

Kafka – Topics & Partitions



Instructor: Ekpe Okorafor

1. Big Data Academy - Accenture
2. Computer Science - African University of Science & Technology

Agenda

- Introduction
- Topics
 - General Principle
 - Ordering Problems
 - When to Split, When to Combine
- Partitions

Introduction

When adopting a streaming platform such as Apache Kafka, some important questions to answer are:

- What topics are you going to use?
- In particular, if you have a bunch of different events that you want to publish to Kafka as messages, do you put them in the same topic, or do you split them across different topics?
- Two extremes – one topic or millions?

Topics – General Principle

- Topic = collection of events of the same type?
 - put all events of the same type in the same topic,
 - and use different topics for different event types.
- What about cases where the ordering of events matters?

Topics – Ordering Problems

Ordering is not preserved across partitions

- Every topic has at least 1 partition
- If the order of event are important, having these events in different topics = different partitions

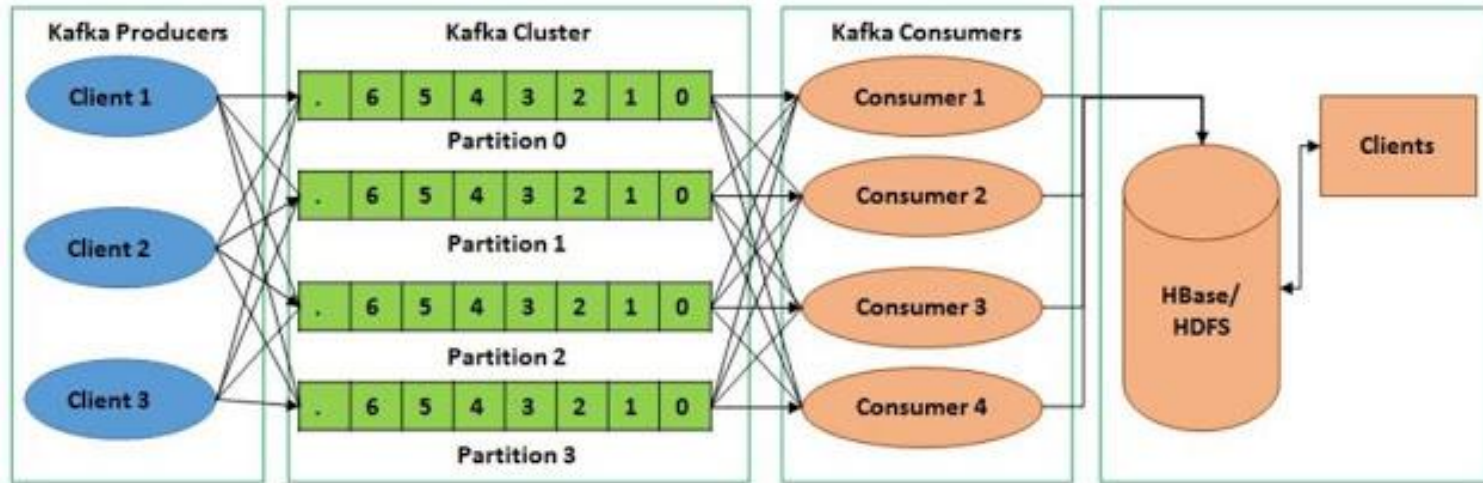
Topics – When to Split, When to Combine

1. The most important rule is that any events that need to stay in a fixed order must go in the same topic
2. If an event entity depends on another, they should be in the same topic. If unrelated the separate
3. Try not to split up events with multiple entities initially
4. If several consumers all read a particular group of topics, this suggests that maybe those topics should be combined

Partitions

- More Partitions Lead to Higher Throughput
 - topic partition is the unit of parallelism in Kafka
- More Partitions Requires More Open File Handles
- More Partitions May Increase Unavailability
- More Partitions May Require More Memory In the Client

Partitions



- The more partitions the greater the Zookeeper overhead
 - With large partition numbers ensure proper ZK capacity
- Message ordering can become complex
 - Single partition for global ordering
 - Consumer-handling for ordering
- The more the partitions the longer the leader fail-over time

ANY
QUESTIONS
?