



APACHE
APISIX



基于 Apache APISIX 与 RocketMQ 构建云原生一体化架构

杜恒

ASF Member, Apache RocketMQ PMC 成员

APACHE APISIX CONNECTS THE WORLD



APACHE
APISIX



Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Alibaba's products remains at the sole discretion of Alibaba.



APACHE
APISIX



CONTENT

- 01 Apache RocketMQ 简介
- 02 云原生时代的 RocketMQ
- 03 借力 APISIX 构建云原生接入体系



APACHE
APISIX



01 Apache RocketMQ 简介



POWERED BY
APACHE
APISIX
APACHE

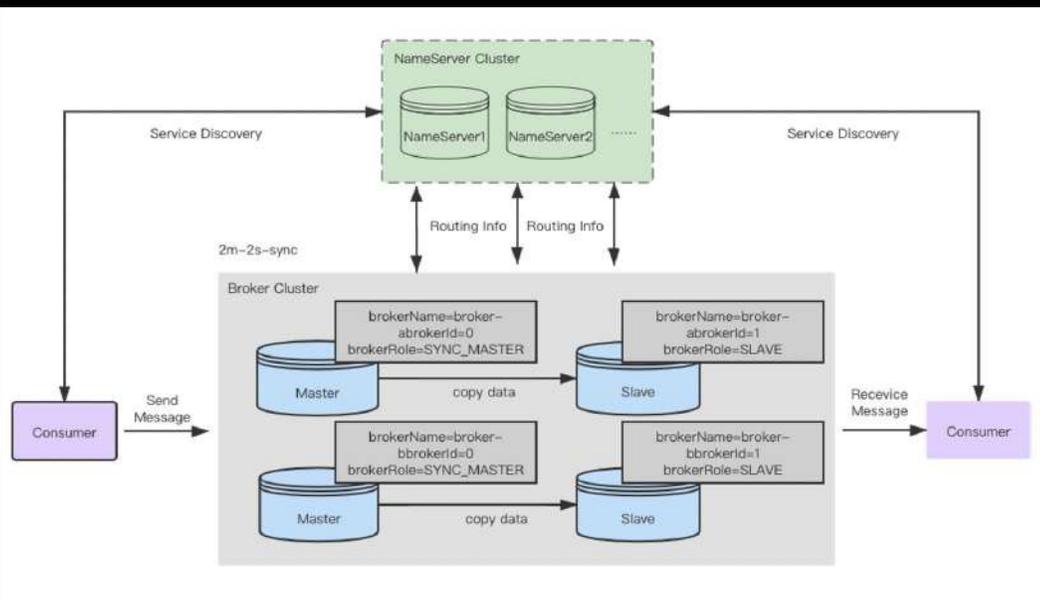
RocketMQ 简介

业务消息领域挑战

- 核心链路，稳定性要求高、时延敏感
- 容量峰值具有随机性，弹性要求高
- 业务场景复杂、集成要求尽可能简单
- 运维及流量调拨要求高

设计思想：

1. 消息不丢、高可靠是架构的基础
2. 时延优先，兼顾吞吐
3. 收敛业务共性问题，提供丰富的业务消息类型
4. 注重可运维性、弹性扩缩、流量调拨能力建设



打造业务消息领域首选

普通消息

顺序消息

延迟消息

事务消息

重试消息

死信消息

零依赖

可扩展

极简架构

低延迟

高吞吐

高性能

强同步刷盘

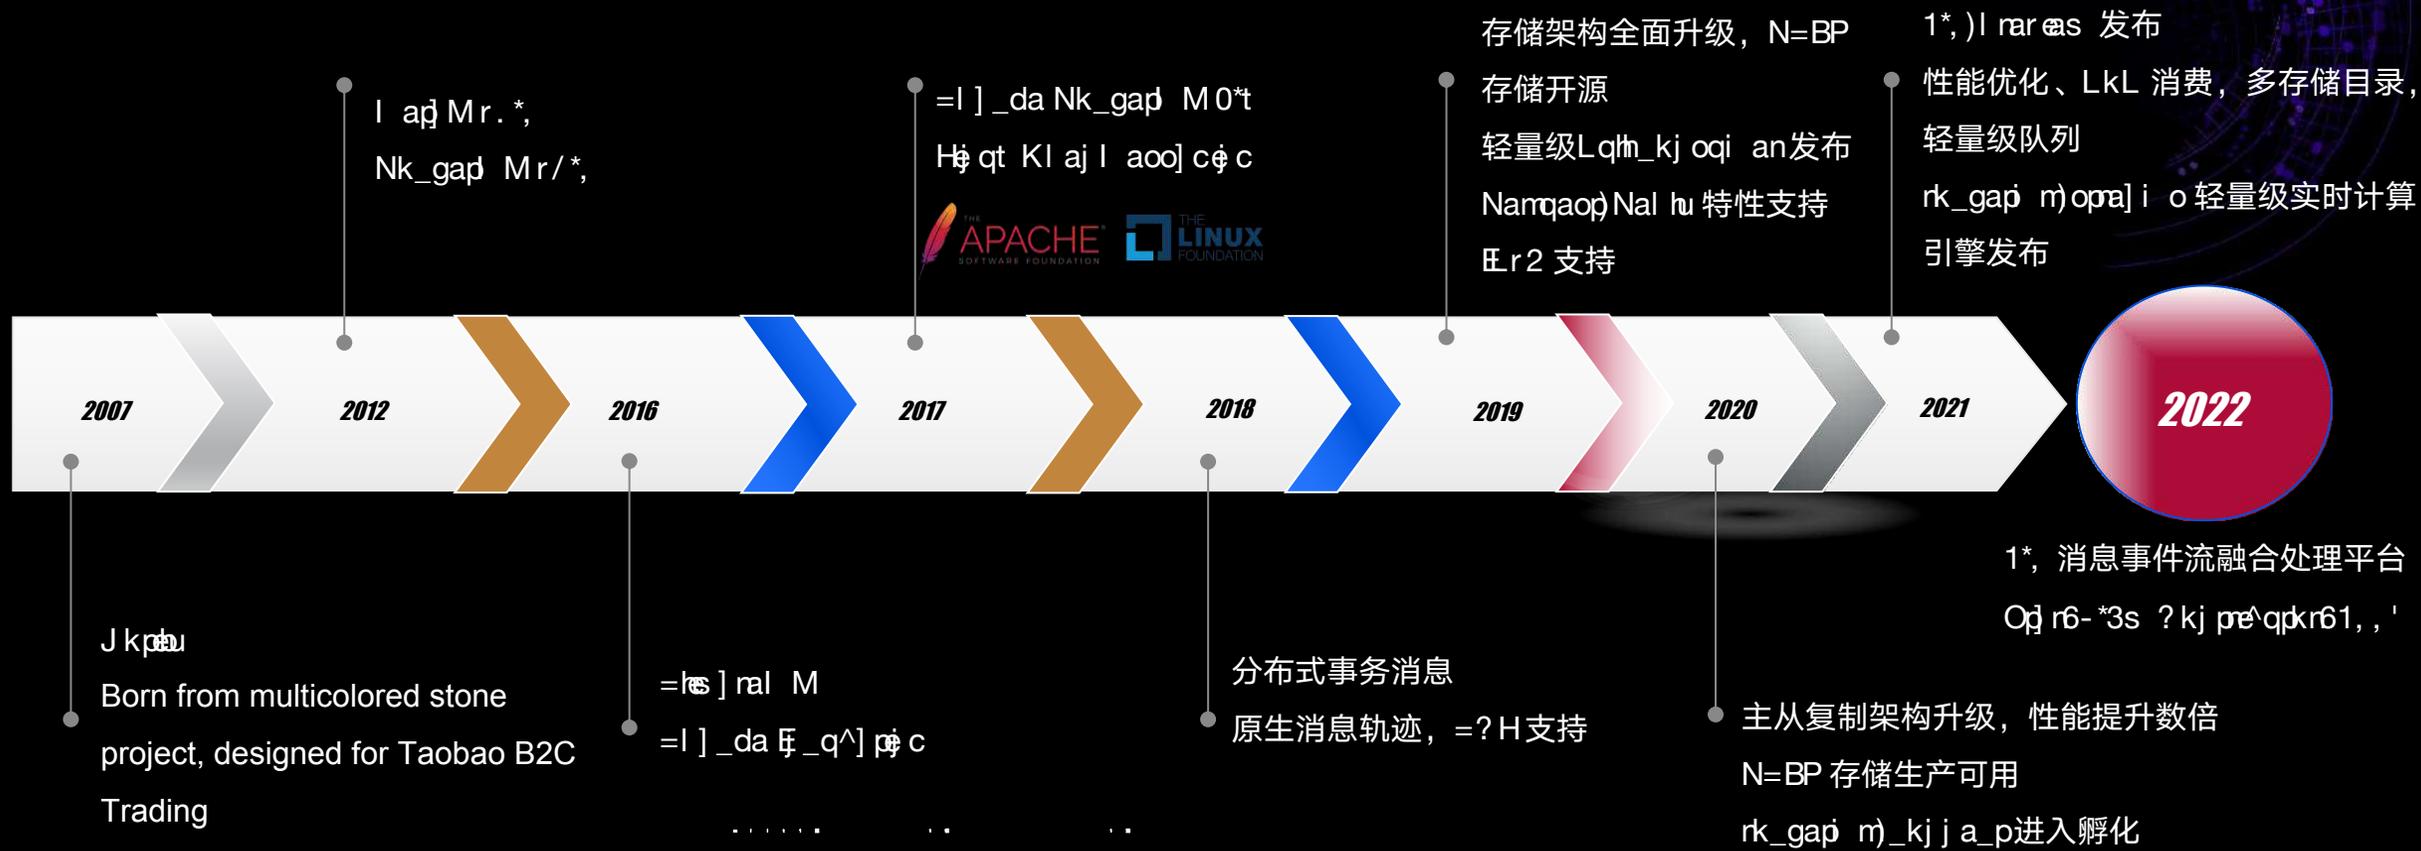
ACK 机制

金融级高可靠



APACHE
APISIX

RocketMQ 发展历程





APACHE
APISIX

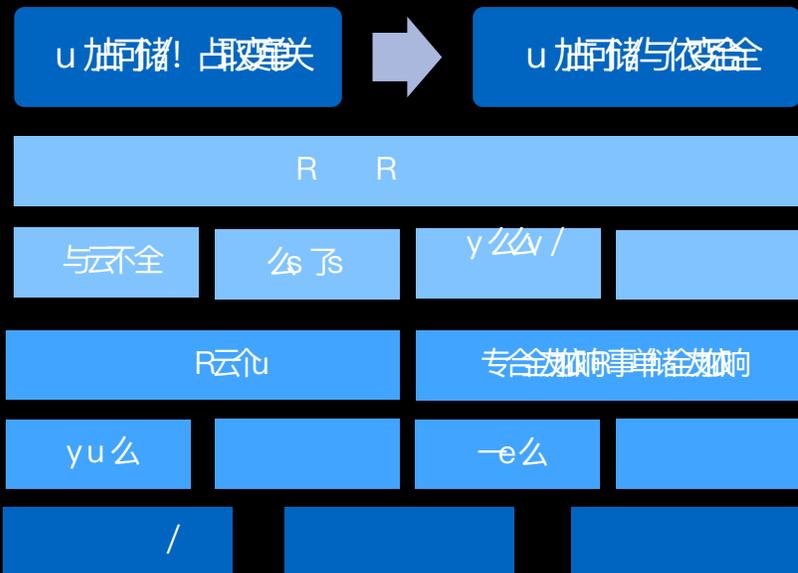


02 云原生时代的 Apache RocketMQ

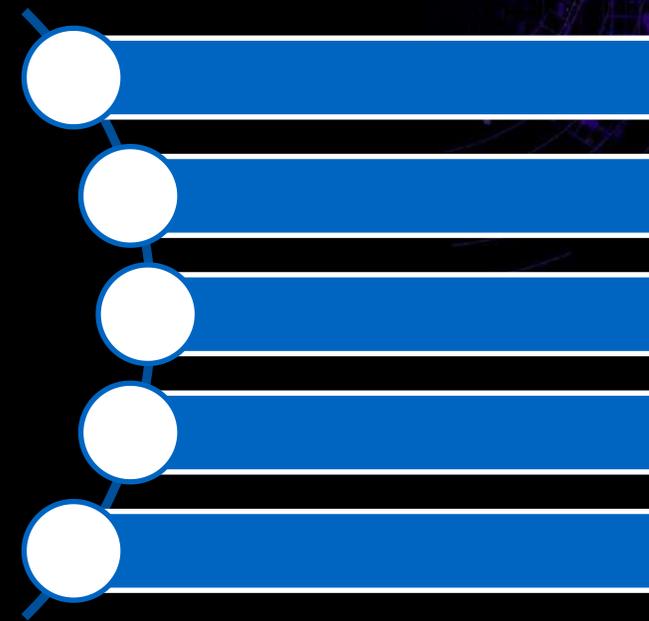


POWERED BY
APACHE
APISIX

云原生时代的挑战



-
-
- 依么
-
- /合取
-
- Q



客户价值:



高 SLA

X



低成本

X

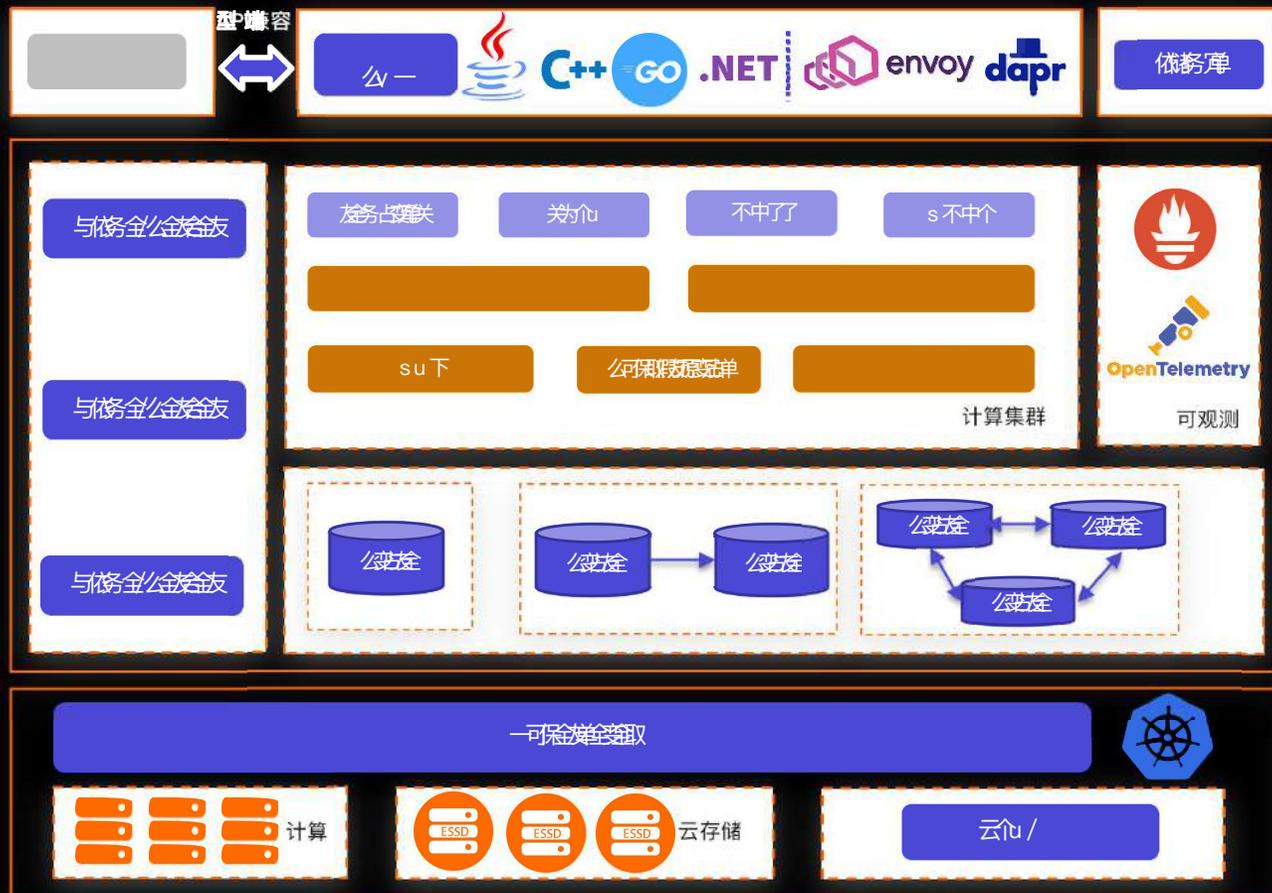


多场景



POWERED BY
APACHE
APISIX

云原生时代的 RocketMQ



高可用架构升级：

- 秒级故障转移，多场景容灾支持
- 无外部依赖，节点间松散耦合
- 自建及云上异构 IaaS 基础设施支持，降低成本

轻量级SDK：

- 全面支持云原生通信标准 gRPC 协议
- 无状态 Pop 消费模式，多语言友好，易集成

从业务走向数据：

- 事件流场景支撑
- 面向 SQL 的轻量级实时计算引擎

可分可合的存储计算分离：

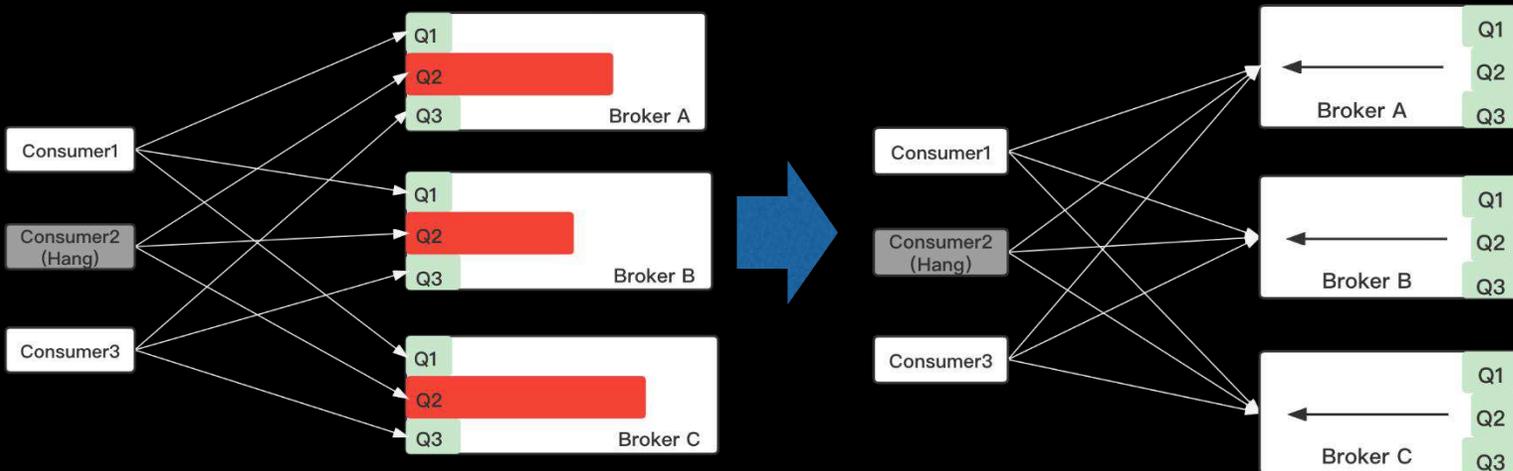
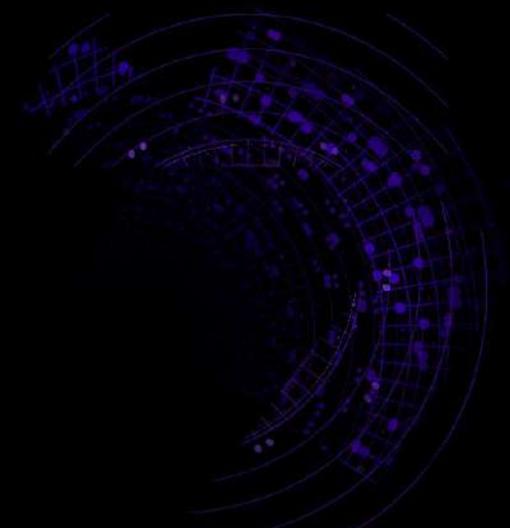
- Broker 升级为真正的无状态服务节点，无 binding
- Broker 和 Store 节点分离部署、独立扩缩
- 可分可合，适应多种业务场景，降低运维负担

云原生基础设施：

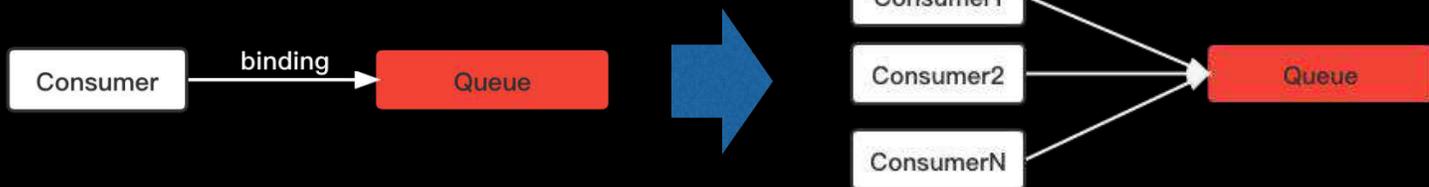
- 可观测性能力云原生，OpenTelemetry 标准化
- Kubernetes 一键式部署扩容交付



APACHE APISIX POP



- Consumer Queue Binding
- Queue
- rebalance
-
-



- Streaming
 - batch
- consumer



APACHE APISIX

无损弹性扩缩

固定分区使用场景

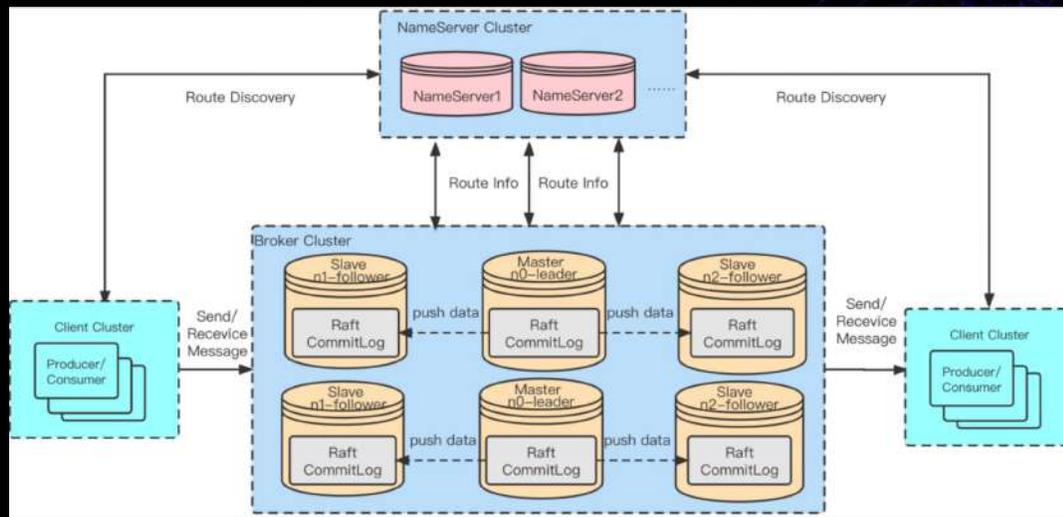
- 任务计算过程中，会将同一个业务类型的数据发到同一个队列
- Binlog 等数据同步过程中，需要保证严格顺序

固定分区面临的挑战

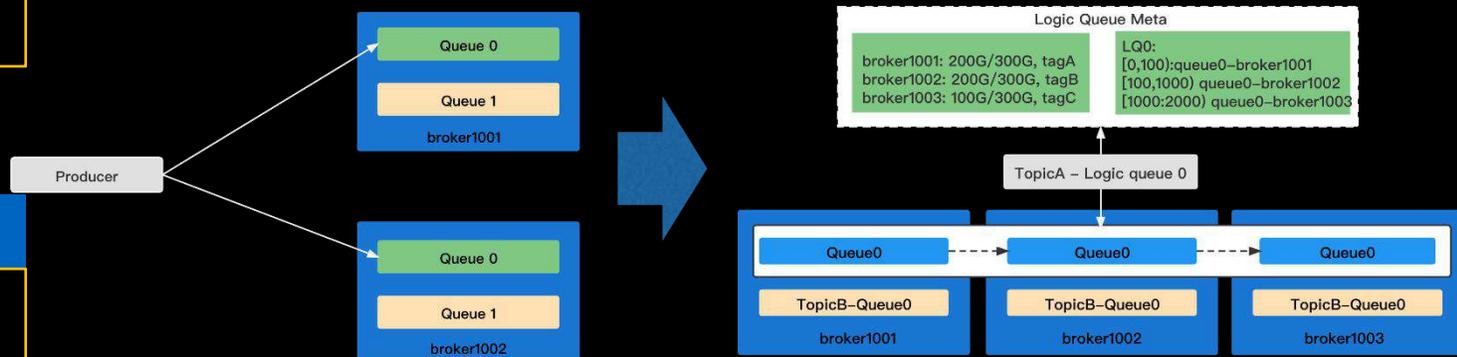
- 无切换架构中，主节点宕机，备节点不可写，分区数量减少
- 容量不足，Kafka 分区迁移时，会导致网络风暴，耗时极长

问题重点

- 在主节点宕机时，备节点要有自动切换为主的能力
- 容量调整时，不能产生数据迁移，且要在秒级完成



RAFT 存储支持：自动主从切换，强一致性保证



逻辑队列：秒级无损弹性扩缩，无数据复制，流量精准调度



APACHE
APISIX

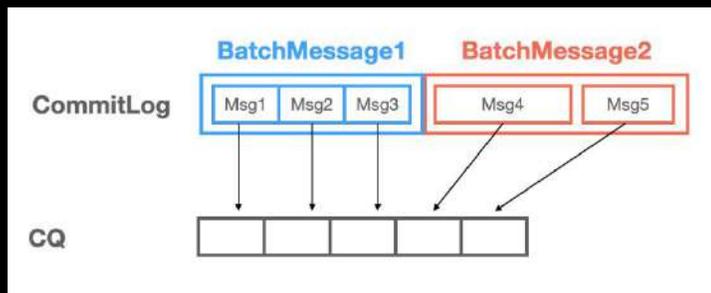
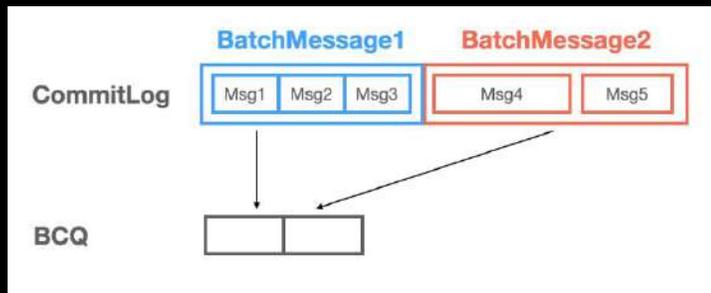
消息与流融合索引支持

核心问题

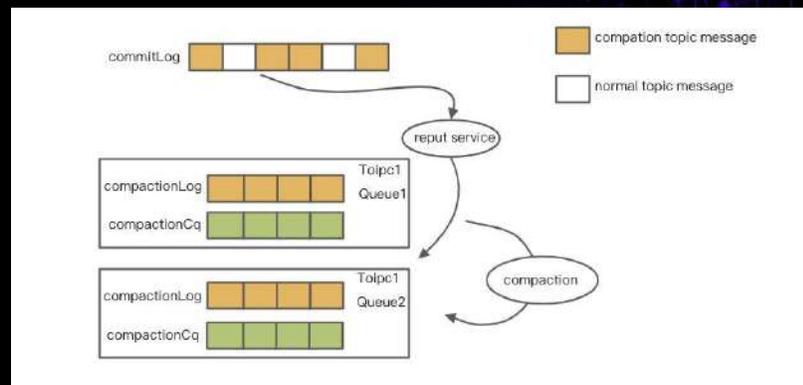
- 消息体小且存储结构面向单条消息，导致提升吞吐困难
- 在进行状态存储时，无 KV 语义支撑

挑战

- Commitlog 格式存储，统一复制算法
- 不借助外部依赖，提供状态存储能力



- 单一数据文件，支持面向流的索引及面向批的索引
- 批量发，批量存，批量读，吞吐提升十倍



Offset	0	1	2	3	4	5	6	7	8	9	10	
Key	K1	K2	K1	K1	K3	K2	K4	K5	K5	K2	K6	Log Before Compaction
Value	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	
Compaction												
				3	4	6	8	9	10			Log After Compaction
Keys				K1	K3	K4	K5	K2	K6			
Values				V4	V5	V7	V9	V10	V11			

- 统一消息，KV 语义，统一 API，提供状态存储能力
- 流表二象性支持



APACHE
APISIX



03 借力 APISIX 构建云原生接入体系

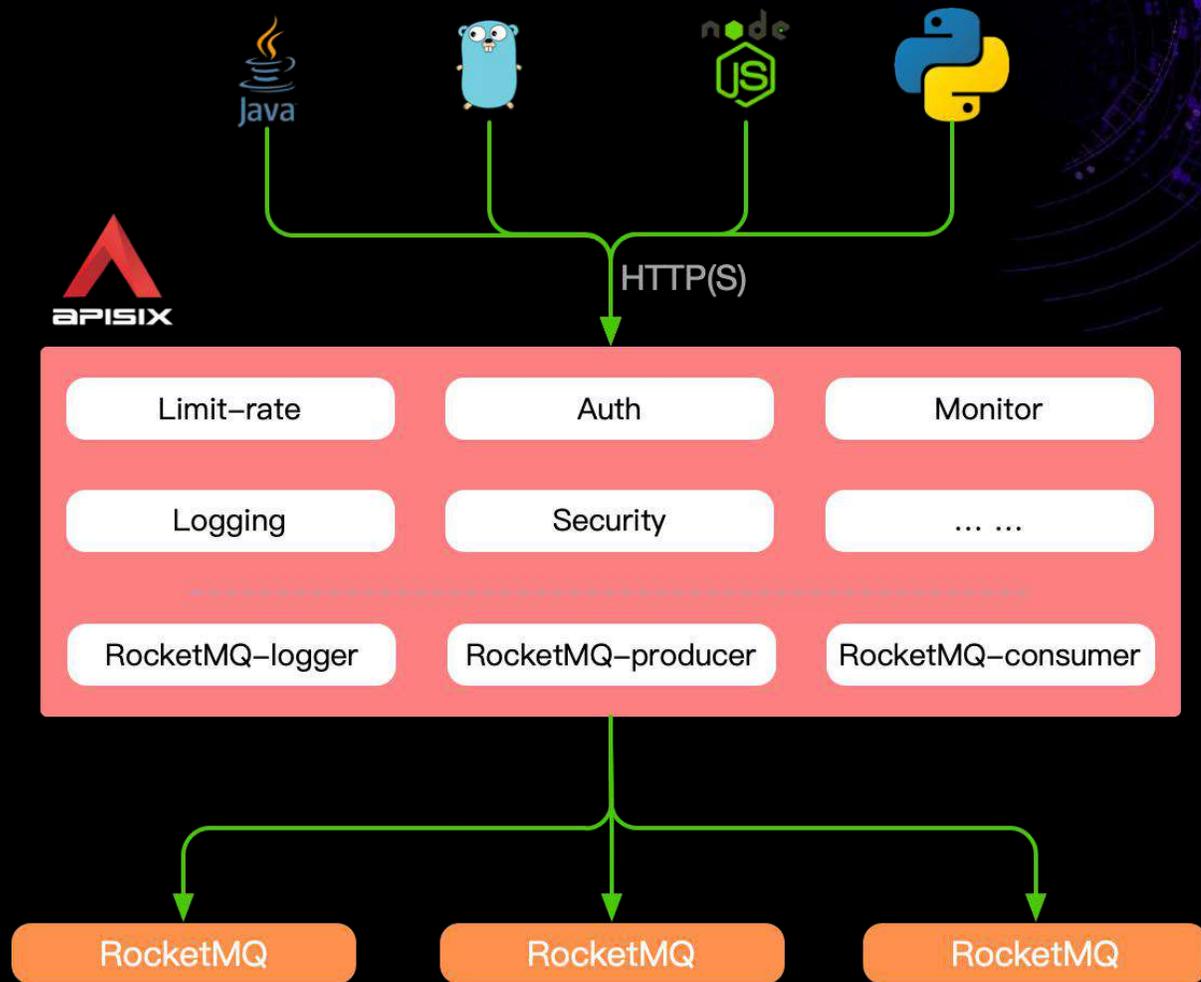


APACHE
APISIX

借力 APISIX 构建云原生接入体系

RocketMQ Gateway

- logger
- limit-count
- 安全





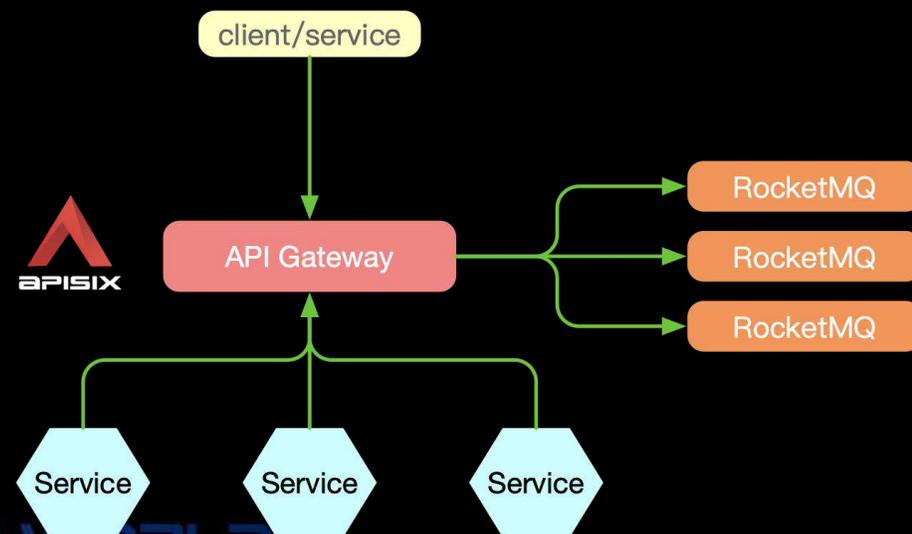
APACHE
APISIX

借力 APISIX 构建云原生接入体系

```
# Producer
```

```
$ curl http://127.0.0.1:9080/apisix/admin/routes/1 -H 'X-API-KEY: YOUR_KEY' -X PUT -d '{
  "plugins": {
    "rocketmq-producer": {
      "nameserver_list": ["127.0.0.1:9876"]
    }
  },
  "uri": "/rocketmq/producer/*"
}'
```

```
$ curl http://127.0.0.1:9080/rocketmq/producer/topic1 \
-H 'Content-Type: application/json' \
-d '{"orderId": "100"}'
```





POWERED BY
APACHE
APISIX

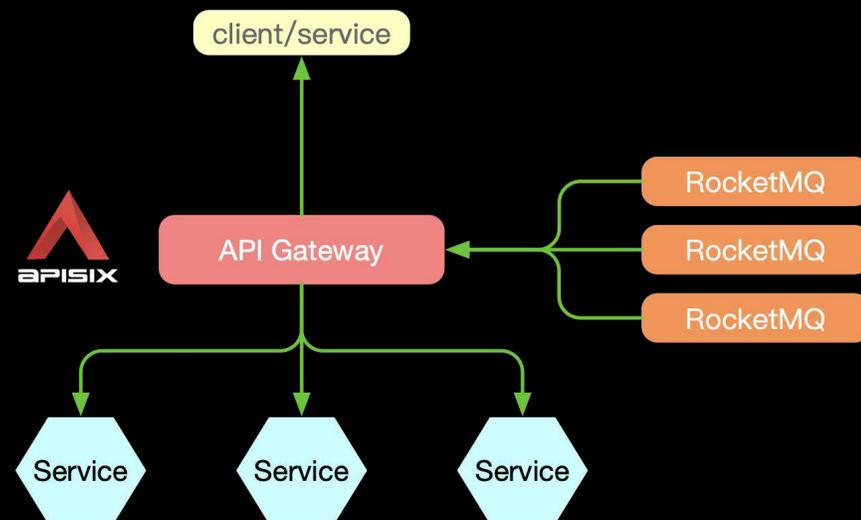
借力 APISIX 构建云原生接入体系

```
# Consumer
```

```
$ curl http://127.0.0.1:9080/apisix/admin/routes/1 -H 'X-API-KEY: YOUR_KEY' -X PUT -d '{  
  "plugins": {  
    "rocketmq-consumer": {  
      "nameserver_list": ["127.0.0.1:9876"],  
      "consumer_group": "test_group"  
    }  
  },  
  "uri": "/rocketmq/consumer/*"  
}'
```

```
$ curl http://127.0.0.1:9080/rocketmq/consumer/topic1 \\  
-H 'Content-Type: application/json'
```

```
{"orderId": "100"}
```



APACHE APISIX CONNECTS THE WORLD



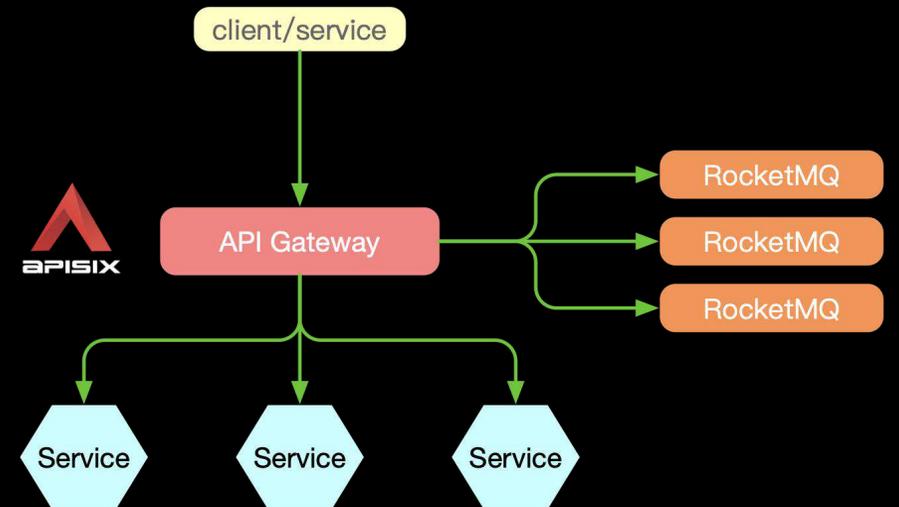
APACHE
APISIX

借力 APISIX 构建云原生接入体系

```
# logger
```

```
$ curl http://127.0.0.1:9080/apisix/admin/routes/1 -H 'X-API-KEY: YOUR_KEY' -X PUT -d '{  
  "plugins": {  
    "rocketmq-logger": {  
      "nameserver_list": ["127.0.0.1:9876"],  
      "topic": "log"  
    }  
  },  
  "upstream": {  
    "nodes": {  
      "127.0.0.1:1980": 1  
    },  
    "type": "roundrobin"  
  },  
  "uri": "/foo/bar/*"  
}'
```

```
$ curl http://127.0.0.1:9080/foo/bar/abc
```



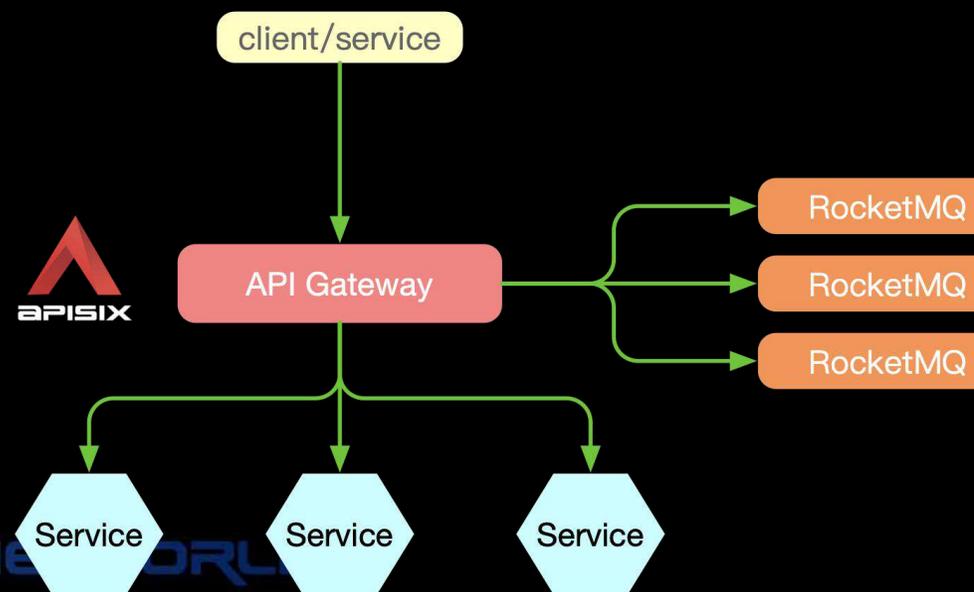


APACHE
APISIX

借力 APISIX 构建云原生接入体系

```
# logger + limit-count
```

```
$ curl http://127.0.0.1:9080/apisix/admin/routes/1 -H 'X-API-KEY: YOUR_KEY' -X PUT -d '{  
  "plugins": {  
    "rocketmq-logger": {  
      "nameserver_list" : ["127.0.0.1:9876"],  
      "topic": "log"  
    },  
    "limit-count": {  
      "count": 2,  
      "time_window": 60,  
      "rejected_code": 503,  
      "key_type": "var",  
      "key": "remote_addr"  
    }  
  },  
  "upstream": {  
    ... ..  
  },  
  "uri": "/foo/bar/*"  
}'
```



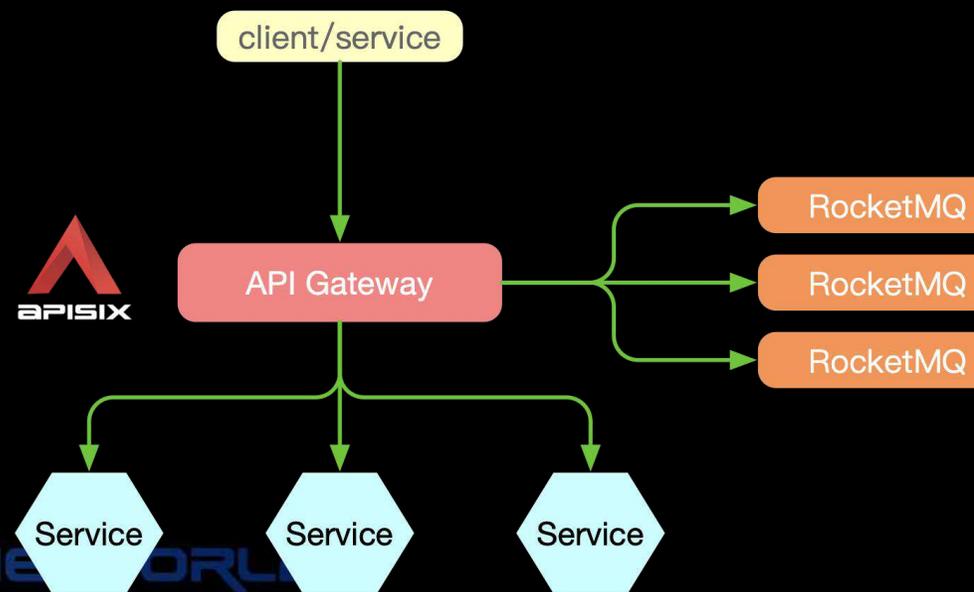


APACHE
APISIX

借力 APISIX 构建云原生接入体系

```
# logger + limit-count
```

```
$ curl http://127.0.0.1:9080/apisix/admin/routes/1 -H 'X-API-KEY: YOUR_KEY' -X PUT -d '{  
  "plugins": {  
    "rocketmq-logger": {  
      "nameserver_list" : ["127.0.0.1:9876"],  
      "topic": "log"  
    },  
    "limit-count": {  
      "count": 2,  
      "time_window": 60,  
      "rejected_code": 503,  
      "key_type": "var",  
      "key": "remote_addr"  
    }  
  },  
  "upstream": {  
    ... ..  
  },  
  "uri": "/foo/bar/*"  
}'
```





APACHE
APISIX

借力 APISIX 构建云原生接入体系

作者：余洲

RocketMQ committer

Apache APISIX contributor

<https://github.com/yuz10/lua-resty-rocketmq>



小火箭

浙江 杭州



扫一扫上面的二维码图案，加我为朋友



APACHE
APISIX



感谢聆听
THANKS

APACHE APISIX CONNECTS THE WORLD