



Apache RocketMQ – Trillion Messaging in Practice

冯嘉(Von Gosling)





Von Gosling, Alibaba Architect , Master of Computer Science, Over 7 years of industrial software experience, especially in distributed tech., high performance web architecture and performance tuning. Possesses many patents in distributed system, recommender etc.

Open source fanatic. Focuses on distributed systems, cloud computing and big data fields. Apache RocketMQ PPMC member & committer.

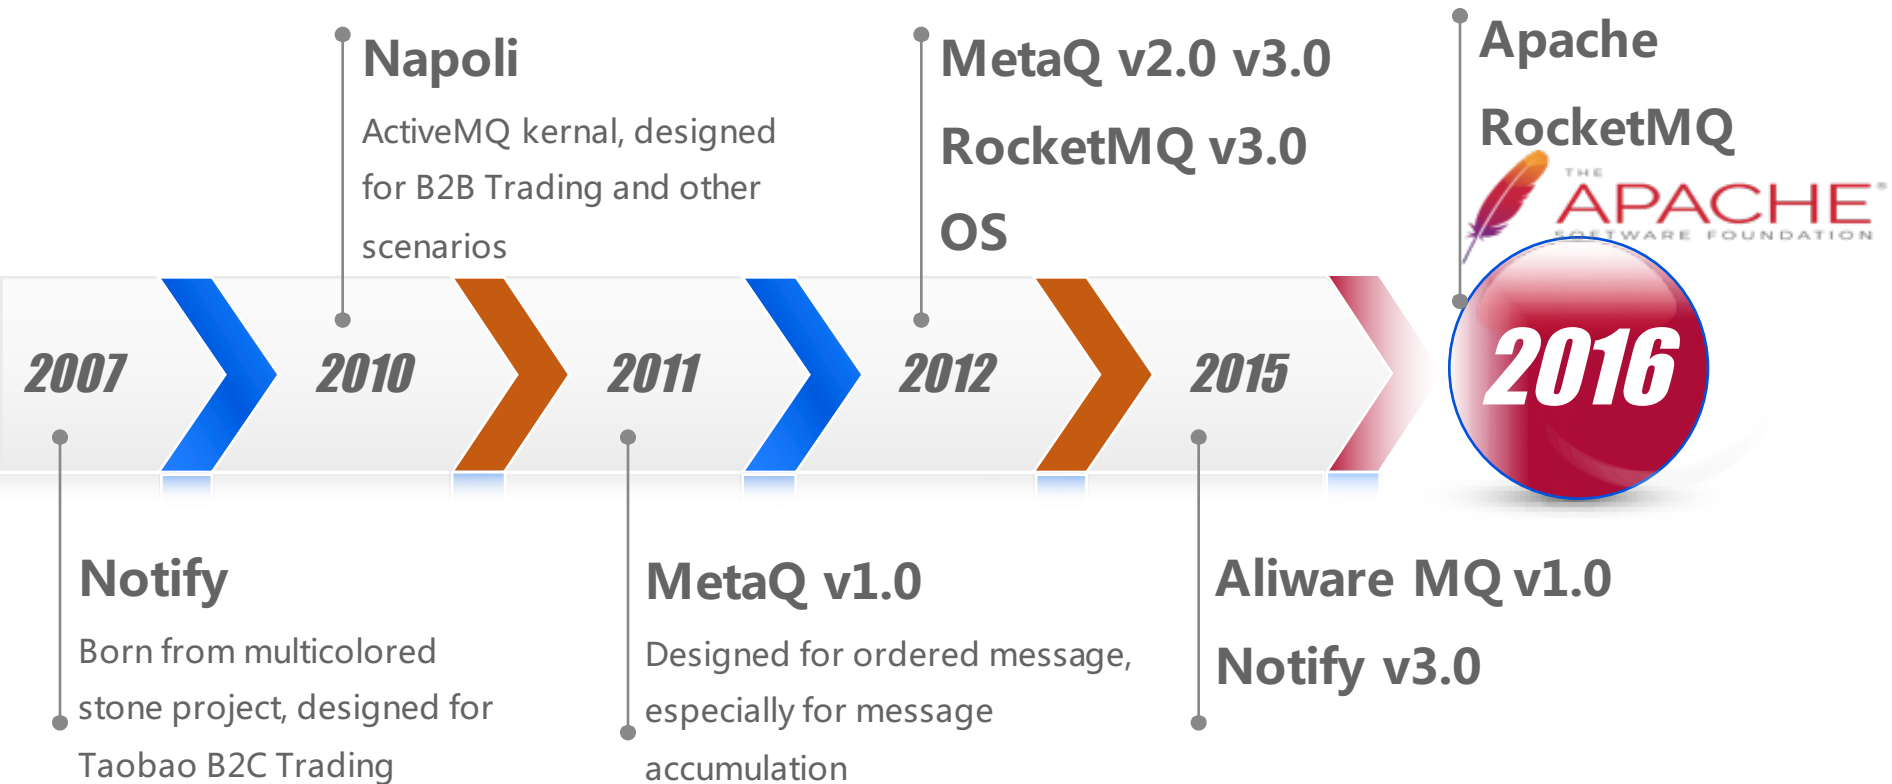
Currently in charge of [Aliware MQ](#). Email: vongosling@apache.org





Part I

From Alibaba to Apache





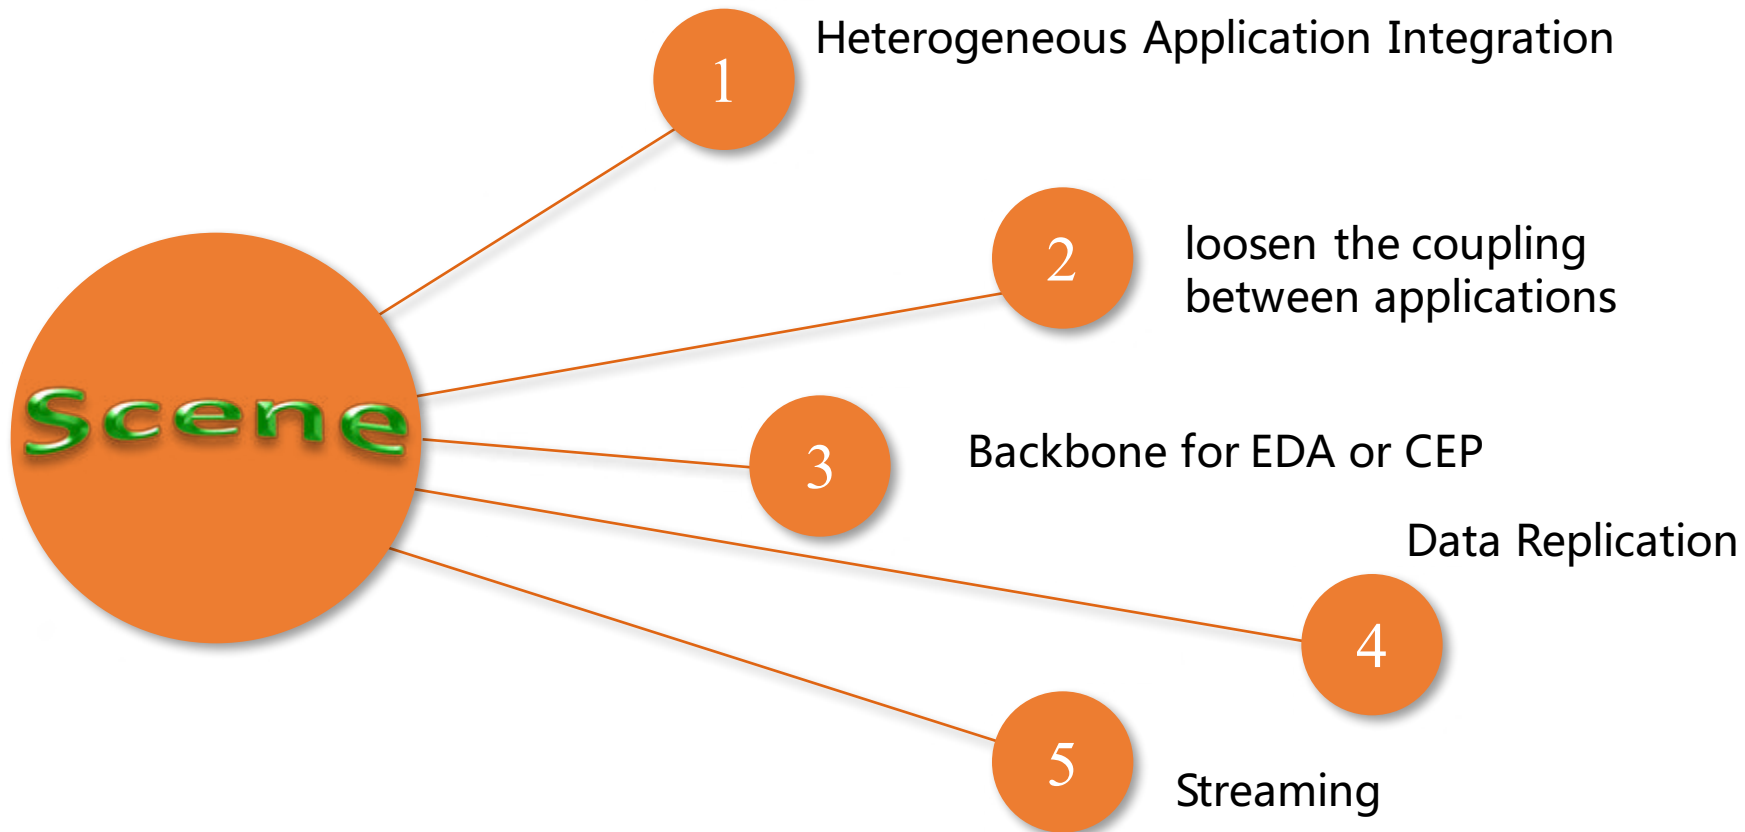
Ordered messaging , Pull model



Transactional messaging,
Push model

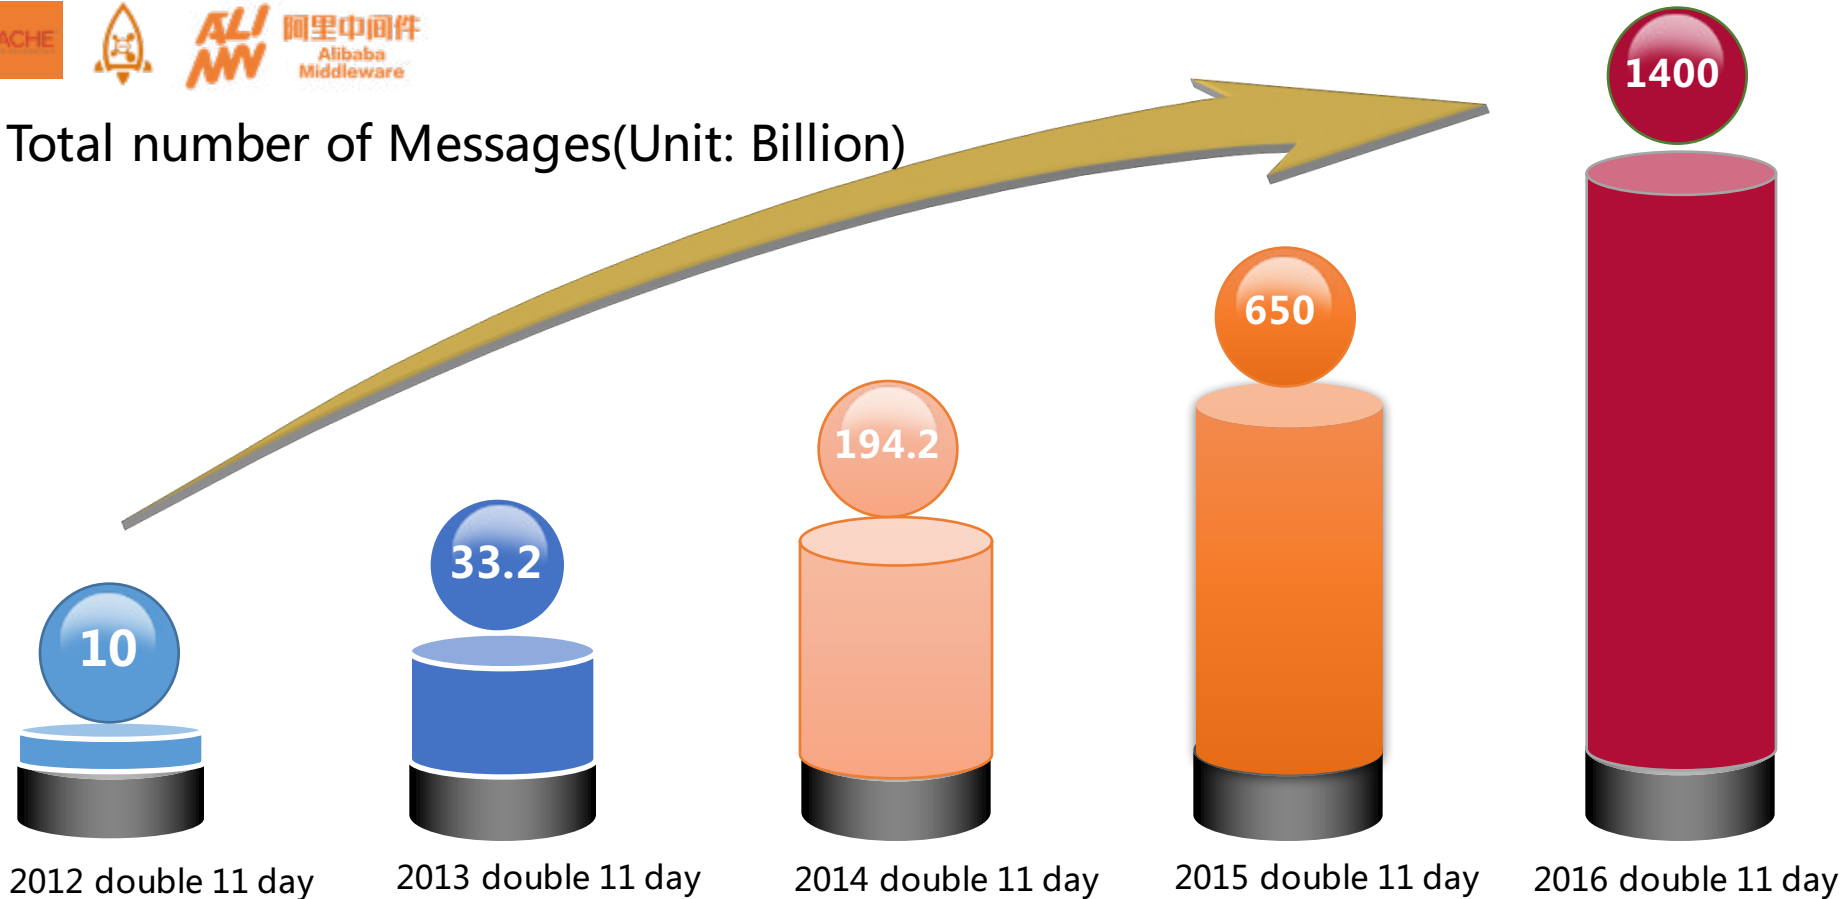


Commercial Distribution ,
Cloud MQ





Total number of Messages(Unit: Billion)





What's Next ↗

Apache RocketMQ

in Practice



01/Architecture and Features

02/Monitoring and Administering

03/Performance Tuning

04/Apache RocketMQ 4.X



阿里中间件
Alibaba
Middleware

01

Architecture and Features

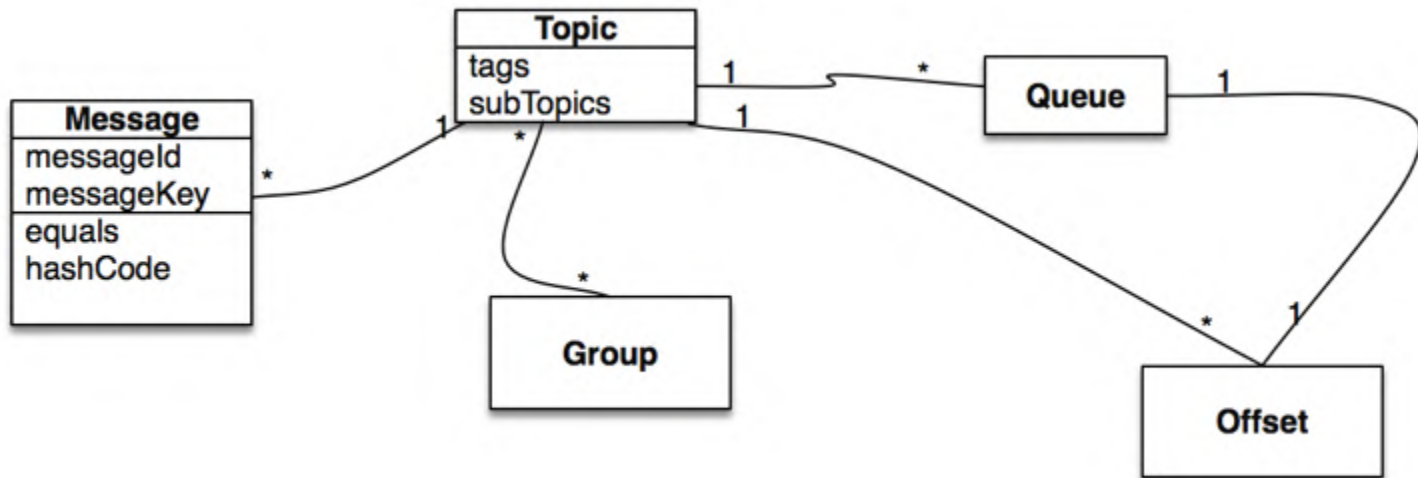


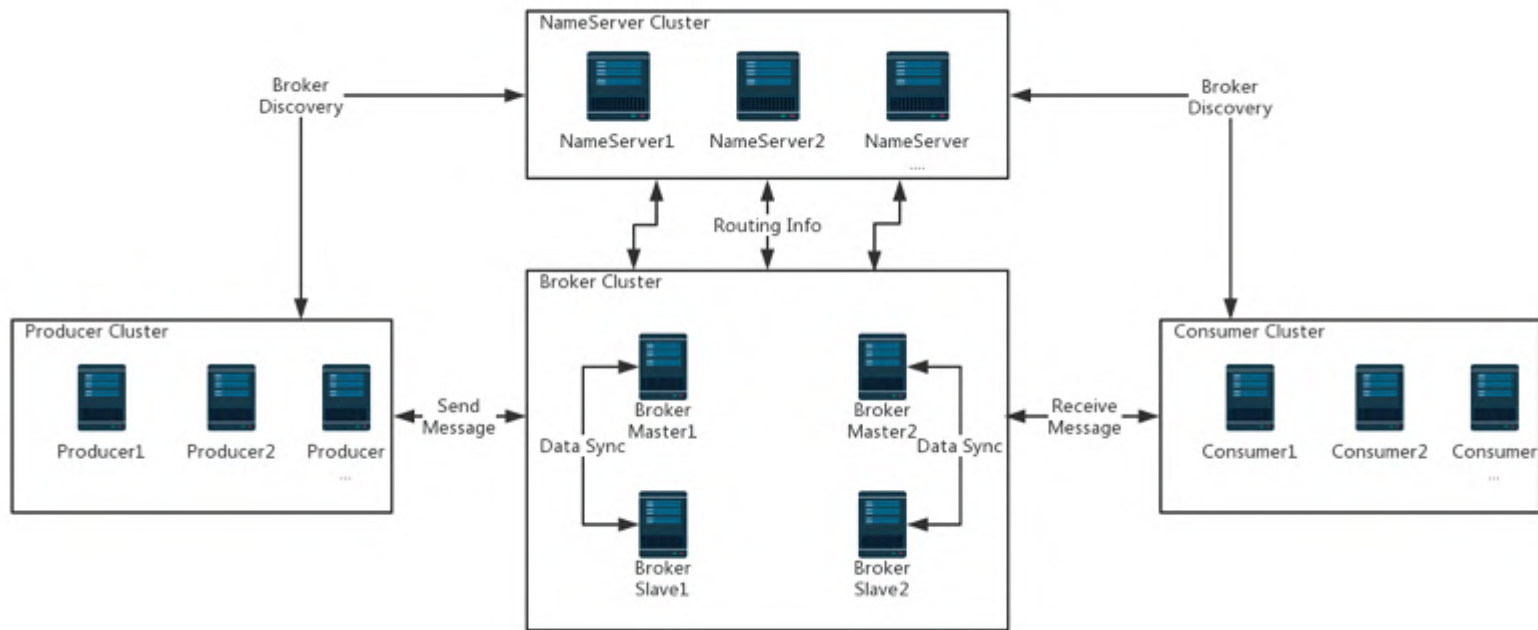
..... 合抱之木生于毫末；九层之台起于累土；千里之行始于足下.....

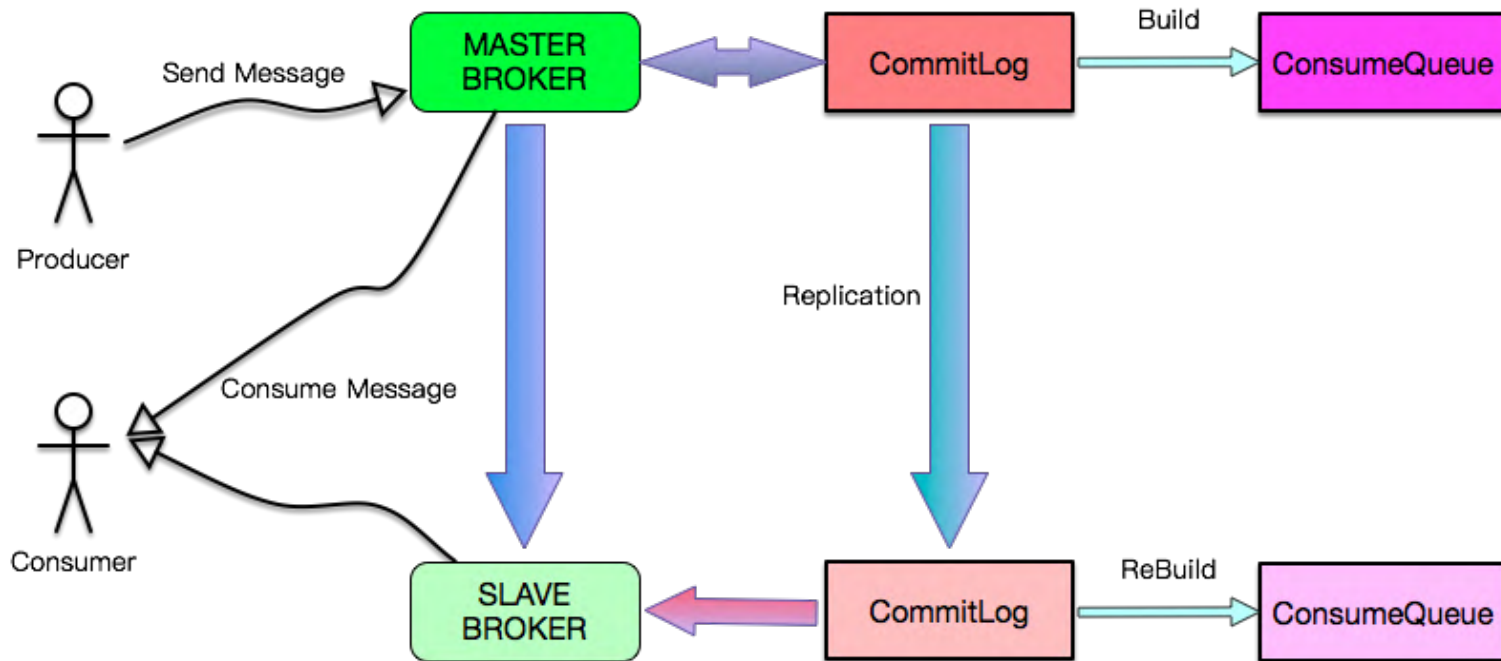


阿里中间件
Alibaba
Middleware

Architecture











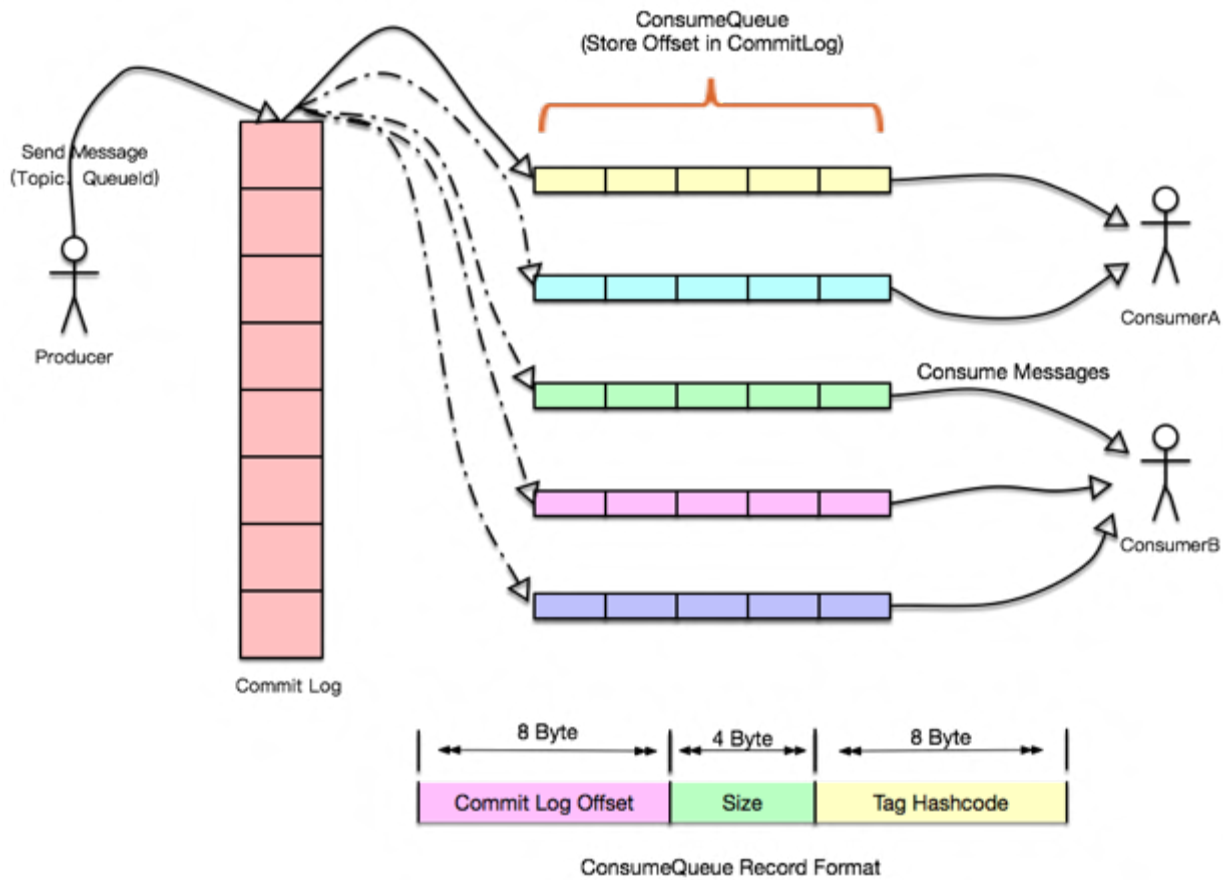
Slot Table

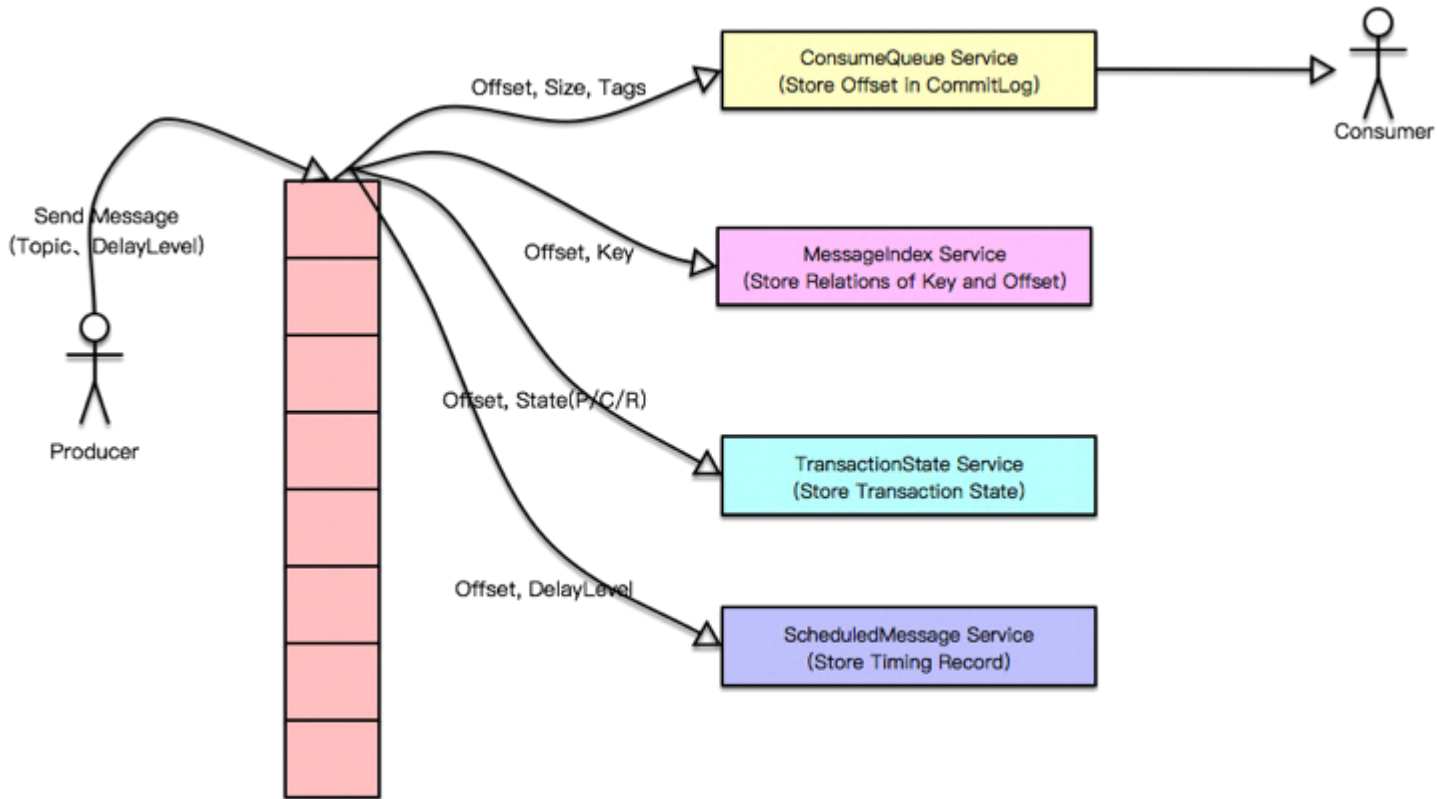


Index Linked List

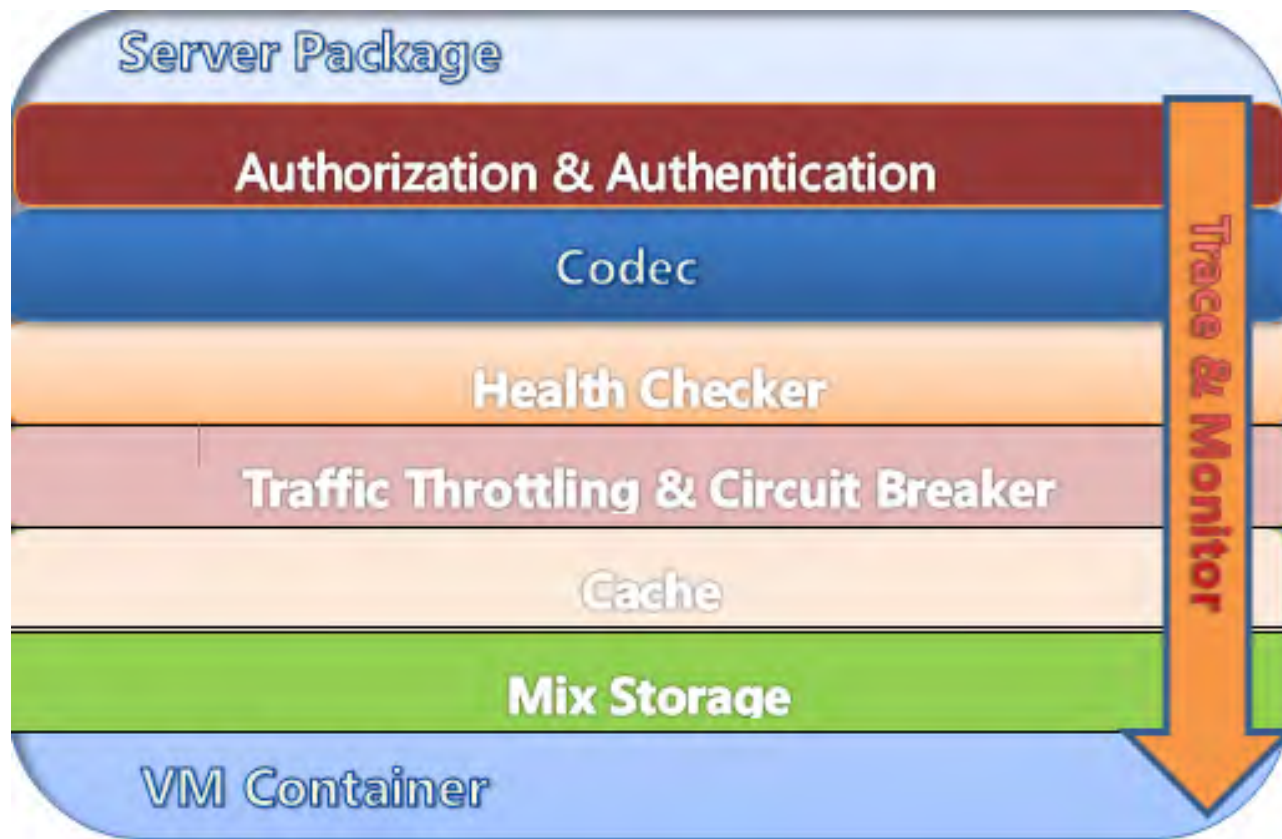


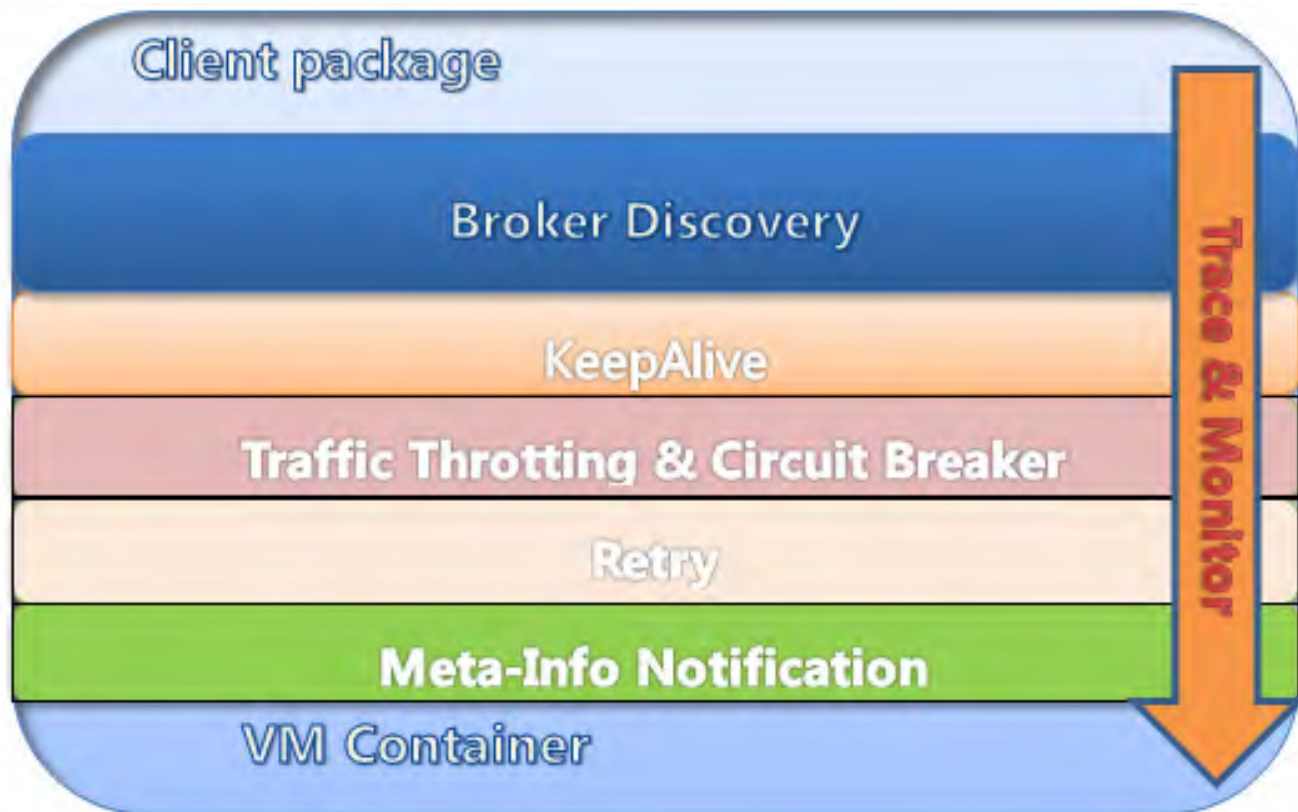
Index File Format





Commit Log







阿里中间件
Alibaba
Middleware

Community Features

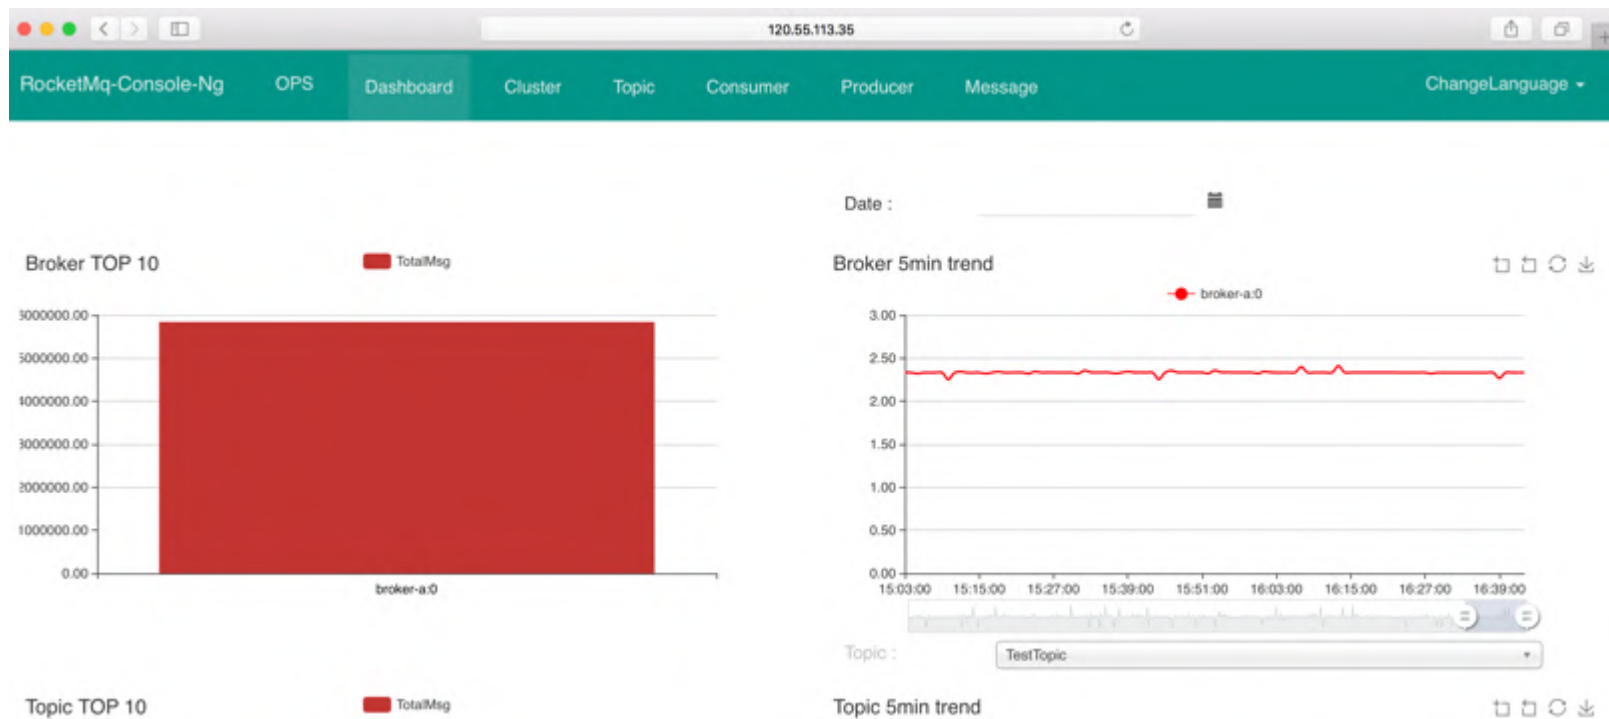


- ✓ Ordered and Non-Ordered message
- ✓ Broadcast message
- ✓ Scheduled message
- ✓ Transactional message
- ✓ Batched message
- ✓ One-way message
- ✓ Traceable message
- ✓ Retroactive message
- ✓ Message Filtering
- ✓ Massive Accumulation
- ✓ Backoff Strategy
- ✓ Delivery QoS
- ✓ Community features – console, docker, streaming etc.

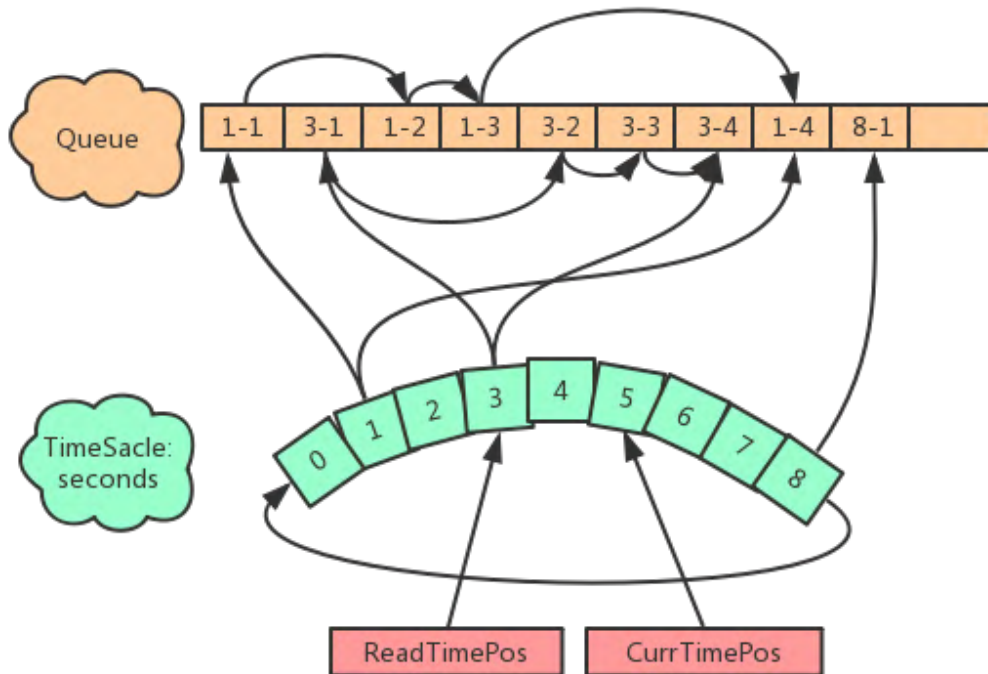




Console

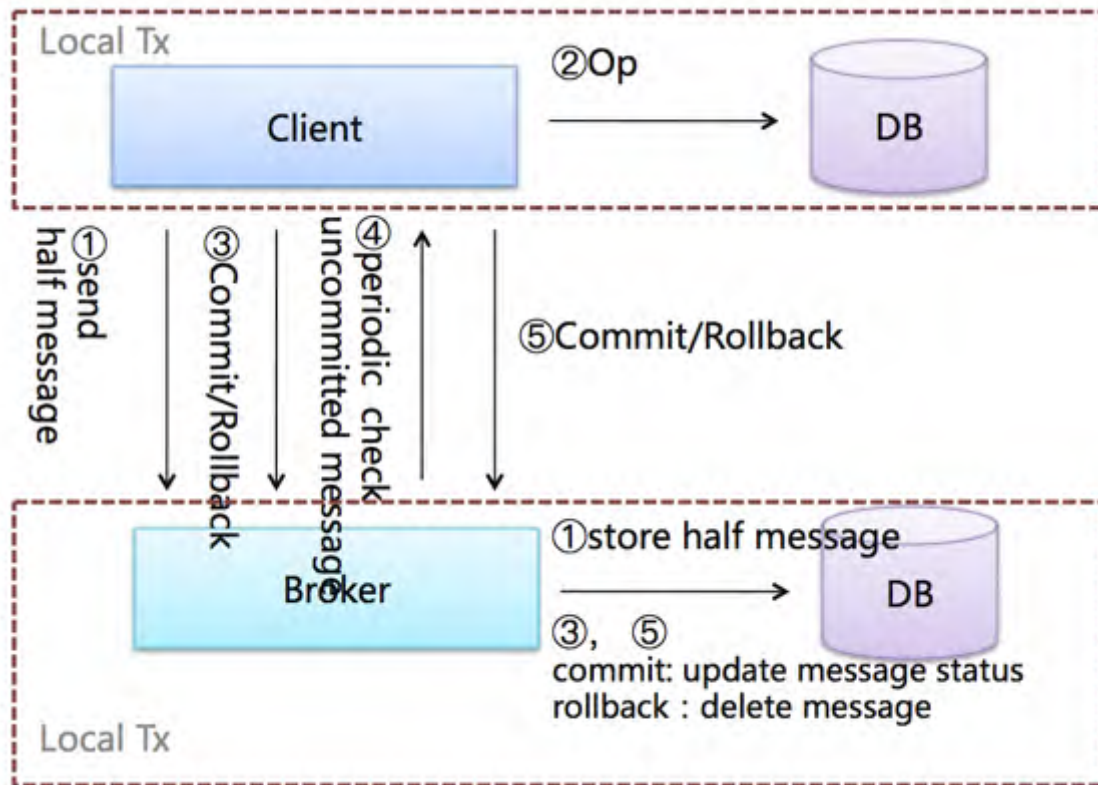


Scheduled Message





Transactional Message – Eventually Consistent

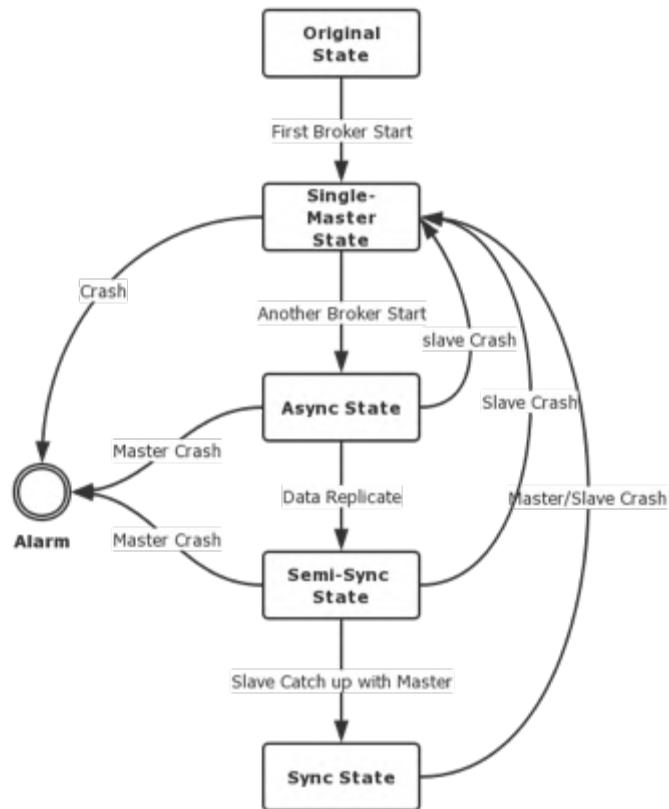
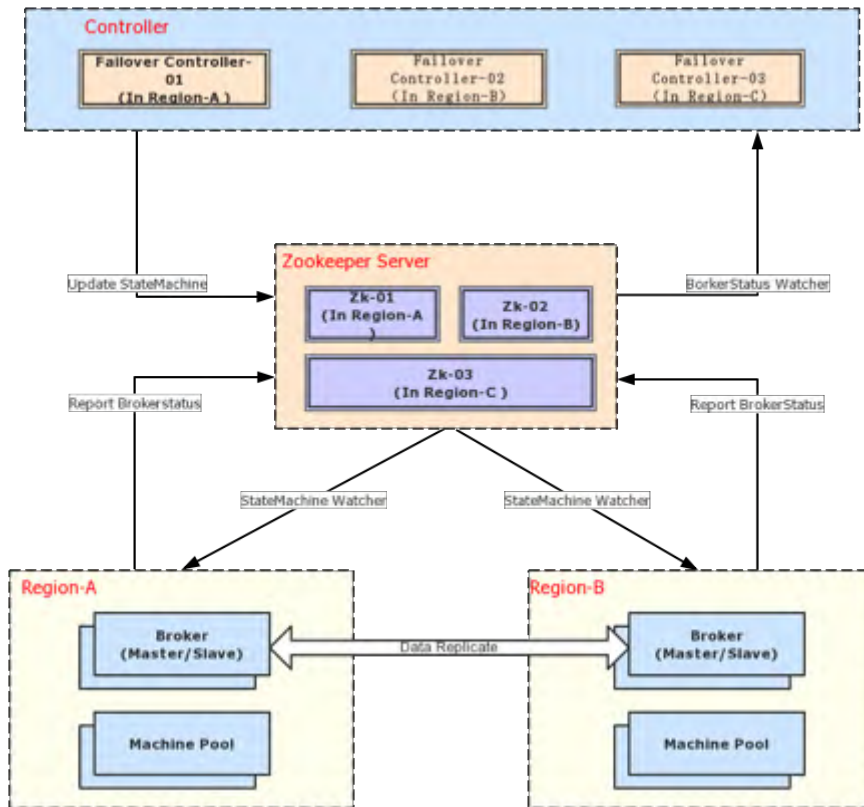




阿里中间件
Alibaba
Middleware

Commercial Features

HA





Message Tracing Dashboard





Kafka Service

发布Topic



【商用Region】：提供可靠、稳定的服务，详细计费价格信息请点击 [\[计费价格信息\]](#)

【测试Region】：只提供公网测试时使用，已商用

Topic:

1.只能包含字母，数字，短横线(-)和下划线(_)

于64字节将被自动截取

消息类型: 无顺序消息 分区顺序消息 全局顺序消息 ✓ kafka消息(公测中)

该类消息可以兼容 MQ 及 Kafka 协议(kafka-client-0.10 及以上版本)，具体详情请参考[\[接入文档\]](#)

备注:

确定

取消

Topic管理 搜索 发布Topic

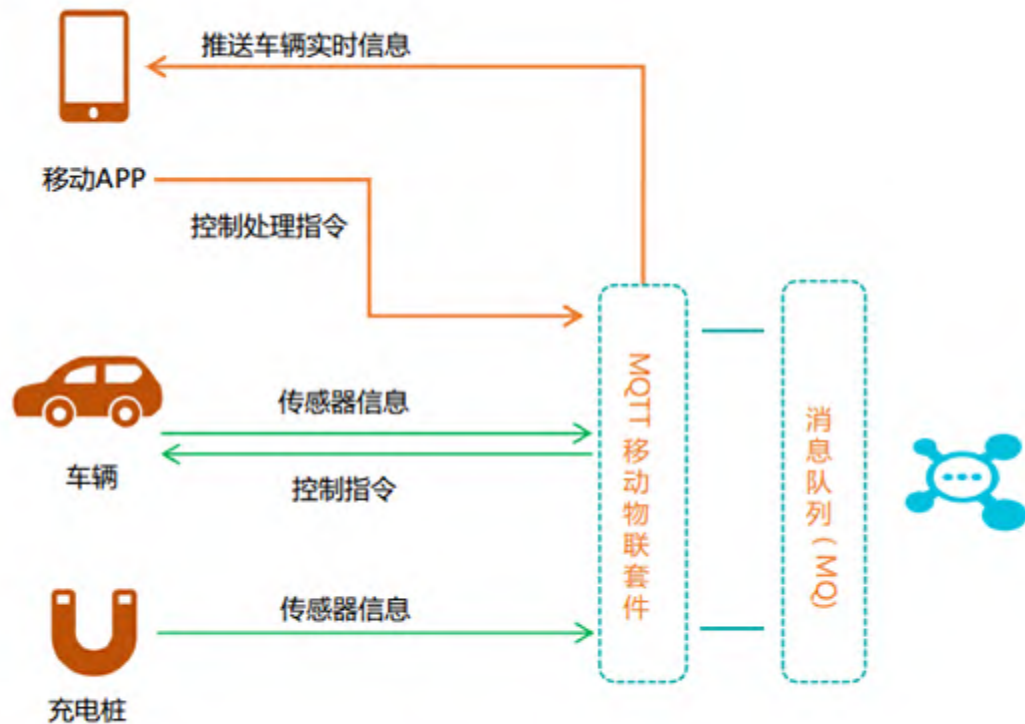
Topic	类型	发布状态	权限	创建时间	操作
kafka-one-test	kafka消息(公测中)	发布中	所有者	2017-03-20 17:45:23	申请发布 申请订阅 申请MQTT GroupID 删除 授权 查看授权 资源列表
kafka-proxy-test	无序消息	发布中	所有者	2017-02-09 14:54:50	申请发布 申请订阅 申请MQTT GroupID 删除 授权 查看授权 资源列表

共 2 条，每页显示：20条 1



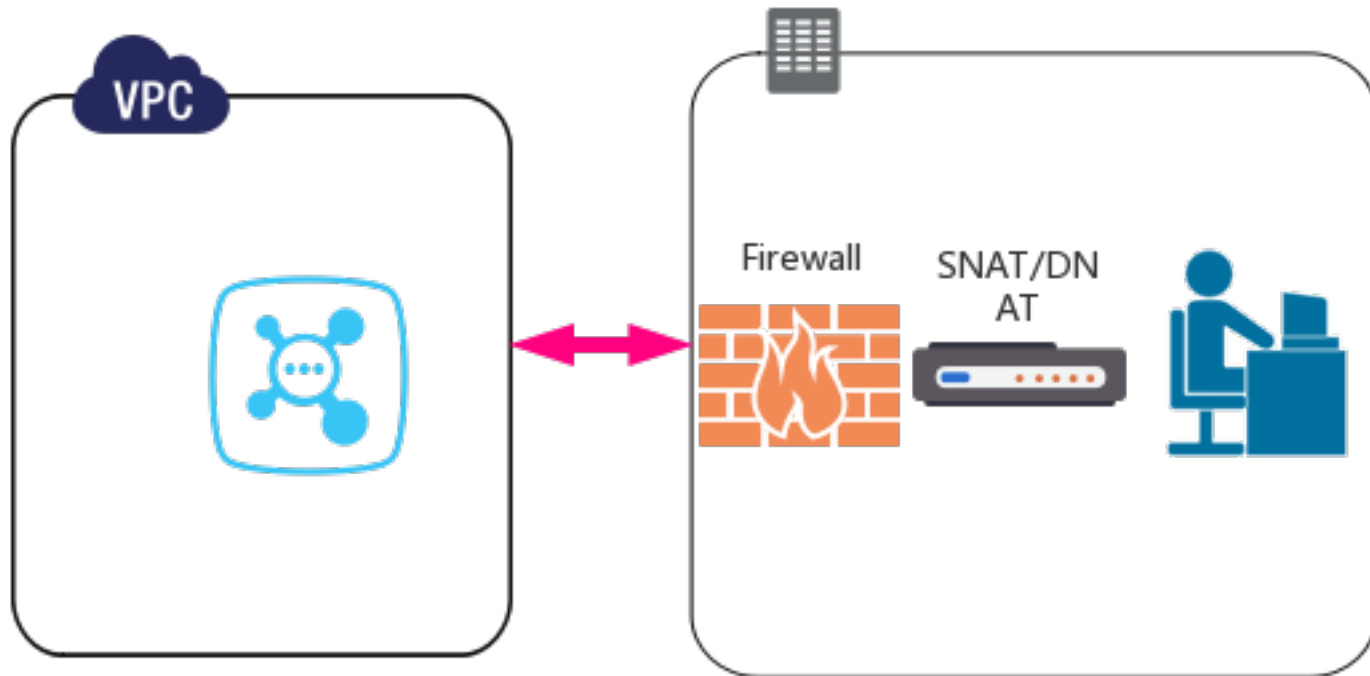
阿里中间件
Alibaba
Middleware

IoT - MQTT





Relay Service






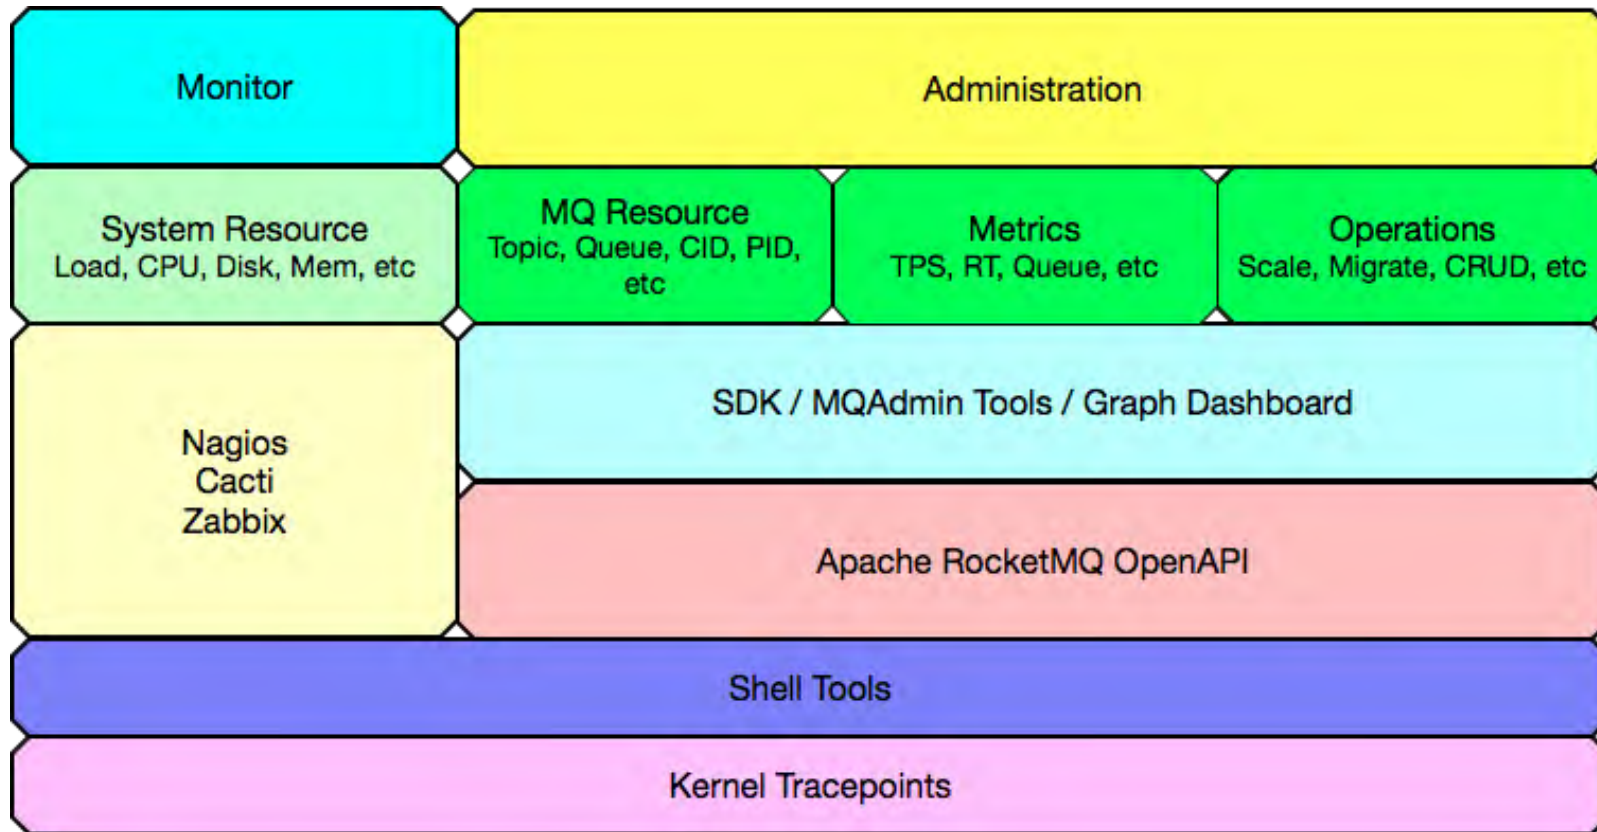
阿里中间件
Alibaba
Middleware

02

Monitoring and Administering

千磨万击还坚韧，任尔东西南北风







```
$sh mqadmin
```

The most commonly used mqadmin commands are:

healthCheck health check

deleteSubGroup Delete subscription group from broker.

updateBrokerConfig Update broker's config

getBrokerConfig Get broker config by cluster or special broker

updateTopicPerm Update topic perm

topicRoute Examine topic route info

topicStatus Examine topic Status info

queryMsgById Query Message by Id

queryMsgByKey Query Message by Key

consumerProgress Query consumers's progress, speed

consumerStatus Query consumer's internal data structure

clusterList List all of clusters

resetOffsetByTime Reset consumer offset by timestamp(without client restart).

...

...



阿里中间件
Alibaba
Middleware

```
$sh mqadmin clusterlist
```

#Cluster Name	#Broker Name	#BID	#Addr	#Version	#InTPS(LOAD)
share	share-01	0	10.10.10.10:10911	V4_0_0	3840.73(0,0ms)
42.31 0.8394					7758.35(0,0ms)
share	share-02	1	10.10.10.11:10911	V4_0_0	3764.42(1,0ms)
82.97 0.3245					7734.55(1,0ms)



```
$sh mqadmin consumerstatus -g CID_SHOPPING -s
```

```
001 10.10.10.4@29077 V4_0_0 1493971139224/10.10.10.4@29077#-1244857759#-  
163453988#14074607828995785  
002 10.10.10.2@28953 V4_0_0 1493971139224/10.10.10.2@28953#-1244857759#-  
163453988#14074505387423377  
003 10.10.10.1@11865 V4_0_0 1493971139224/10.10.10.1@11865#-1244857759#-  
163453988#14074508601484056  
004 10.10.10.3@23180 V4_0_0 1493971139224/10.10.10.3@23180#-1244857759#-  
163453988#14074400788700508
```



```
$sh mqadmin querymsgbykey -k Query_Key -t ORDER_TOPIC
```

```
#Message ID
```

```
707C8CC435FE3FD00E77180D7FCB0B24
```

```
#QID
```

```
0
```

```
#Offset
```

```
1
```



```
$sh mqadmin querymsgbyuniquekey -i 707C8CC435FE3FD00E77180D7FCB0B24 -t  
ORDER_TOPIC  
Topic:          ORDER_TOPIC  
Tags:           [ApacheCon]  
Keys:           [Query_Key]  
Queue ID:       0  
Queue Offset:   1  
CommitLog Offset: 148053682794  
Reconsume Times: 0  
Born Timestamp: 2017-05-05 16:05:37,867  
Store Timestamp: 2017-05-05 16:05:37,880  
Born Host:      10.143.33.157:42775  
Store Host:     10.151.203.124:10911  
System Flag:    0  
Properties:     {TRACE_ON=true, KEYS=Query_Key, MSG_REGION=cn-hangzhou,  
UNIQ_KEY=707C8CC435FE3FD00E77180D7FCB0B24, TAGS=ApacheCon}  
Message Body Path: /tmp/rocketmq/msgbodies/707C8CC435FE3FD00E77180D7FCB0B24
```



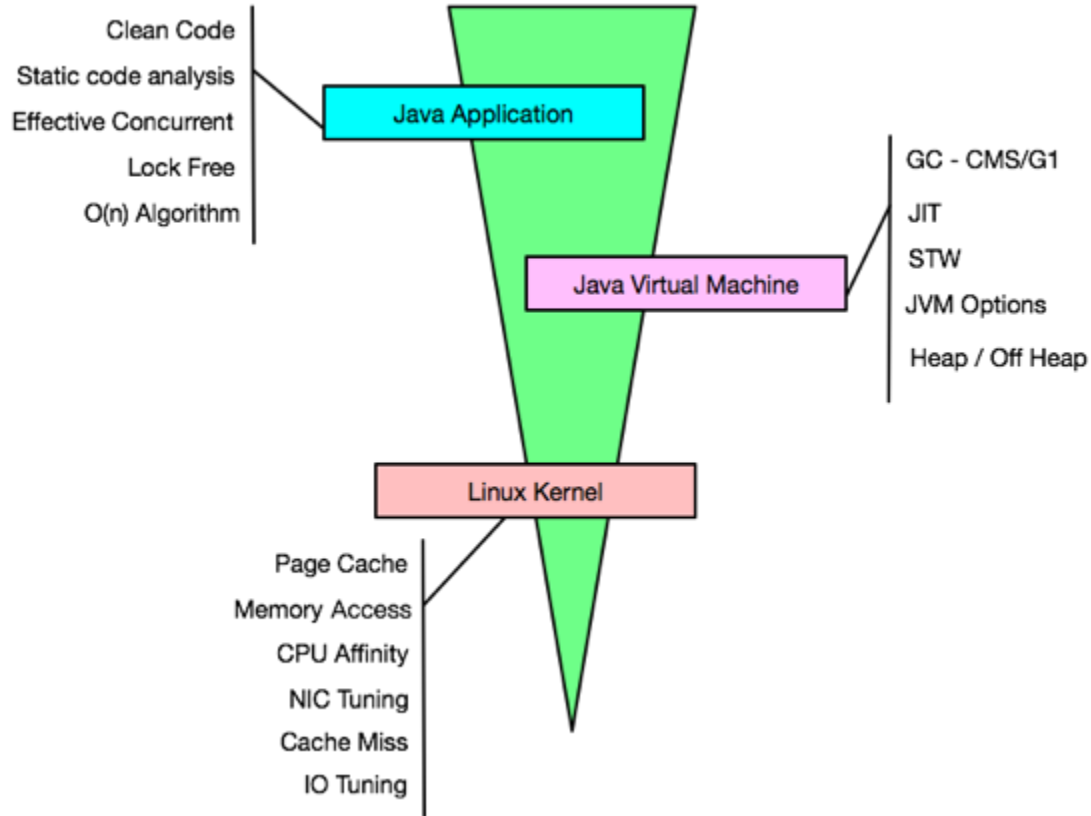
阿里中间件
Alibaba
Middleware

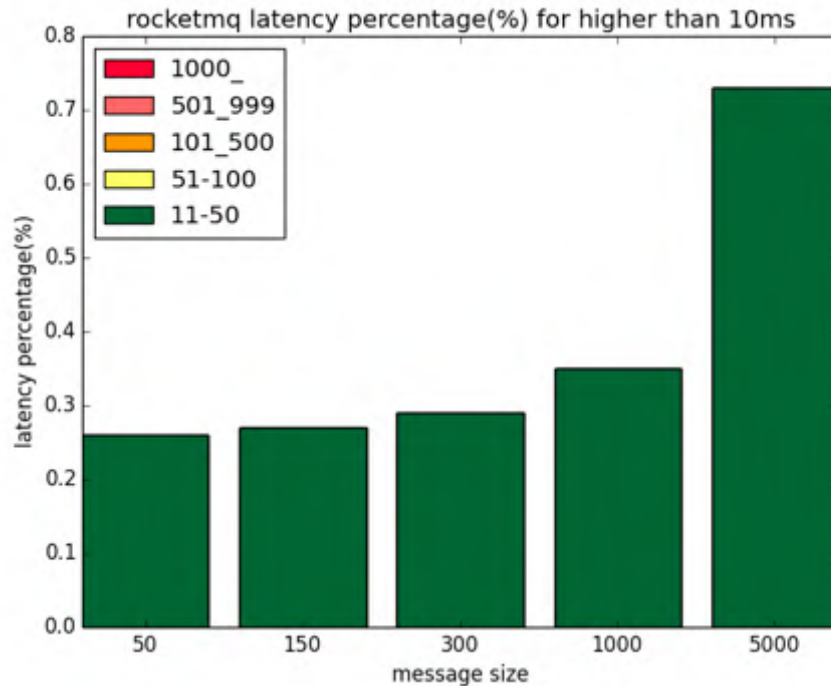
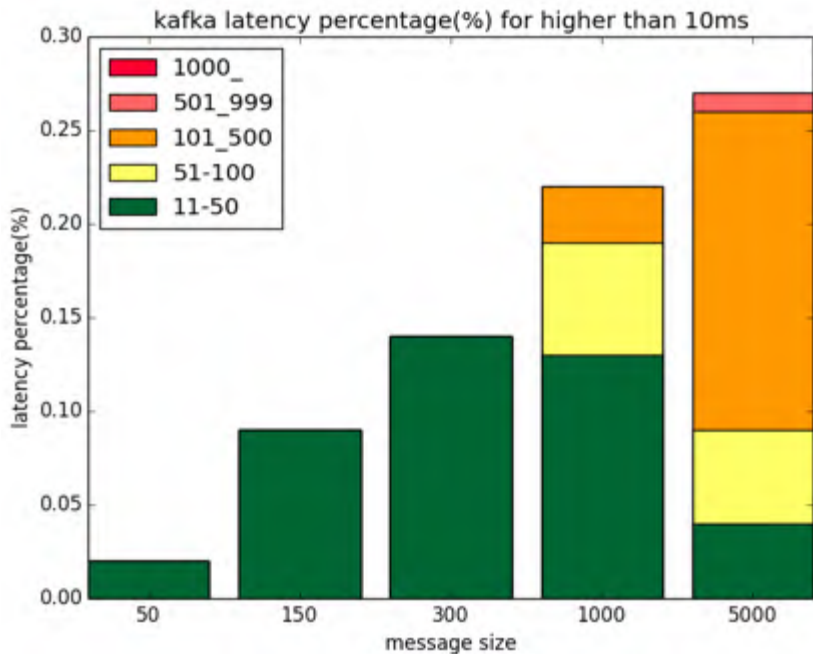
03

Performance Tuning

行路难，行路难，多歧路，今安在。
长风破浪会有时，直挂云帆济沧海。








RocketMQ vs. Kafka



RocketMQ latency distribution heat map





阿里中间件
Alibaba
Middleware

04

RocketMQ 4.X



咬定青山不放松，立根原在破岩中



E-Commerce



High Concurrency
High Reliability



Finance



IoT

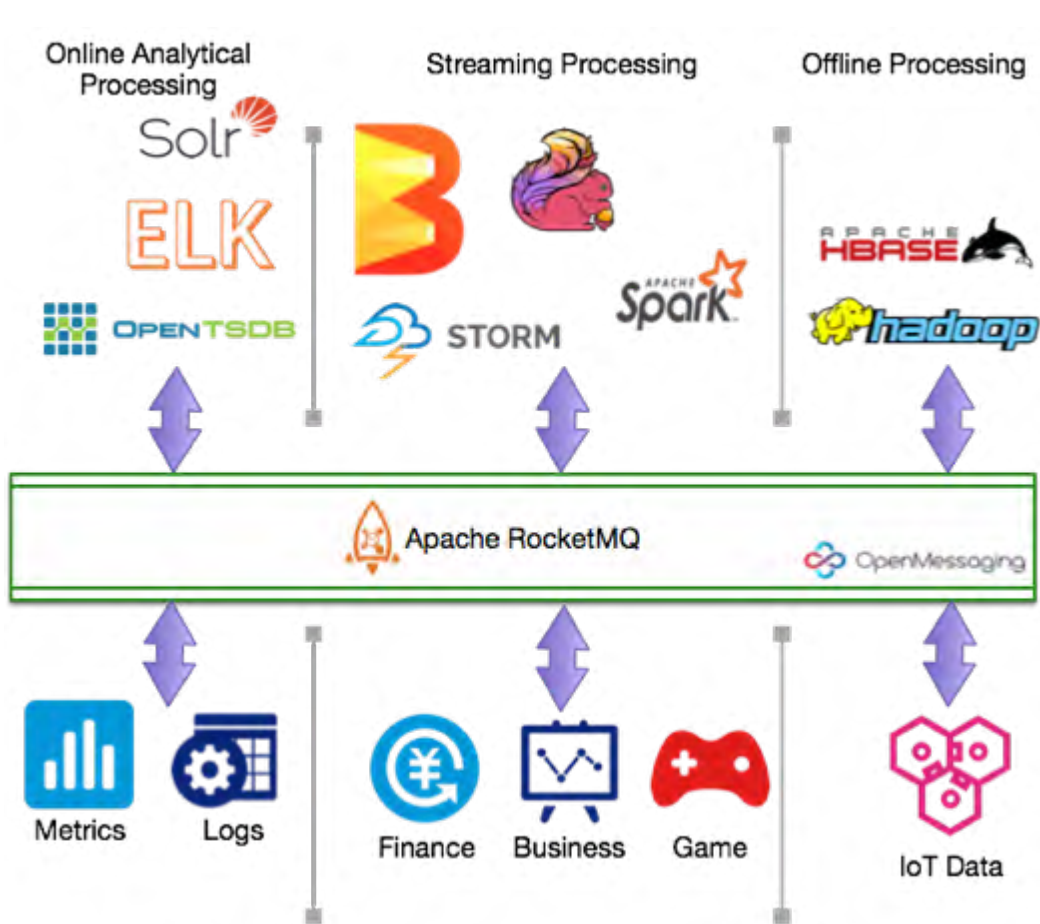


Massive Online Devices
High Throughput



BigData







OpenMessaging

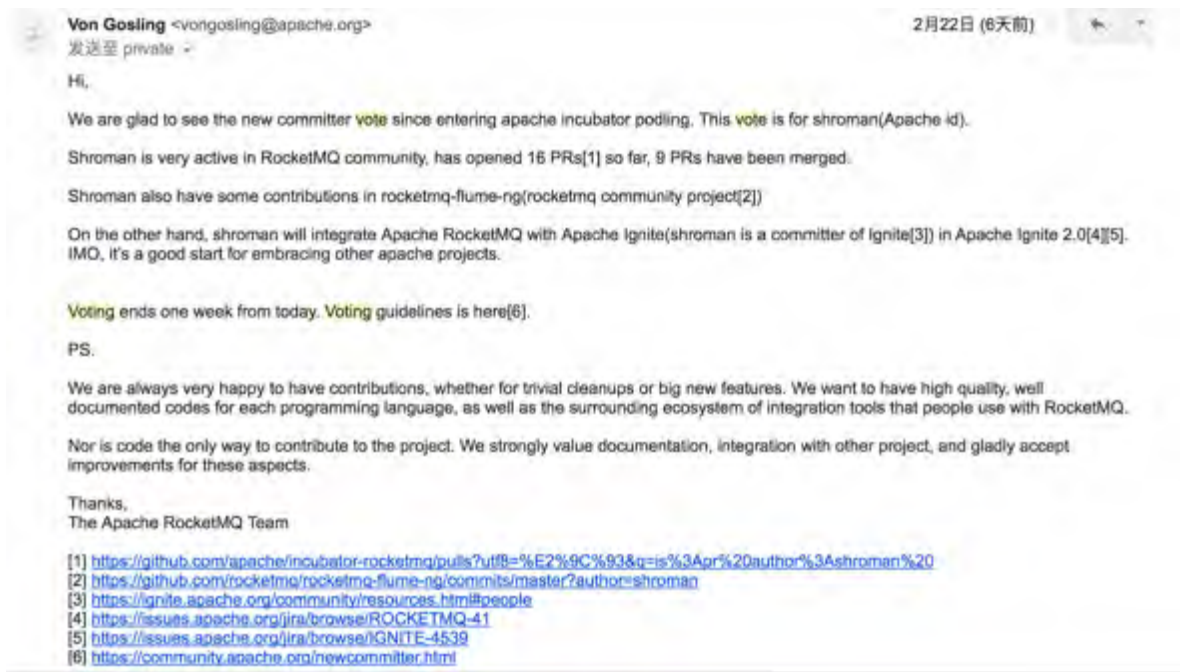
A vendor-neutral open standard for distributed messaging and streaming

OpenMessaging, which includes the establishment of industry guidelines and messaging, streaming specifications to provide a common framework for finance, e-commerce, IoT and big-data area. The design principles are the cloud-oriented, simplicity, flexibility, and language independent in distributed heterogeneous environments. Conformance to these specifications will make it possible to develop a heterogeneous messaging applications across all major platforms and operating systems.



Apache RocketMQ Community

- ✓ Workshop, Meetup
- ✓ Code Marathon
- ✓ Community Merit





谢谢 THANKS!

Alibaba Middleware

Recommended Latest Articles :

1. From Alibaba to Apache: RocketMQ' s Past, Present, and Future. [InfoQ](#)
2. Apache RocketMQ背后的设计思路与最佳实践. [云栖社区](#)
3. 专访RocketMQ联合创始人：项目思路、技术细节和未来规划. [InfoQ](#)
4. 万亿级数据洪峰下的分布式消息引擎. [电子书节选章节](#)
5. 捐赠Apache背后那些鲜为人知的故事. [中间件博客节选](#)



沙加
浙江 杭州



扫一扫上面的二维码图案，关注微博