



# Flink: 构建下一代大数据处理引擎

巴真 @FlinkChina Meetup



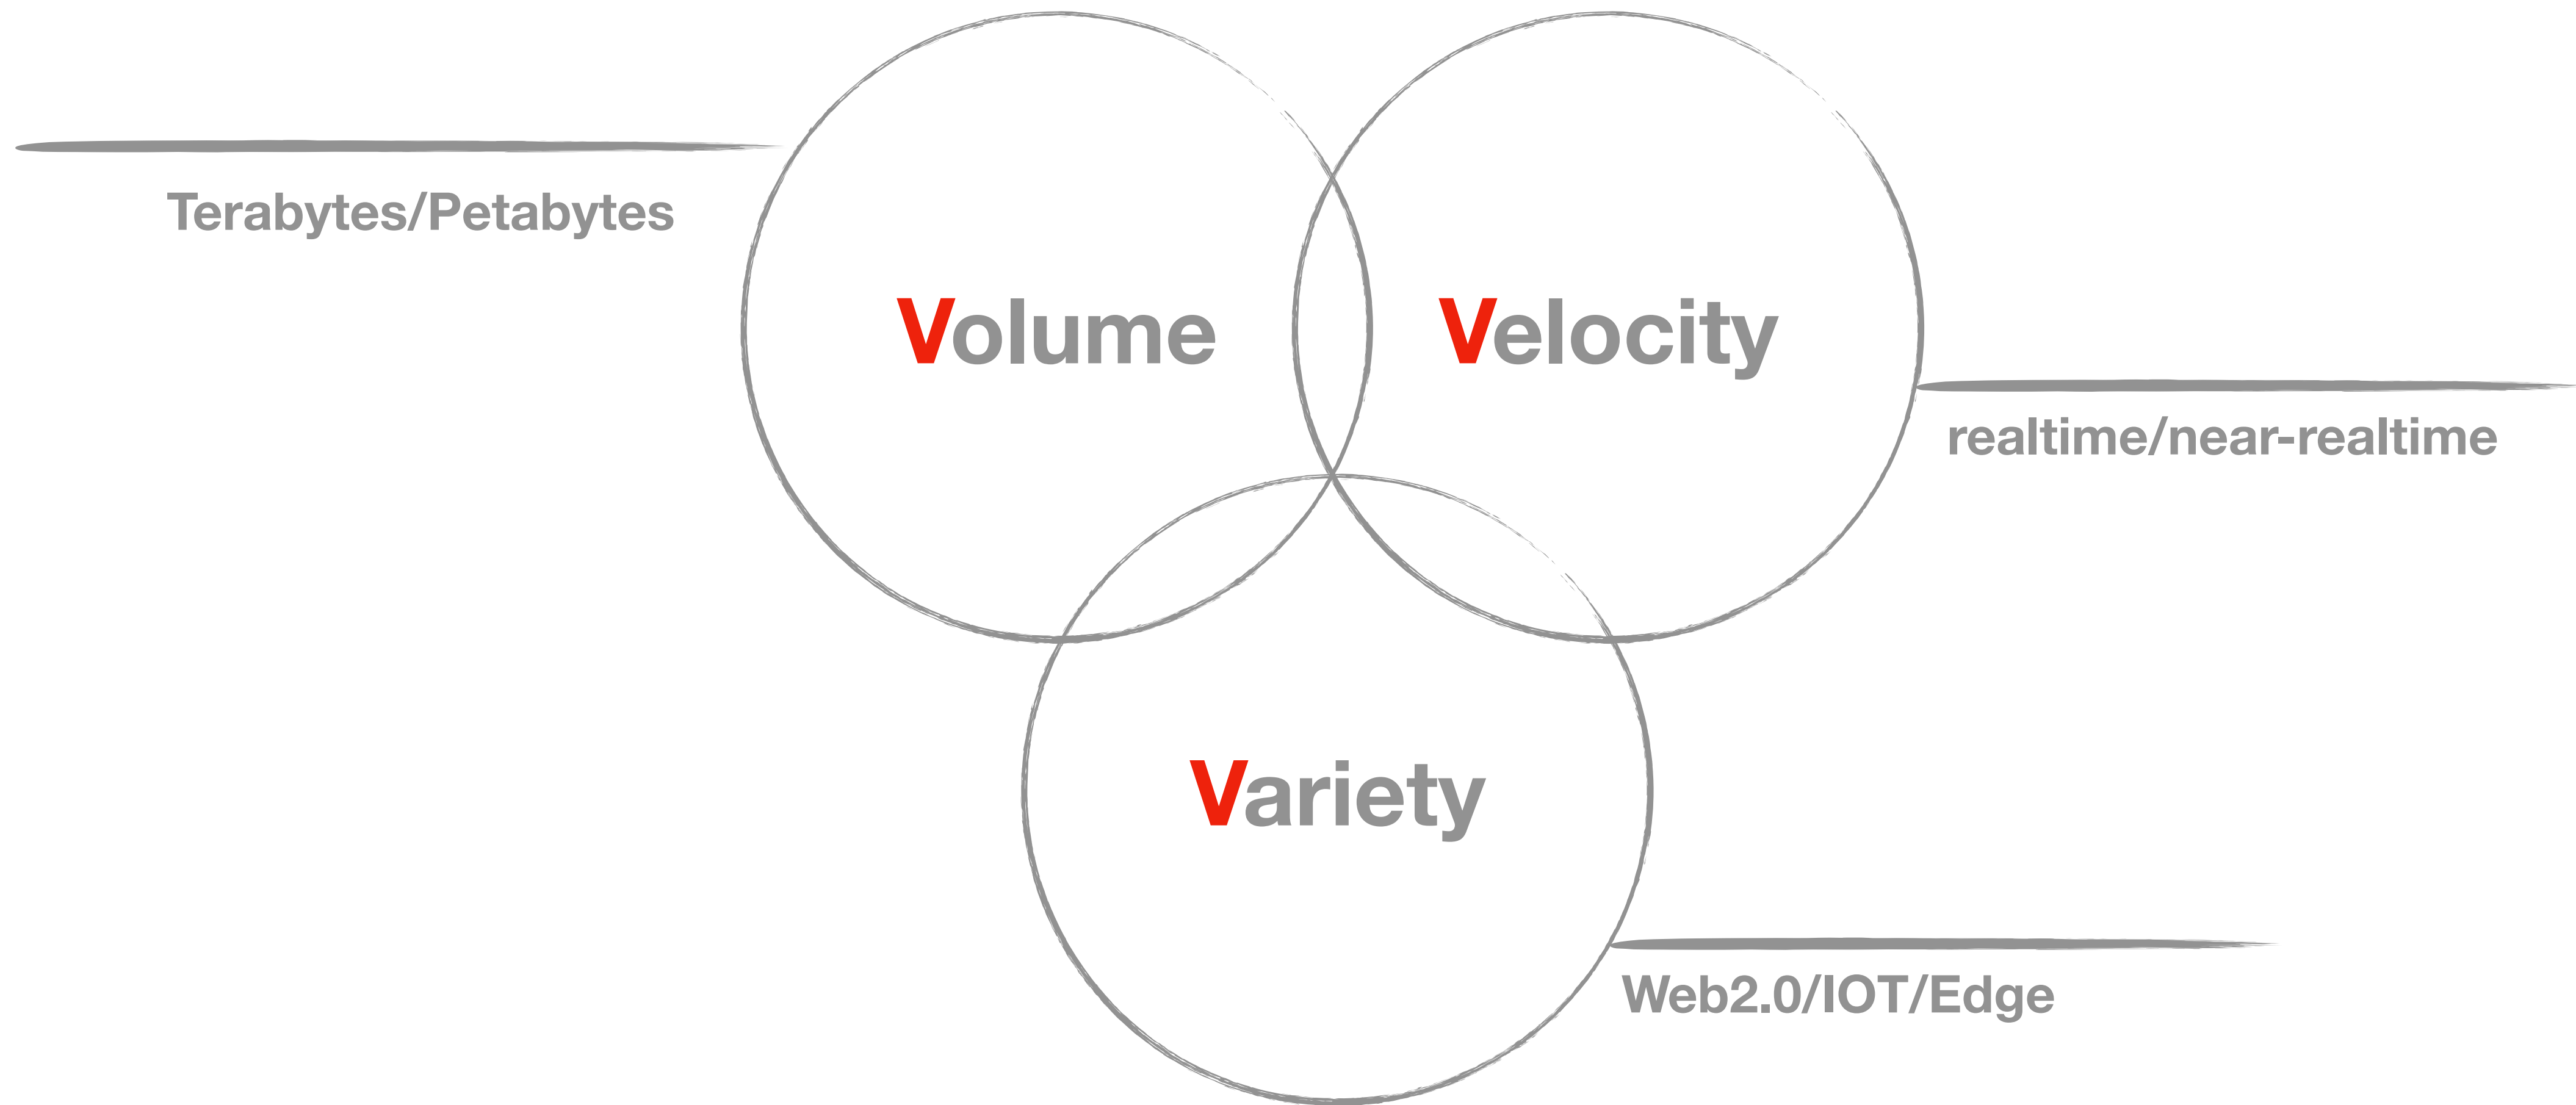
# Agenda

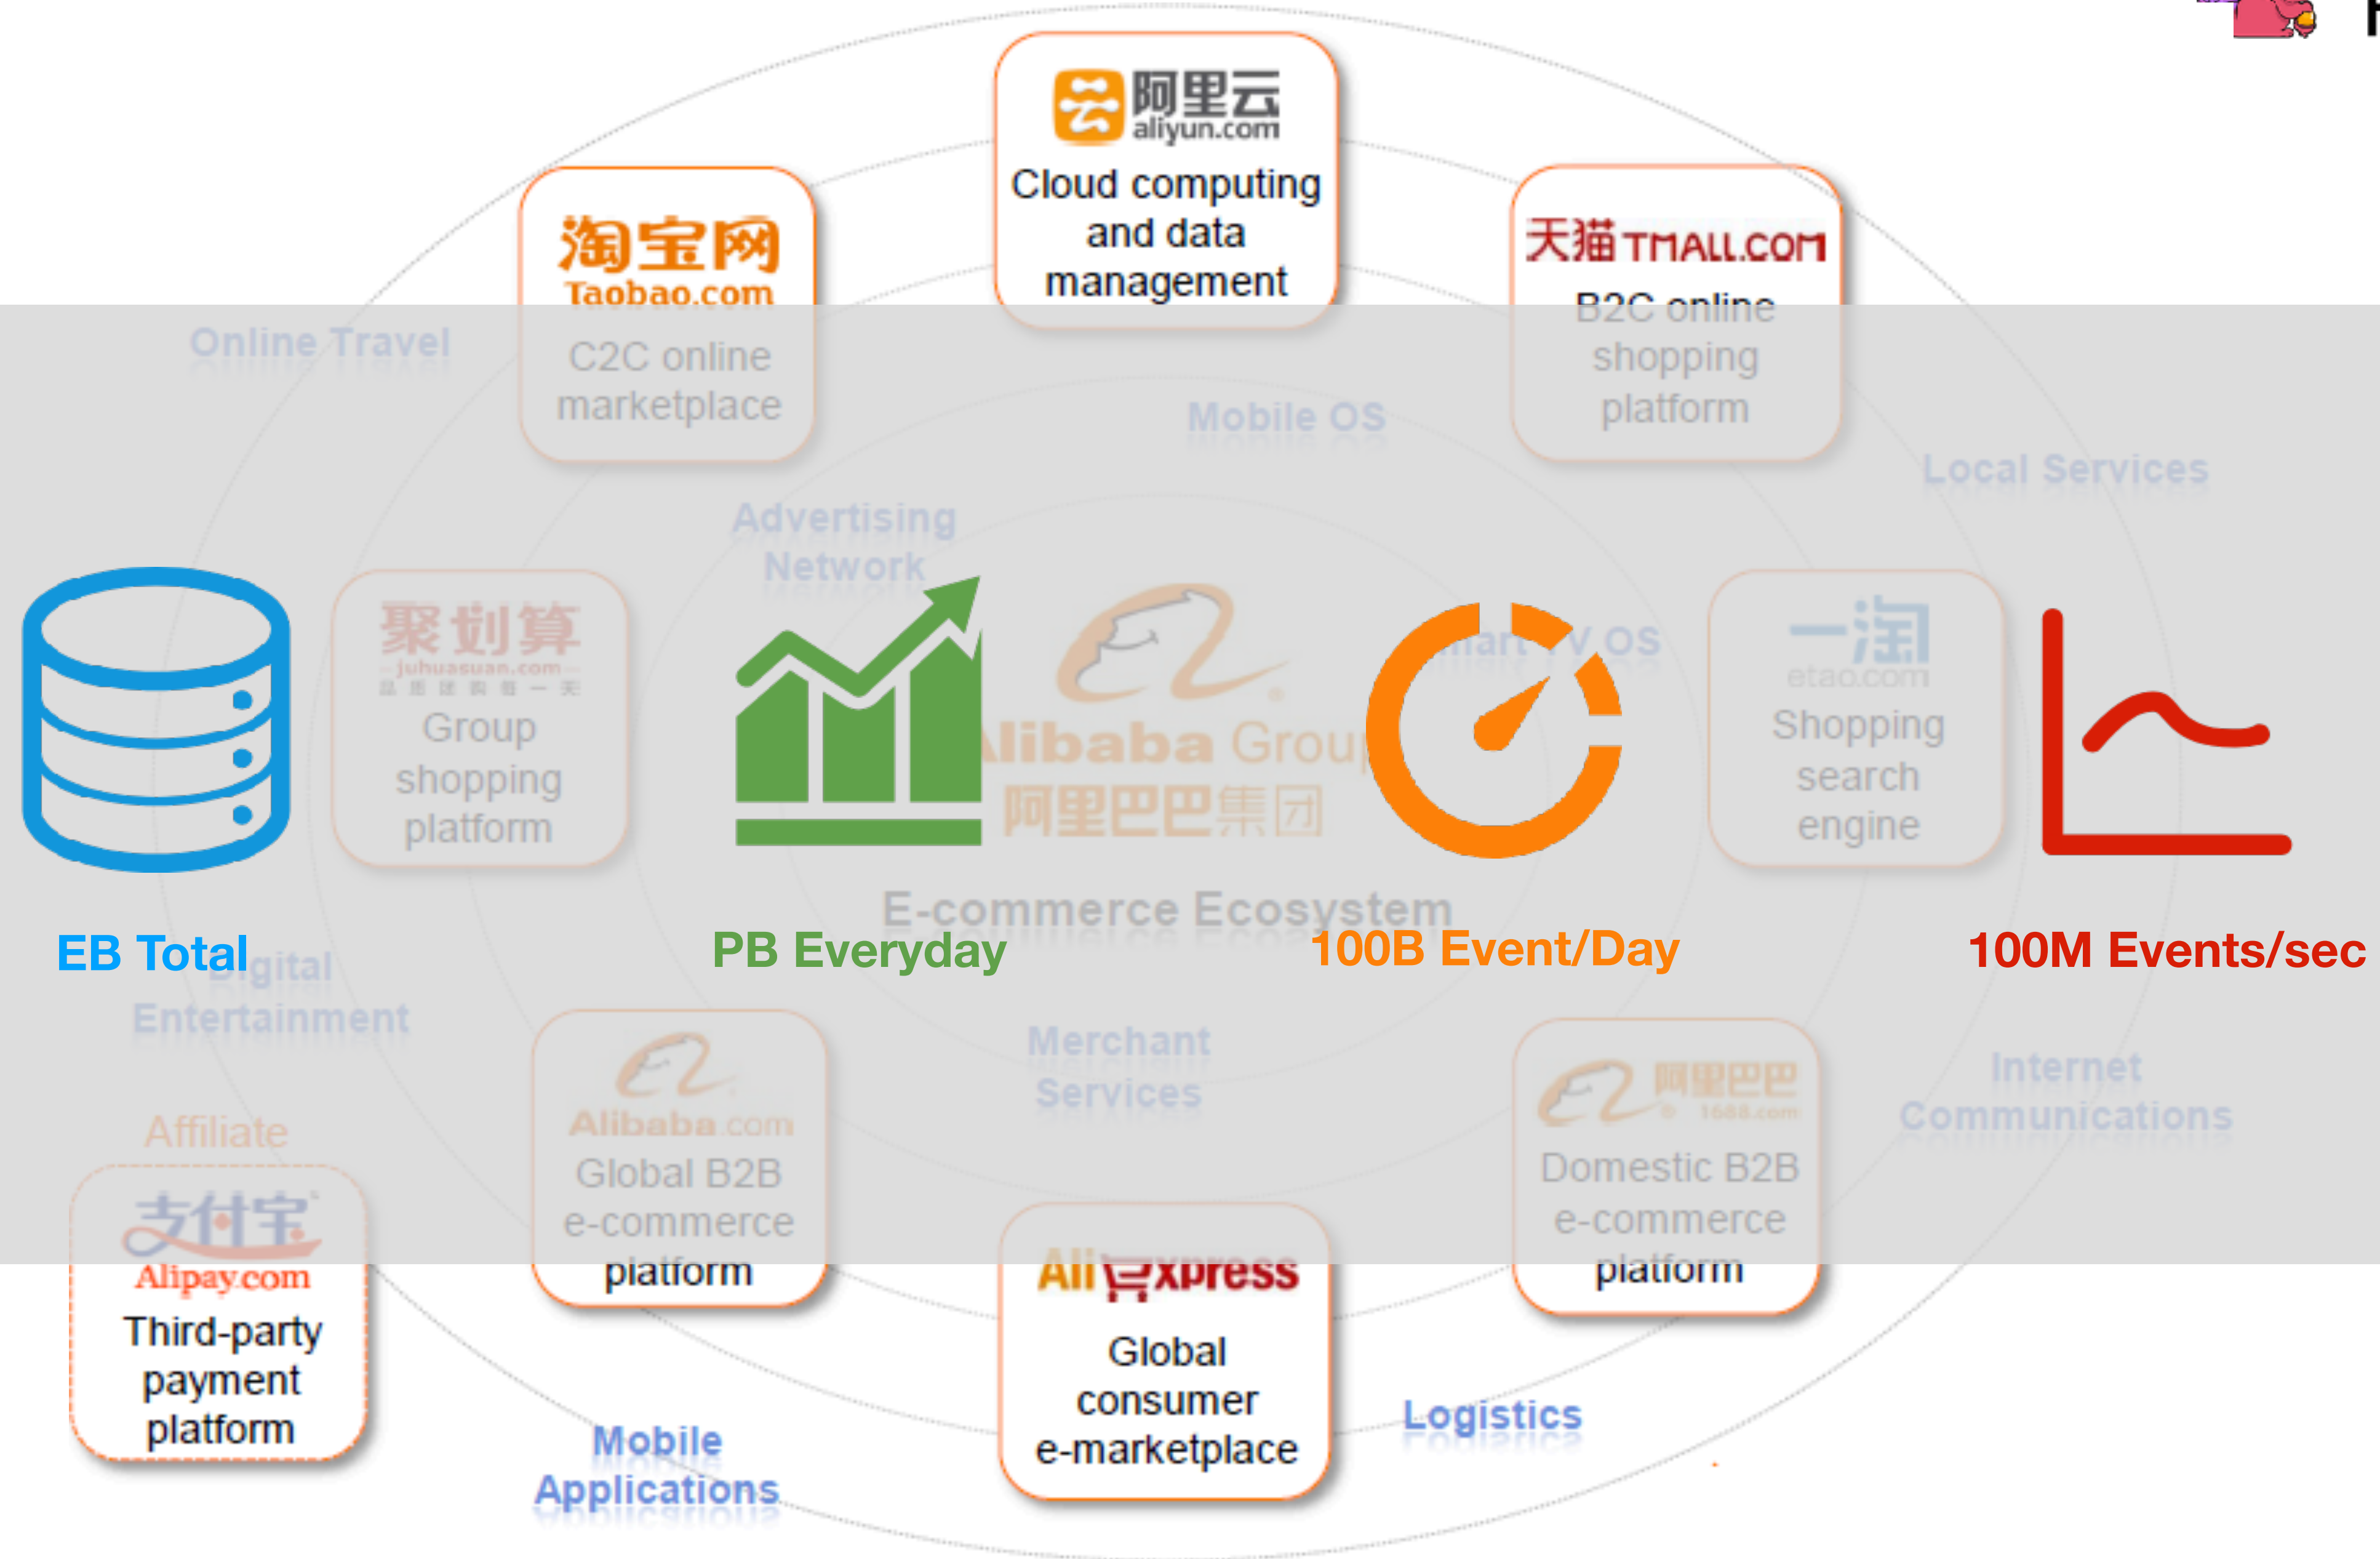
数据趋势

业界案例

阿里思考

Flink@阿里







**AWS S3**



**阿里云OSS**

**Semi-Structured Unstructured**



472M r/s



Sub-second  
latency



Exactly-once



Highly Available

24:00:00 2016 11.11 Global Shopping  
Festival total GMV is RMB 120.7 billion.



计算碎片化



模型多样化



BIG DATA LANDSCAPE 2017

INFRASTRUCTURE

**HADOOP ON-PREMISE**  
cloudera Hortonworks  
MAPR Pivotal  
IBM InfoSphere  
bluedata jethro

**HADOOP IN THE CLOUD**  
amazon Microsoft Azure  
Google Cloud Platform  
IBM InfoSphere  
altiscale  
CAZENA CenturyLink

**STREAMING**  
amazon databricks  
confluent stream  
GridGain METAMARKETS  
dataArtisans  
DATATORRENT

ANALYTICS

**DATA ANALYST PLATFORMS**  
Microsoft pentaho alteryx  
guavus AYASDI  
WATTIVO Detameer Quid  
ClearStory OrigamiLogic Interano  
Bottlenose ARIMO ENDOR MODE

**DATA SCIENCE PLATFORMS**  
IBM KNIME data iku  
DOMINO that rapidminer  
CONTINUUM Alpine  
ALGORITHMIA Anqoss

APPLICATIONS - ENTERPRISE

**SALES**  
CHORUS INSIDESALES.COM  
conversica clari AVISO TACT  
fusemachines TRAPP

**MARKETING - B2B**  
RADIUS App Annie  
EVERESTING Lattice  
infer HENTIGO  
Genseo tubular

**MARKETING - B2C**  
Zeta bloomreach  
blueyonder persado  
Olahuna ACTIONIO  
SALTHRU BLUECORE  
QUANTIFIND  
Amplero

**CUSTOMER SERVICE**  
MEDALLIA zendesk  
CLEARBRIDGE Gainsight  
CLICKFOX NGDATA  
DigitalGenius  
AUTOMAT  
magal

**NOSQL DATABASES**  
Google Cloud Platform  
ORACLE amazon dynamo  
Microsoft Azure MarkLogic  
mongoDB DATASTAX  
KEROSPIKE Couchbase  
redislabs Aflawless

**NEWSQL DATABASES**  
SAP Clustrix  
Pivotal nuodb  
Cockroach Labs  
memsql splice  
MongoDB VOLTDB  
clouddata Tachyon  
druid paradigm4

**GRAPH DBS**  
neo4j  
IBM ORACLE

**MPP DBS**  
TERADATA  
VERTICA  
Microsoft Azure  
Pivotal  
snowflake  
Infoworks

**CLOUD EDW**  
amazon  
Google Cloud Platform  
Microsoft Azure  
Pivotal  
snowflake  
Infoworks

**BI PLATFORMS**  
Microsoft  
amazon  
Google Cloud Platform  
Wave Analytics  
Microsoft Dynamics  
GoodData

**VISUALIZATION**  
Tableau  
Google Cloud Platform  
Qlik Looker  
COLONIS  
CHARTIO

**VERTICAL ANALYTICS**  
PREDIX  
Cylo CAPE  
UPTAKE  
TACHYUS  
Abikam

**STATISTICAL COMPUTING**  
sas  
SPSS  
MATLAB

**DATA SERVICES**  
Palantir  
CASA  
kaggle  
FF

**HUMAN CAPITAL**  
entelo  
hiQ  
Rookit  
Stella  
pymetrics

**LEGAL**  
RAVEL  
Seal  
Everlaw  
Brevia

**FINANCE**  
anaplan  
uoro  
tidemark  
LUSHERA

**ENTERPRISE PRODUCTIVITY**  
slack  
facebook  
ORACLE  
butterfly  
KASISTO

**BACK OFFICE AUTOMATION**  
Captricity  
AppZen

**SECURITY**  
TANUM  
CYLANCE  
StackPath  
DARKTRACE  
Vectra  
ANOMALI  
sift science  
Recorded Future  
Securix

**DATA TRANSFORMATION**  
talend pentaho  
alteryx TRIFACTA  
tome Paxata  
StreamSets

**DATA INTEGRATION**  
Informatica  
Marten  
snoplogic  
segment  
podium data  
Import Skitch

**DATA GOVERNANCE**  
informatica  
IBM skyhigh  
collibra  
Alation

**MGMT / MONITORING**  
amazon New Relic  
APPROXIMATIONS octio  
WAVEFRONT  
splunk  
Yoocon Numenty

**MACHINE LEARNING**  
Google Cloud Platform  
H2O  
DataRobot  
VIZENI SKYTREE

**HORIZONTAL AI**  
IBM Watson Cortana  
servient  
Voyager  
Affective  
OSARO

**SPEECH & NLP**  
Google Cloud Platform  
NarrativeScience  
semantic machines  
IDIBON  
snips

**ADVERTISING**  
AppNexus  
criteo  
Integral  
MOAT  
drawbridge  
Liveness  
DataXu gumgum  
Optix

**EDUCATION**  
NEWTON  
Clever  
Declara  
kidaptive  
KNOWLEDGE

**GOVERNMENT**  
Socrata  
OPENGOV  
mark43  
enigma  
FacalNote  
OpenText

**FINANCE - LENDING**  
OnDeck  
affirm  
Kreditech  
INSIKT  
MoneyLion  
TrueAccord  
cignifi aire

**FINANCE - INVESTING**  
Dataminr  
KENSHC  
Quantopian  
NUMERA  
iSENTIUM  
ALGOZ

**REAL ESTATE**  
Opendoor  
VTS  
CREDEN  
reonomy  
COMPSTAK

**INSURANCE**  
Thomson  
Lemonade  
CYENCE  
Shift Technology  
Tractable

APPLICATIONS - INDUSTRY

**STORAGE**  
amazon  
Google Cloud Platform  
Microsoft Azure  
ALLUXIO  
ramblockstorage  
Datto COHO  
paragon

**CLUSTER SERVICES**  
amazon  
docker  
MESOSPHERE  
Core OS

**APP DEV**  
Lightbend  
rainforest  
CSK

**CROWDSOURCING**  
amazon  
Upwork  
Workfusion

**HARDWARE**  
Google TPU ARM  
nvidia Graphcore  
MYTHIC  
nvidia  
Movidius SCORTEX

**SEARCH**  
elastic  
ORACLE  
Thoughtworks  
Lucidworks  
swiftype  
MAANA  
alphasense  
Searchlink

**LOG ANALYTICS**  
splunk  
sumologic  
loggly  
kibana  
logz.io

**SOCIAL ANALYTICS**  
Hootsuite  
NETBASE  
DATA SIFT  
synthesio  
smallreach  
bitly predata

**WEB / MOBILE / COMMERCE ANALYTICS**  
Google Analytics  
mixpanel  
sumo.i  
retention  
granify  
custora

**HEALTHCARE**  
FLATIRON  
HealthTop  
Gingerio  
COTA  
iACure  
Imagin  
METABOTA  
Glow  
zebra  
ovig  
enric  
Quantu

**LIFE SCIENCES**  
color  
oeo zymogen  
BenevolentAI  
ZEPHYRUS  
Citrine  
twoAR

**TRANSPORTATION**  
UBER  
TESLA  
CLEANPATH  
drive.ai  
OPTIMUS  
iHexar  
ZODIAC  
nauto  
pilot.ai  
OTTO

**AGRICULTURE**  
FARMERS  
FarmLogs  
mavrX  
prosporo

**COMMERCE**  
Instacart  
RECALLBOX  
HowGood

**OTHER**  
eHarmony  
spherix  
KNOX  
santry

CROSS-INFRASTRUCTURE/ANALYTICS  
amazon Google Cloud Platform Microsoft IBM SAP Hewlett Packard Enterprise sas data vmware TIBCO TERADATA ORACLE NetApp

OPEN SOURCE

**FRAMEWORK**  
HADOOP YARN TEZ  
Flink MESOS  
Spark COOP

**QUERY / DATA FLOW**  
Spark SQL presto  
SLAMDATA  
DWH

**DATA ACCESS**  
cassandra mongoDB  
COUCHDB  
FARIKEL  
Spanner

**COORDINATION**  
talend  
Apache Zookeeper  
Apache Ambari

**STREAMING**  
Spark  
Flink  
kafka  
druid  
STORM

**STAT TOOLS**  
ScalaLab  
SciPy

**AI / MACHINE LEARNING / DEEP LEARNING**  
theano  
Caffe  
neon  
FeatureFu  
mllib  
DL4J  
Aerosolve

**SEARCH**  
elasticsearch  
Solr  
Lucene


**LOG ANALYSIS**  
elasticsearch  
kibana  
logstash

**VISUALIZATION**  
BLAKEY  
Rodeo

**COLLABORATION**  
superset  
ANACONDA

**SECURITY**  
Apache Ranger  
KNOX  
santry





数万Dataer



百万Job

---

在阿里，人人都是数据分析师



BI Engineer



AI Engineer

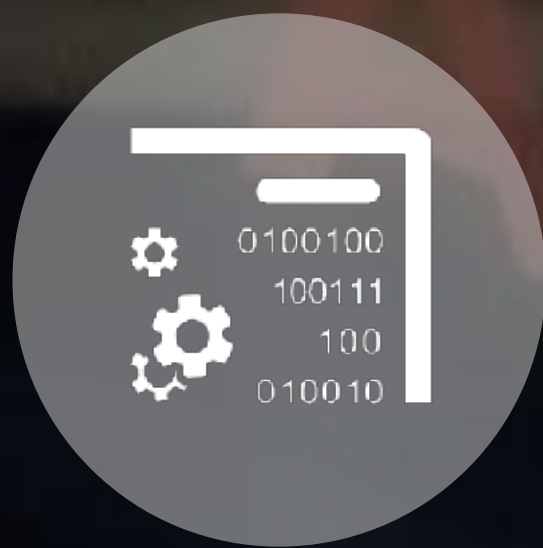
---

在阿里，BI工程师在转型AI工程师

# 阿里在思考



| 计算类型多，但核心模型可列举



| 描述方式多，但核心抽象可列举

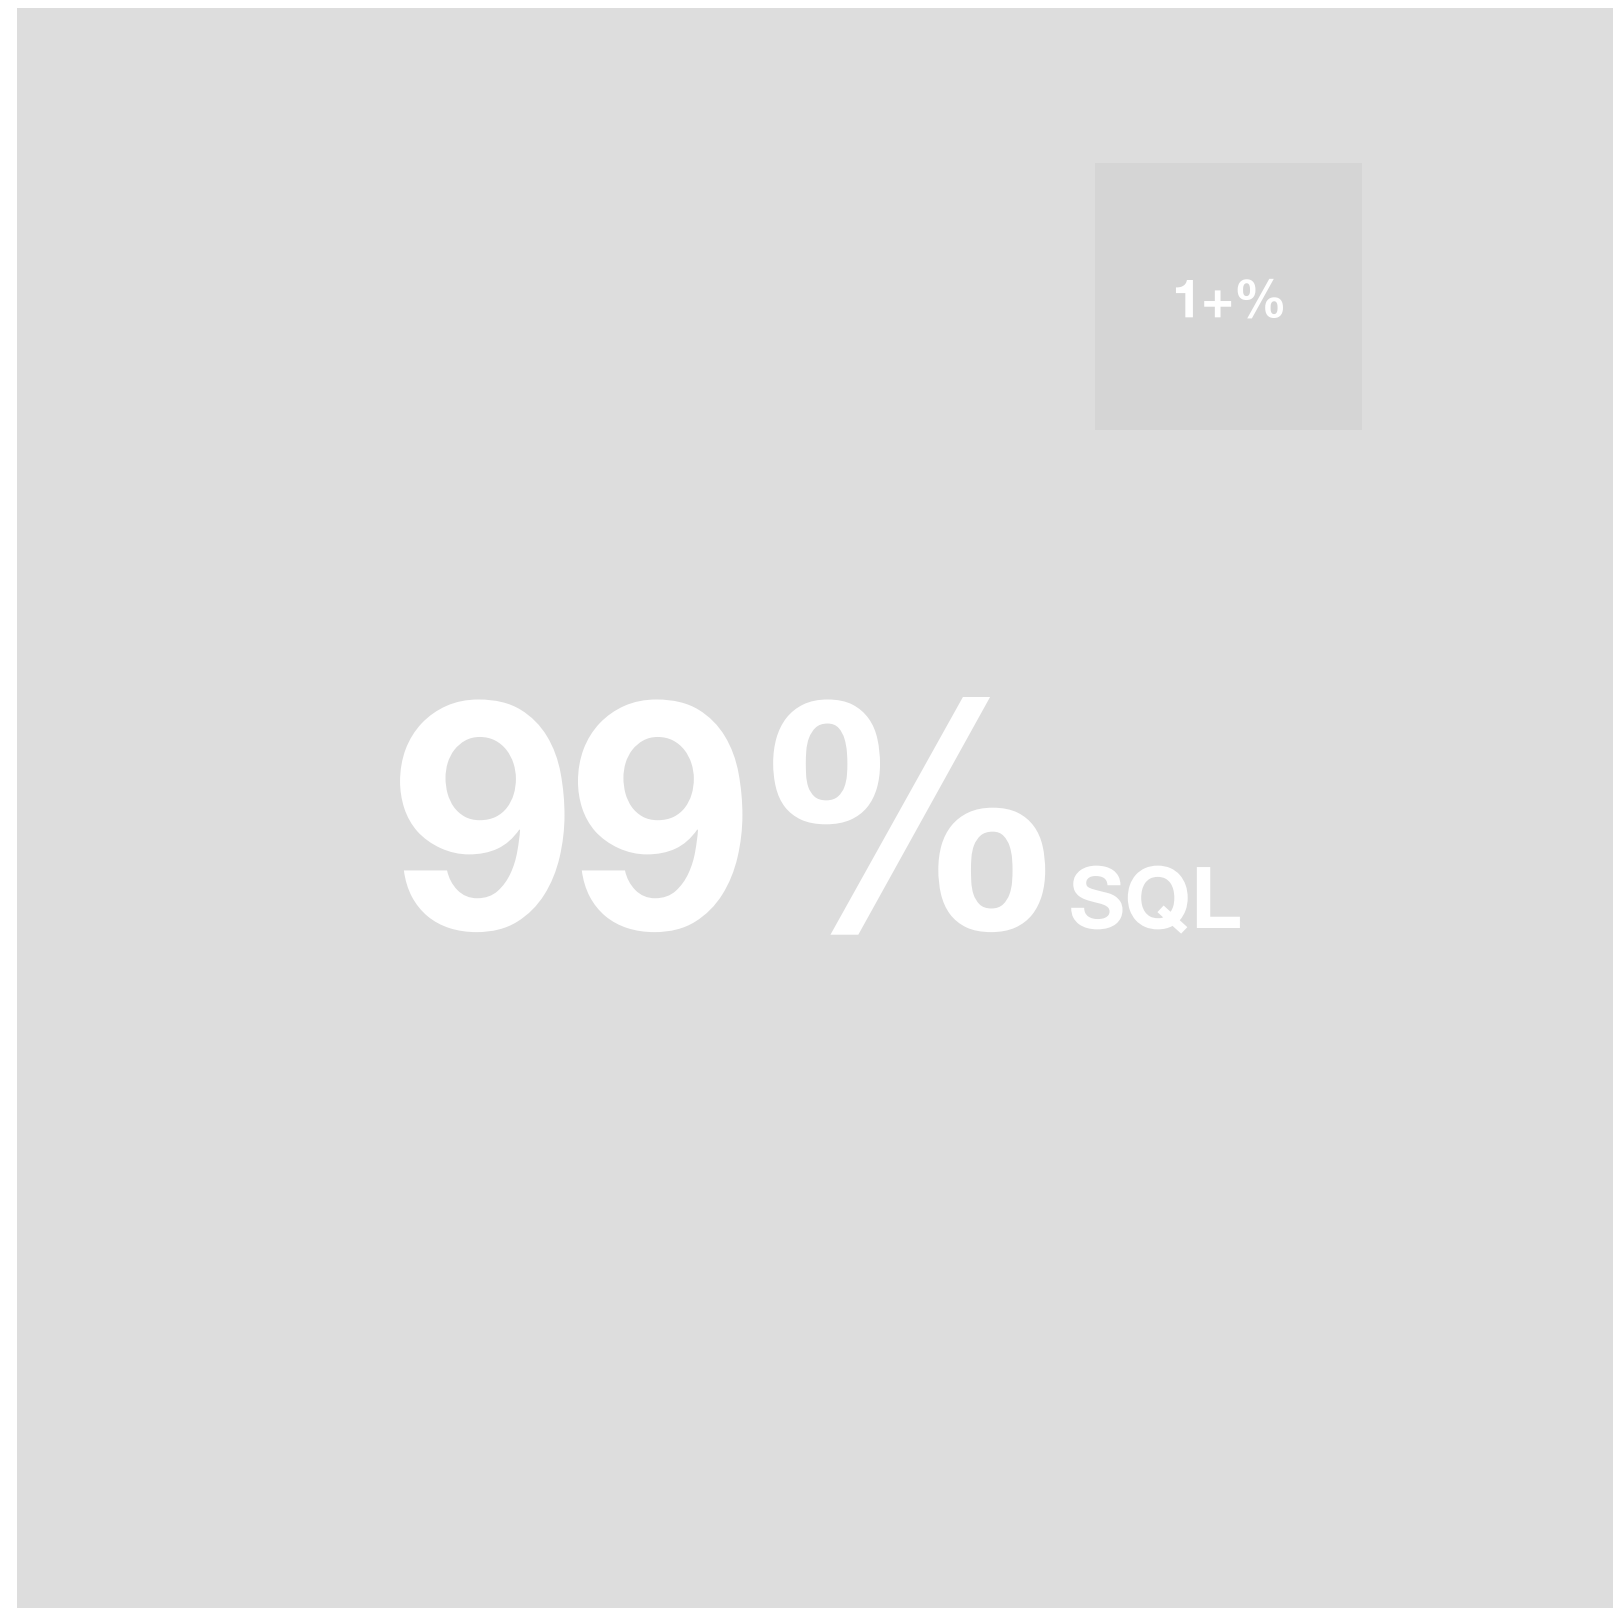


| 人参考数据决策 -> 机器计算数据决策

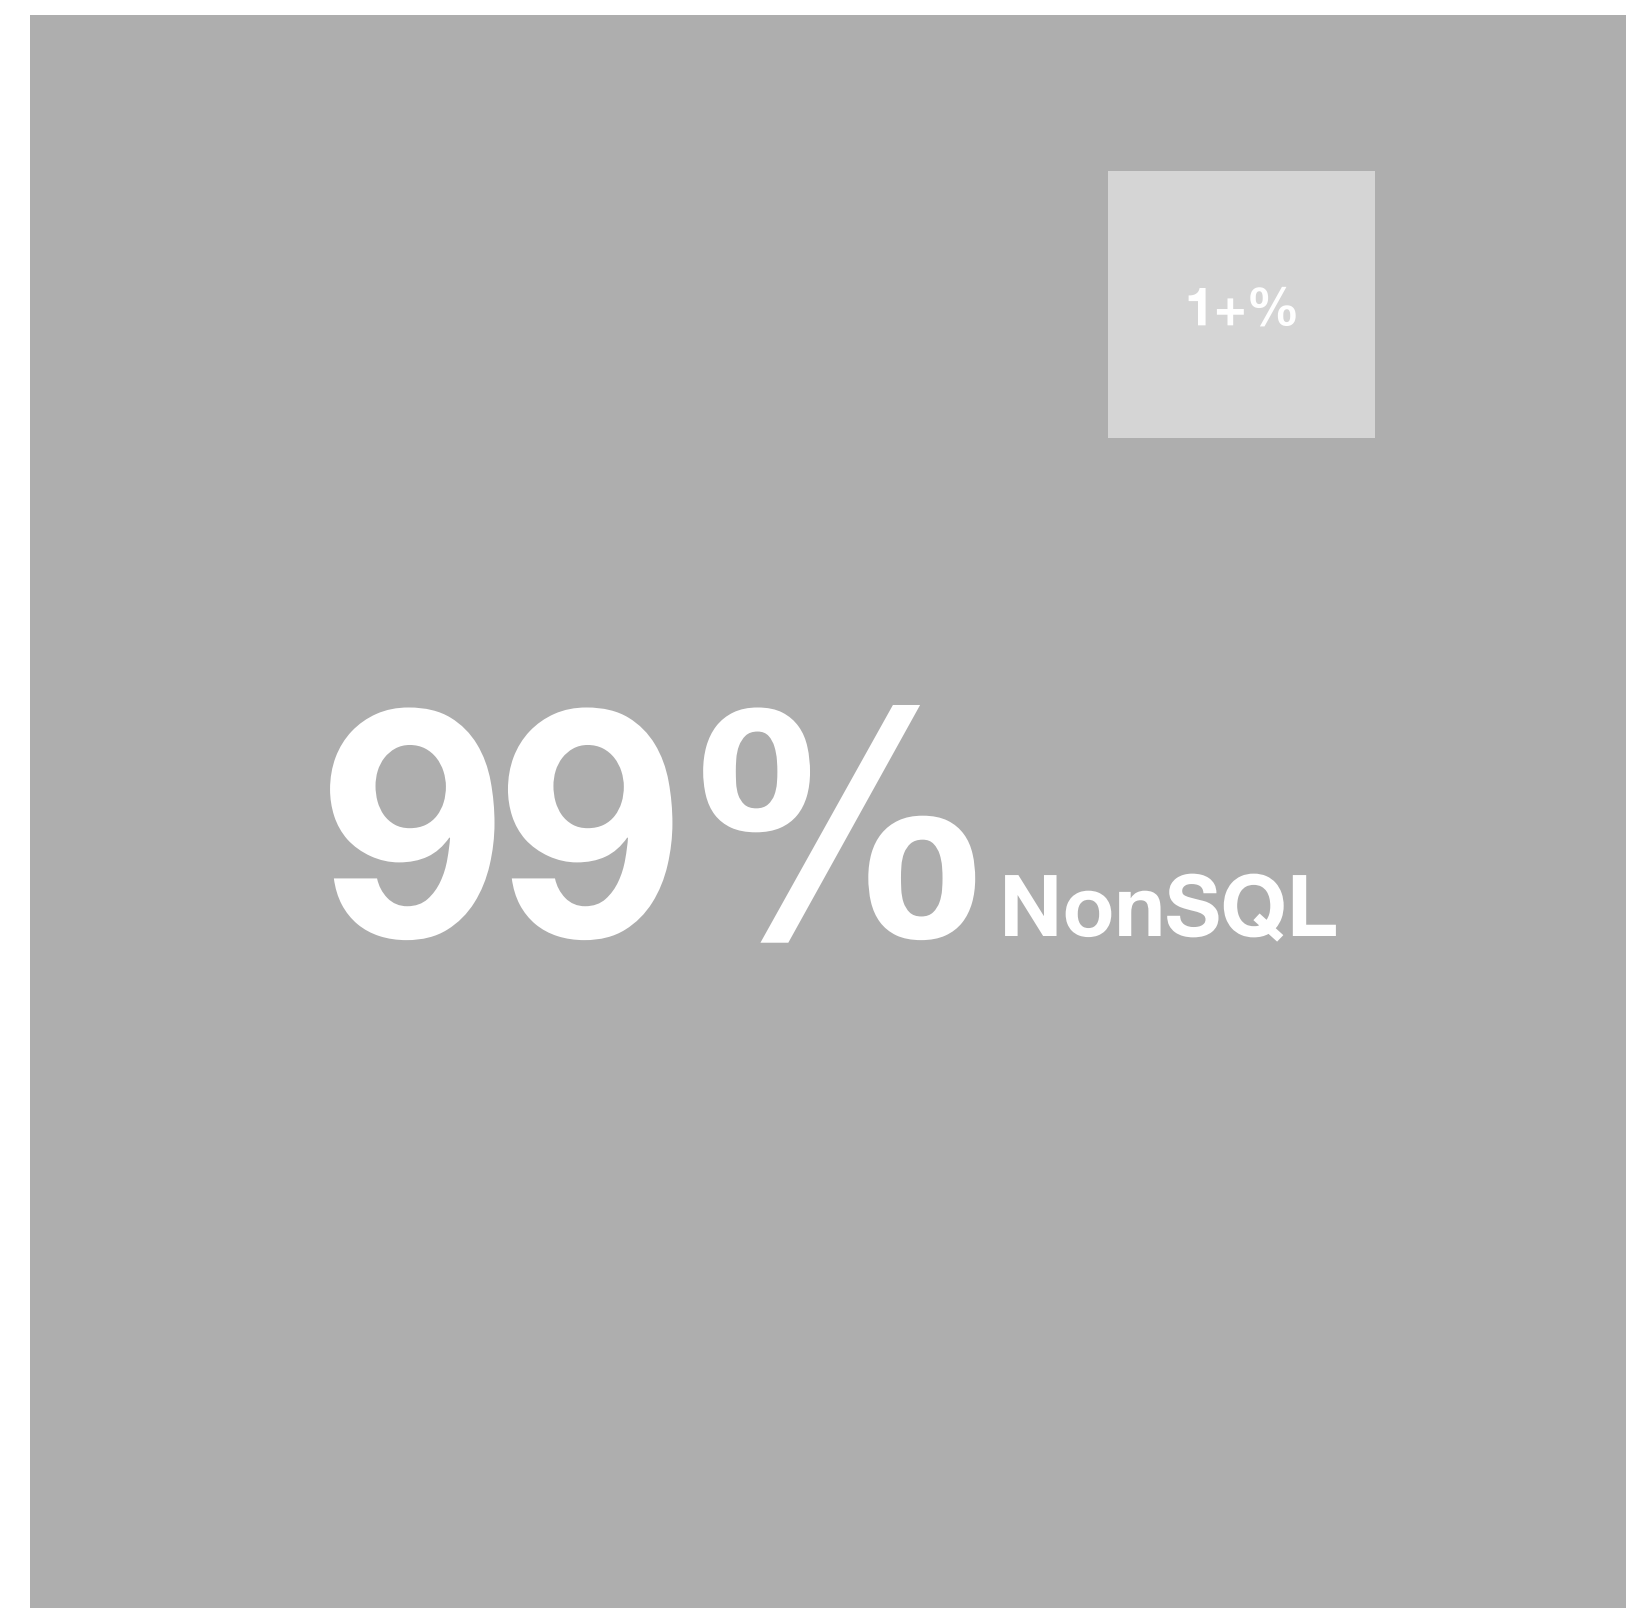


---

计算类型多种多样，但入口业务仅此数种



Batch

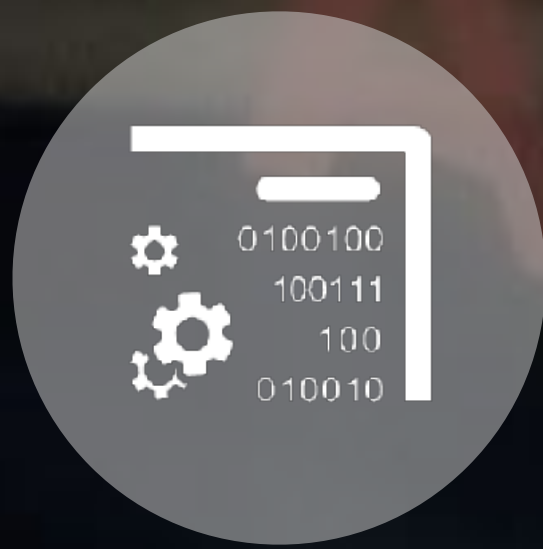


Stream

阿里的思考



| 统一计算引擎: Batch/Stream/OLAP/ML/Graph...



| 统一抽象方式: Unified SQL & API



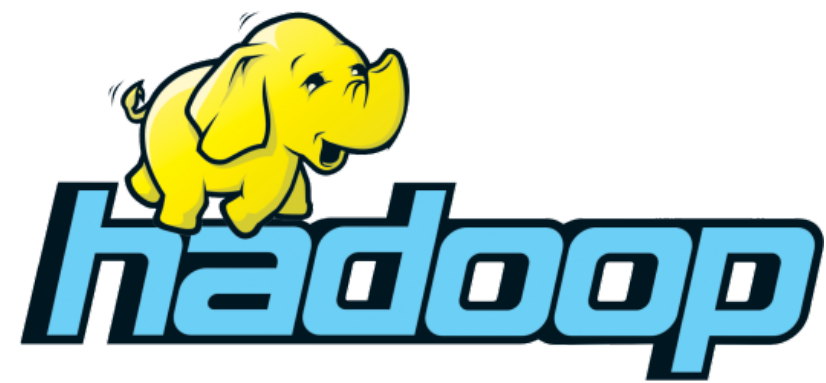
| 统一BI+AI引擎: 数据清洗&数据训练

阿里的思考



如何构建下一代大数据处理引擎？

Alibaba.com



第一代

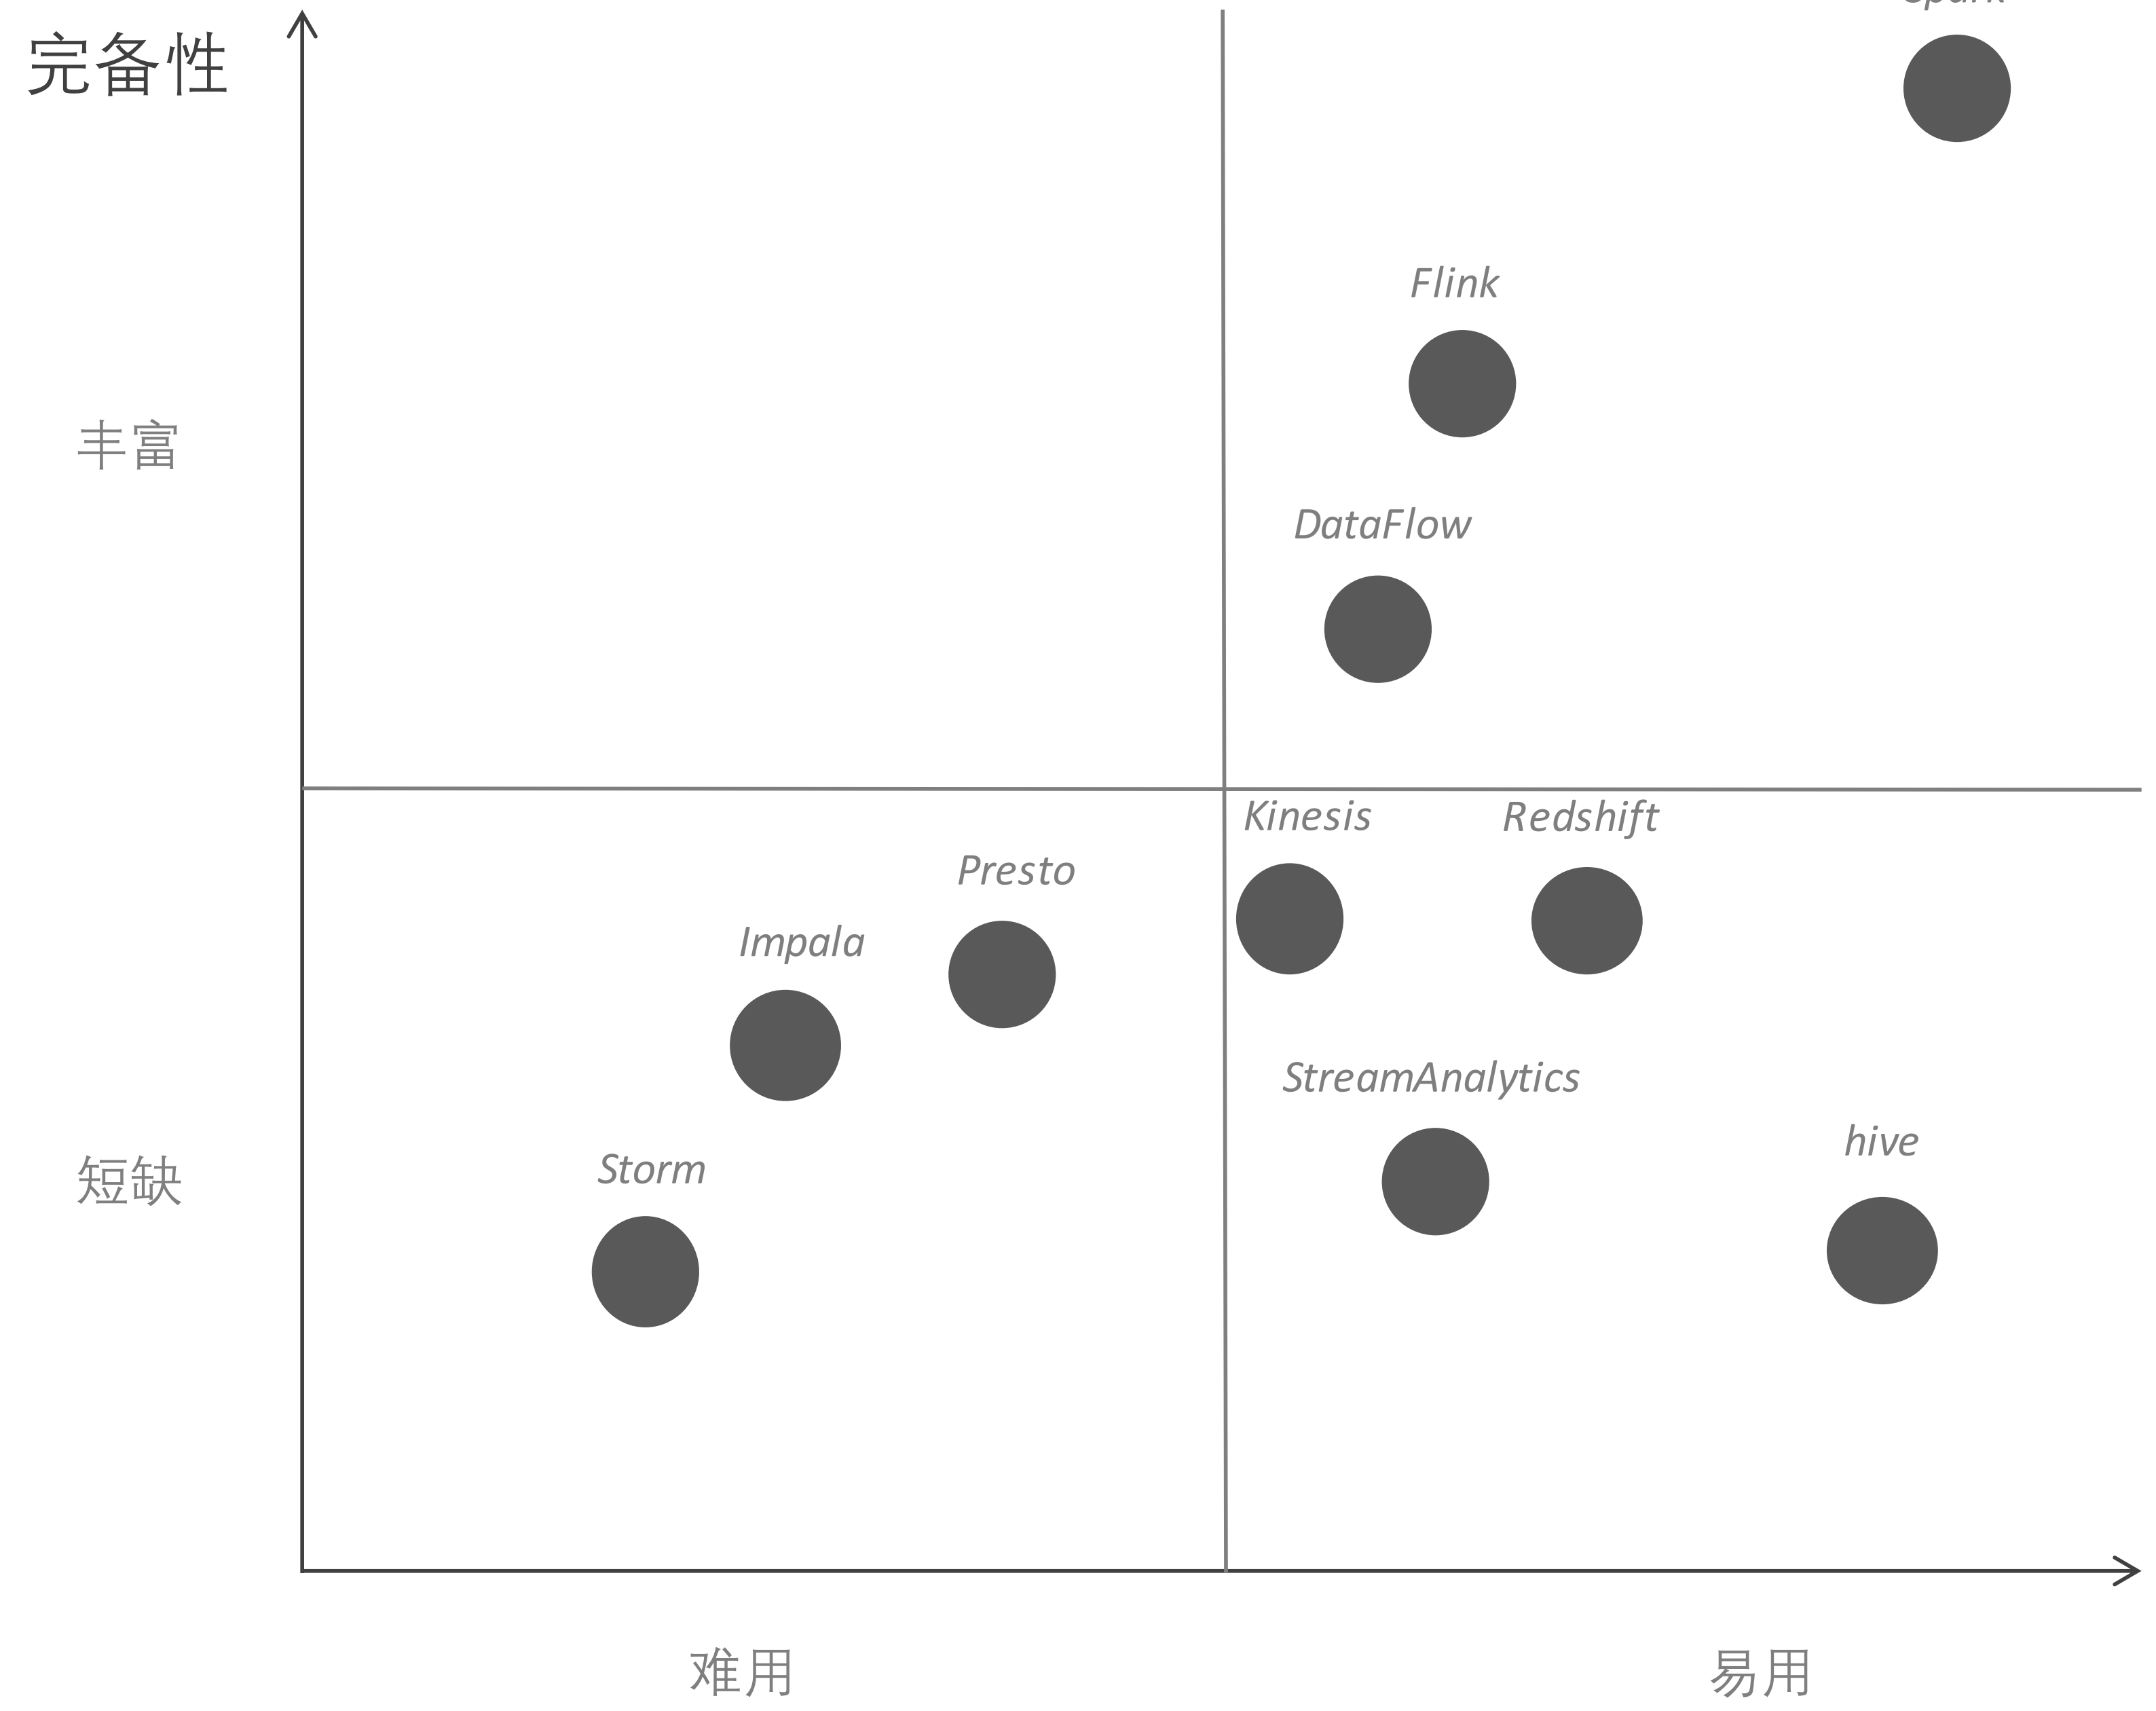


第二代



第三代





### Spark优势明显

功能完备: Batch/Stream/ML  
简单易用: API丰富/文档较丰富/生态对接  
社区活跃: 运营出彩、大数据最活跃社区

### 云产品系列

产品化: 产品化成熟, 易于上手  
稳定性: 稳定性好, 商业托管  
社区小: 产品封闭, 不易形成社区

### 开源软件

社区活跃: 产品开放, 易于形成社区  
稳定性差: Bug多, 不稳定  
产品粗糙: 产品化初级, 上手较难



Spark期望一套软件覆盖主要计算模型，但实际覆盖不完整

**21%**用户认为Spark Streaming在功能(集中在窗口)和时延(亚秒)等比不上Flink，增量流式业务考虑使用Flink

**17%**用户认为Spark ML部分落后，包括提供更多算法、对接TF，部分业务迁移到TF框架运行

**25%**的用户认为当前缺乏好用、内置的上层平台，包括开发界面、工作流调度，用户使用Spark同样需要重新搭建平台系统

稳定性/调优/排错 仍未解决

**31%** 用户吐槽Spark集群不稳定，经常性OOM导致业务产出不稳定

另外，几乎同样客户群体(说明都是深入生产使用Spark用户)都认为Spark作业排错、调优困难，易用性不够

中文资料/社区严重缺乏，未能形成有效组织

**30%**用户吐槽当前文档、案例过少，特别在调优、排错方面，用户往往不知所措  
用户同时认为相关中文资料相比更少，时效性也落后英文社区太多  
中美语言差异导致中国市场更加空白



# Flink: 下一代流式处理系统

——为什么Flink要比Storm/Spark更加优秀？



Apache Flink

+



Alibaba's Improvements

=

**Alibaba Blink**

Alibaba Blink Platform



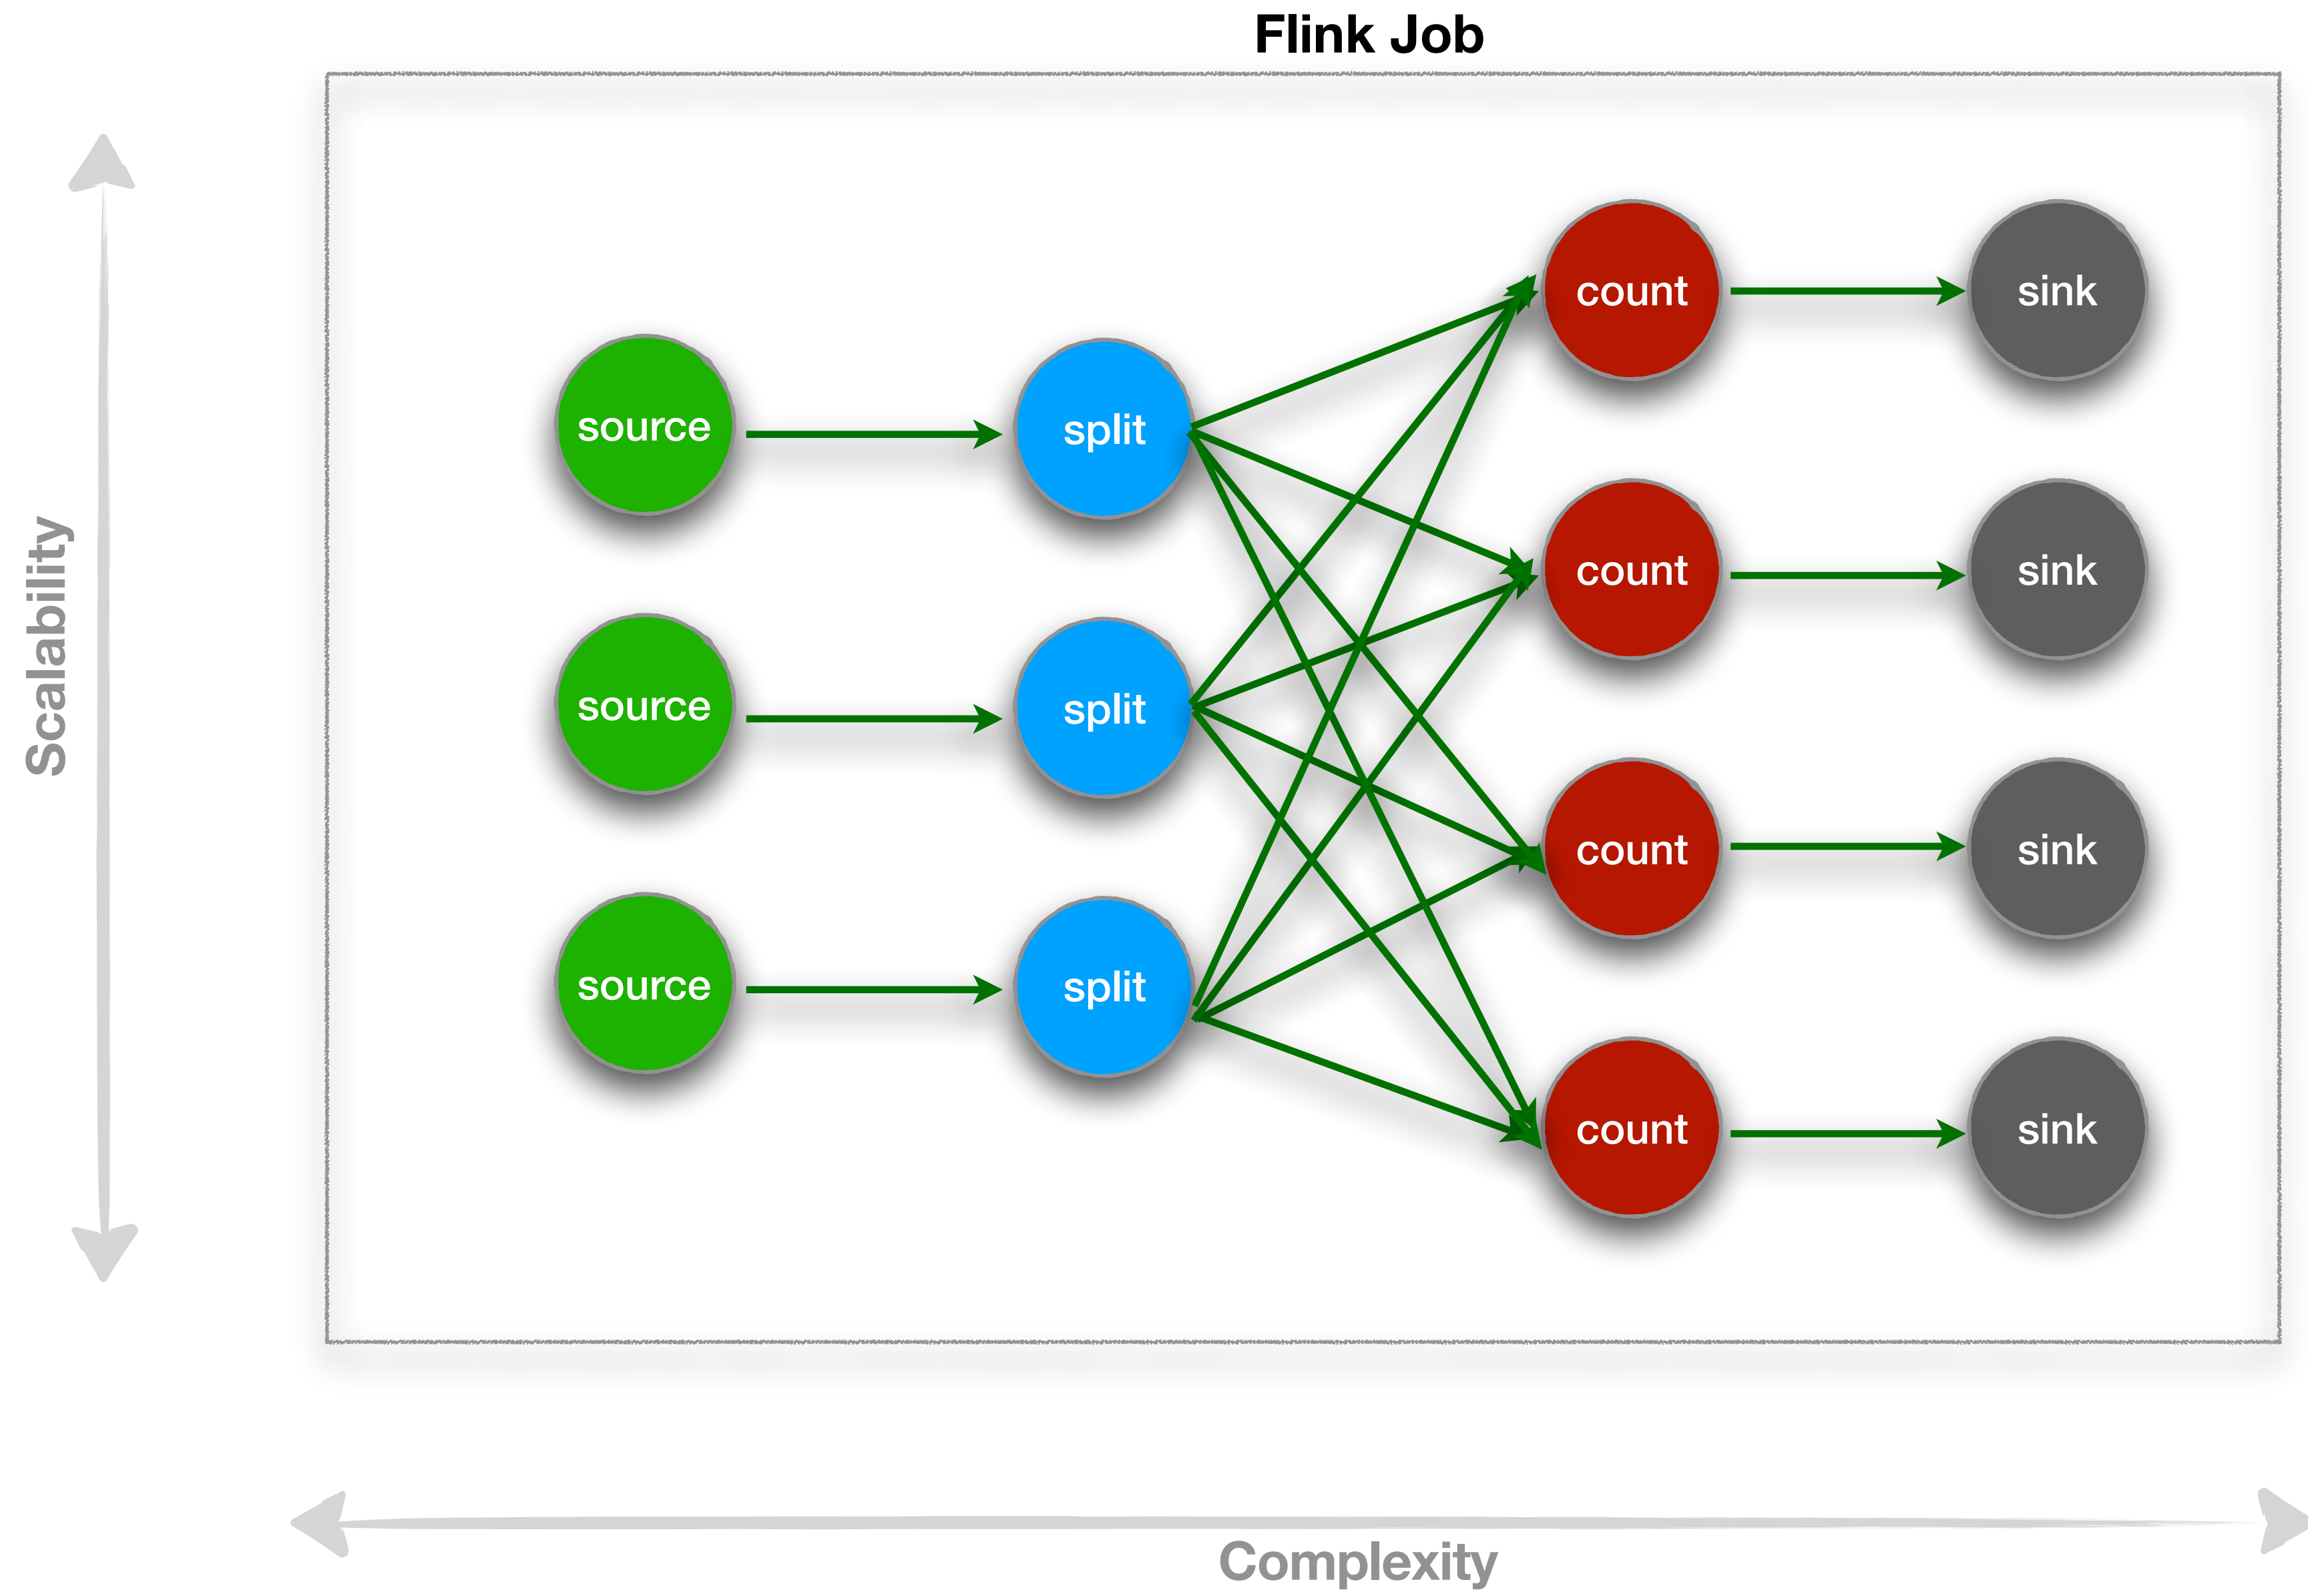
Scalability



Unified



API&SQL





## Scalability&Complex

### **FLIP-6: Deployment and Process Model**

解决大规模部署+云计算隔离问题

### **FLIP-12: Asynchronous IO**

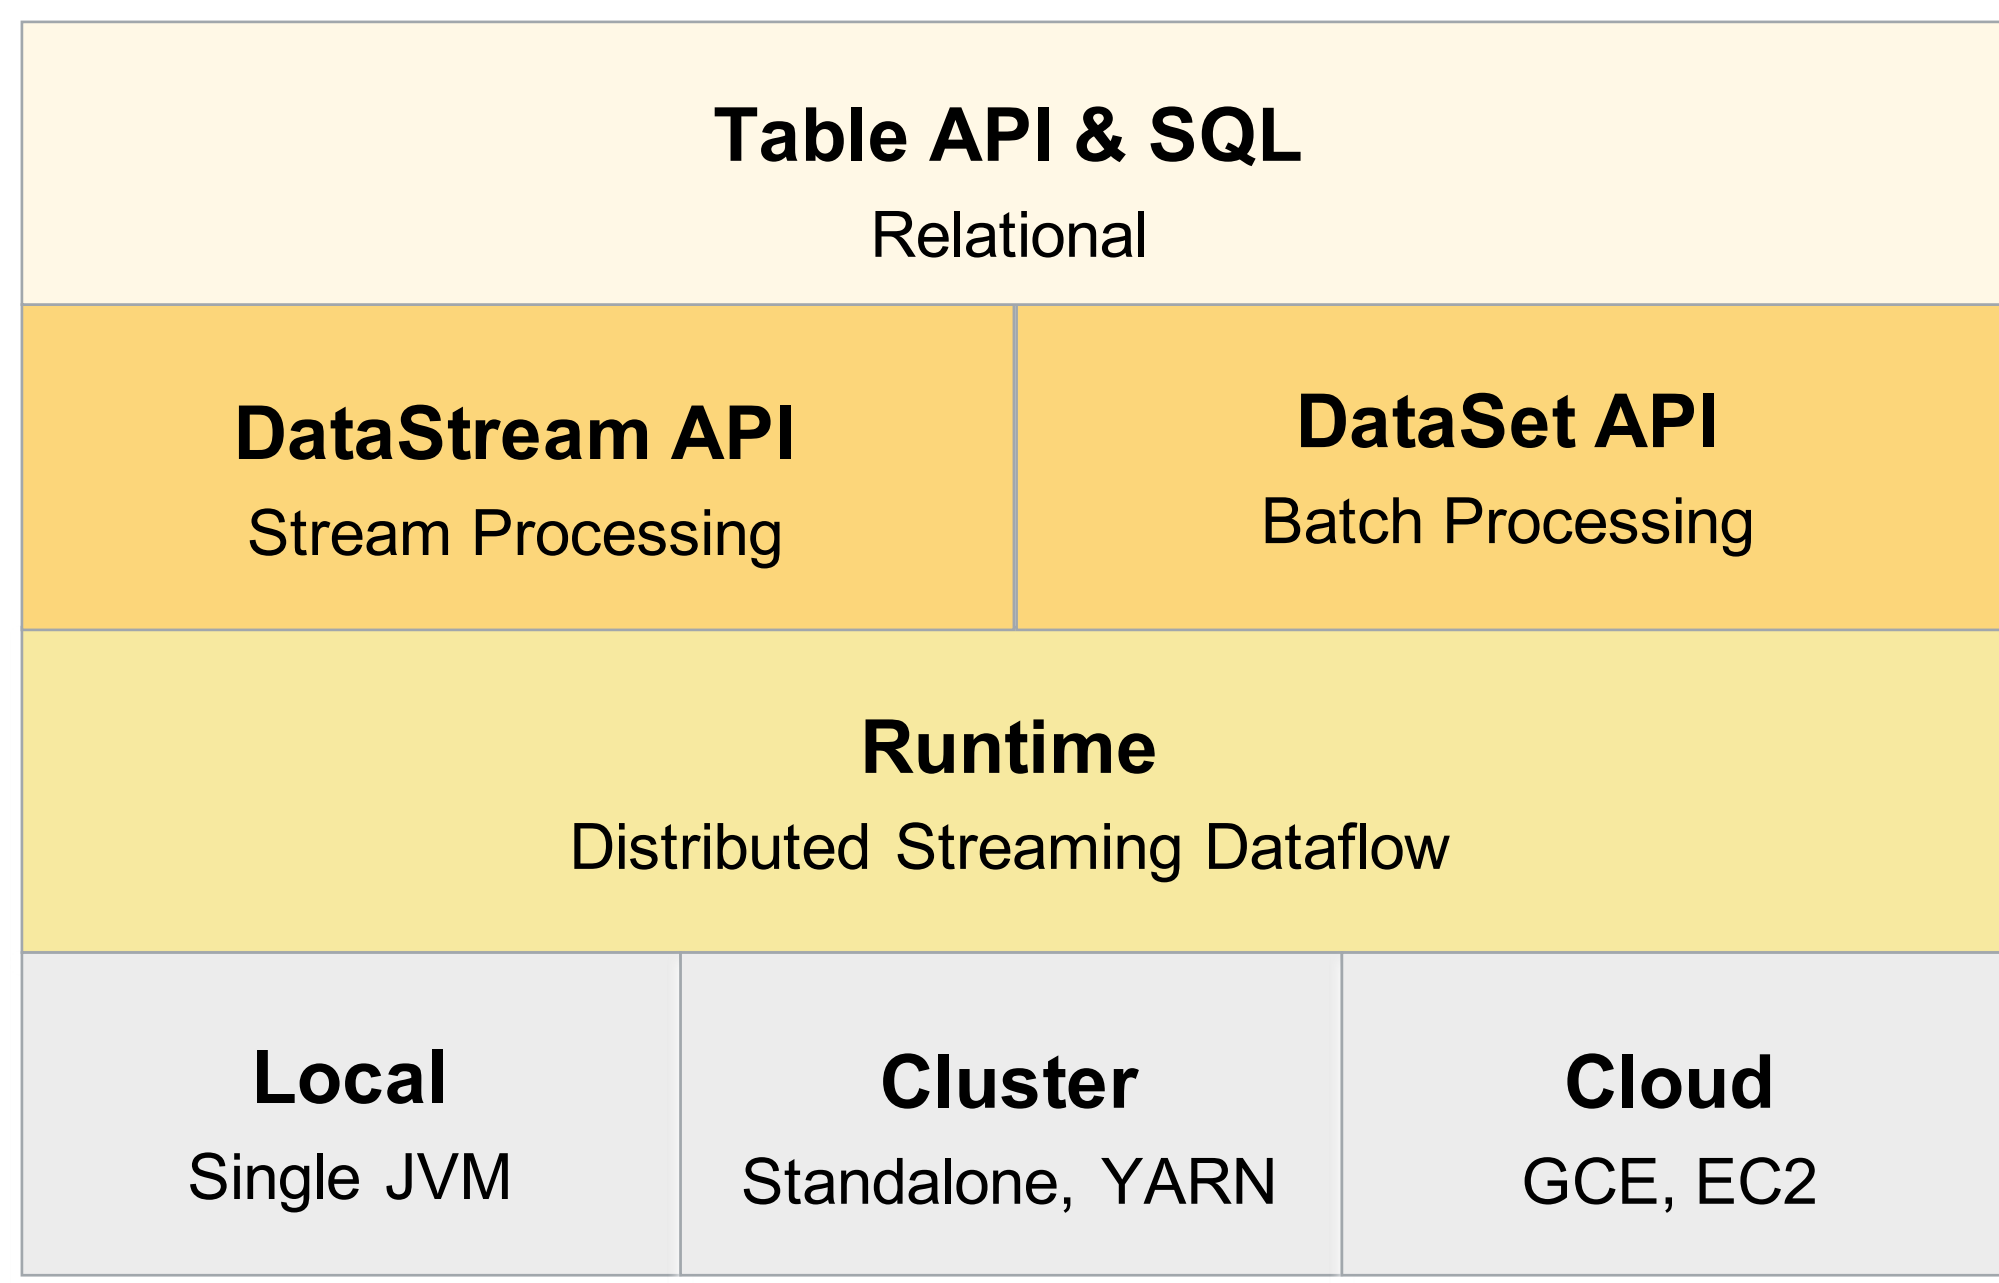
解决流计算时延问题

### **Incremental Checkpoints**

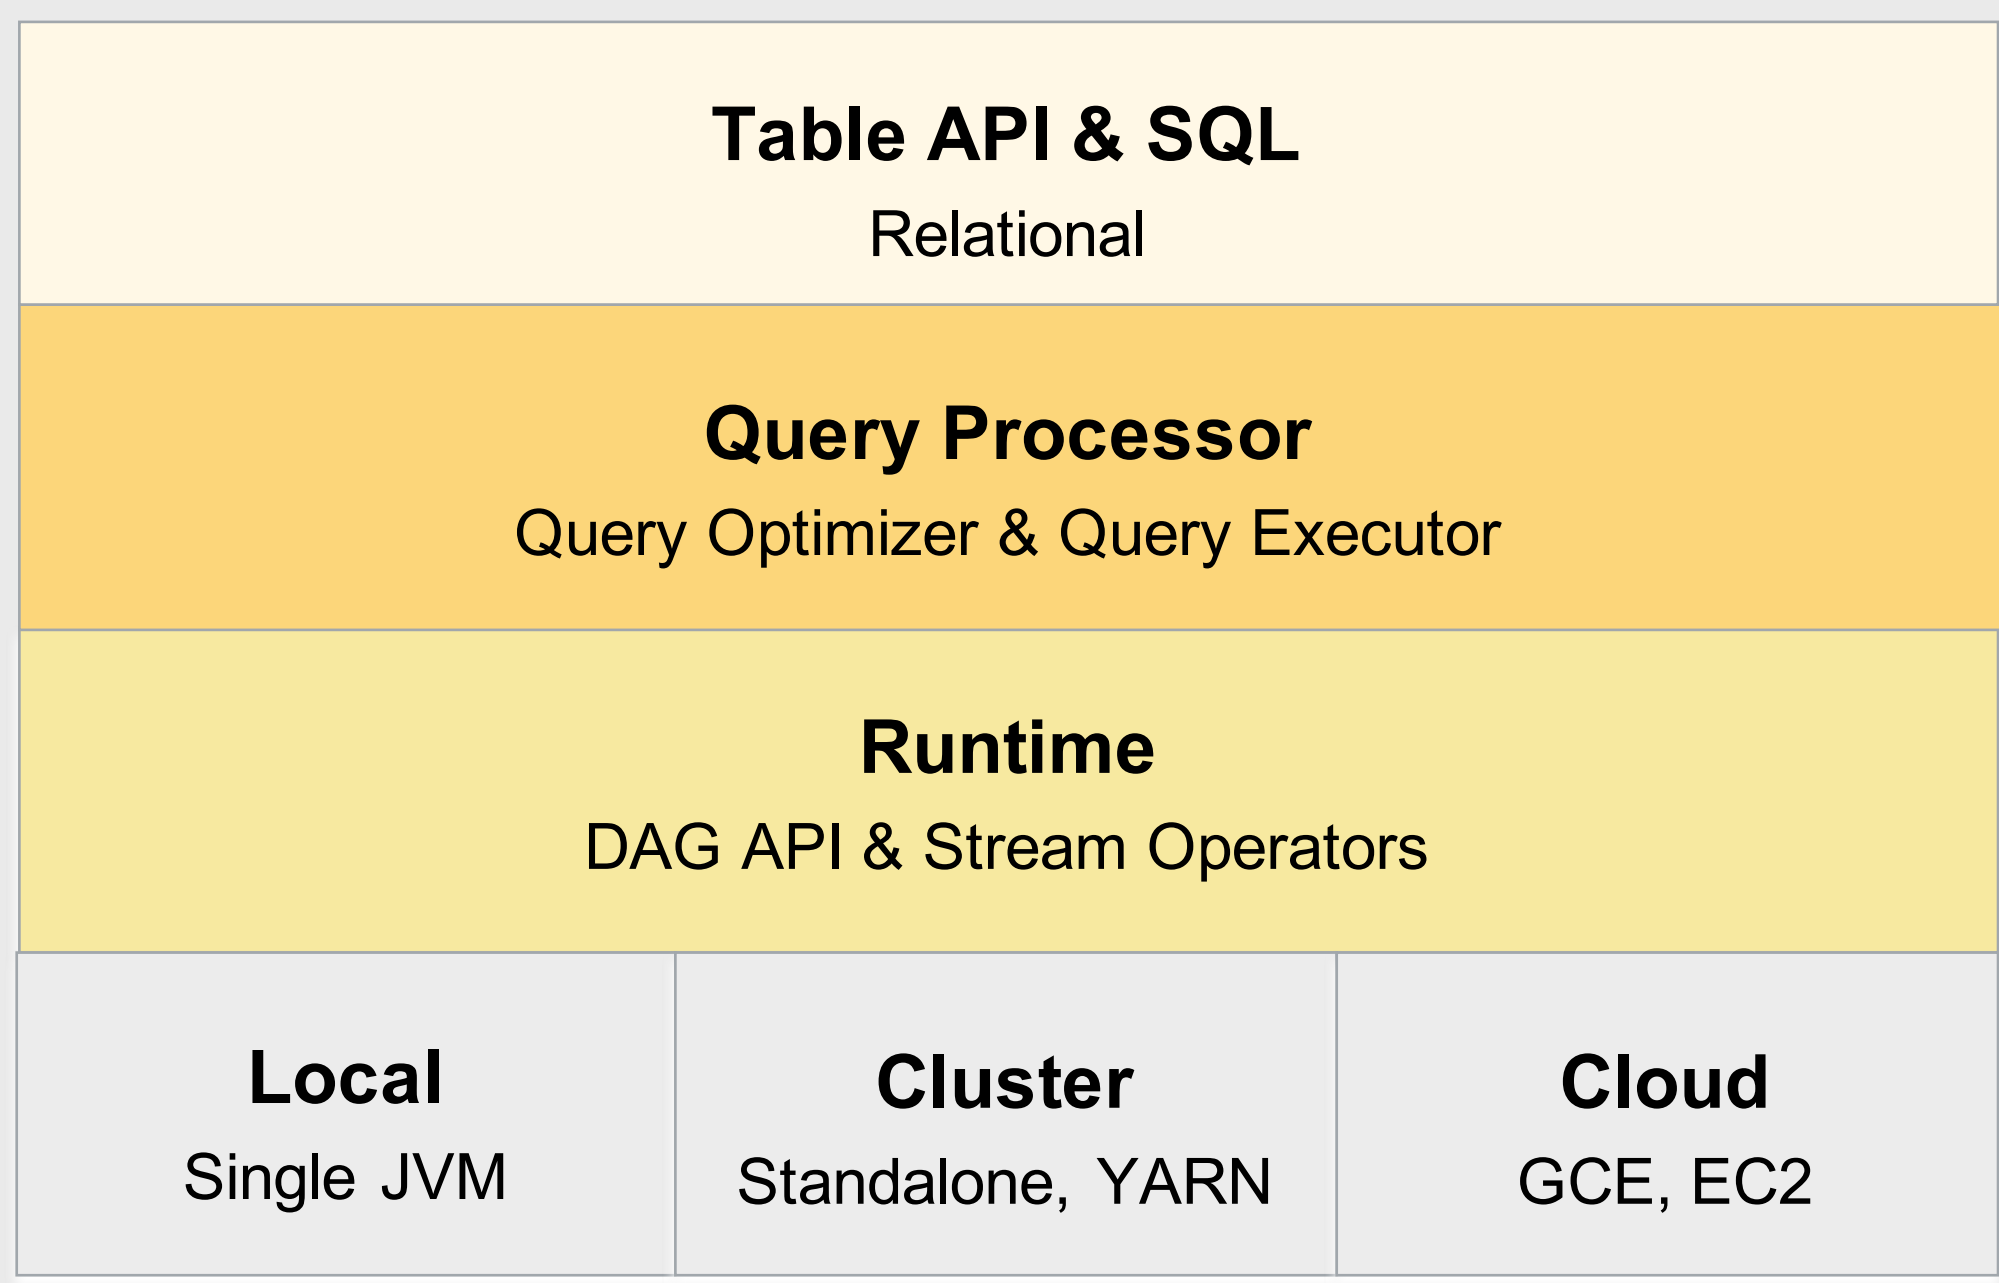
减少CP磁盘IO, 解决流计算时延问题

### **Dynamic Balance&Local-Global AGG**

减少Data Skew, 避免木桶效应



**Flink Stack**



**Blink Stack**

**Unified Engine Unified API**



Unified SQL & Unified API

Batch

Stream

OLAP

ML



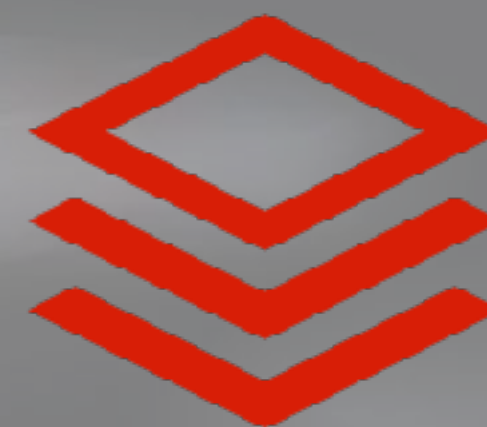
**Declarative**



**Optimizable**



**Understandable**



**Stable**



**Unify**

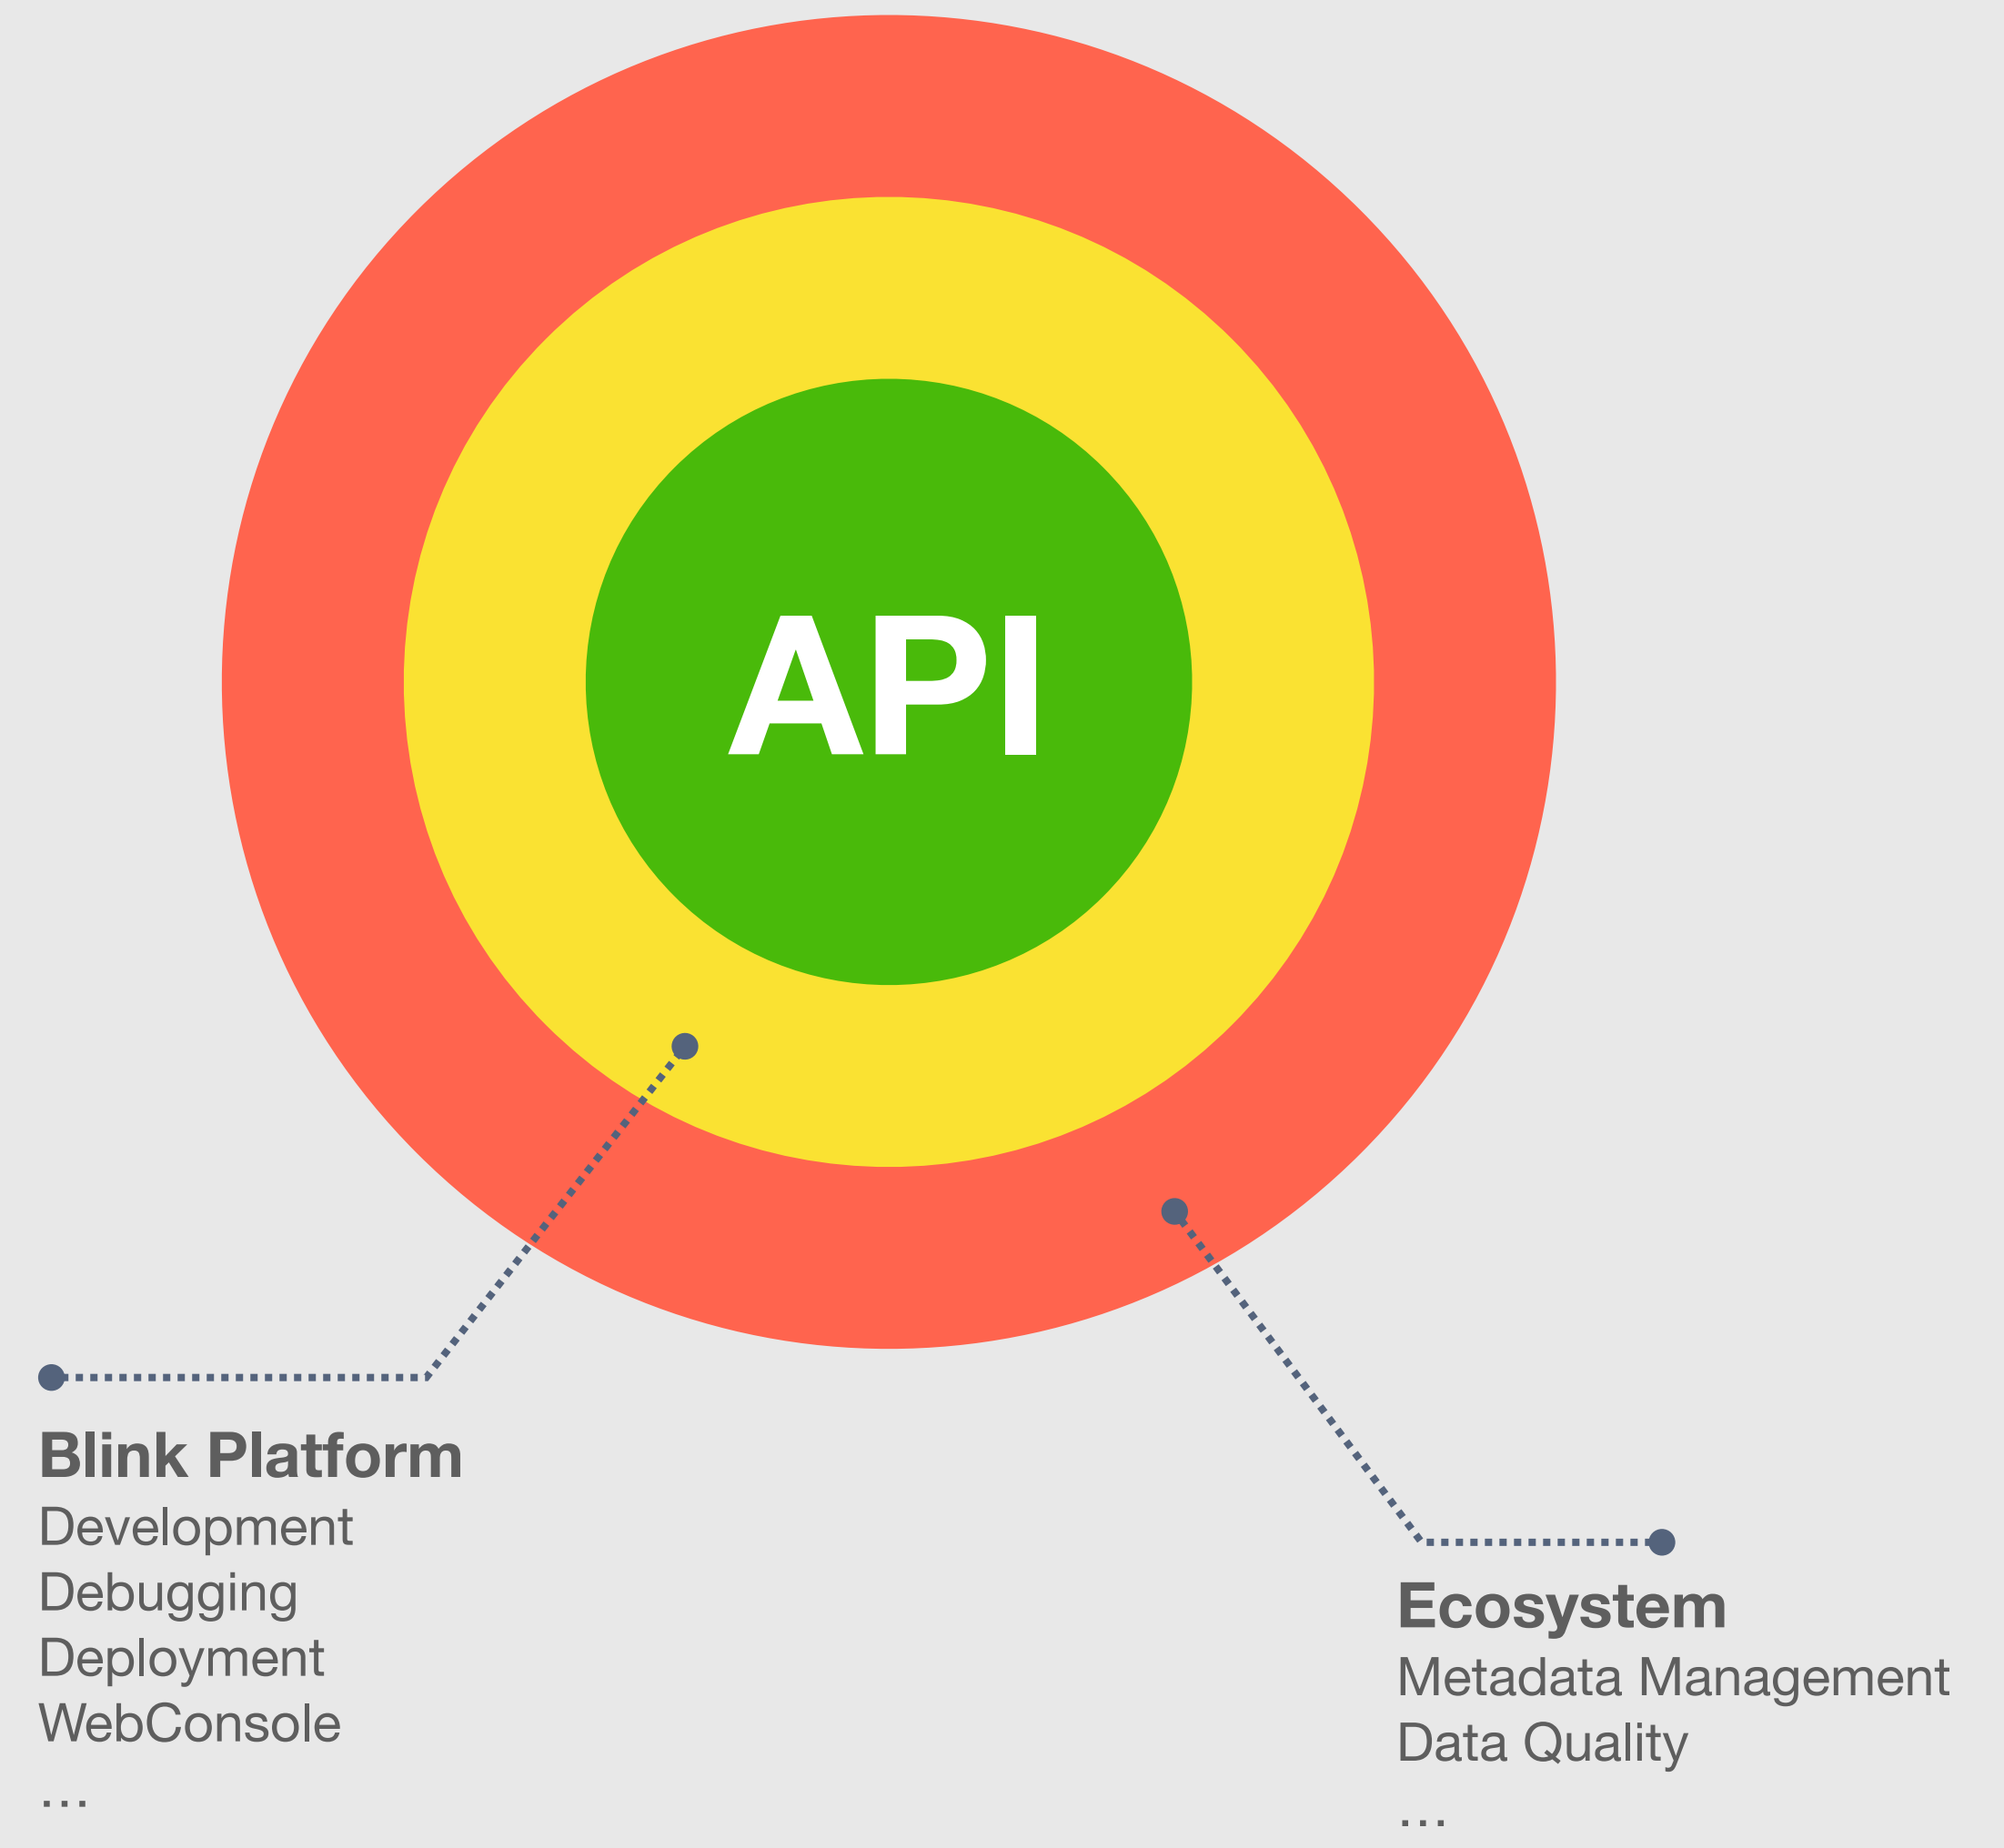


Just API?



SQL/API Only Solves Part of the Problem

We Provide More!



开源计划



(约)一个季度内Blink分支贡献给Apache Flink



**Flink Forward China**

**12月20-21日 @北京 国家会议中心**

亚洲区第一场Flink官方大会，预计3000+参会人员

**Flink中文社区大牛&公司悉数到场**

华为、腾讯、滴滴、饿了么、阿里等共同举办

**Call For Keynotes&Sponsors**

召集keynotes&Sponsors





Thanks

巴真 



 在钉钉上扫一扫加我

 巴真



扫一扫上面的二维码图案，加我微信