

**DataStax**

# **Starlight-for-RabbitMQ, powered by Apache Pulsar**

**Christophe Bornet, Apache Pulsar committer**  
**ApacheCon North America 2022**

October 6, 2022


# Why Starlight-for-RabbitMQ ?

- RabbitMQ is popular in the Message Queue world
- Common troubles with RabbitMQ
  - Scaling
  - Resilience and durability
  - Throughput and predictability
  - Erlang
- Migration from RabbitMQ to Apache Pulsar

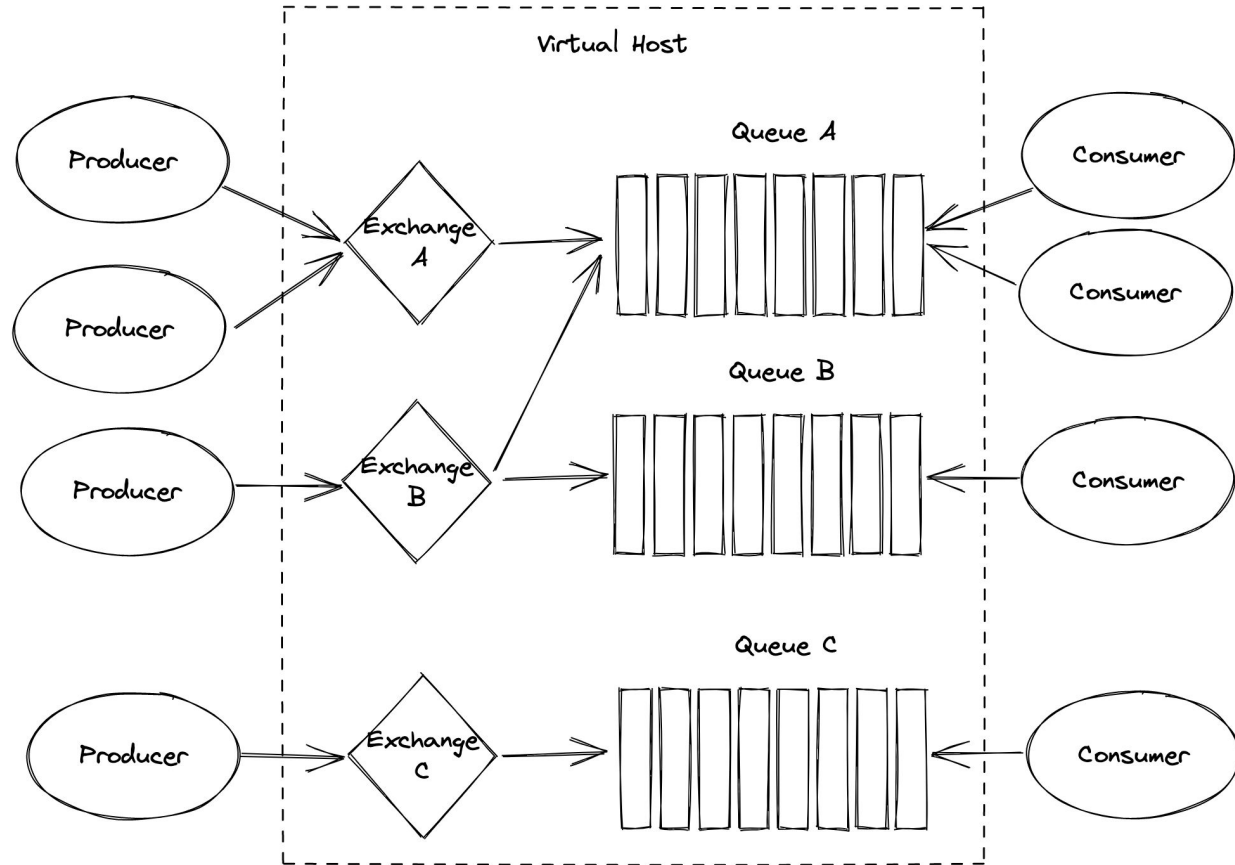


# AMQP 0.9.1 concepts (RabbitMQ)

# Exchanges and queues

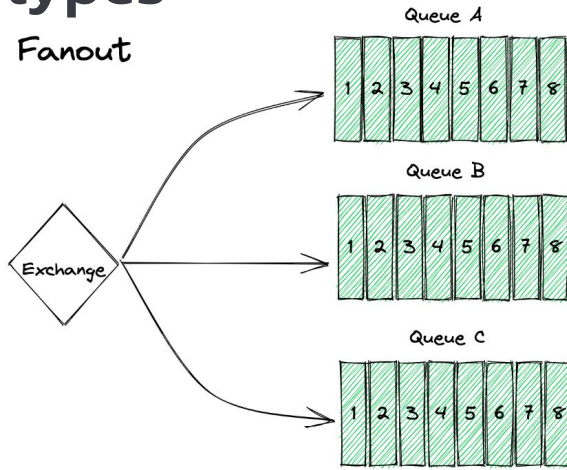


# Virtual Hosts

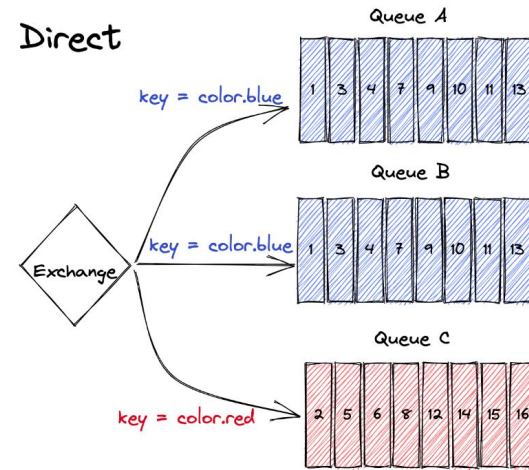


# Exchange types

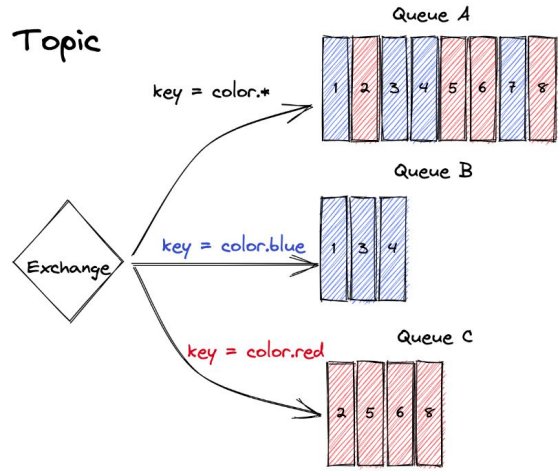
## Fanout



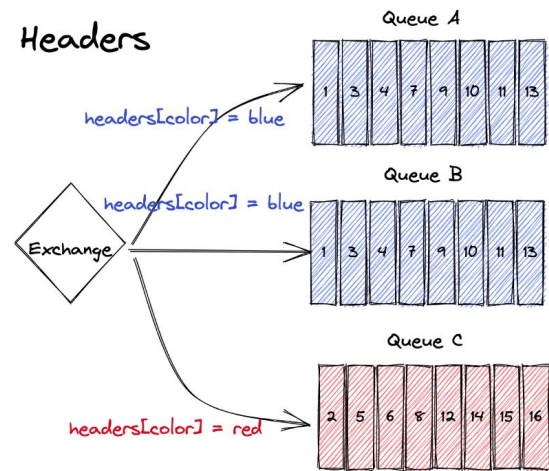
## Direct



## Topic



## Headers



# At least once delivery

- Durable exchanges and queues
- Publisher confirms
- Consumer acknowledgements
- Persistent messages
  - Messages are removed from queues as soon as they are acknowledged
  - Disk writes are batched and flushed every 200ms
  - A persistent message might not go to disk if it's acknowledged before the batch flush.





DS

# Pulsar concepts



# Pulsar distributed architecture

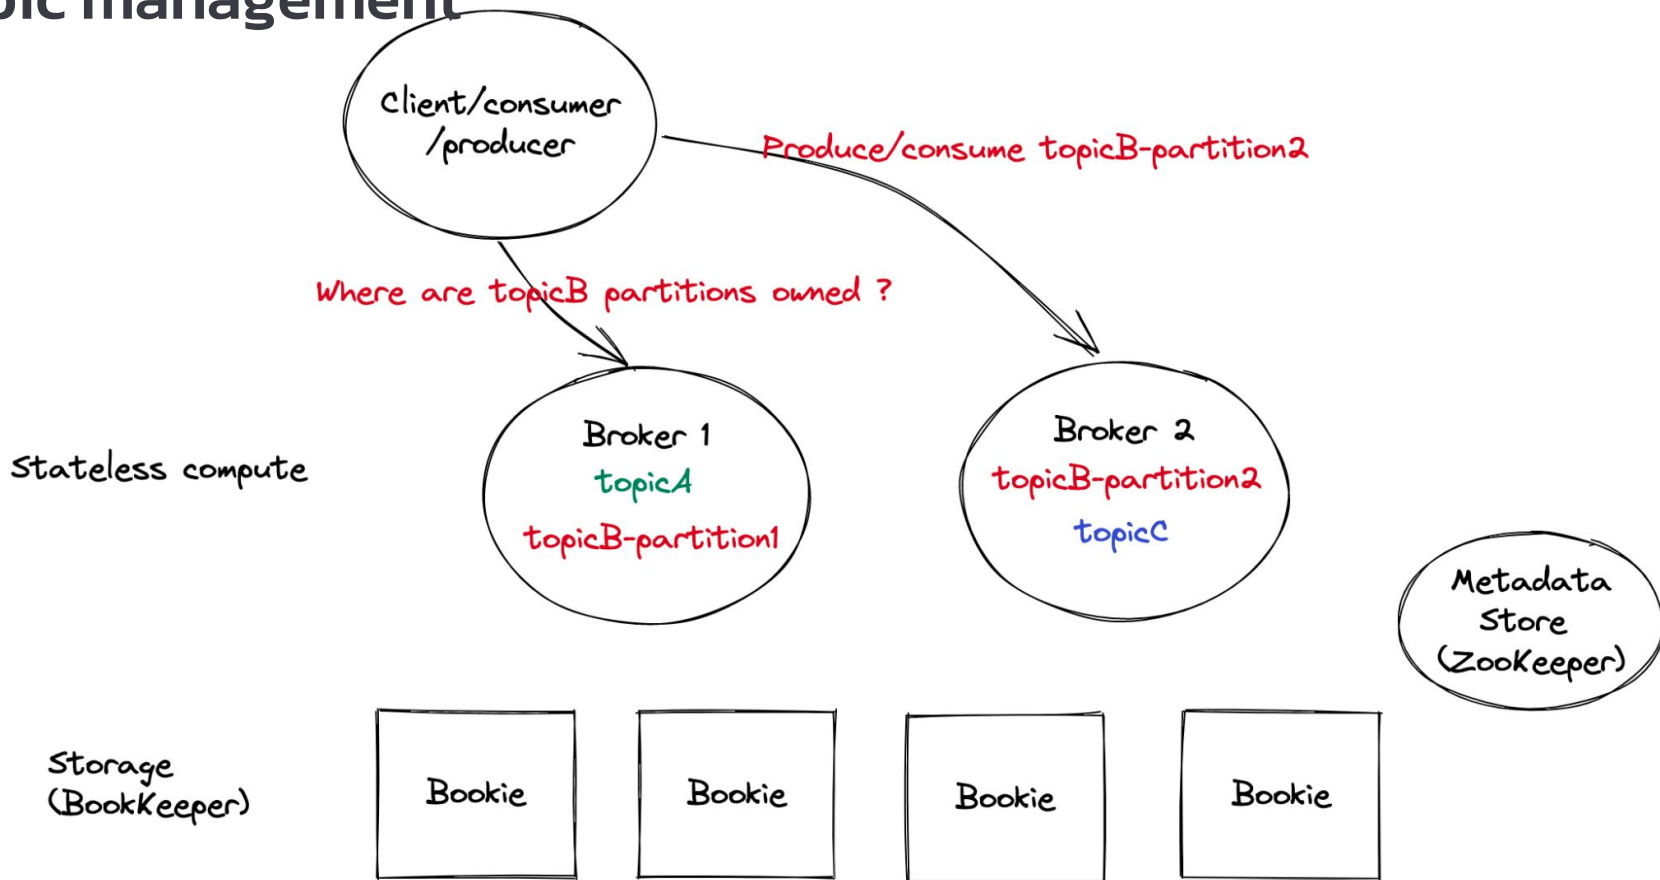
Stateless compute



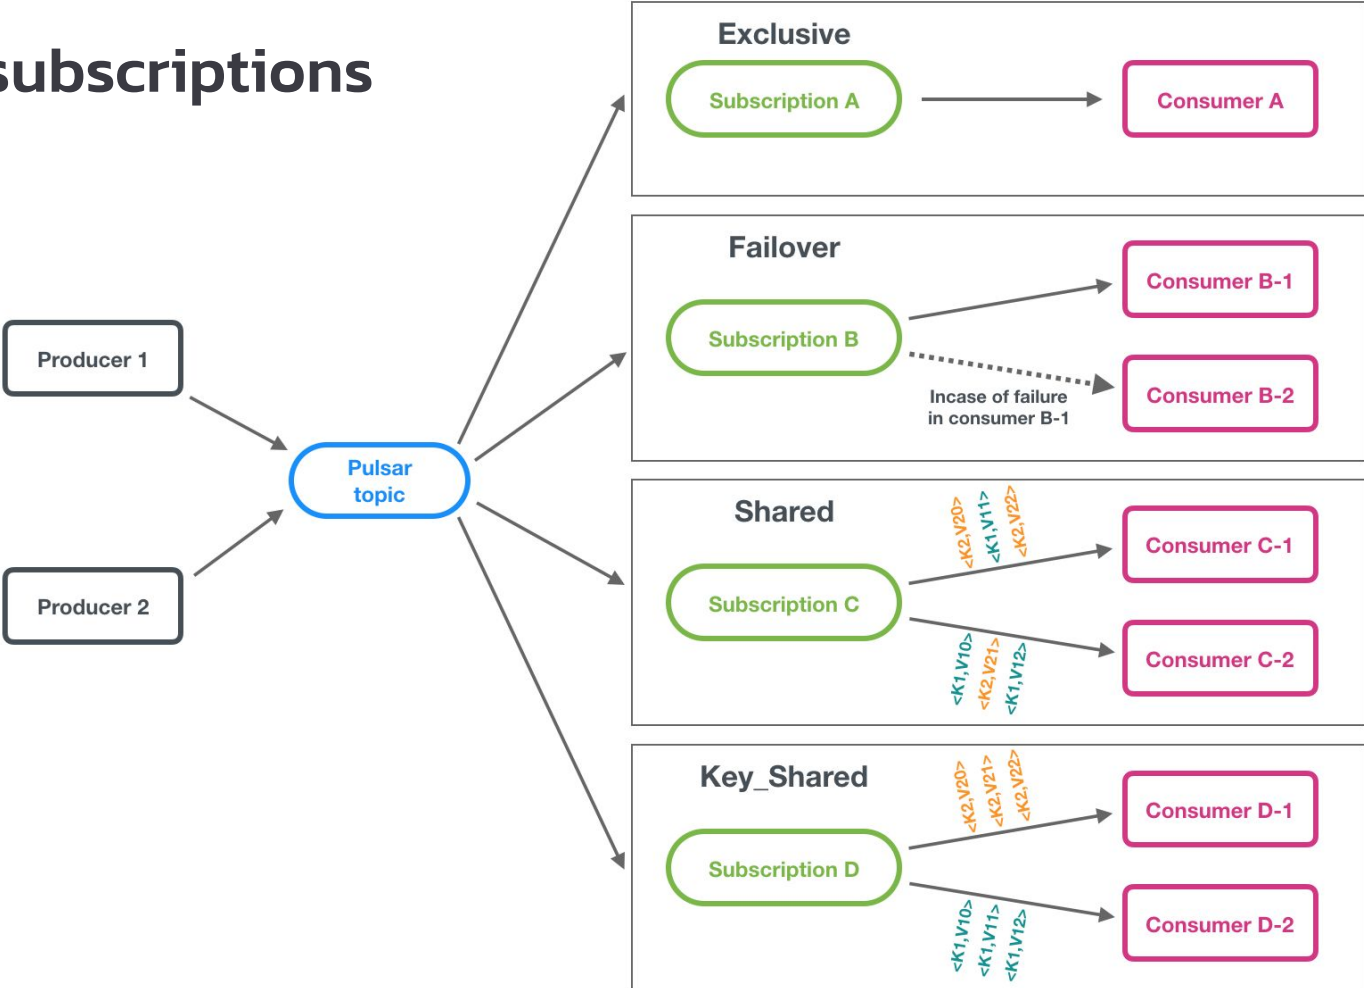
Storage (BookKeeper)



# Topic management



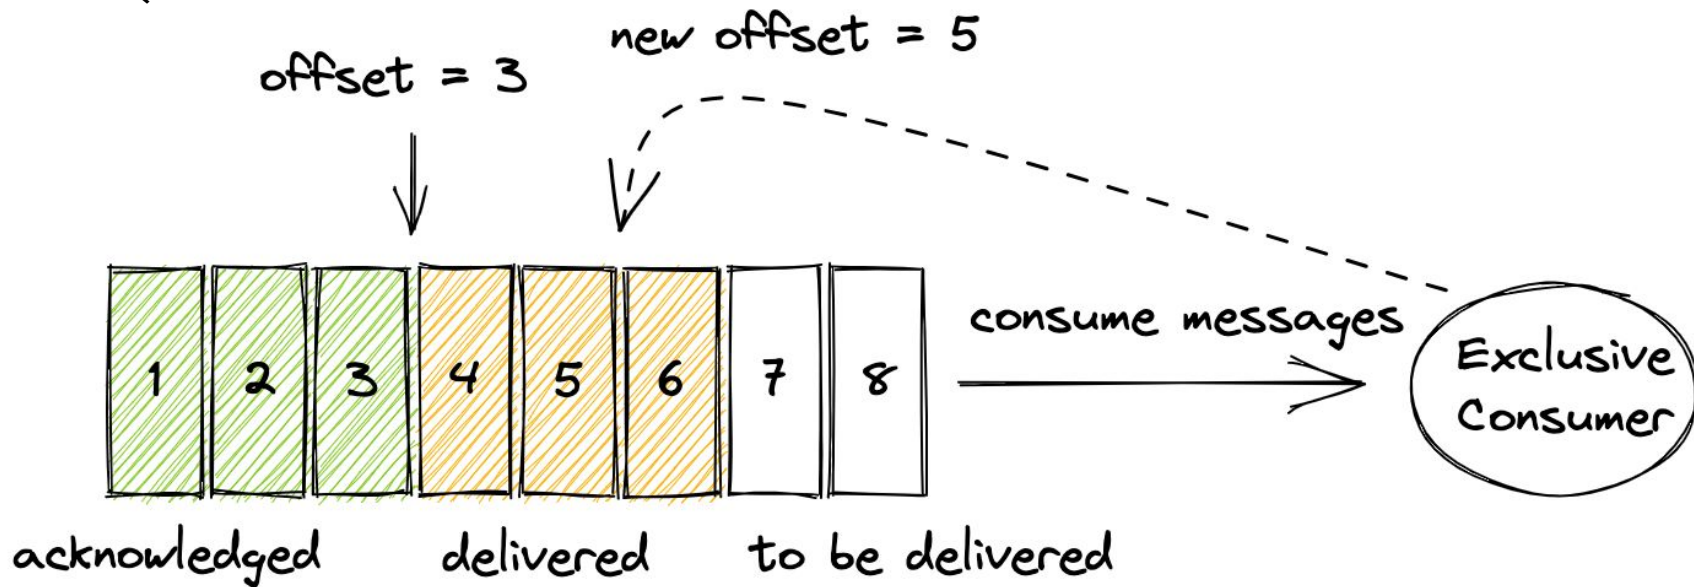
# Topic subscriptions



# Message Queue or Streaming ?

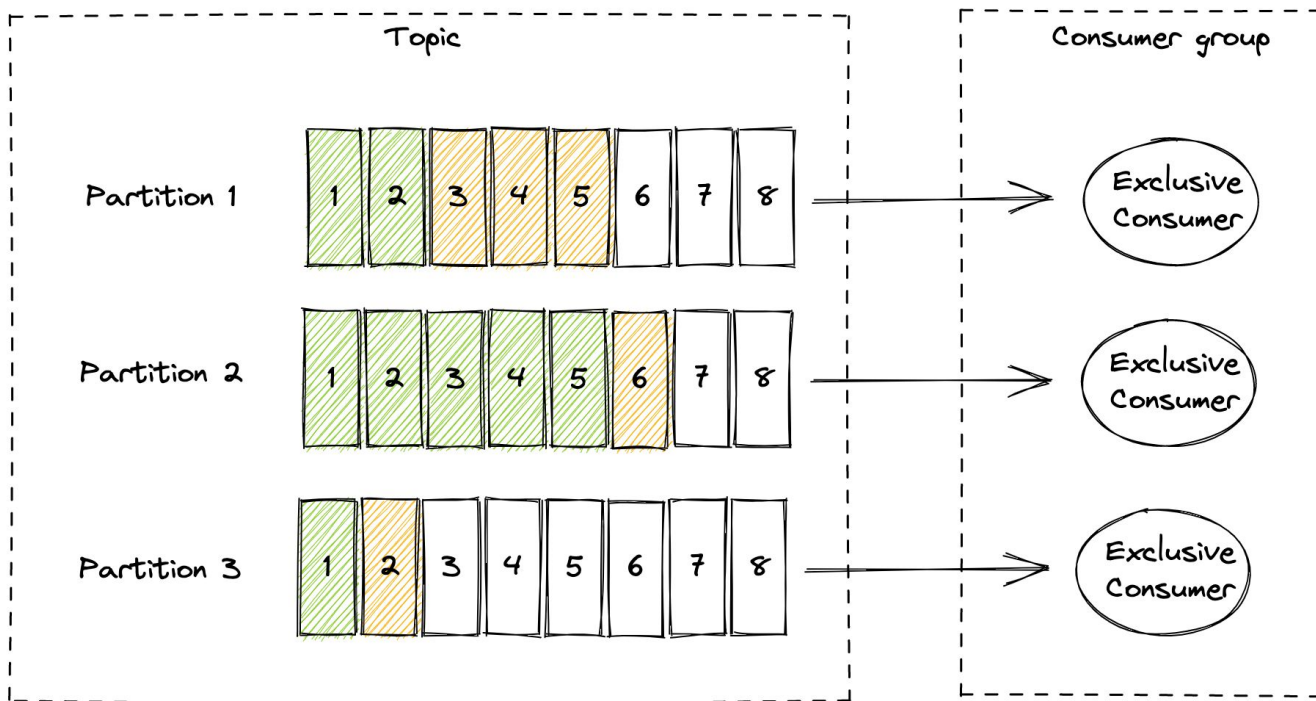
# Streaming = in-order exclusive consumer

- Pulsar Exclusive subscription
- Kafka
- RabbitMQ



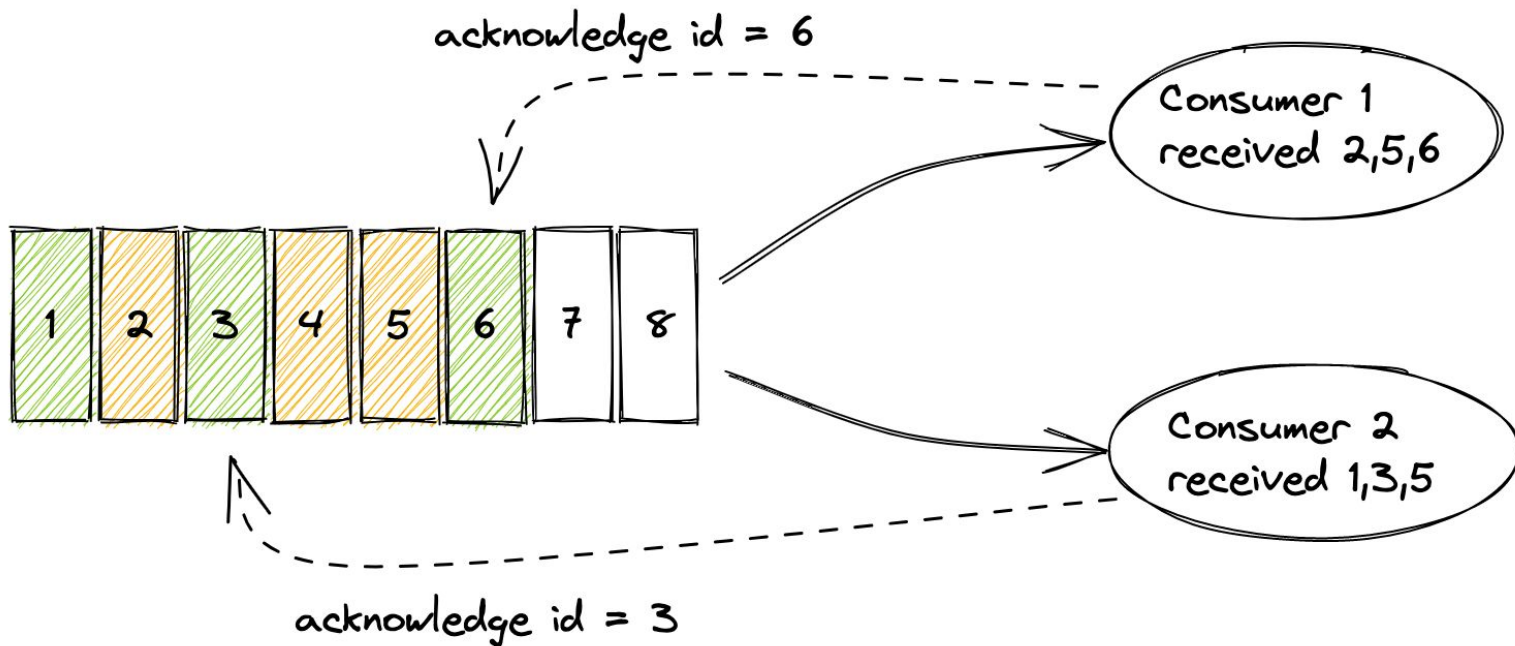
# Streaming with partitioned topics

- Pulsar Failover subscription
- Kafka

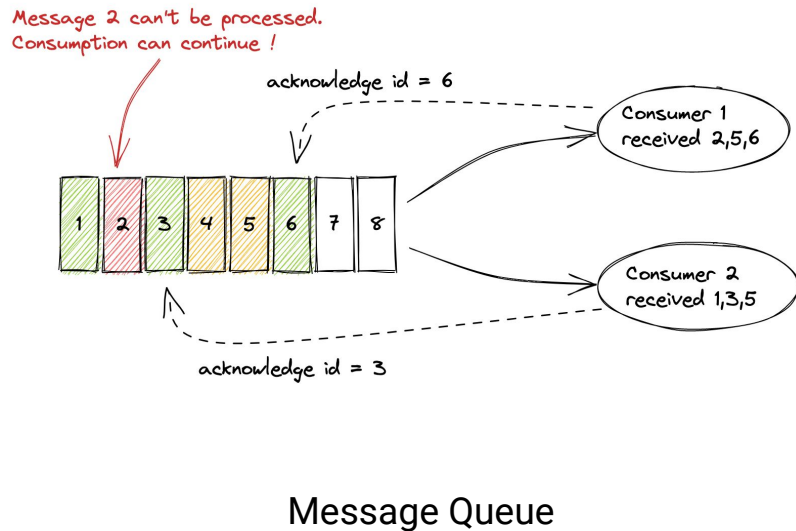
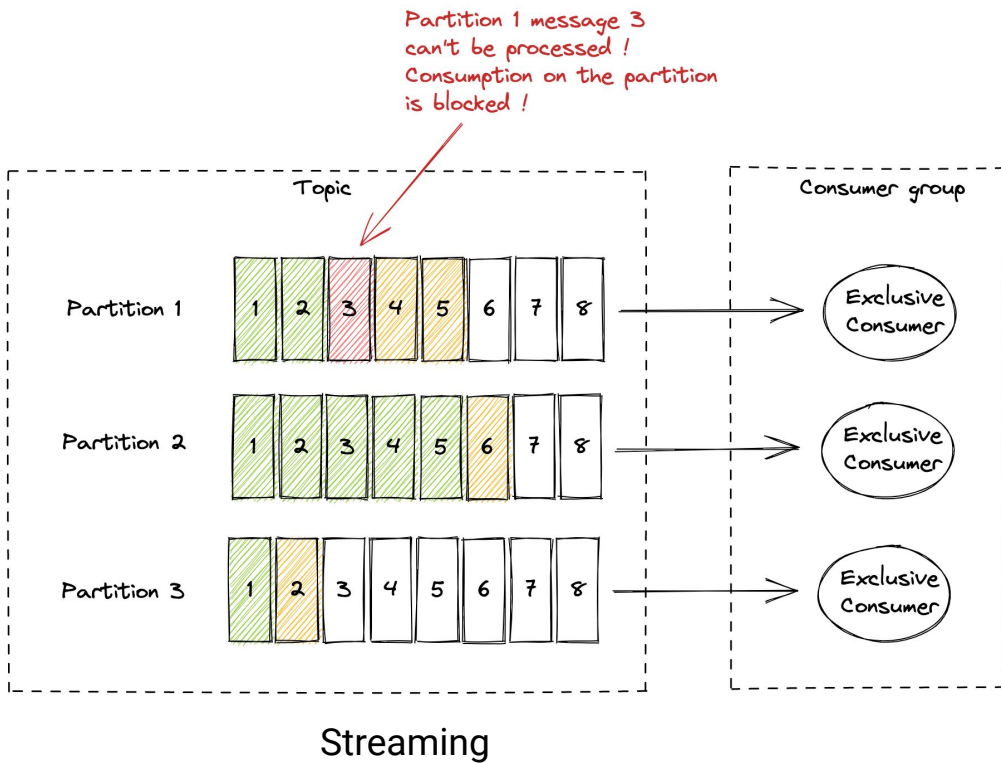


# Message Queue = unordered competing consumers

- Pulsar Shared subscription
- RabbitMQ



# Individual acknowledgements are essential for competing consumers





# How do we map Pulsar to RabbitMQ ?

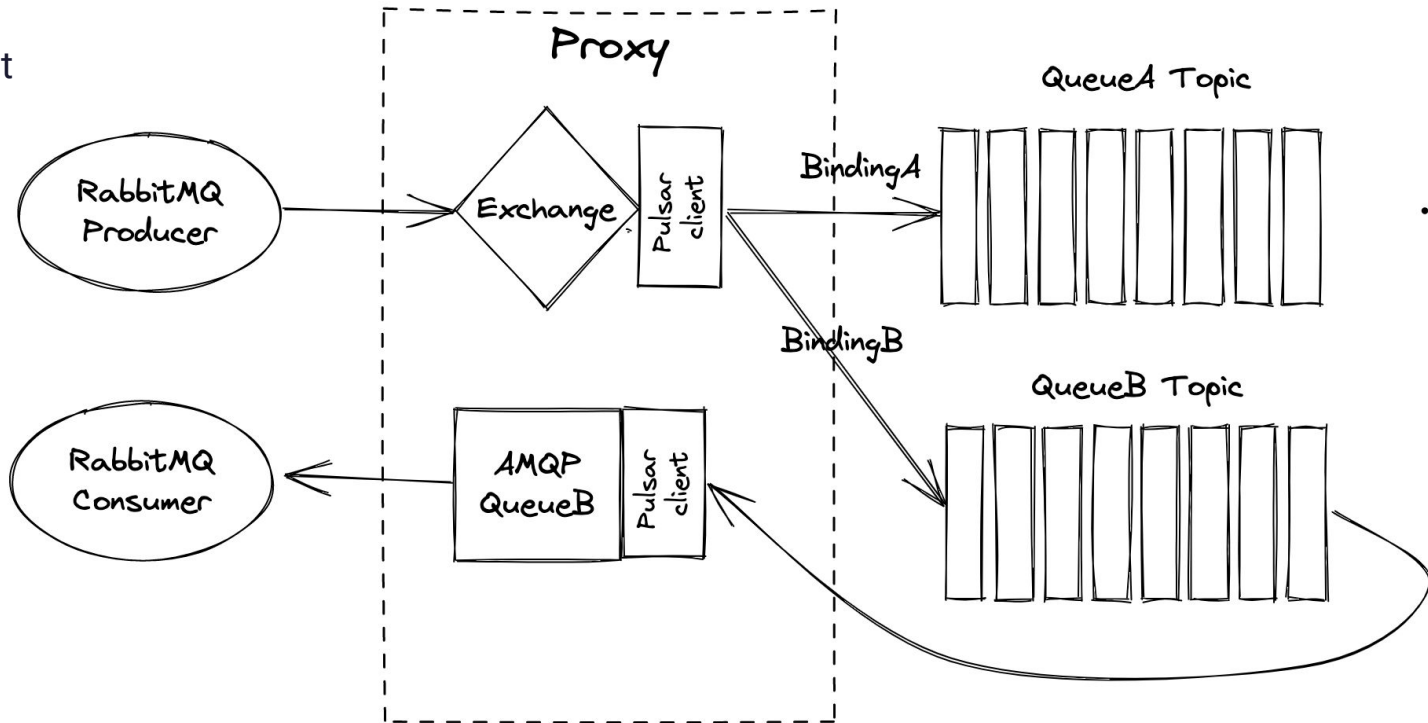
# Solution #1: copy full message to all bound queues

## Pros

- “Simple” to implement
- Scales with the number of brokers
- Pulsar client optimizations (batching, ...)

## Cons

- Not efficient in I/O and disk usage
- Network hop due to proxy



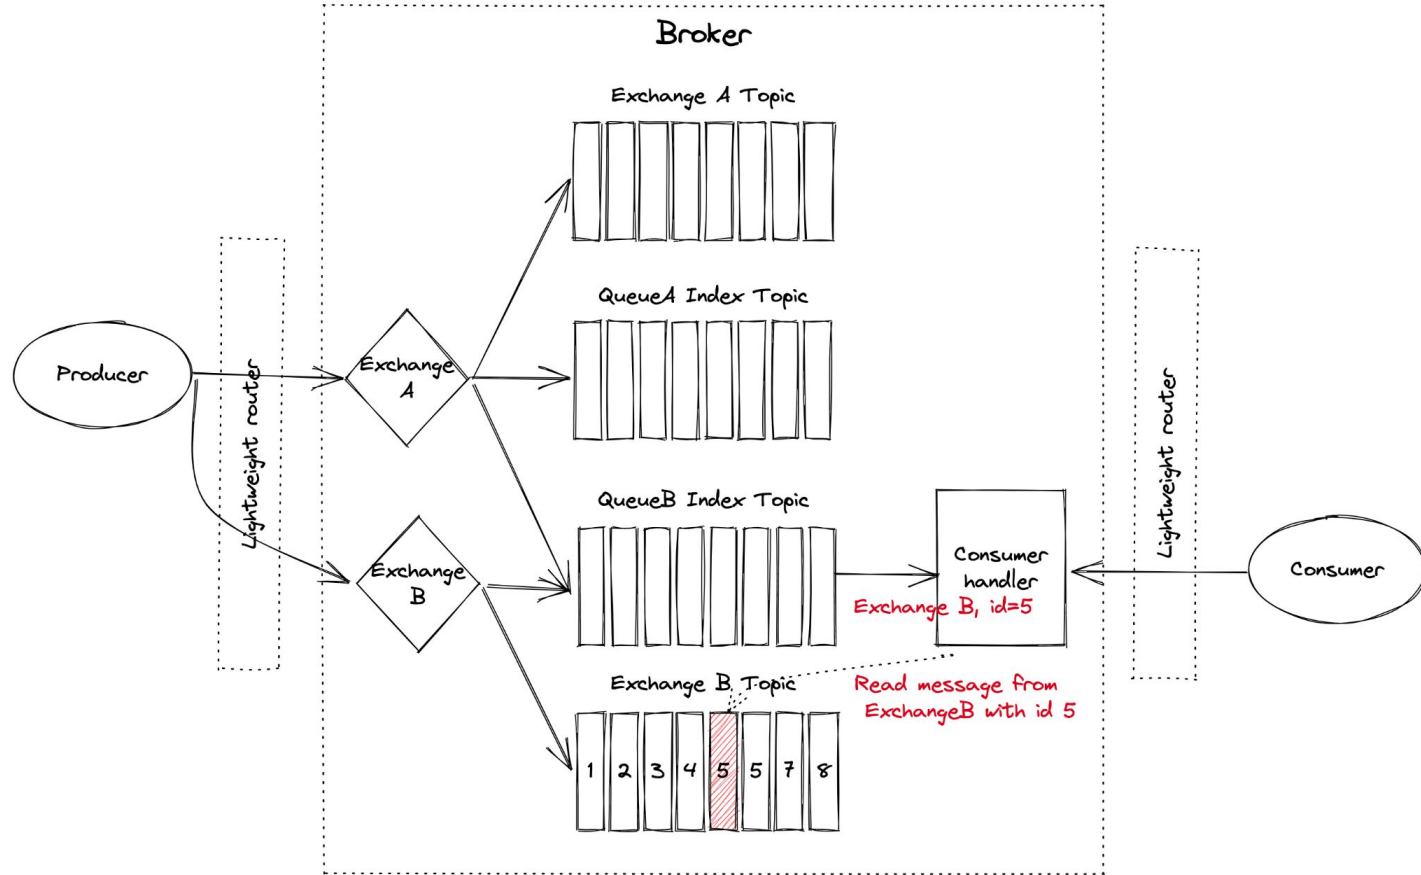
# Solution #2: copy message index to all bound queues (AoP)

## Pros

- Direct read/write to ledger is performant
- Zero-copy techniques can be used

## Cons

- Can't scale a VHost as all Exchanges and Queues for a given VHost must be owned by the same broker
- Network hop due to proxy/router



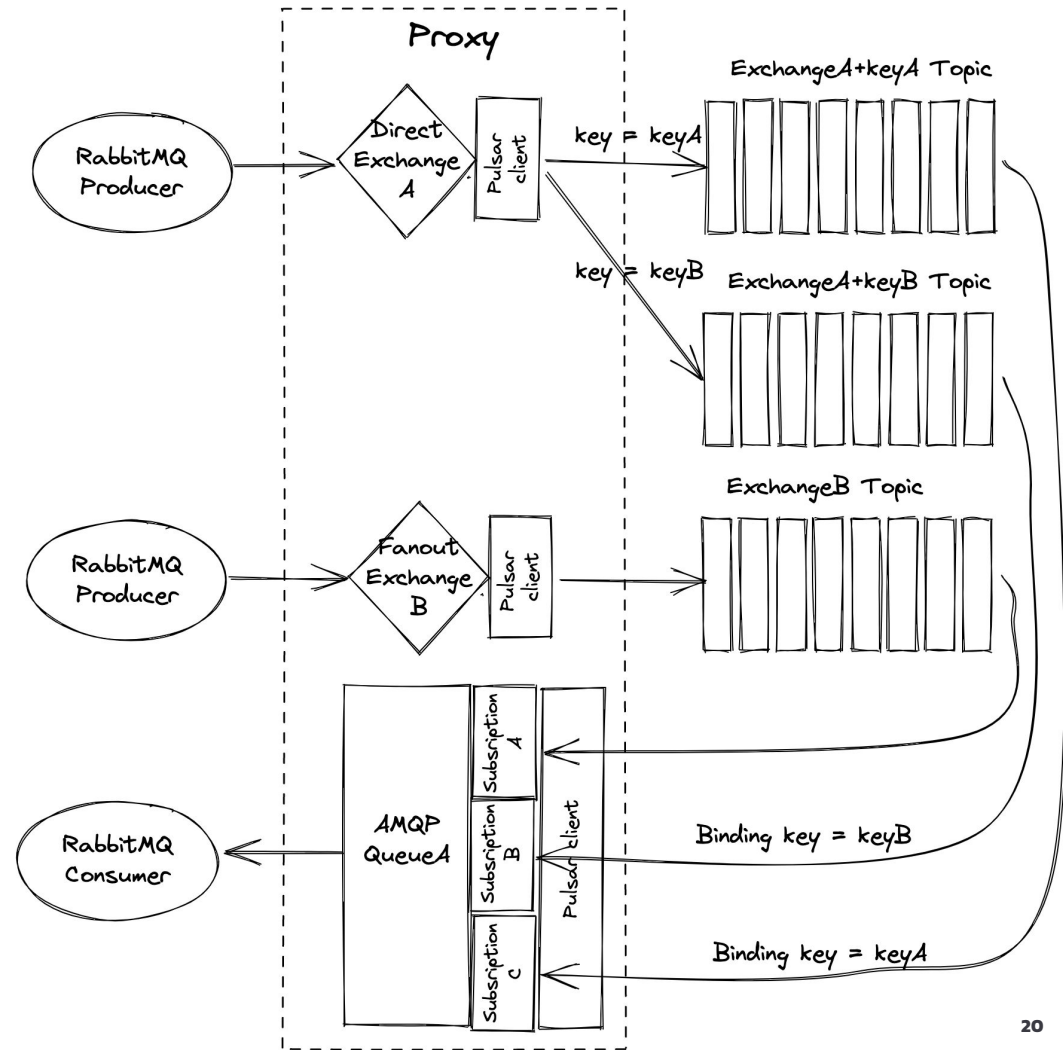
# Solution #3: use subscriptions as bindings (Starlight-for-RabbitMQ)

## Pros

- Scales with the number of brokers thanks to Pulsar Shared subscriptions
- Pulsar client optimizations (batching, ...)

## Cons

- More memory copies than reading/writing directly from the ledger
- Network hop due to proxy



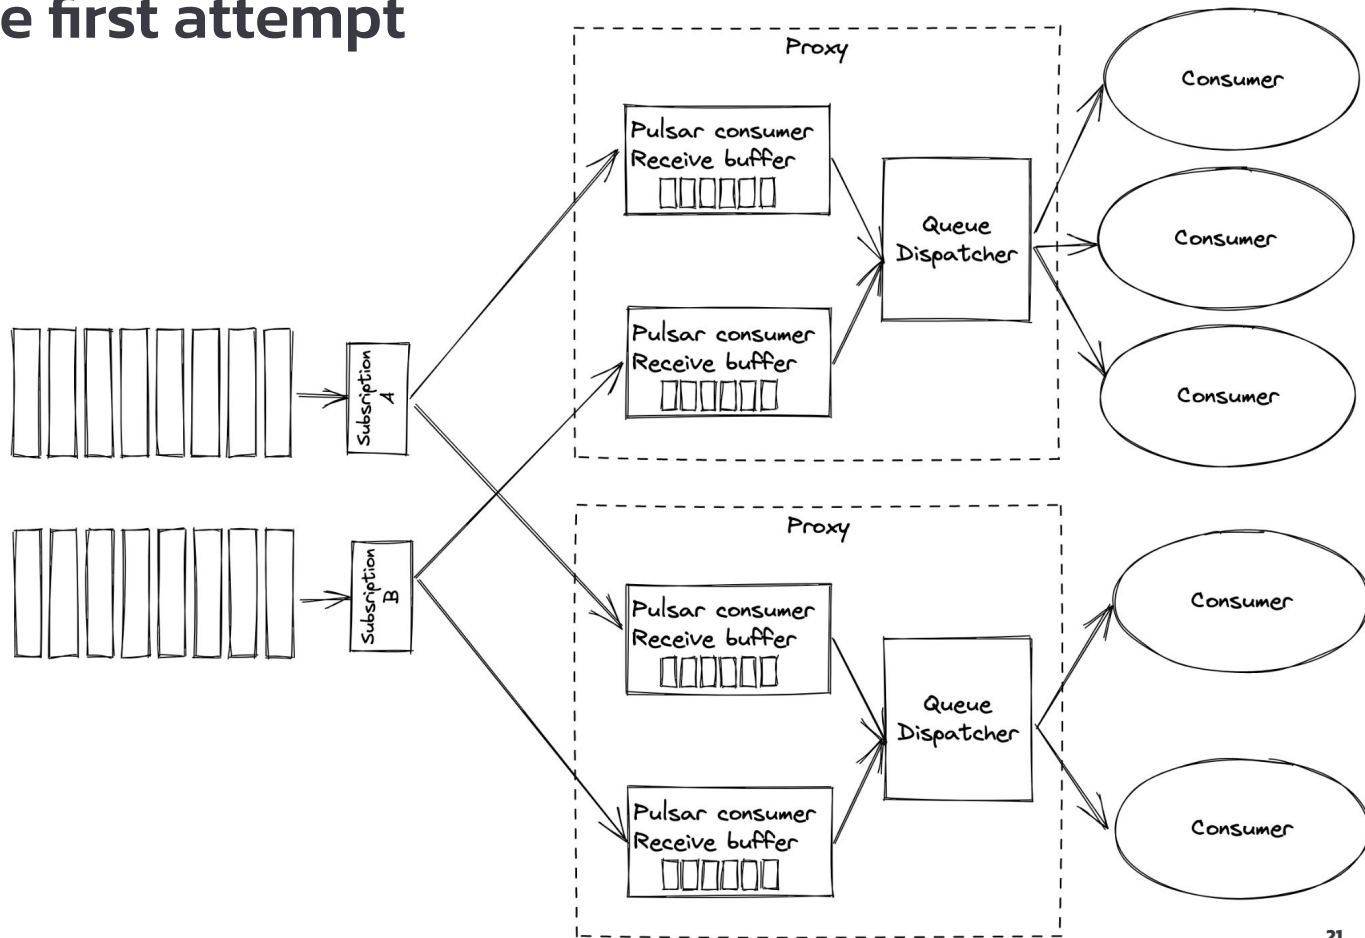
# Consuming-side first attempt

## Pros

- The lifecycle of the channel and consumer can be very AMQP 0.9.1 spec compliant.

## Cons

- Uneven distribution of messages per consumer.
- Messages can get stuck in the proxy receive buffers if there are no consumers.



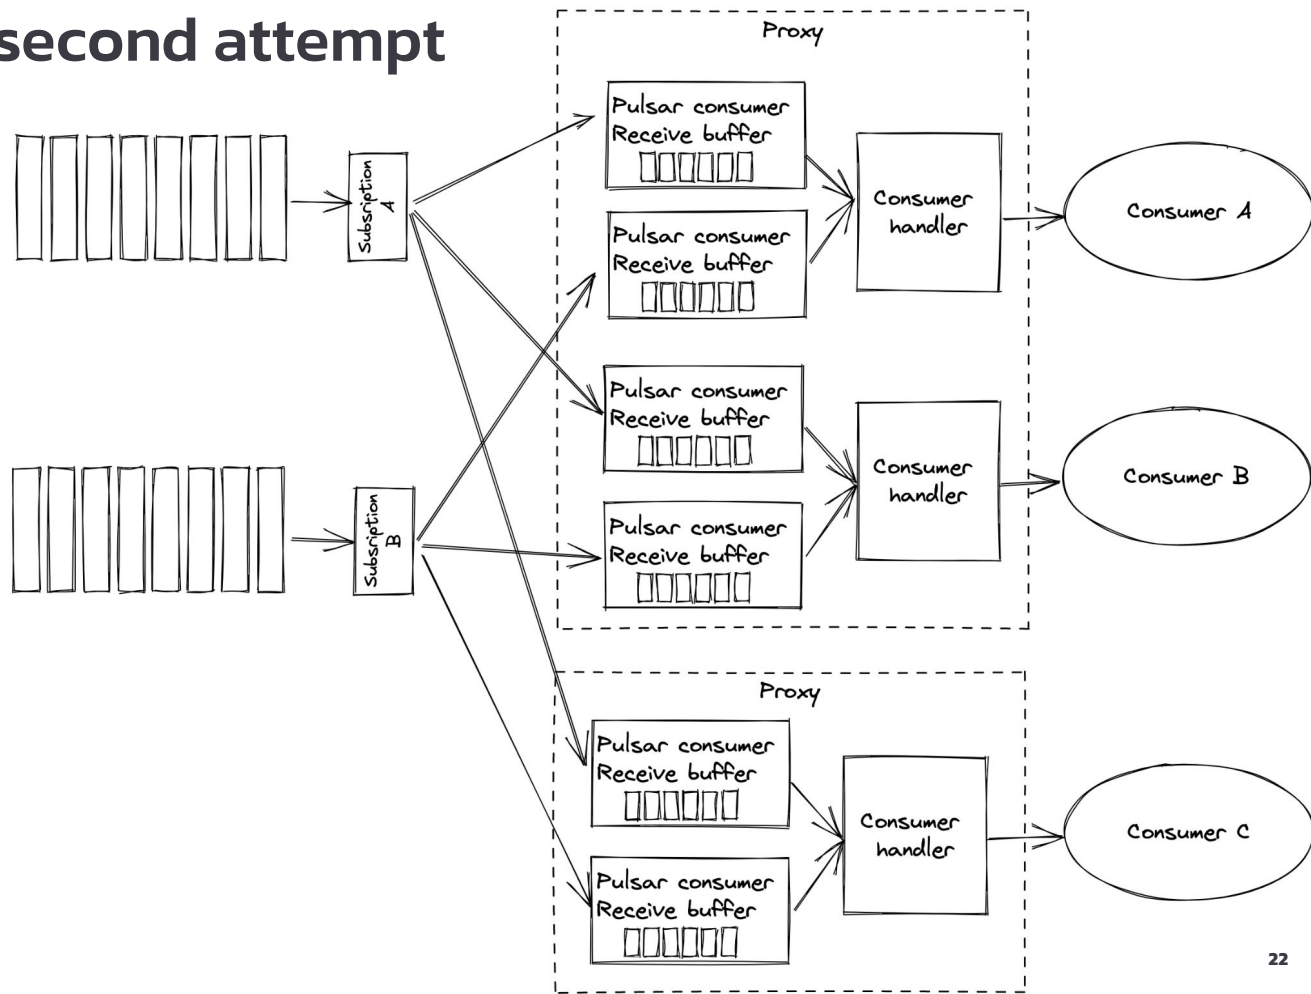
# Consuming-side second attempt

## Pros

- Messages distributed evenly among the consumers
- If a consumer is canceled, the messages in the receive buffer are sent to another consumer.

## Cons

- Divergence from AMQP spec as messages are resent when the consumer is closed instead of when the channel is closed



# Authentication and authorization

- JWT authentication
  - In RabbitMQ: rabbitmq\_auth\_backend\_oauth2 plugin + AMQP “PLAIN” auth mechanism
    - Username ignored
    - JWT passed in password
  - Reuse Pulsar Token authentication implementation
- TLS authentication
  - In RabbitMQ : rabbitmq-auth-mechanism-ssl + AMQP “EXTERNAL” auth mechanism
  - Reuse Pular TLS authentication implementation
- Authorization
  - AMQP Virtual Host mapped to Pulsar tenant+namespace
  - Check that the namespace exists
  - Check that the authenticated role is admin of the Pulsar tenant with the PulsarAdmin client.

# Development

Fully open-source  
under Apache  
Software Licence

In Java, so most  
people can  
understand the code  
and contribute

Use the Apache  
QPid protocol library  
to implement the  
AMQP protocol layer

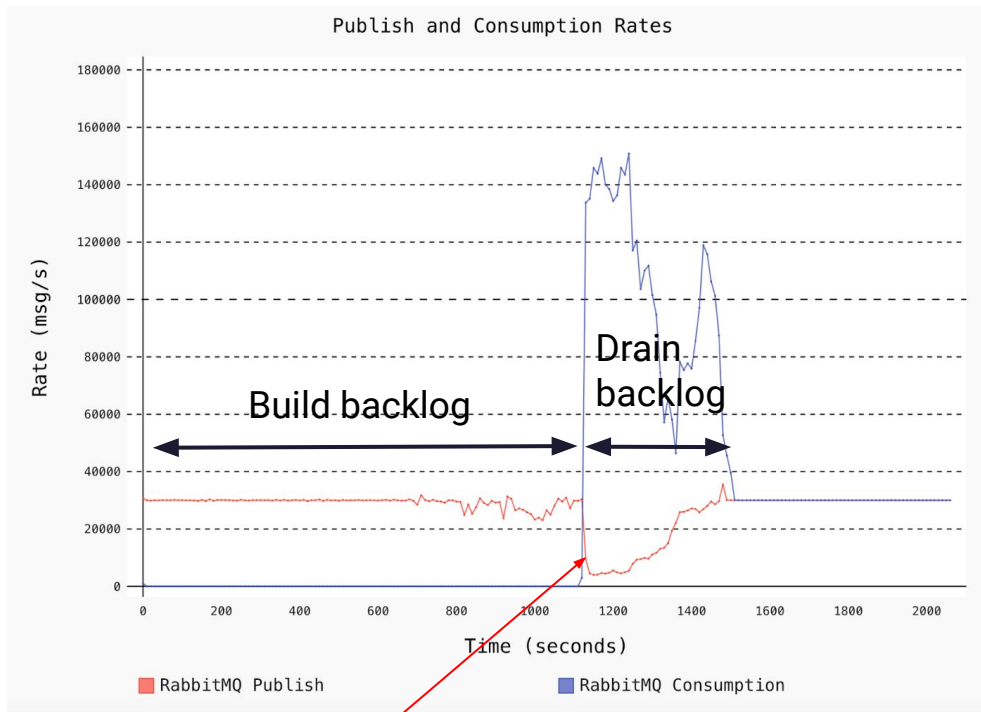
Integration tests  
using the tests from  
the  
rabbitmq-java-client



# Performance tests

- Done with the OpenMessaging Benchmark Framework
- Distributed on a cluster of 3 nodes
- Setup for at-least-once delivery
  - Use of “mirrored queues” for RabbitMQ
- Message size: 1kB
- Consumer stopped, build a backlog of 30 GB
- Then resume the consumption and drain the backlog

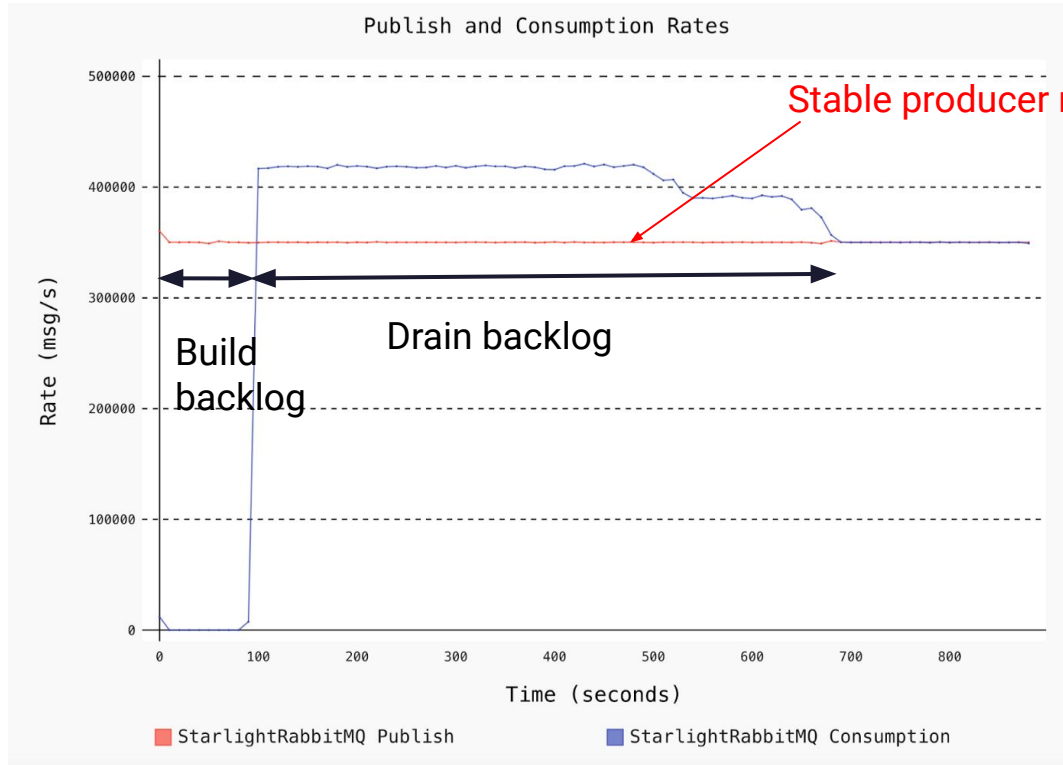
## RabbitMQ



Producer rate of 30 MBps not maintained !

# Performance tests

## Starlight-for-RabbitMQ



# Let's discuss



[@cbornet\\_](https://twitter.com/cbornet_)



[apache-pulsar.slack.com](https://apache-pulsar.slack.com)



[PulsarQuestions@datastax.com](mailto:PulsarQuestions@datastax.com)

**DataStax**

**Thank You!**