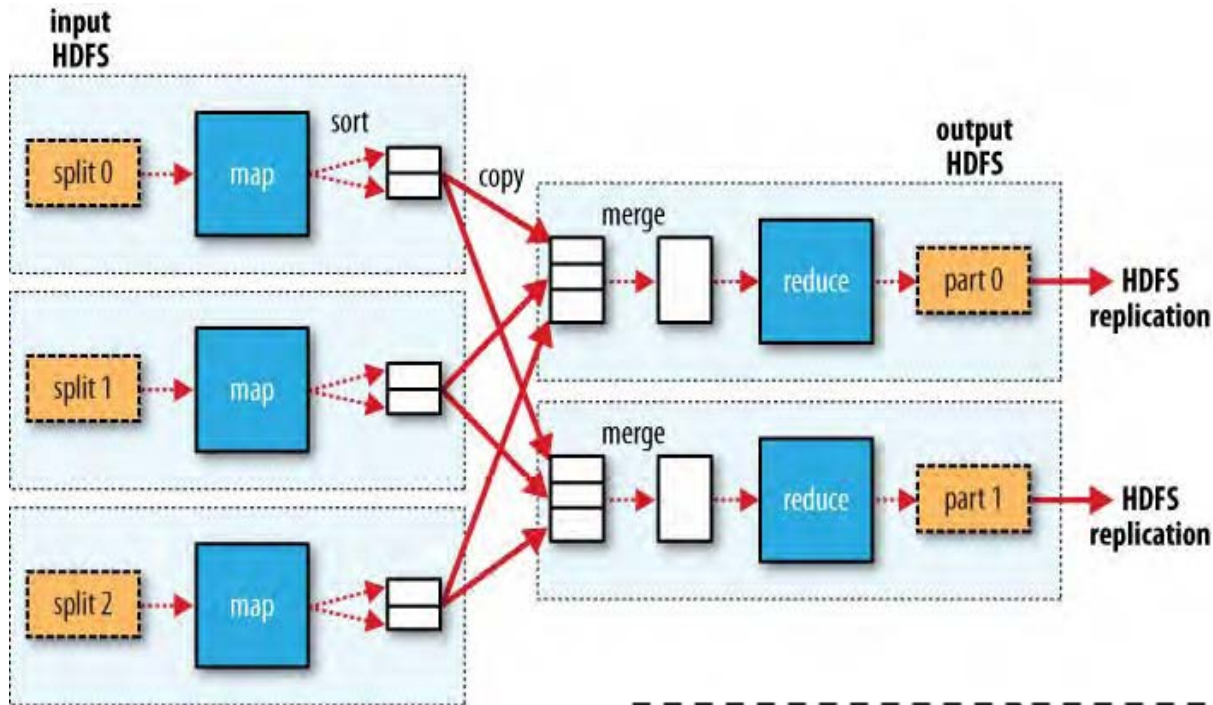
Open Source Big data system



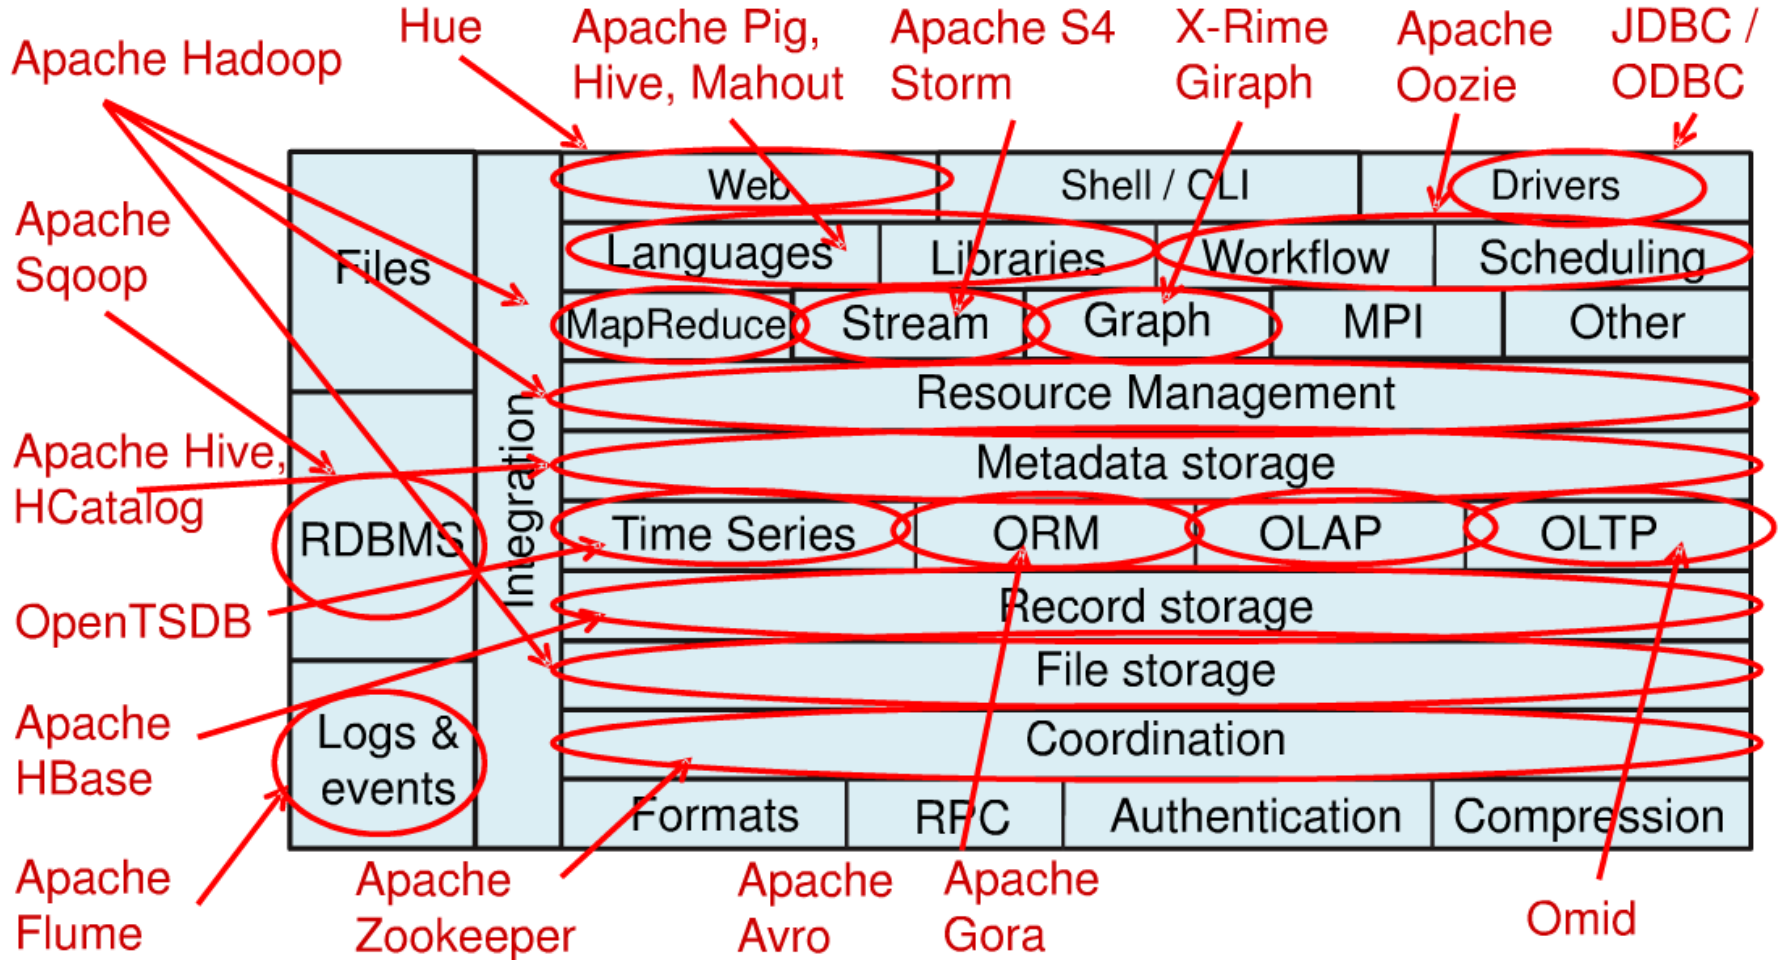
MapReduce

- (1) a **programming paradigm**:
 - it allows users to expressing distributed computations at a massive scale
 - need to provide two functions only: map() and reduce()
- (2) a description of a **processing pipeline and system**:
 - Execution framework for organizing and performing such computations
 - **scales to very large clusters**, > 10,000 nodes
- (3) several implementations of (2):
 - Google's proprietary MapReduce, Hadoop, ...

MapReduce & Hadoop



Hadoop Ecosystem

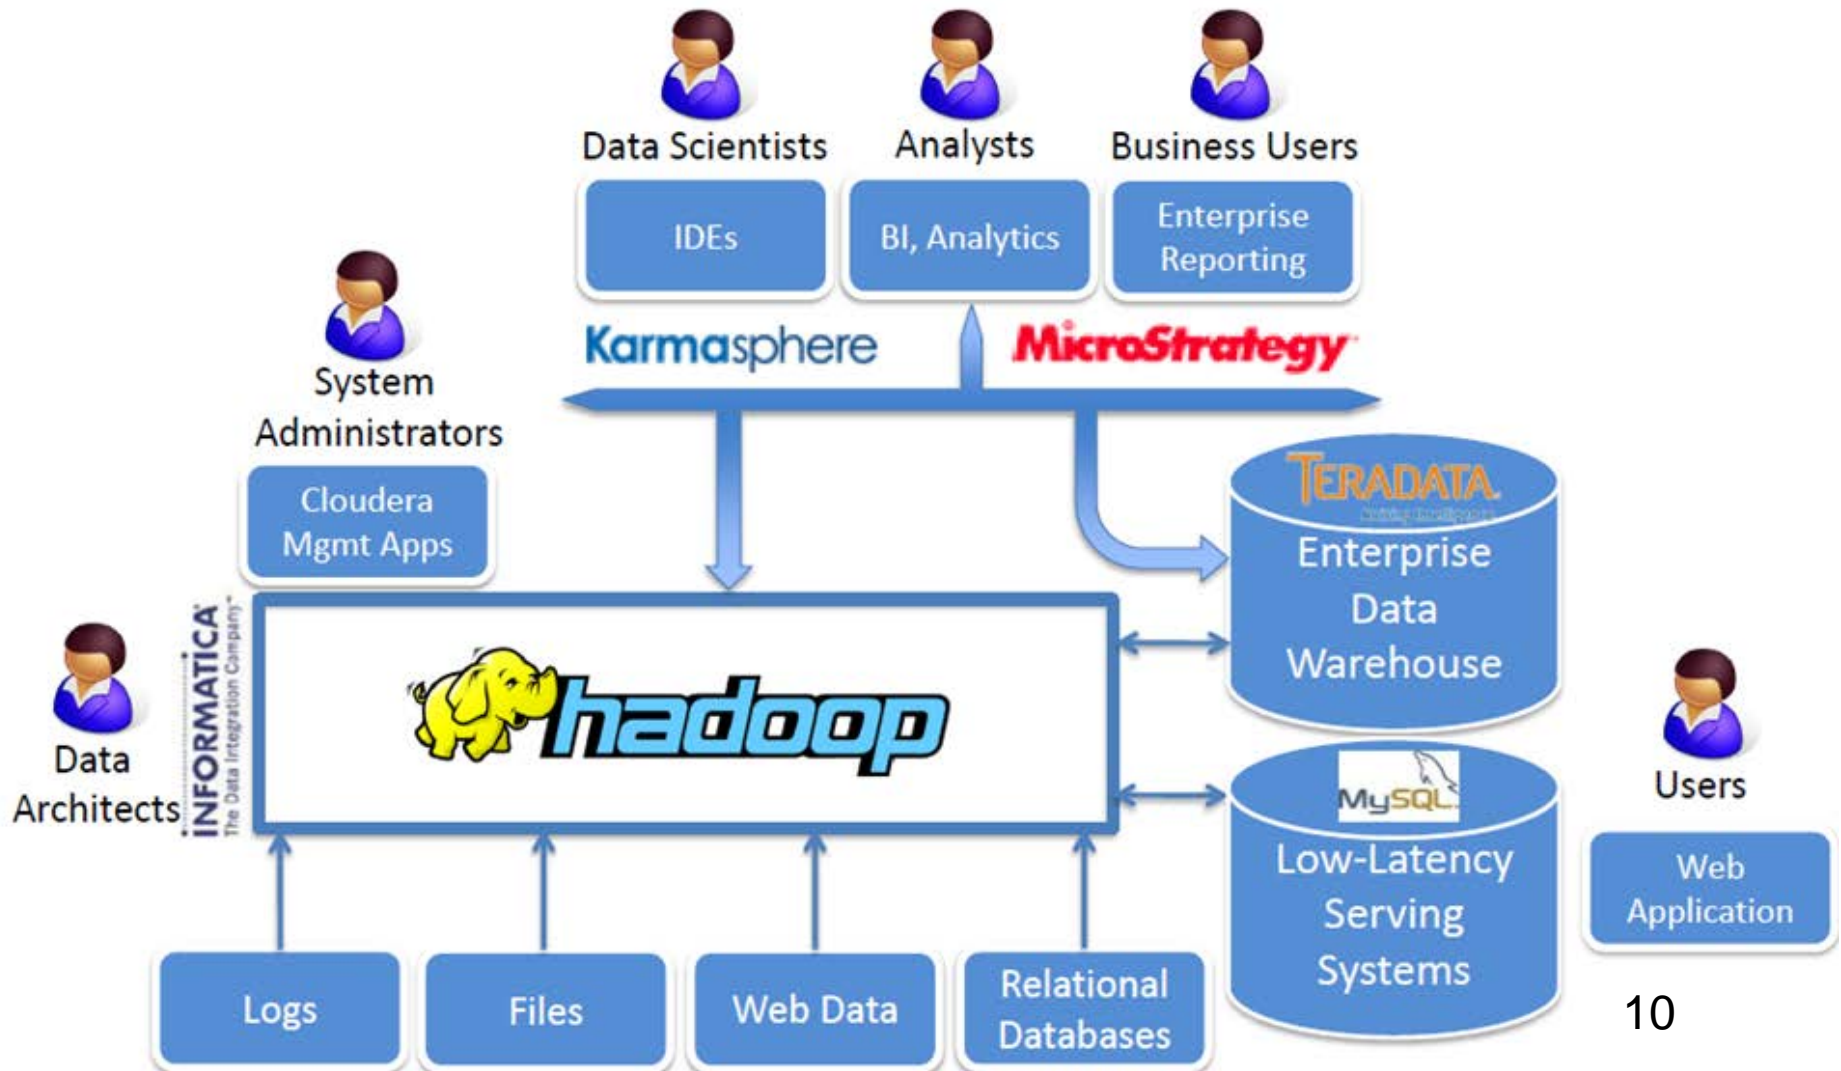


Hadoop-distribution:

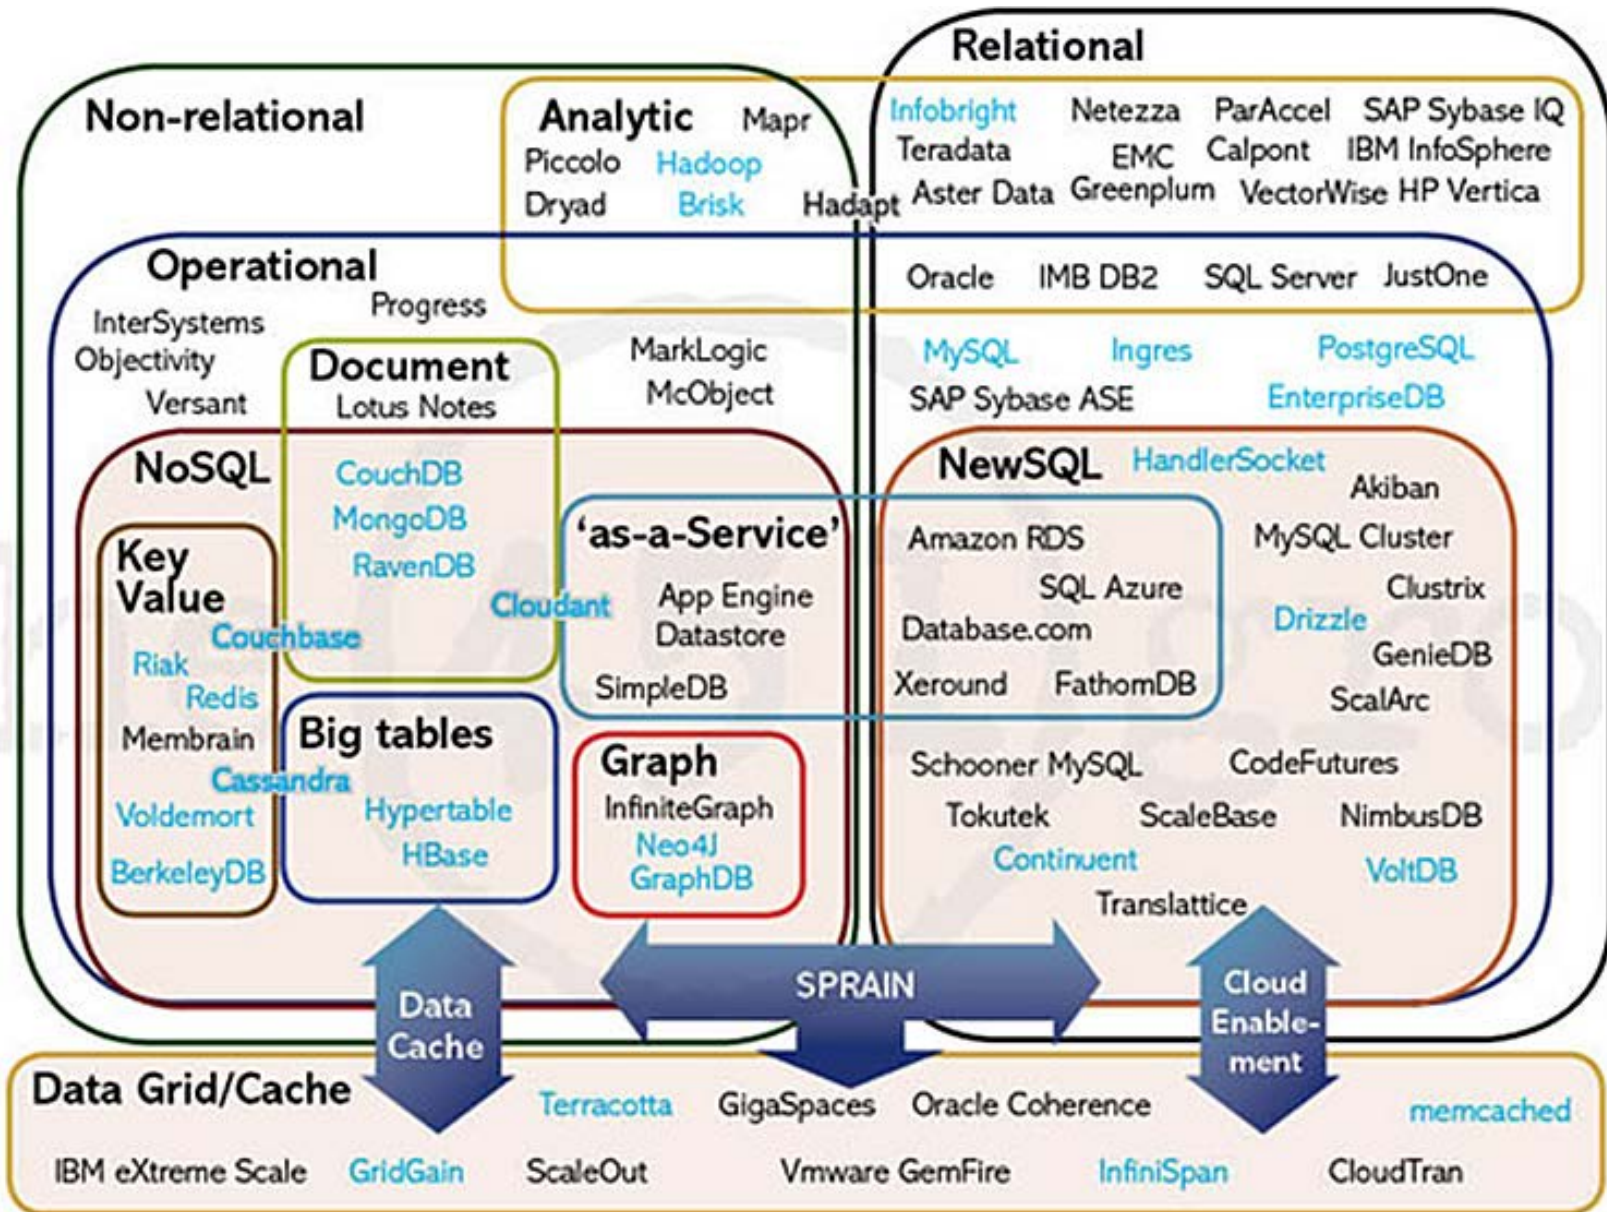
Cloudera , DataStax Brisk , EMC Greenplum HD Family, IBM InfoSphere BigInsights, Mapr M3 and M5, Platform

Hadoop Application

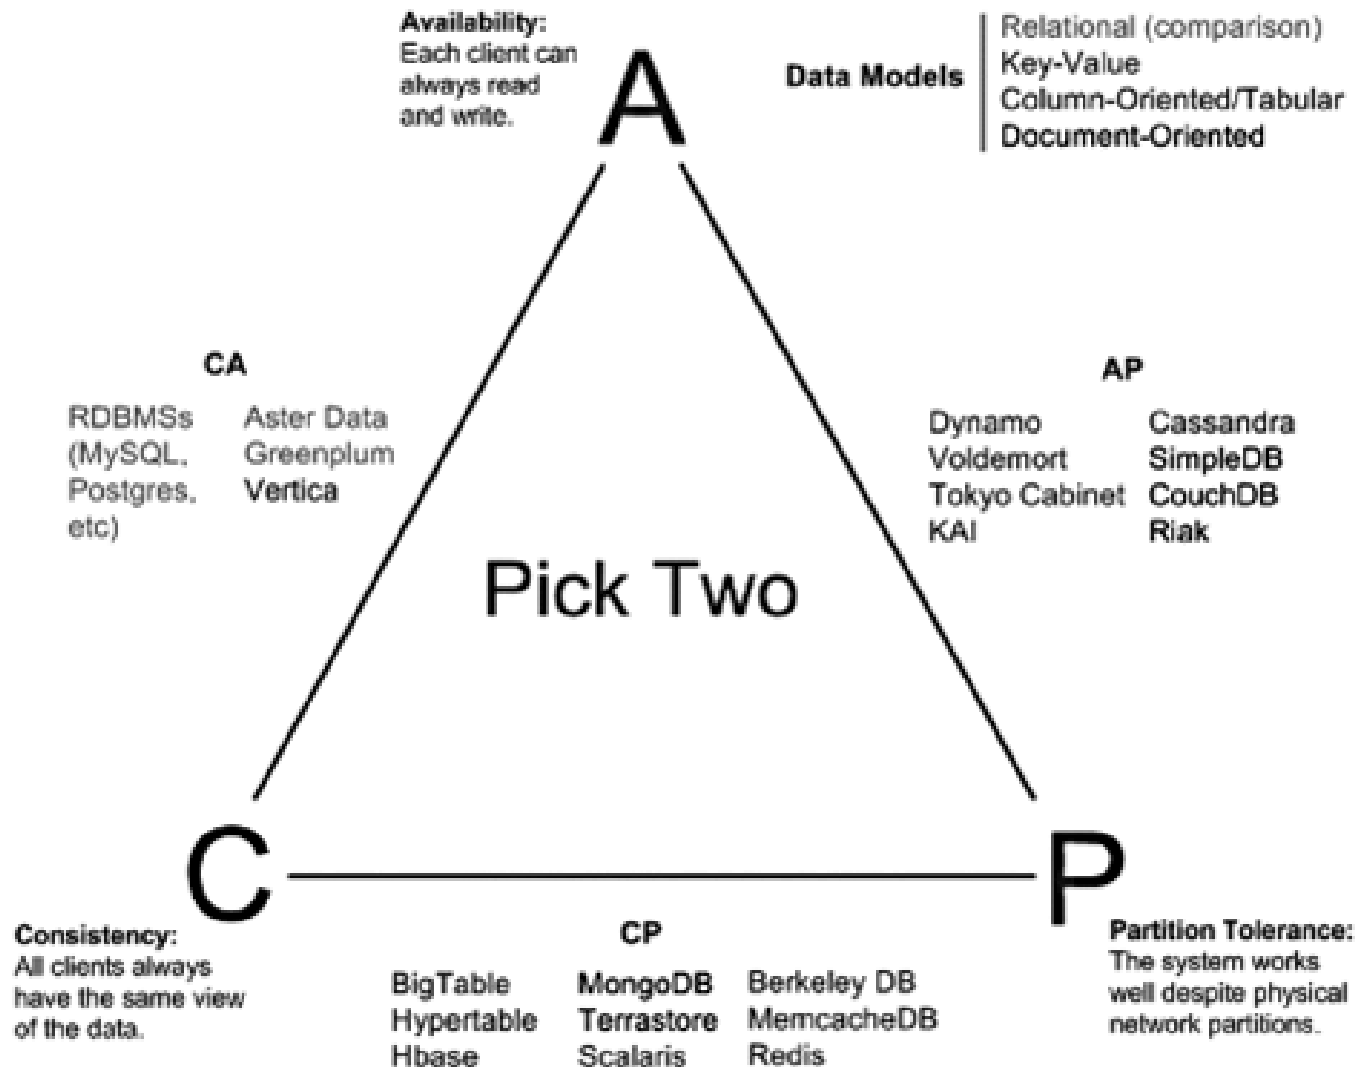
Where Does Hadoop Fit in the Enterprise Data Stack?



NoSQL



Distributed storage system



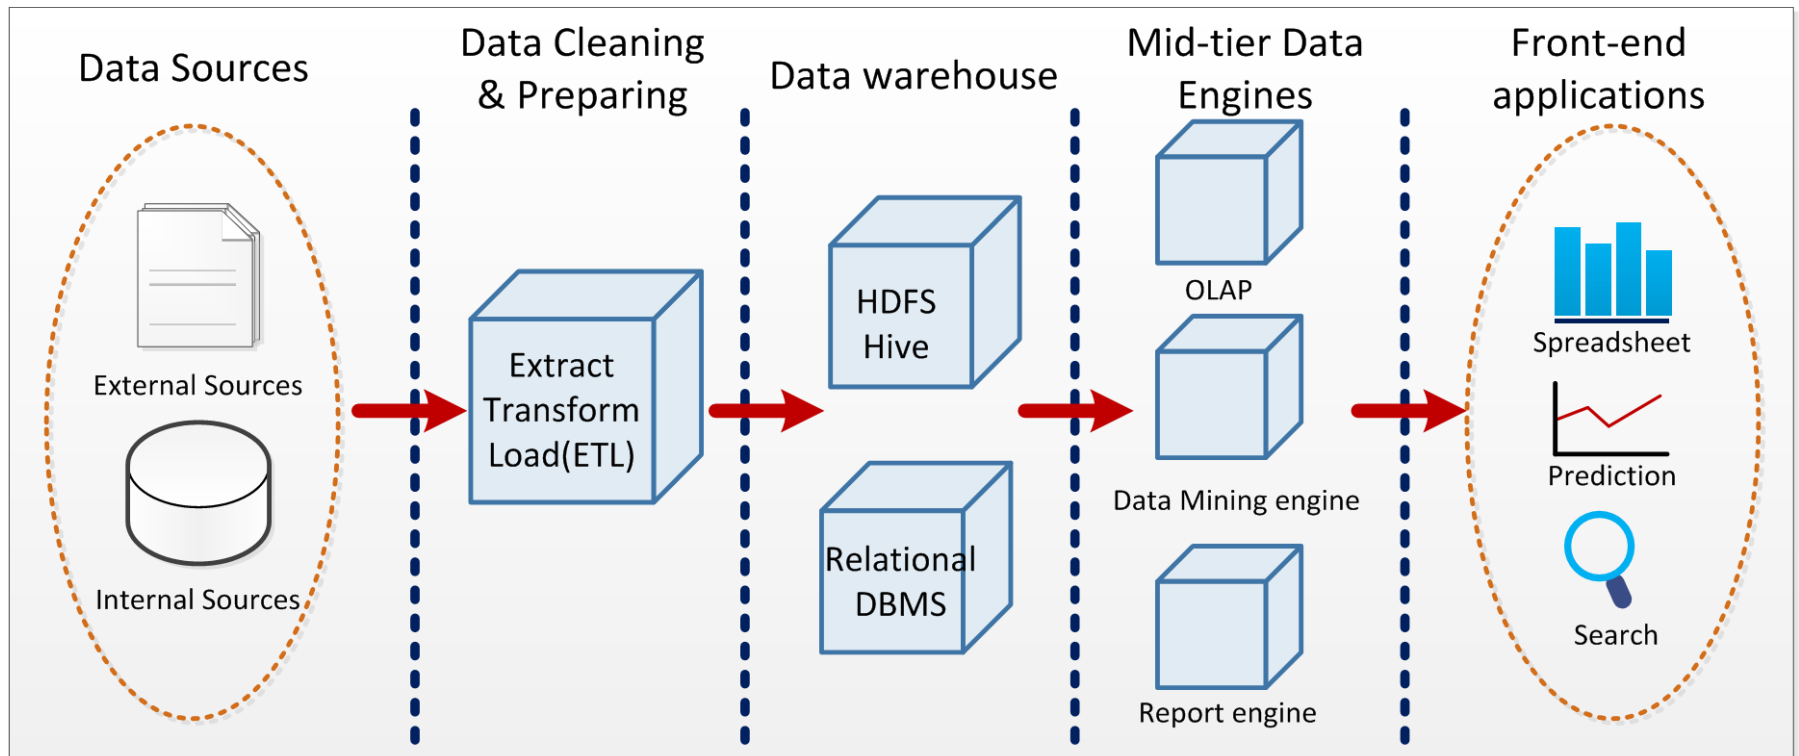
What can we do with these Open Source?

MDD-MR :

Model Driven Development framework for MapReduce Applications

Big data analysis Application

- Data warehousing and reporting
 - ETL、OLAP
- Deep analysis
 - Predicate analysis



(1) Hard to develop & maintain

- MapReduce DAGs: hand coded, high cost
- Not easy to handle changes

Write in Pig Latin:

or java code more complicated

```
-----  
visits      = load '/data/visits' as (user, url, time);  
gVisits     = group visits by url;  
visitCounts = foreach gVisits generate url, count(visits);  
  
urlInfo     = load '/data/urlInfo' as (url, category, pRank);  
visitCounts = join visitCounts by url, urlInfo by url;  
  
gCategories = group visitCounts by category;  
topUrls     = foreach gCategories generate top(visitCounts,10);  
  
store topUrls into '/data/topUrls';
```


Hard to monitor the progress, only the CMD & results

```
root@ubuntuOnceDQ: ~
s to: /root/pig_1295599472760.log
2011-01-21 16:44:32,889 [main] INFO org.apache.pig.backend.hadoop.executionengi
ne.HExecutionEngine - Connecting to hadoop file system at: file:///
grunt> visits = load '/data/visits' as (user, url, time);
2011-01-21 16:44:59,954 [ma
00: Error during parsing. L
018" (24), after : ""
Details at logfile: /root/p
grunt> gVisits = g
2011-01-21 16:44:59,957 [ma
00: Error during parsing. U
Details at logfile: /root/p
grunt> visitCounts = forea
2011-01-21 16:44:59,959 [ma
00: Error during parsing. U
Details at logfile: /root/p
grunt>
grunt> urlInfo = l
2011-01-21 16:44:59,973 [ma
00: Error during parsing. L
018" (24), after : ""
Details at logfile: /root/p
```

Queue details for default - Firefox - 火狐中国版

文件(F) 编辑(E) 查看(V) 历史(S) 书签(B) 工具(T) 帮助(H)

http://133.133.133.90:50030/jobqueue_details.jsp?queueName=default

Gmail - 收件箱 (1887) - ... x Google 日历 hadoop - Google 搜索 MapReduce Tutorial Queue details for def... x

Hadoop Job Queue Scheduling Information on **ubuntuOnceDQ**

Scheduling Information : Queue configuration
Capacity Percentage: 50.0%
User Limit: 100%
Priority Supported: NO

Map tasks
Capacity: 3 slots
Used capacity: 0 (0.0% of Capacity)
Running tasks: 0

Reduce tasks
Capacity: 3 slots
Used capacity: 0 (0.0% of Capacity)
Running tasks: 0

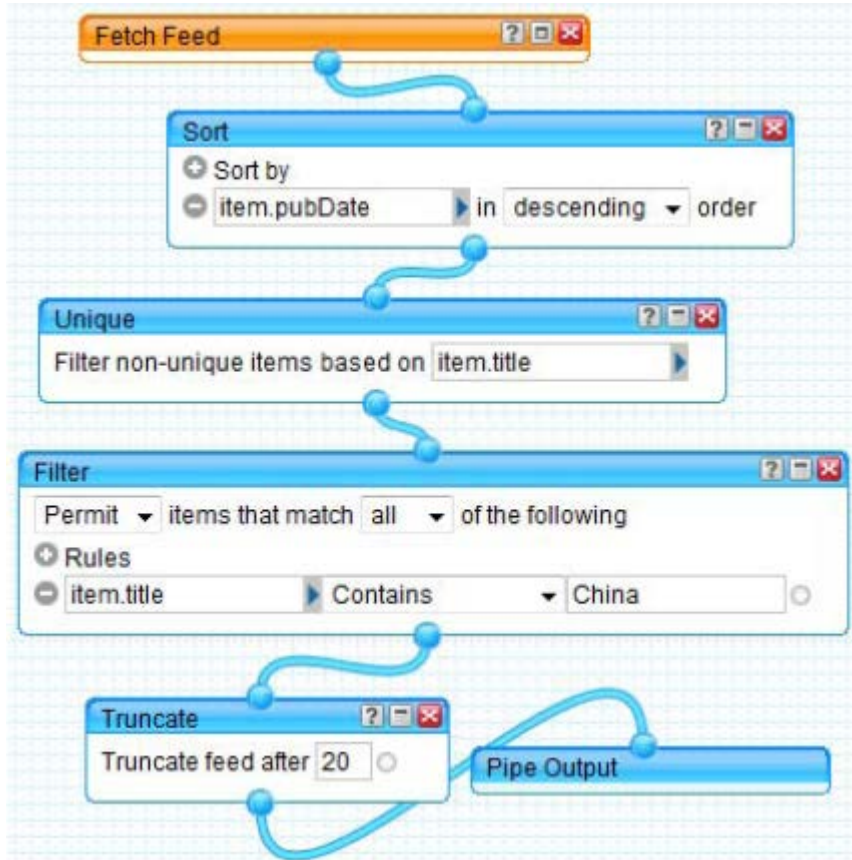
Job info
Number of Waiting Jobs: 0
Number of users who have submitted jobs: 0

No Jobs found for the Queue :: default

[Hadoop](#), 2011.

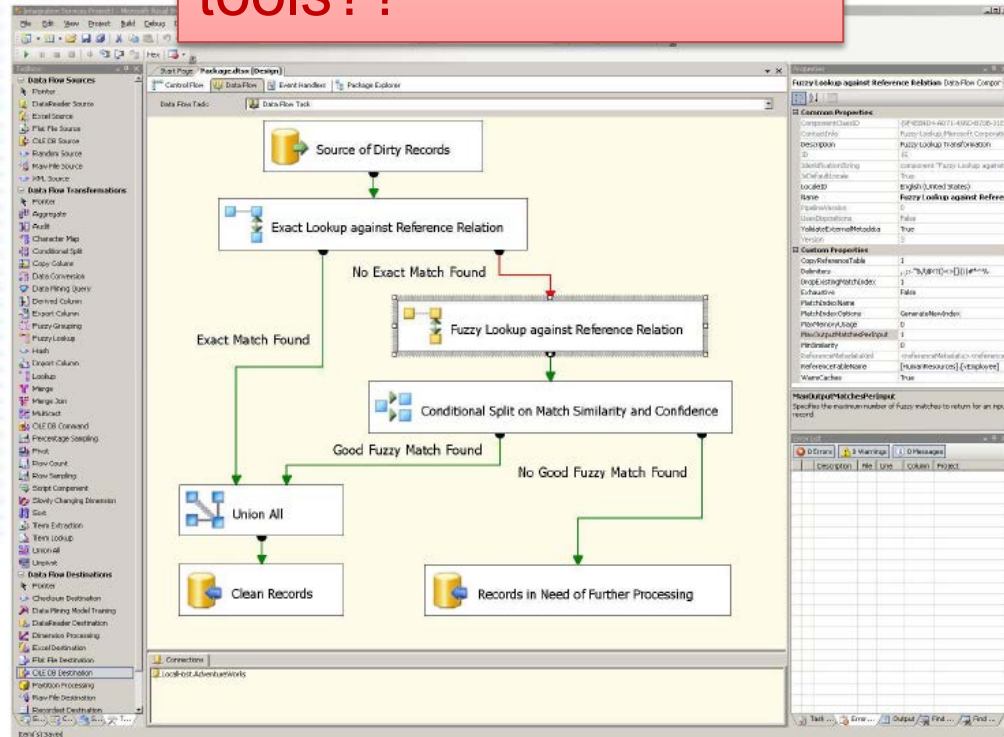
完成 zotero

Better approach in Existing platform —visually design and maintain



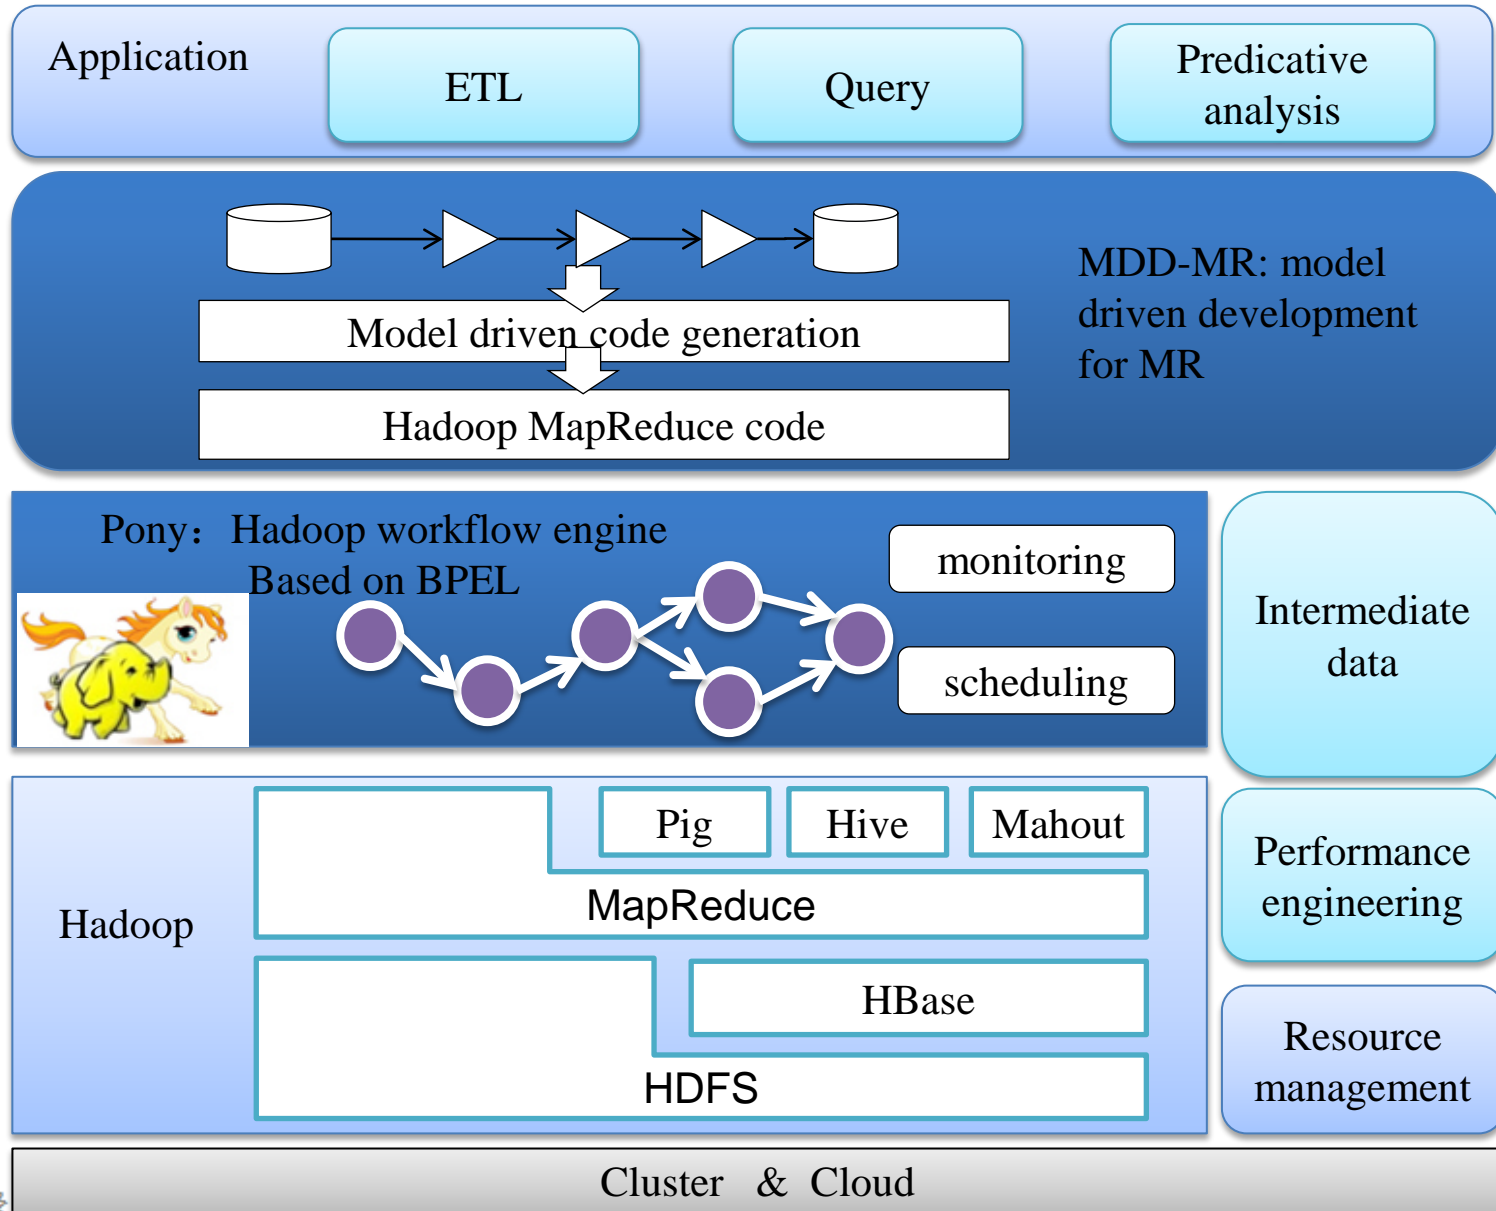
Mashups:
Yahoo Pipes, et al

Can we design
MapReduce DAGS
Just like in these
tools??



ETL:
MS SQL Server (and IBM, Oracle...)

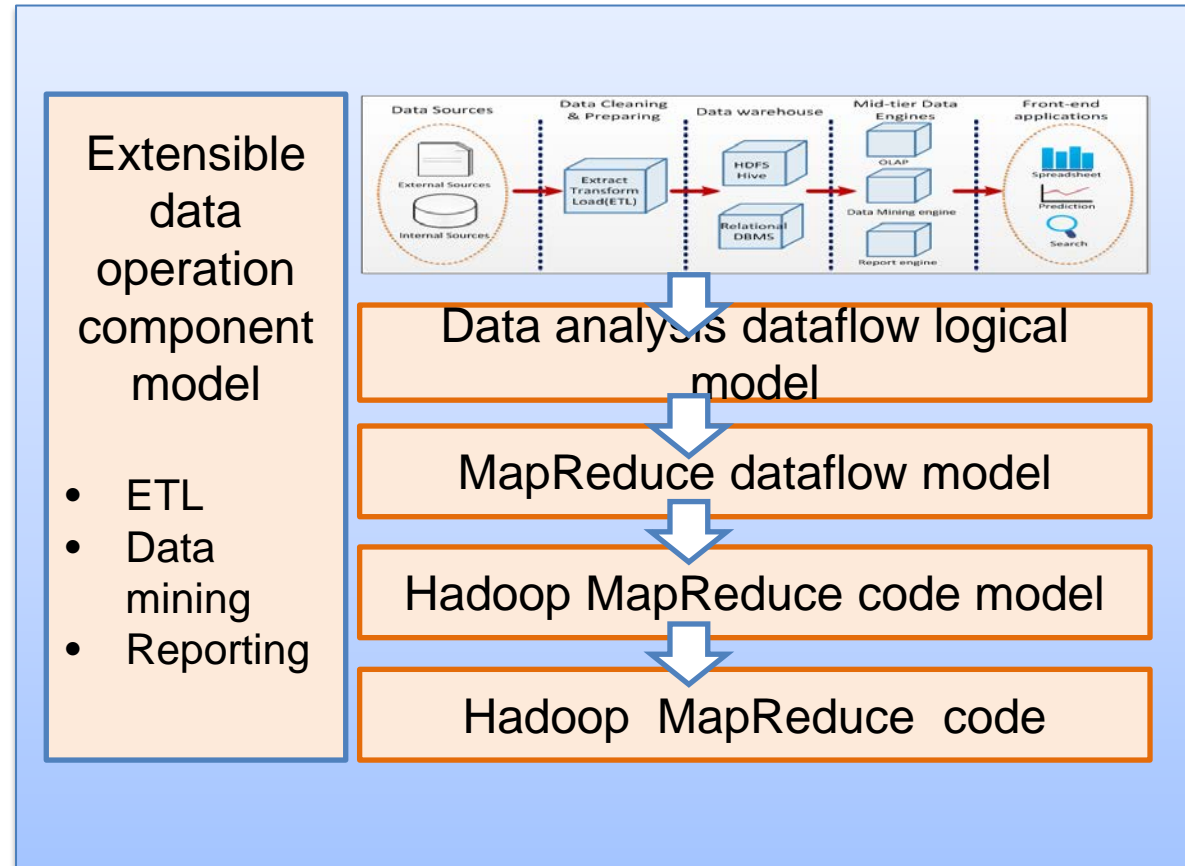
Our work



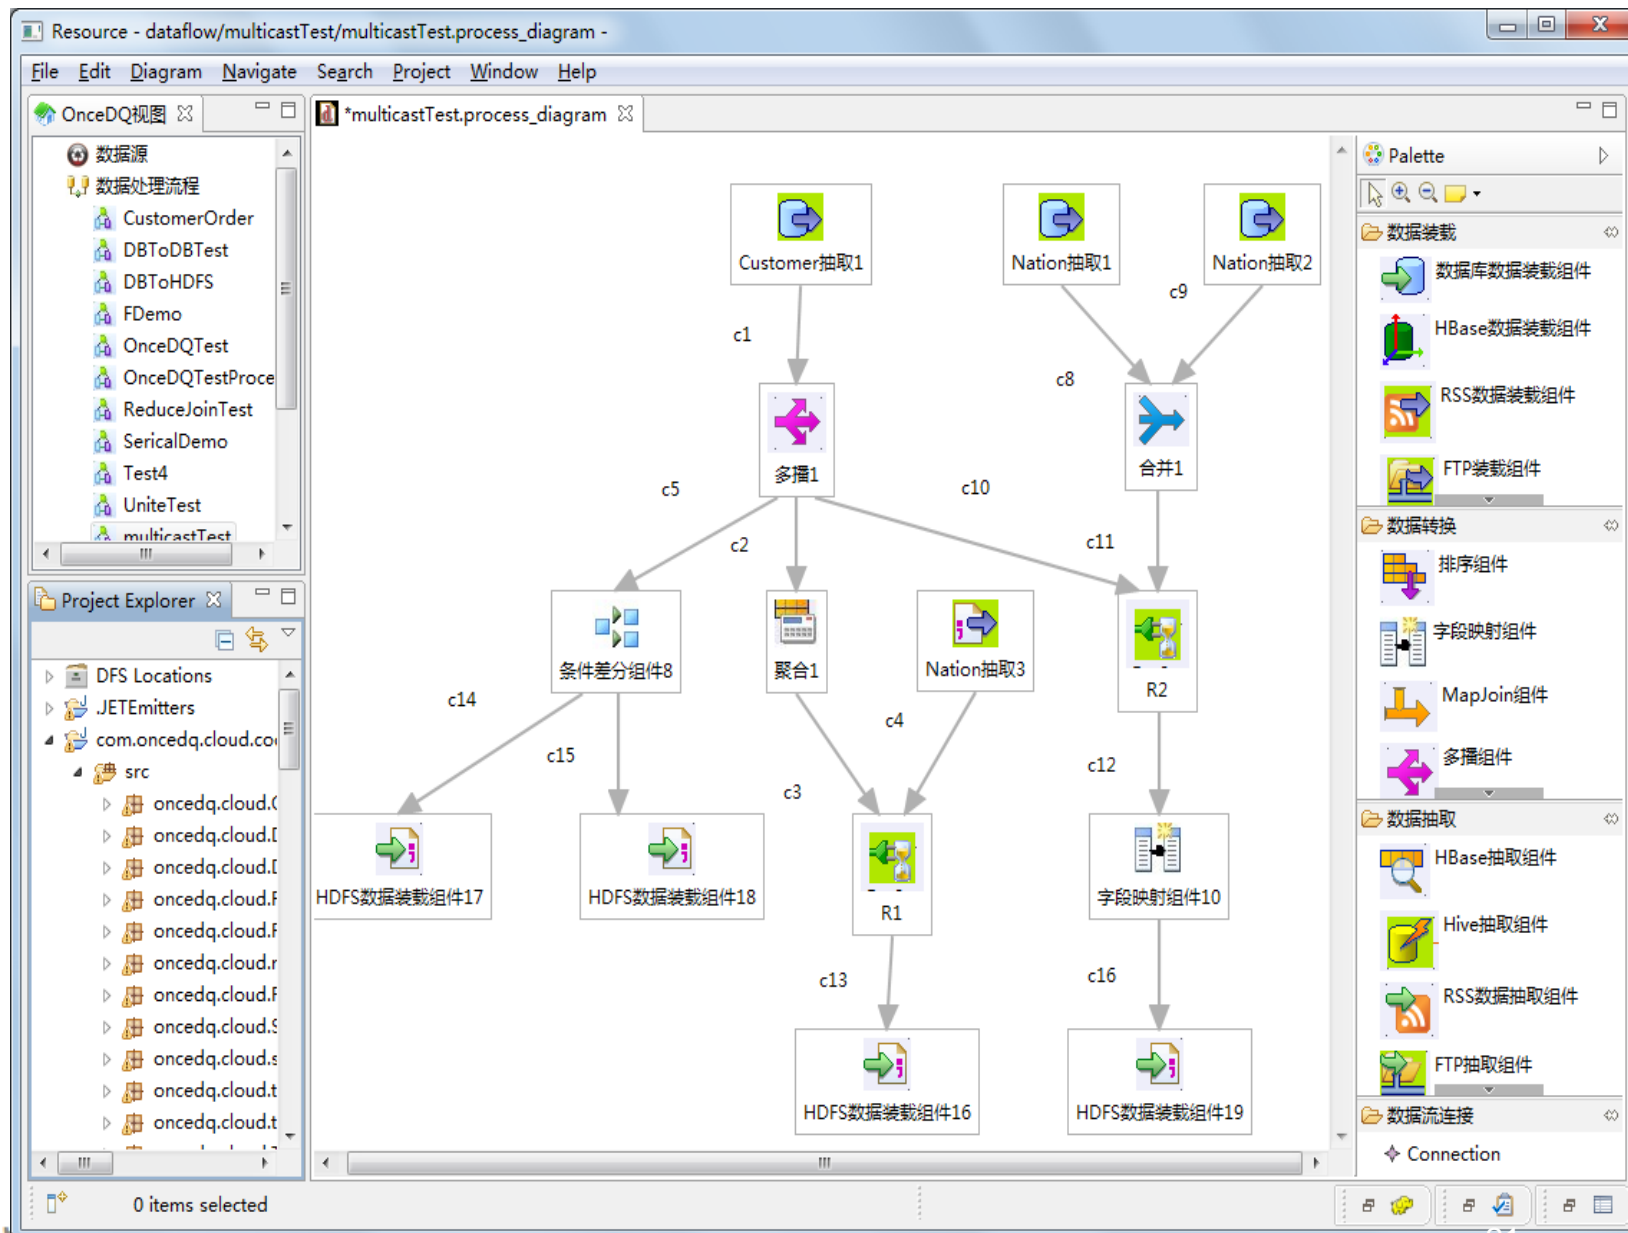
MDD-MR

○ Features

- **Visually design dataflow application for cloud computing platform**
- Support **data sources any where**, in the cloud or not
- **Innovative performance optimization techniques.**
- Tens operation components and can add more



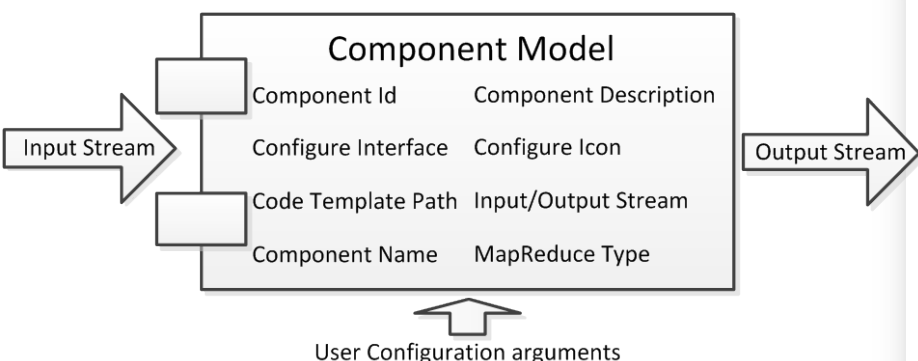
MDD-MR



MDD-MR

○ component model

- Define by OSGi extension point
- User can define new component



```
<?xml version="1.0" encoding="UTF-8"?>
<?eclipse version="3.4"?>
<plugin>
  <extension
    point="com.oncebi.components.component">
    <component
      description="Arima经济学预测模型"
      configInterface="com.oncebi.economy.model.arima.ui.ArimaCo
nfigPage"
      icon="images/arima.png"
      id="com.oncebi.economy.model.arima"
      name="Arima经济学预测模型"
      mapredType="None"
      oType="Model"
      template="Template/"
      templatePath="Template/"
      type="unblocked">
      <input
        description="输入流"
        id="arimaIn"
        name="ARIMA输入"
      >
    </input>
  </component>
</extension>
</plugin>
</pre>
```


MDD-MR

- Data mining components
 - Mahout algorithm
 - R
- Reporting

The screenshot displays the Once BI Platform interface with several views highlighted by red boxes and labels:

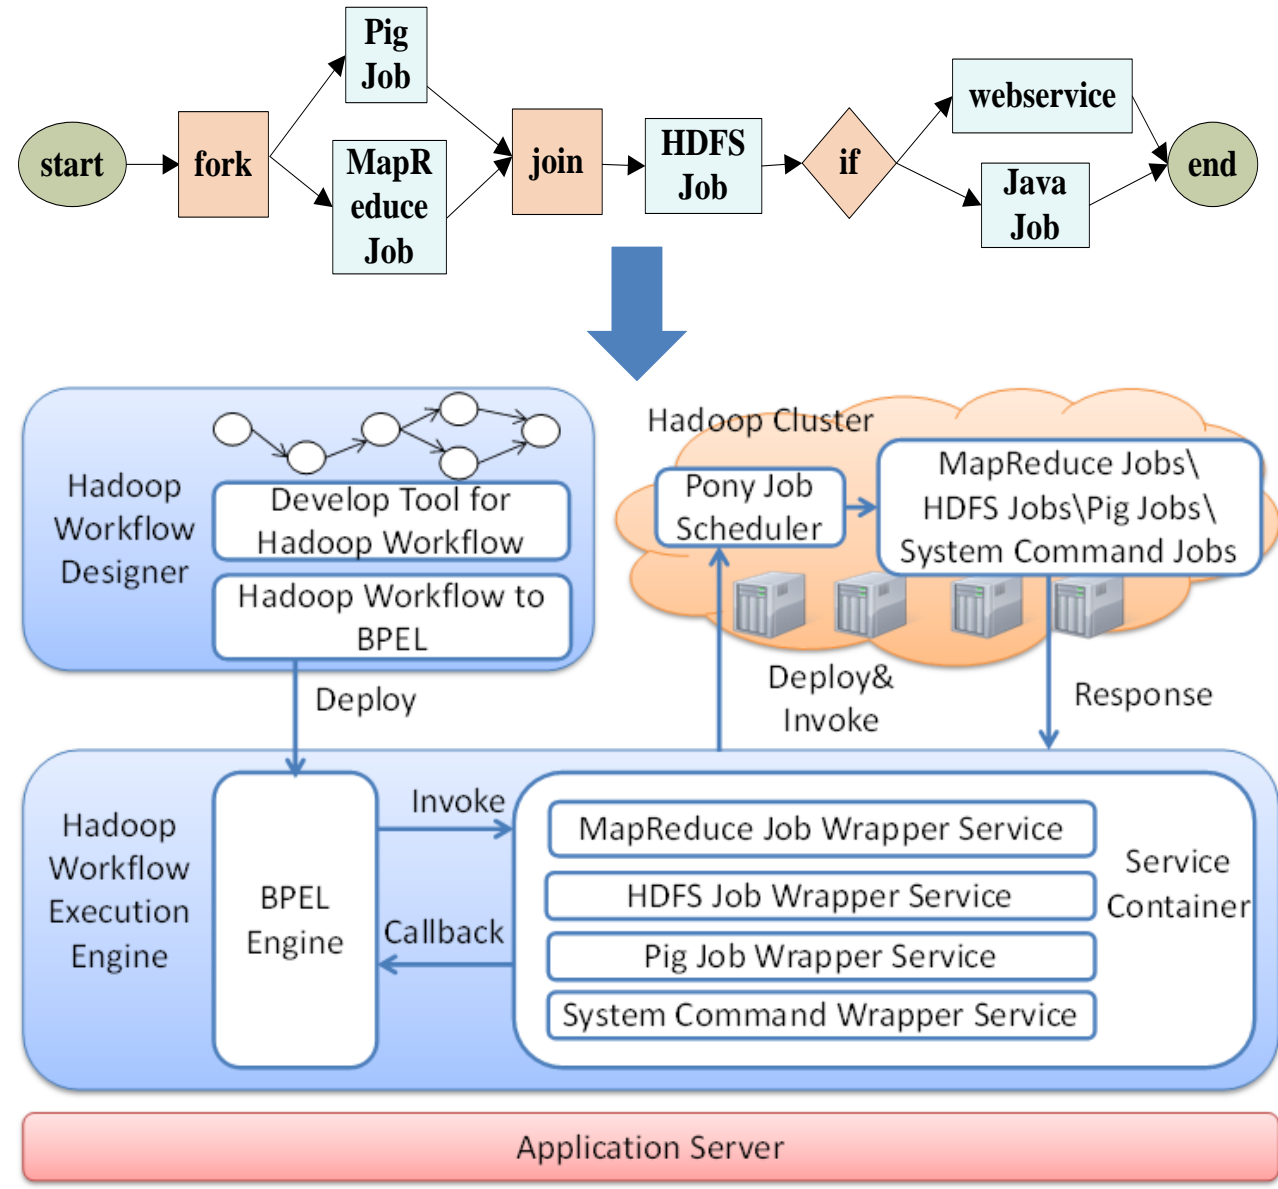
- OnceBi视图**: Located in the top-left pane, showing a project tree with components like 'arima', 'report', 'sot', and 'test'.
- 代码视图**: Located in the middle-left pane, showing a file explorer with code files like 'q.cloud.code'.
- R配置视图**: Located in the bottom-left pane, showing R environment configuration with 'Embedded - R_HOME: C:\R\R-2.14.0'.
- Hadoop集群配置视图**: Located in the bottom-left pane, showing Hadoop cluster configuration with 'Location: m110' and 'Master node: 133.133.2.110'.
- 过程配置视图**: Located in the center pane, showing a process diagram with components like '数据库数据抽取组件21', 'Arima经济学预测模型22', and '时间序列可视化组件23'.
- 组件属性视图**: Located in the bottom-right pane, showing the configuration for a component, including 'JDBC驱动: com.mysql.jdbc.Driver', '数据库连接URL: jdbc:mysql://133.133.134.32:3306/oncebi', and '过滤条件: name='花生' AND statistic_date>'2010-05'.
- 图表视图**: Located in the top-right pane, showing a histogram titled 'Histogram of rnorm(999)' with a frequency axis from 0 to 200 and an x-axis from -3 to 3.
- 报表视图**: Located in the bottom-right pane, showing a data table titled '新发地每日成交量价格统计表' (Xinfadi Daily Volume and Price Statistics Table).

品名	成交量 (吨)	最高价	最低价	均价
花生	47.000	7.00	4.00	5.50
花生	0.000	1.80	1.50	1.65
花生	0.000	2.00	1.80	1.90
花生	0.000	1.80	1.60	1.70
花生	0.000	1.80	1.60	1.70
花生	0.000	1.80	1.60	1.70
花生	0.000	2.00	1.60	1.80
花生	0.000	2.40	2.00	2.20
花生	0.000	2.40	2.00	2.20
花生	0.000	2.40	2.00	2.20
花生	436.580	6.00	5.80	5.80
花生	0.000	0.00	0.00	0.00
花生	85.000	8.00	5.00	6.50
花生	1.000	9.00	3.00	6.00
花生	19.000	8.00	3.00	5.50
花生	5.000	6.00	3.00	4.50

Pony: Workflow engine under MDD-MR

○ Hadoop workflow engine base on BPEL

- Base on OnceBPEL
- Visually deploy and monitor
- Easy to integrate with enterprise application




RR-HBase: Real-time Rational query system on HBase

Real-time OLAP

- Exploration Analysis
- Interactive Analysis
- Real-time statistic

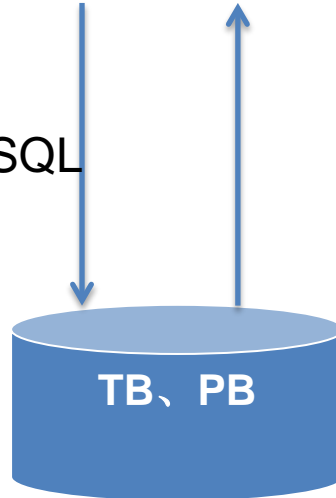
High Concurrency users

1	A	B	C	D	E	F
2	姓名	姓名	地址	电话	购买日期	到货日期
2	T80002722	郭洪涛	19楼21181		2003.6.16	2002.6.14
3	2500017612	周永通	6楼217楼在		2003.6.21	2002.6.21
4	2000017693	段建东	宿舍3楼2号		2003.6.22	2002.6.22
5	T800023439	于北新	聚德1-2-1		2003.6.22	2002.6.22
6	T800023412	南卫民	北京灯厂老办公大楼宿舍2楼2号		2003.6.22	2002.6.22
7	T800023845	董明海	大前门1楼5116号		2003.6.23	2002.6.23
8	2800023993	张永亮	成德机厂宿舍2楼111楼1号		2003.6.23	2002.6.23
9	T800011373	高小敏	中山路255号聚德3楼		2003.6.26	2002.6.26
10	2500018389	李小平	王府井6楼1楼6号		2003.6.26	2002.6.26
11	2000018284	刘刚	北京二环路77号建村村委会4楼5号		2003.6.26	2002.6.26
12	T800022943	刘心	江西南路		2003.6.26	2002.6.26
13	T800023224	文礼平	尚联 银行宿舍1174楼		2003.6.27	2002.6.27
14	2000018388	许福顺	贵德路2楼216楼8号		2003.6.27	2002.6.27
15	2000018385	孙文强	北京中路		2003.6.27	2002.6.27
16	T800023213	李国栋	天津路15号二区宿舍4B-2-9-1		2003.6.27	2002.6.27
17	T800023421	王元才	张发三德士尔街 5楼		2003.6.27	2002.6.27
18	T800023438	潘国华	王府井11楼13-11门		2003.6.28	2002.6.28
19	2000018385	曹应平	牌门楼		2003.6.29	2002.6.29
20	T800023437	高文成	康德路工行宿舍315楼		2003.6.29	2002.6.29
21	T800023442	方光俊	北航宿舍		2003.6.29	2002.6.29



SQL

Data stream



Real Time Analytics Use Cases

- Ecommerce – Auction monitoring, addwards
- Search engines
- Real-time Marketing – Improving conversion rate
- Weather reporting
- Traffic analysis
- Call Center Management
- Supply-Chain Optimization
- Quality Management in Manufacturing
- SLA Monitoring and Maintenance
- Global Shipment & Delivery Monitoring
- Fraud Detection in Financial Companies

Why NoSQL?

solution	Query model	scalability	performance	cost
Relational Database (Oracle,DB2..)	SQL	Manual partition	High cost	Commercial product
Paralleling dataflow programming framework (Hadoop)	MapReduce Hive:SQL subset	good	high latency	open source
NoSQL Database (HBase)	Key-value: Simple model with too simple query interface	good	Fast query	Open source

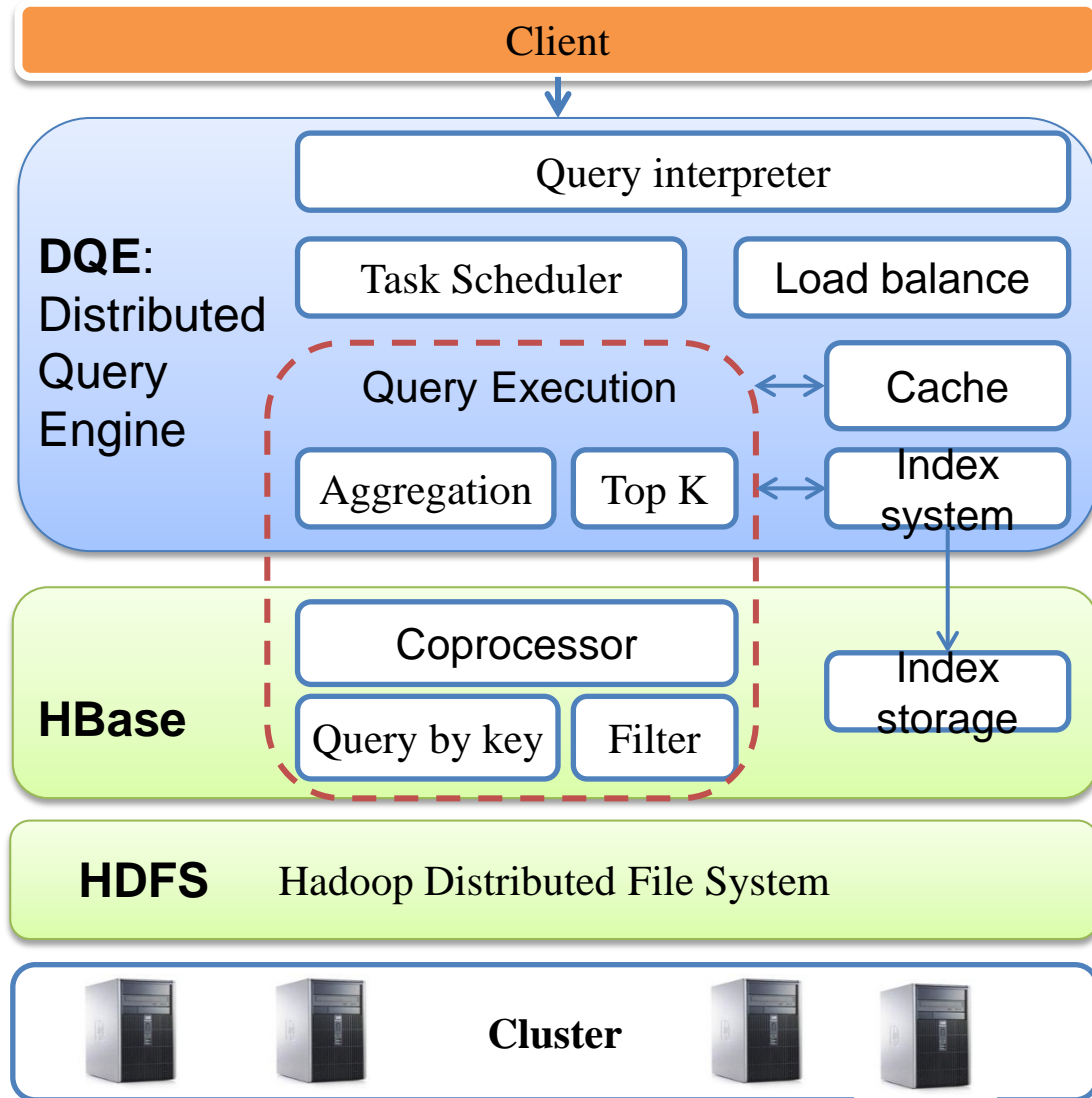
RR-HBase

Function

- High throughput data loading
- High concurrency and low latency point query and range query
- Paralleling Aggregation and join

Features

- HBase as the storage level
- Indexing and caching
- Scalable \ partition \ fault tolerance



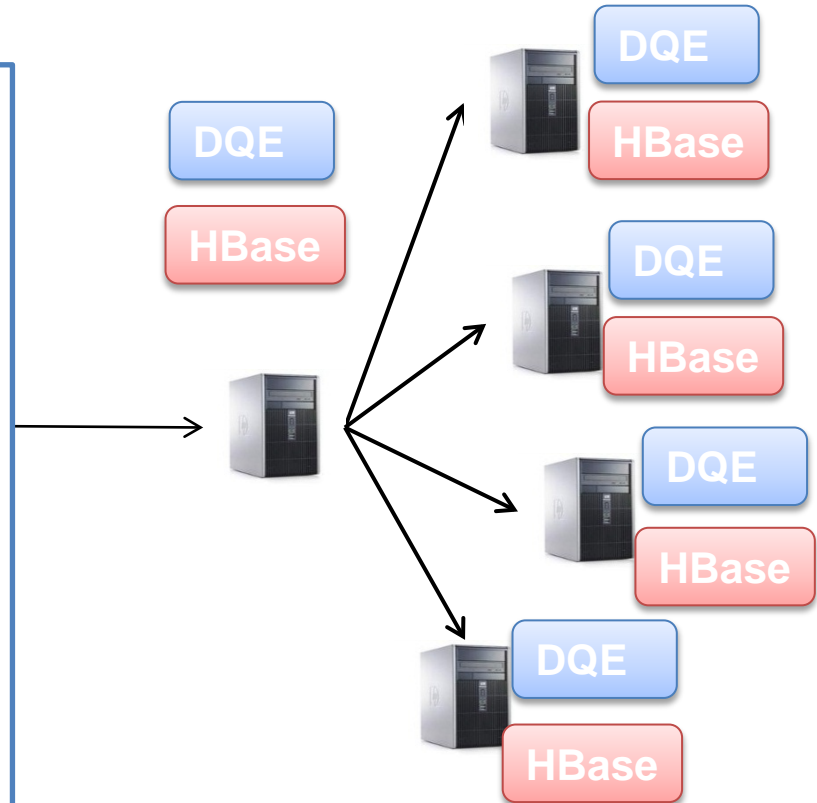
Technology Comparison

	YunTable	Infobright	EMC GreenPlum	Oracle Exadata	Hive	RR-HBase
Query Semantic	SQL	SQL	SQL	SQL	SQL subset	SQL
Technology	Distribute database	MySQL+ column storage	Parallel databae+HDFS	combines massive memory and low-cost disks	MapReduce	Column storage+MR
Performance	Very fast	good	fast	Very fast	slow	Maybe Fast
Data size	PB	TB	Near PB	Near PB	PB	PB
Cost	low	high	high	Very high	low	low
Advantages	Low cost with high performance	MySQL compatible	Very good performance	Enterprise software	Hadoop	Hadoop
Open Source	yes	no	no	no	yes	yes

RR-HBase: case study

- WISE Challenge 2012: Sina Weibo data analysis
 - data: Fellowship network 12.8G、tweet 61.7G
 - Achieve low latency and high throughput for 19 queries
 - ranked the 4th in the Performance Track:
Throughput/Latency/Scalability Measurement.

```
SELECT f1.uid
FROM
(SELECT a.friendID AS uid
FROM FriendList AS a JOIN friendList As b
ON a.friendID= b.uid
WHERE a.uid= b.friendID AND a.uid= "A") AS f,
friendList AS f1 JOIN friendList AS f2
ON f1.friendID = f2.uid
WHERE f2.uid= f.uid AND
f1.friendID = f2.uid AND
f1.uid= f2.friendID AND
f1.uid<>"A" AND
f1.uid<>f.uid
GROUP BY f1.uid
ORDER BY COUNT (f2.uid) DESC
LIMIT 10;
```



Conclusion

Conclusion

○ MDD-MR

- A model driven development framework for MR application
- Provide tens data operation components
- Support ETL, data mining, Reporting
- Can be extended and personalize for certain purposes

○ RR-HBase

- Fast and high concurrency SQL query on HBase
- Fast data loading
- Scalable
- In development

Thank you!

Some pictures in the PPT are obtained from web.