



Agile Integration

PA Agile Integration, Kafka and Event-Streaming

Oliver Berger, Nadja Hagen

Tobias Dehn, Christopher Weiß, Uwe Eisele



Content

- Integration
- Event-Streaming Platform Kafka
- Asynchronous Communication
- Kafka Basics & Components
- Kafka APIs
- Kafka Examples
- Kafka Exercises

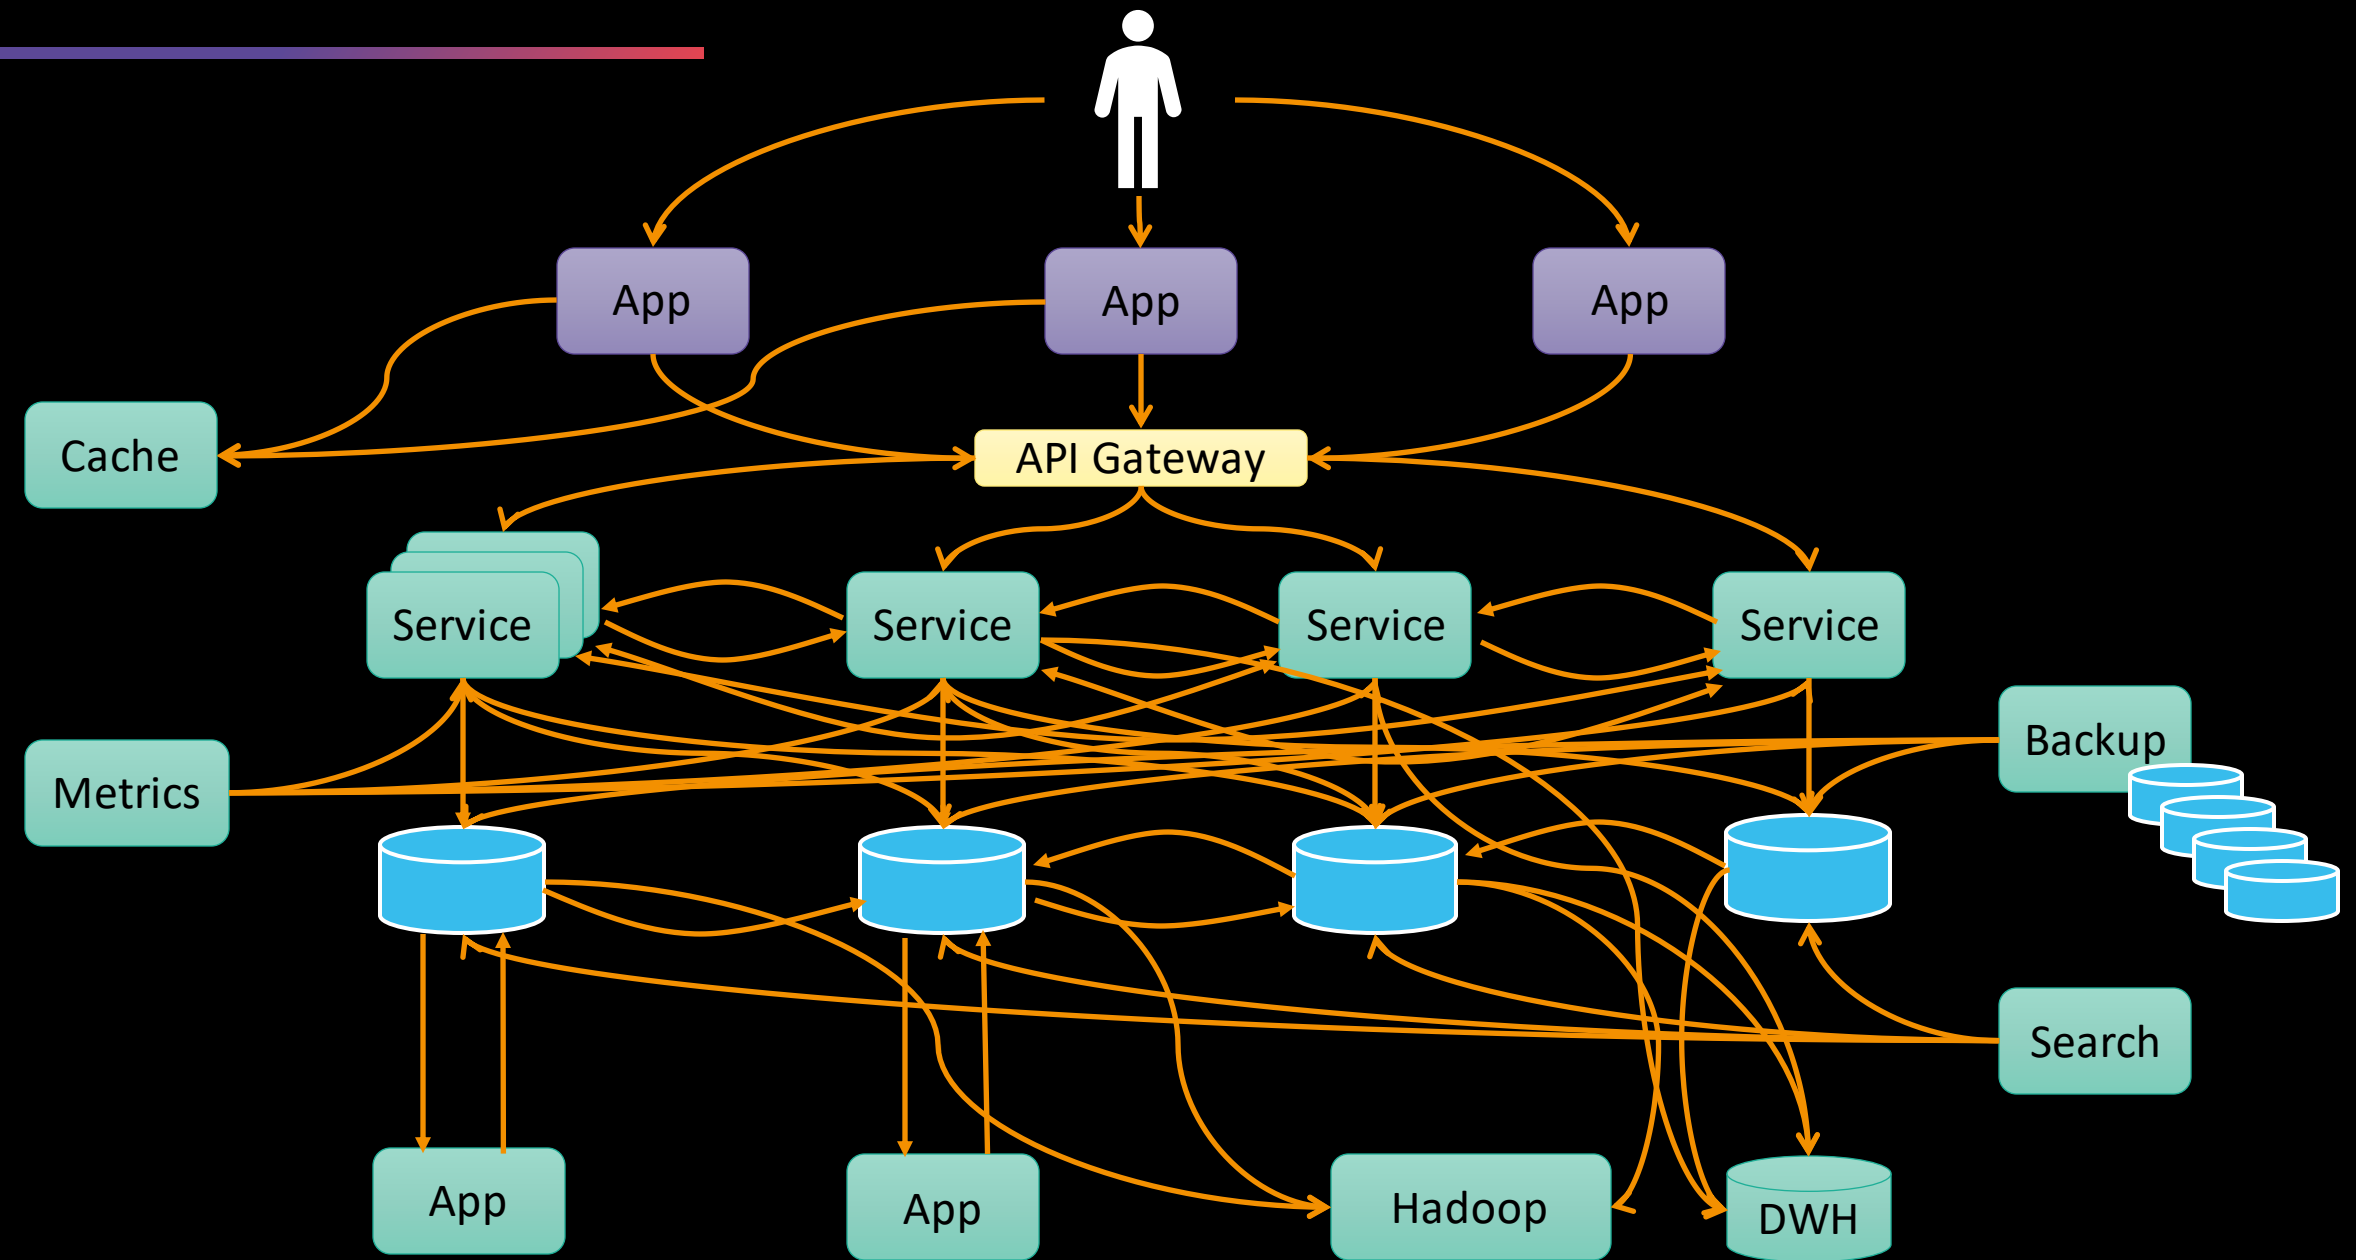


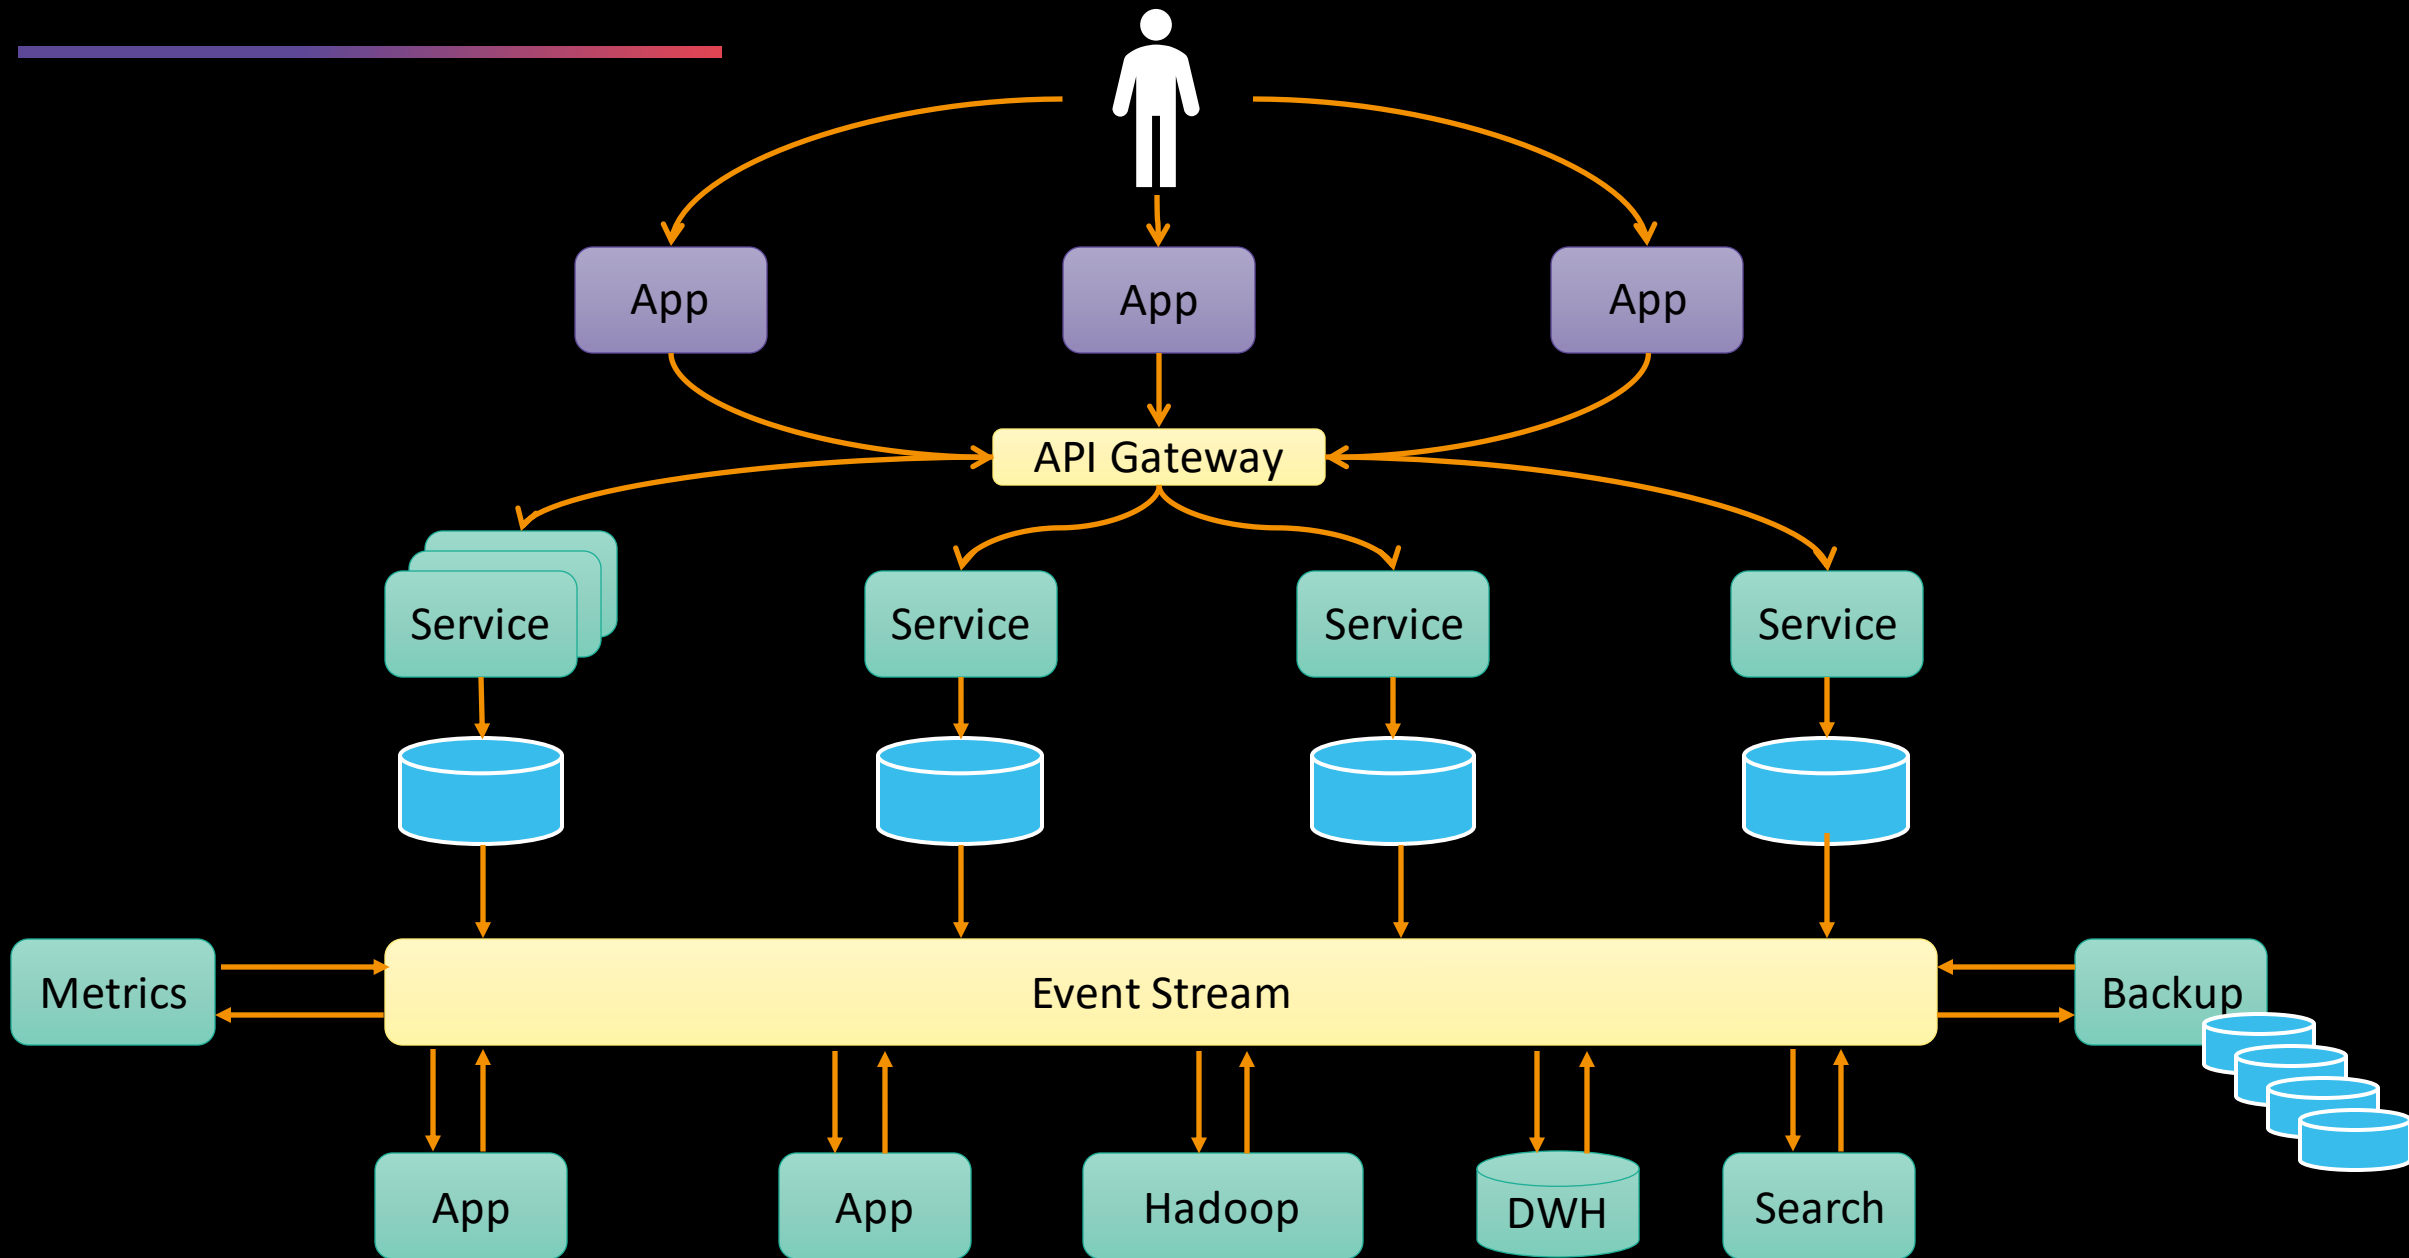
Integration

..of distributed systems is a complex challenge

Integration in Distributed Systems, complex challenge lots of things to consider

Technologies	Programming Language	Application Architecture	Communication Paradigms
<ul style="list-style-type: none">• Standards (SOAP, REST , JMS)• Data formats (JSON, XML, Avro)• Frameworks• Proprietary Interfaces	<ul style="list-style-type: none">• Java• C, C#, .Net• Python• Cobol	<ul style="list-style-type: none">• Client Server• Monolith• SOA• Microservices• Serverless	<ul style="list-style-type: none">• Batch• Realtime• Request-Response• Pub-Sub• Fire & Forget





Content

- Integration
- Event-Streaming Platform Kafka
- Asynchronous Communication
- Kafka Basics & Components
- Kafka APIs
- Kafka Examples
- Kafka Exercises



Event-Streaming Platform Kafka

Why should you be interested?

Apache Kafka: The Event Streaming Platform

- developed at LinkedIn in 2011 and made open source
- can process several trillion (10^{12}) events per day
- originally designed as a messaging queue
- based on an abstraction of a distributed commit log
- evolved from a messaging queue to a full-fledged event streaming platform
- de facto standard for Event-Streaming Platform (> 95% of Event-Streaming-Projects rely on kafka)

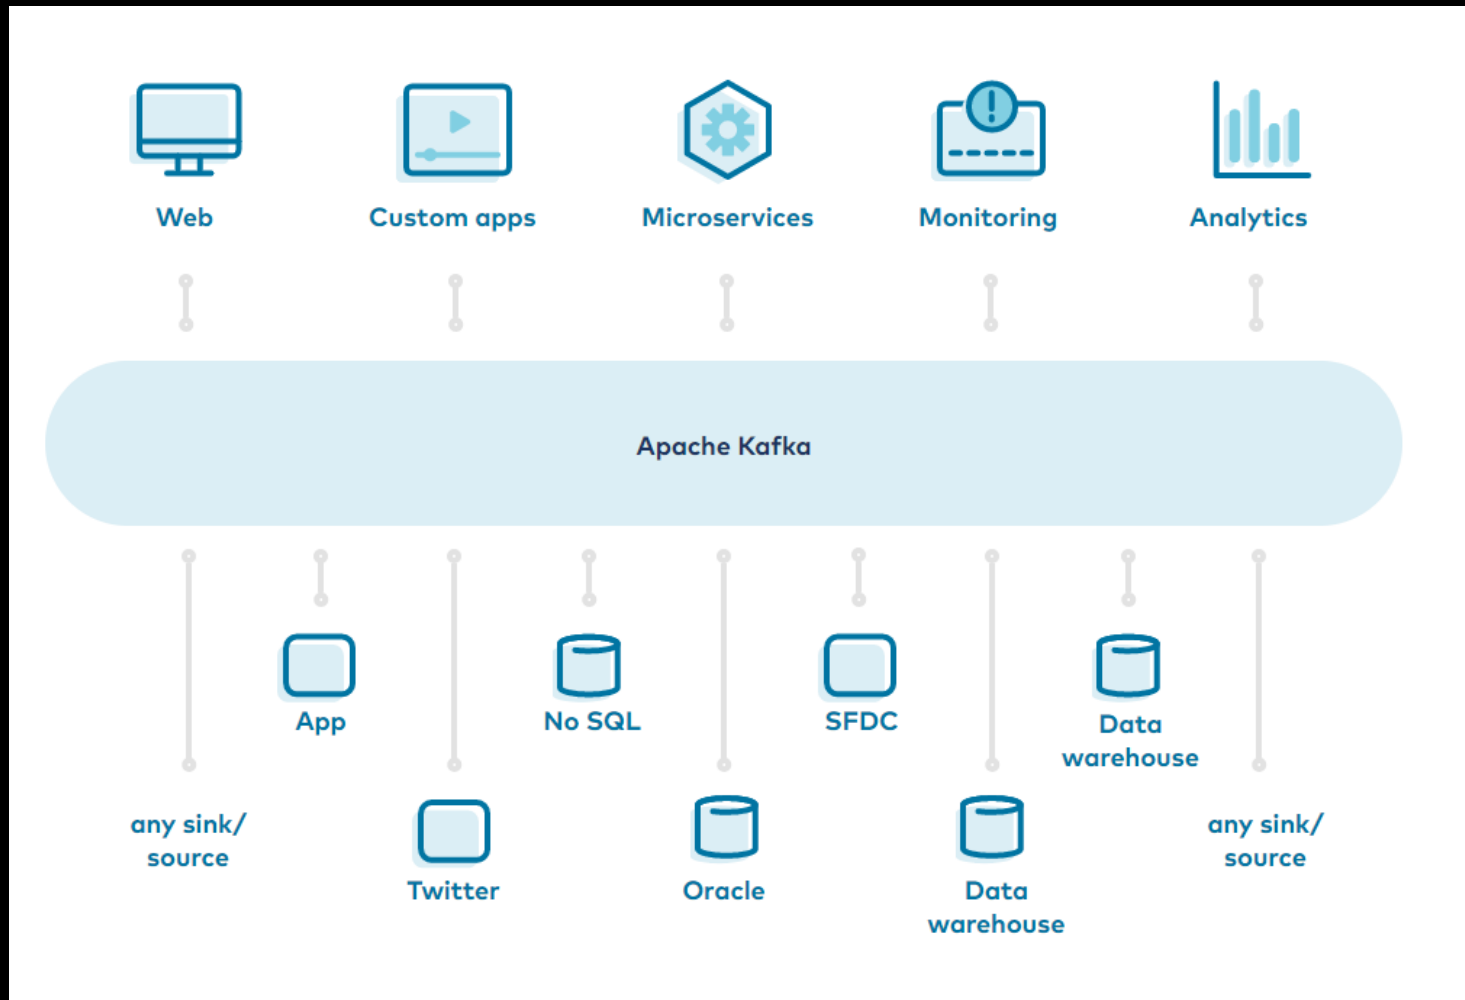


Thousands of Enterprises rely on Kafka and Event-Streaming



Source: <https://kafka.apache.org/powered-by>

Possible Applications of Apache Kafka



3 Key Functionalities of a Streaming Platform

Pub-Sub



Store



Process



Content

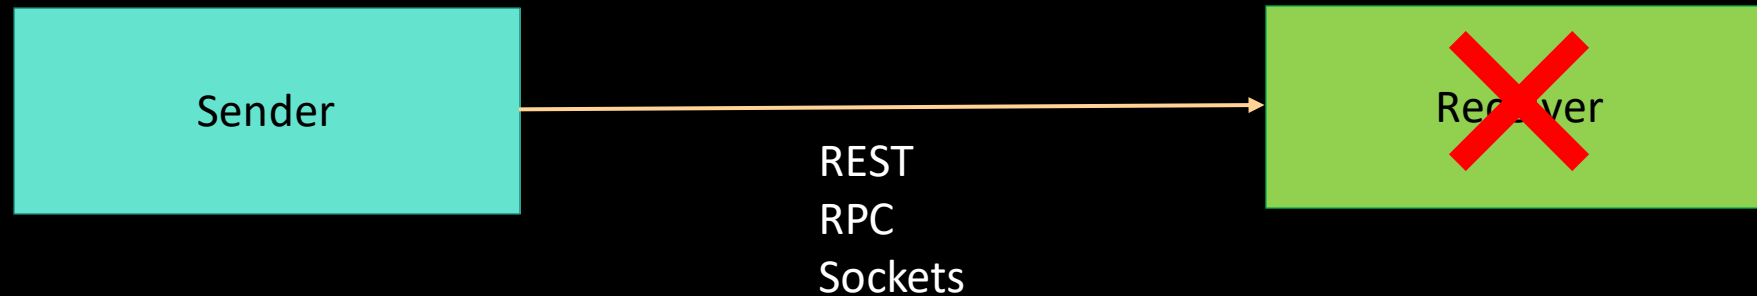
- Integration
- Event-Streaming Platform Kafka
- Asynchronous Communication
- Kafka Basics & Components
- Kafka APIs
- Kafka Examples
- Kafka Exercises



Asynchronous Communication

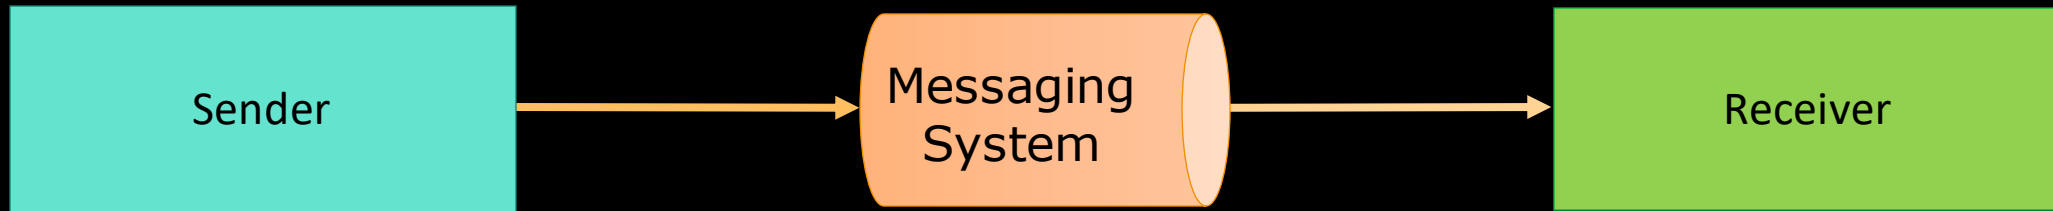
Why do you need a Messaging System?

Why do you need a Messaging System ?



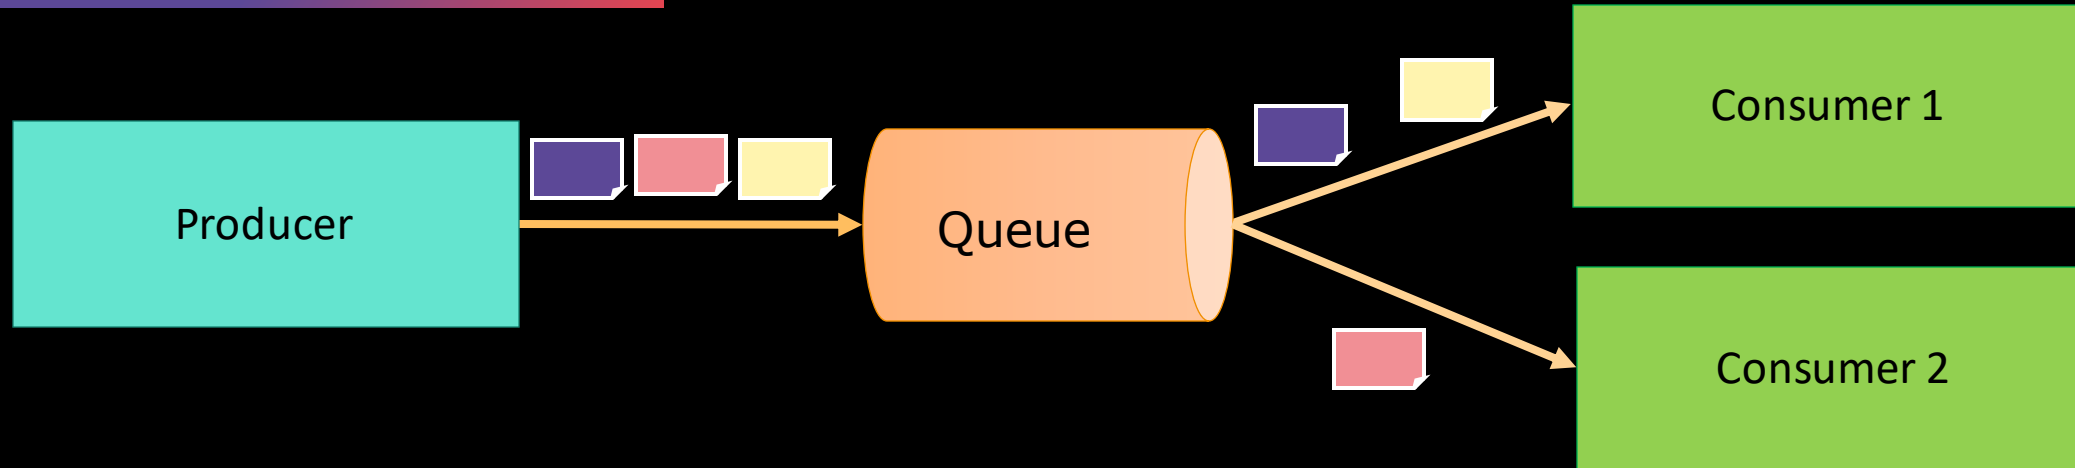
- Challenge 1: Availability
- Challenge 2: Processing Velocity
- Challenge 3: Processing Acknowledgement

Solution: Messaging

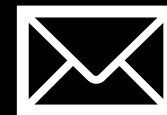


- Decoupling via Messaging
- Examples for Messaging-Systems:
 - MQ-Series,
 - JMS-Messaging (ActiveMQ, Rabbit-MQ),
 - Kafka.
- Transfermode: Queue or Topic

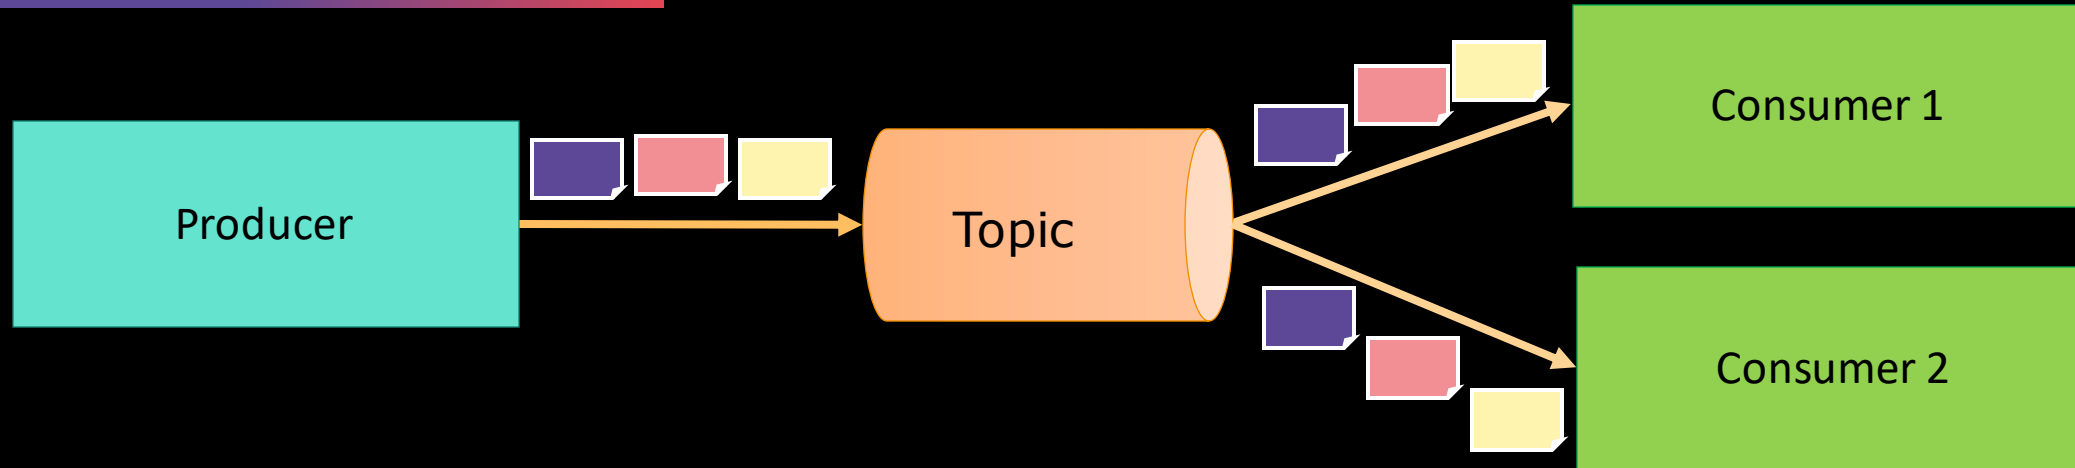
Queue Mode 1 to 1 Topologie



- each message gets processed exactly once



Topic Mode 1 to n Topologie



- each message can be consumed by independent consumers
- each consumer receives all messages after subscription
- sequence of messages is guaranteed
- only new messages are delivered



Content

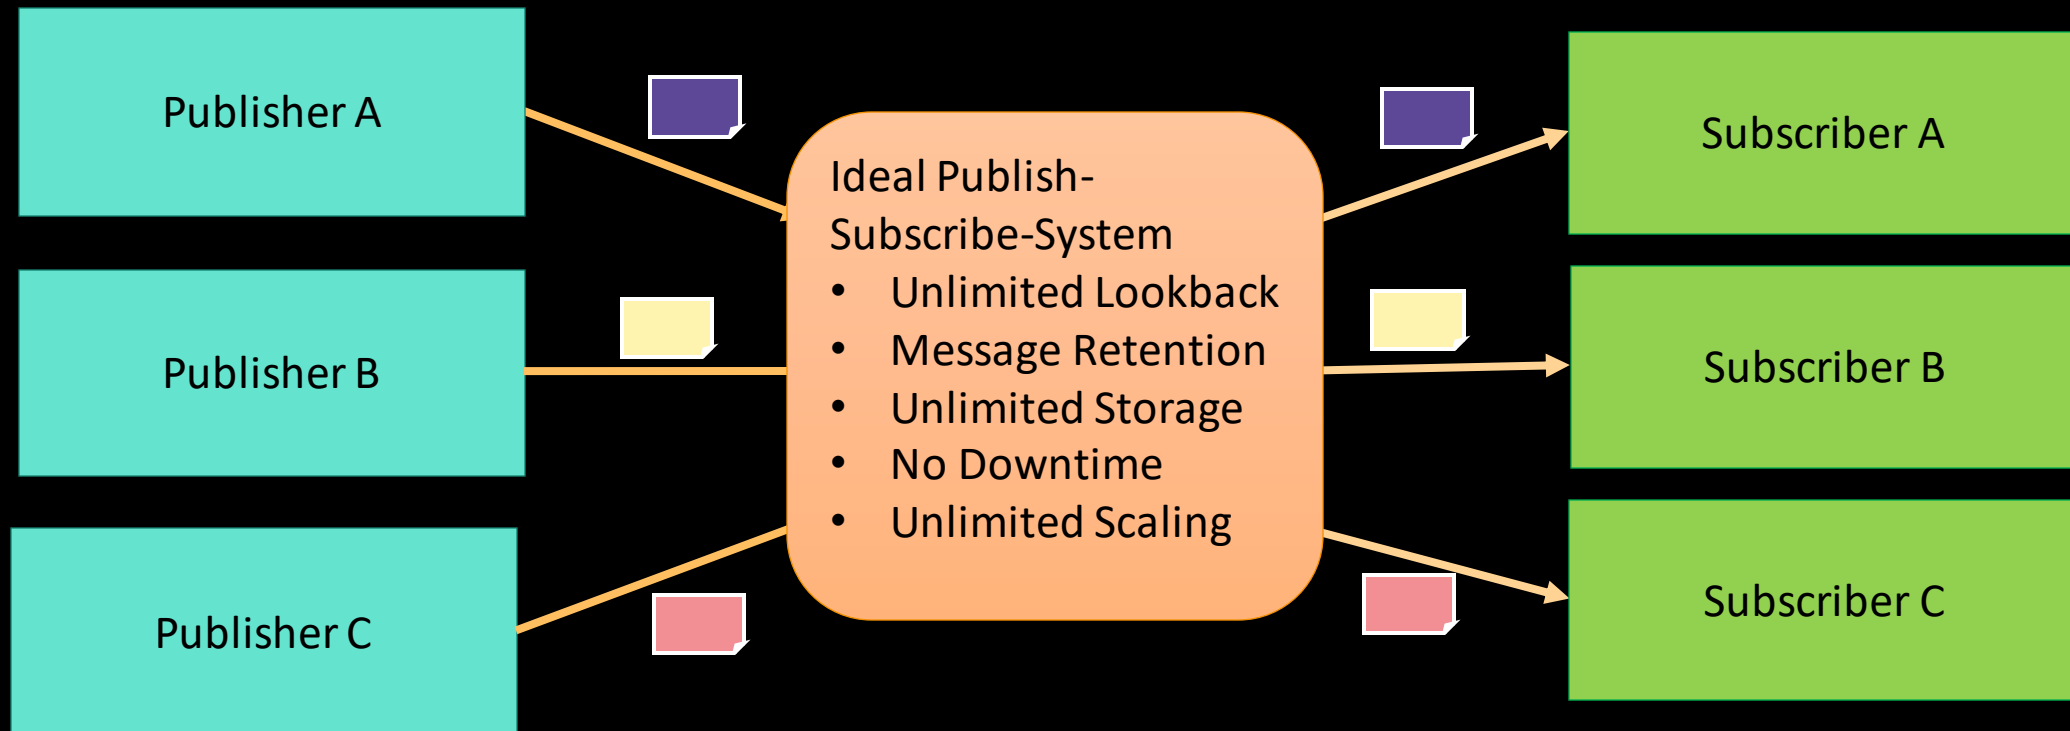
- Integration
- Event-Streaming Platform Kafka
- Asynchronous Communication
- Kafka Basics & Components
- Kafka APIs
- Kafka Examples
- Kafka Exercises



Kafka

Basics & Components

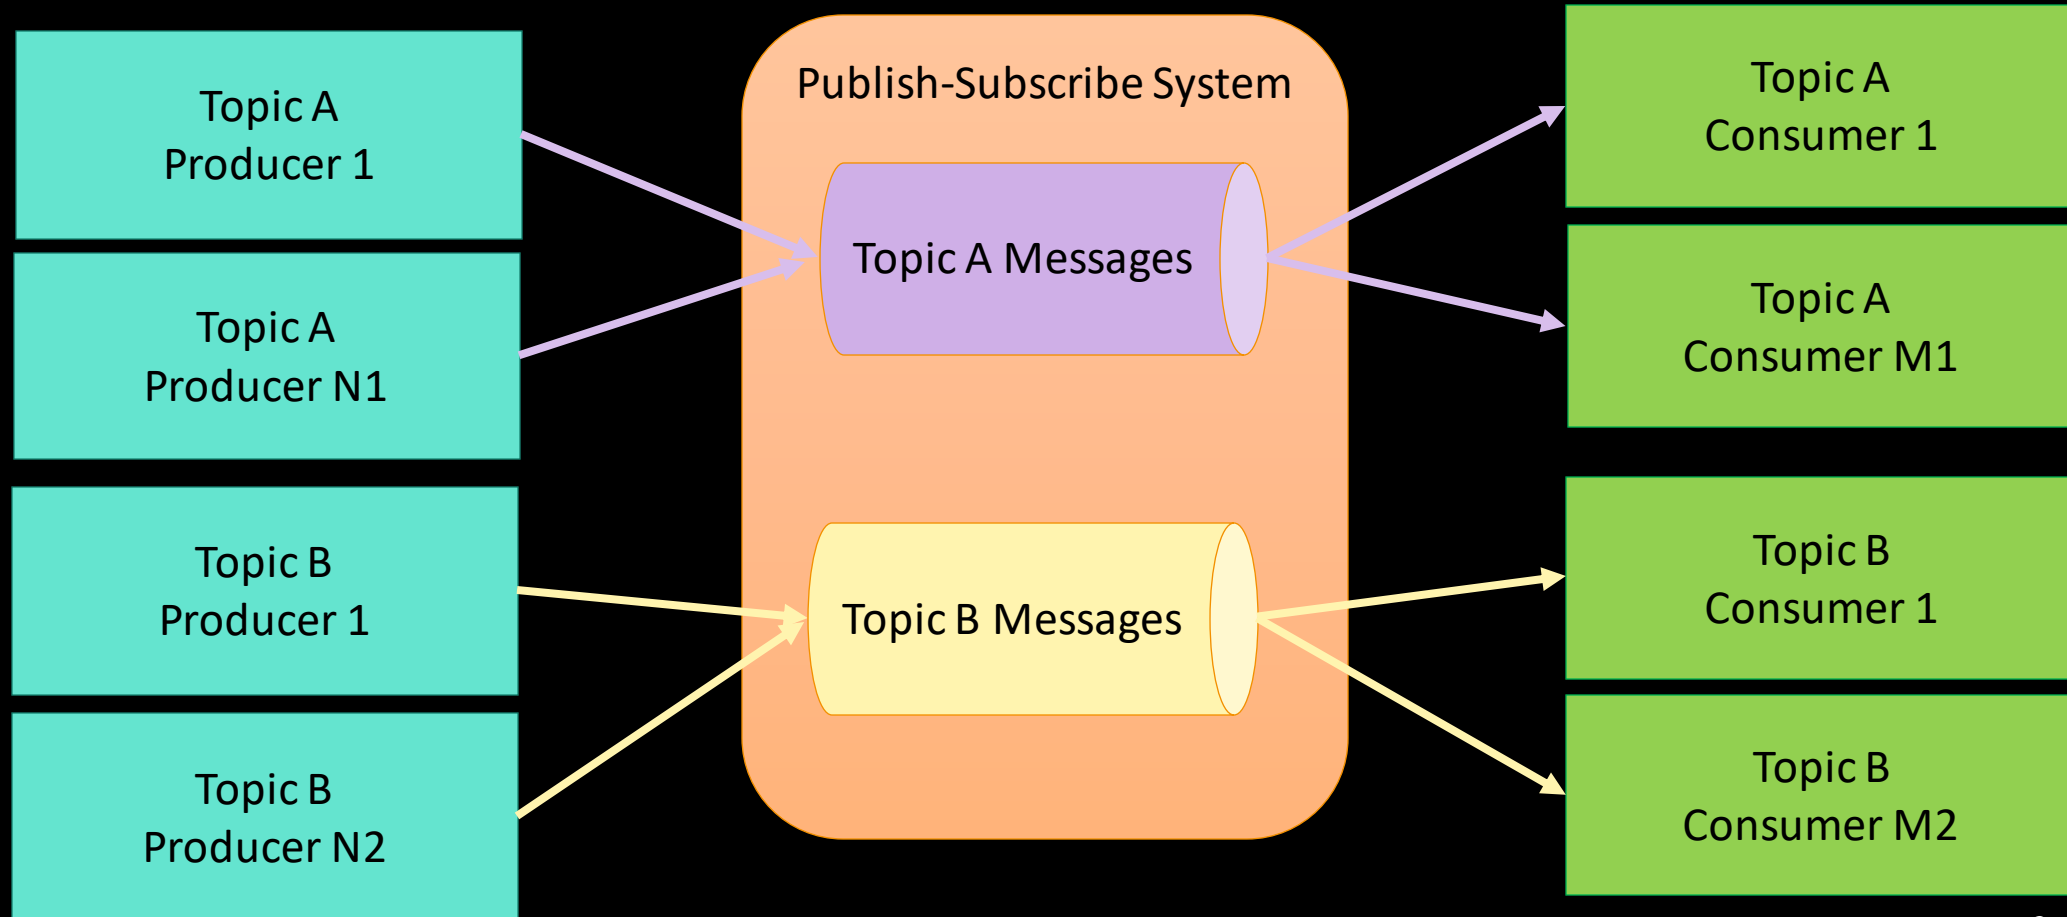
How would an ideal Publish-Subscribe System look like?



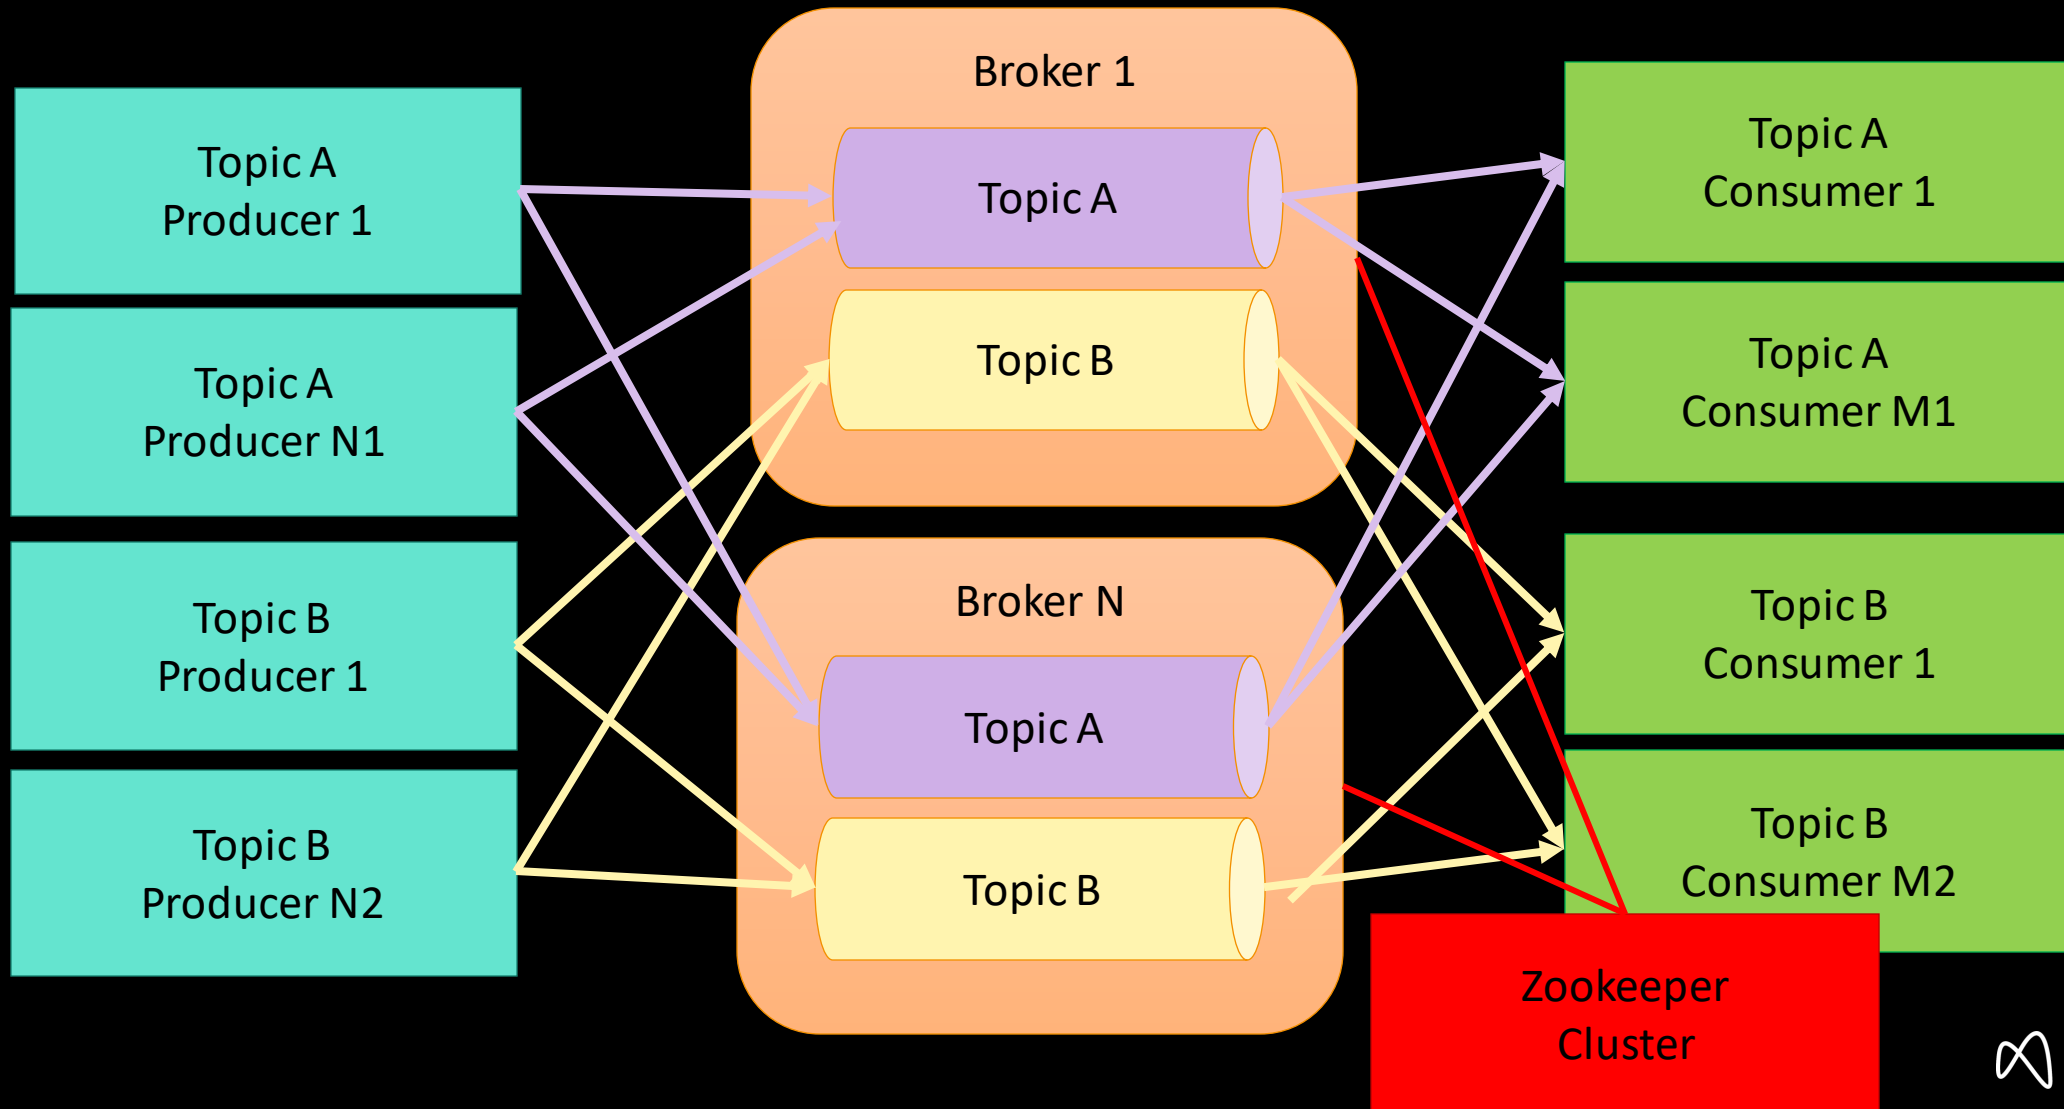
Kafka Architecture in Comparison to the ideal Pub-Sub System

- Key Differences
 - Messaging is implemented on top of a replicated, distributed unmutable commit log.
 - The client has more functionality and, therefore, more responsibility.
 - Messaging is optimized for batches instead of individual messages.
 - Messages are retained even after they are consumed; they can be consumed again.
- Consequences of these Design Decisions
 - extreme horizontal scalability
 - very high throughput
 - high availability
 - but different semantics and message delivery guarantees

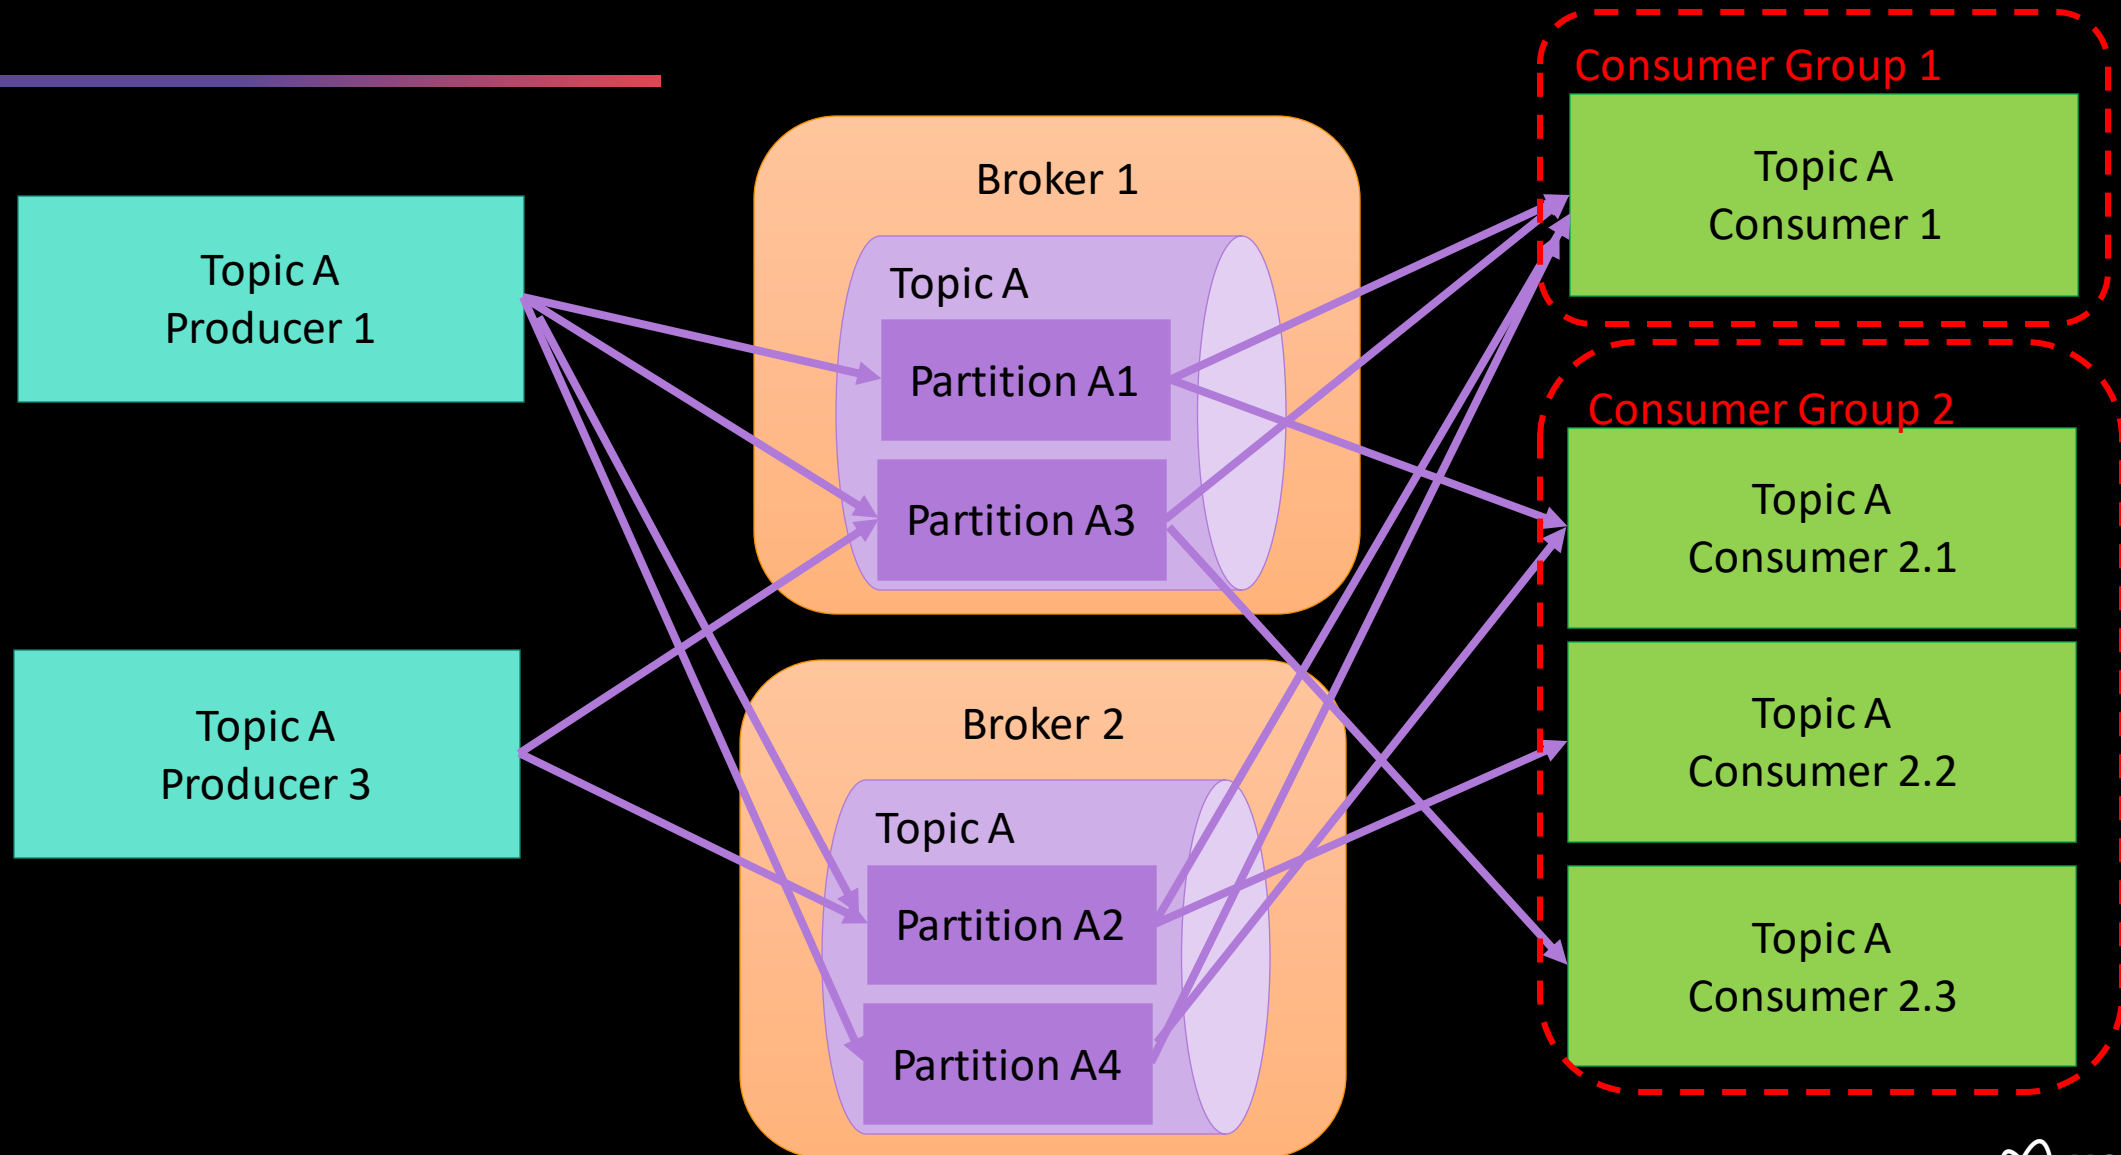
Topics in a Publish-Subscribe System



Broker in a Publish-Subscribe System

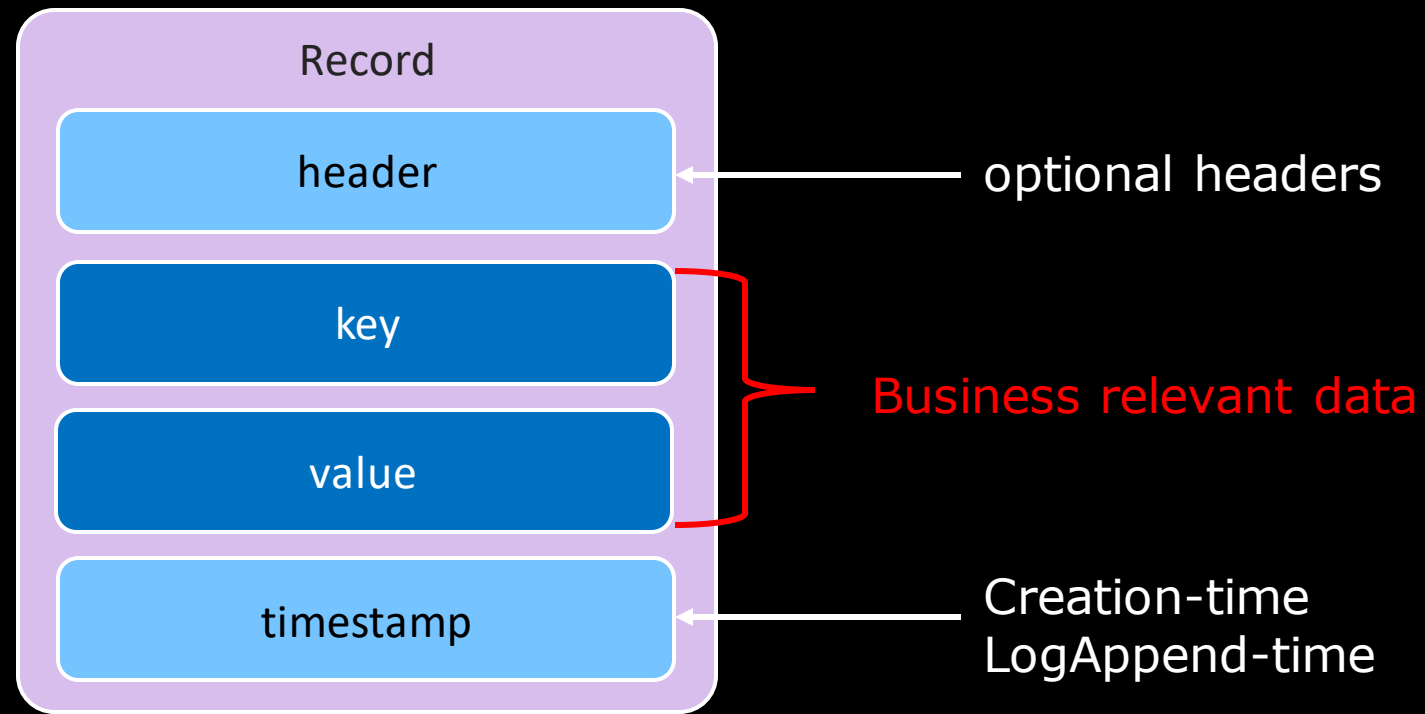


Partitions: Partition Count 4

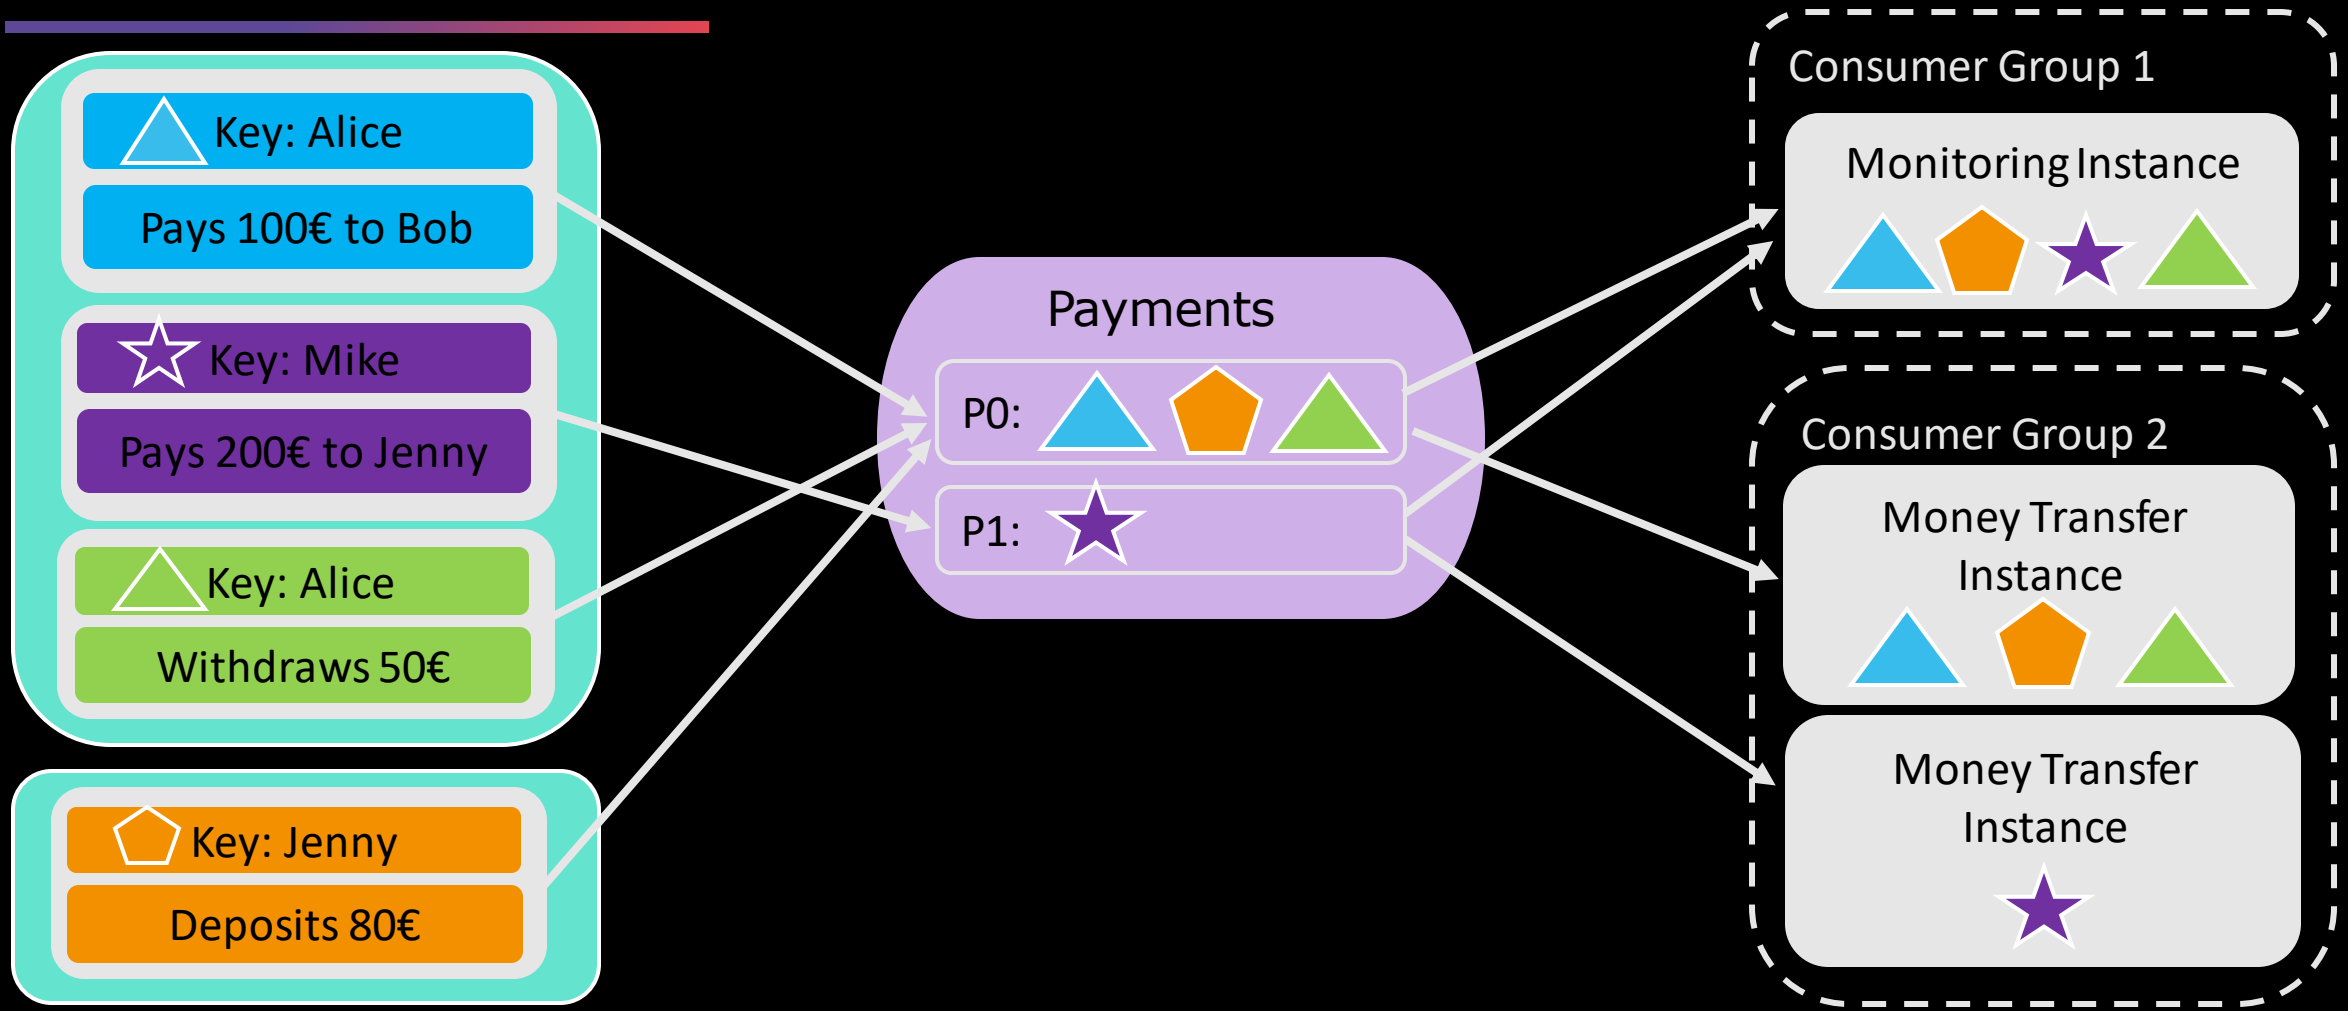


The Record – The Atomic Unit of Kafka

Synonyms: Message or Event

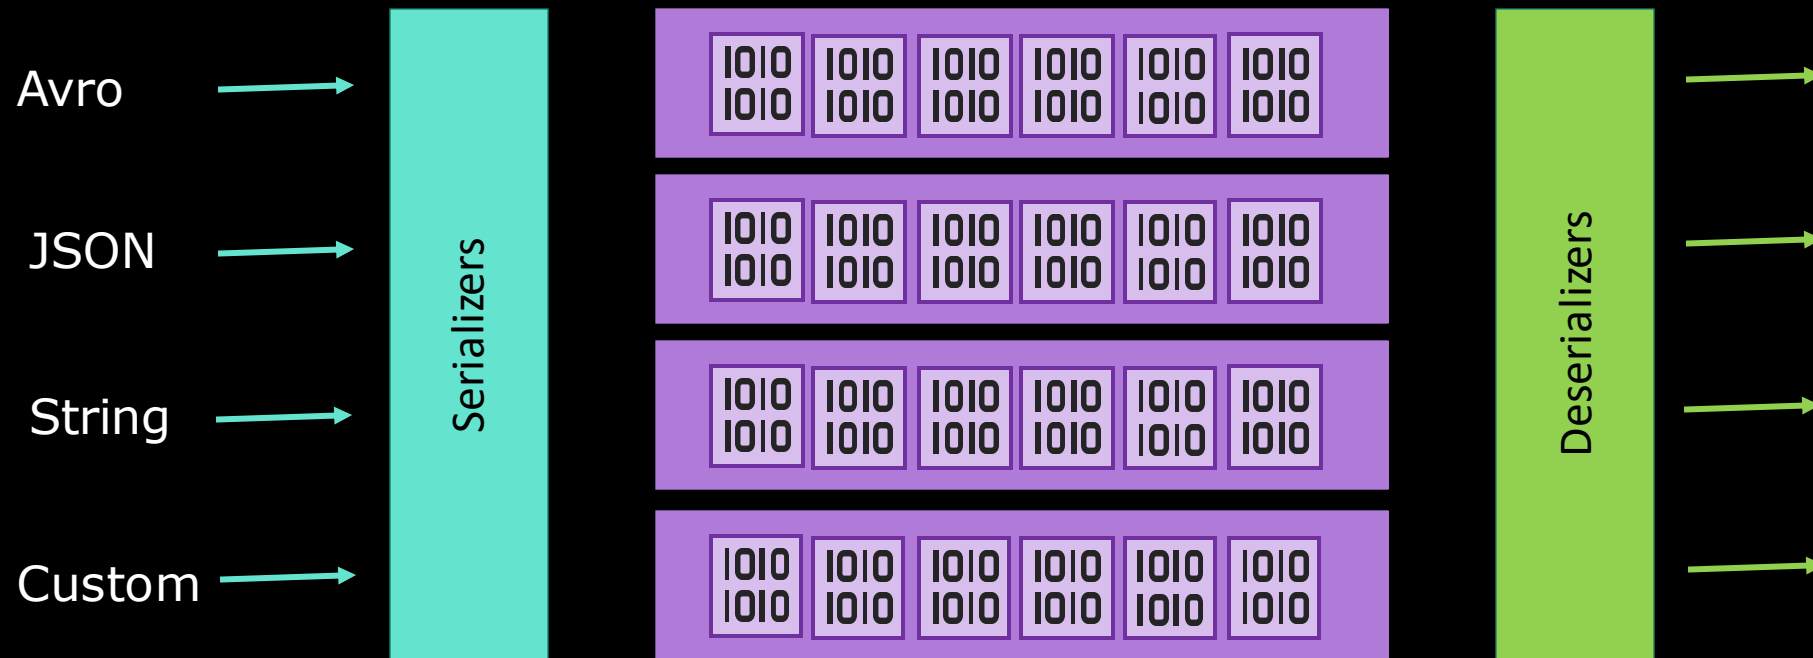


Example: Payment Processing



Serialisation

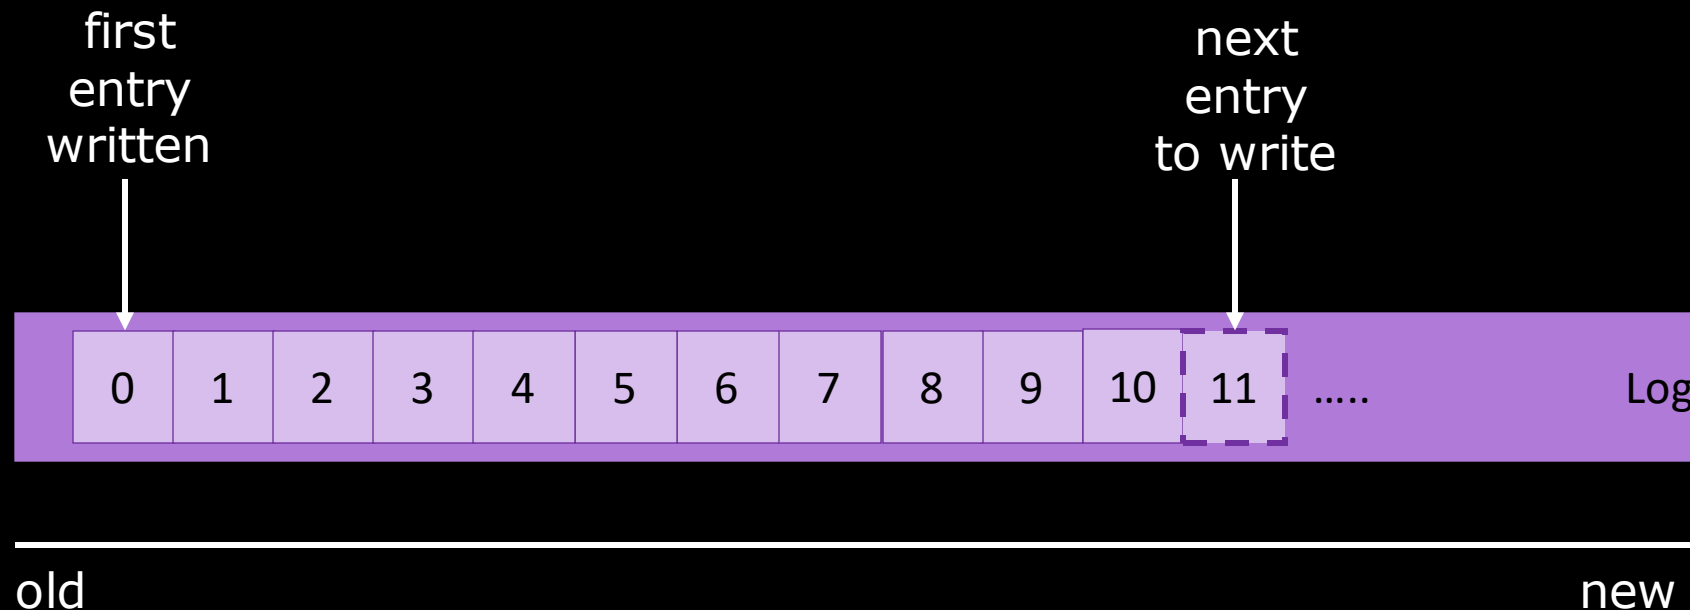
Kafka stores Byte Arrays



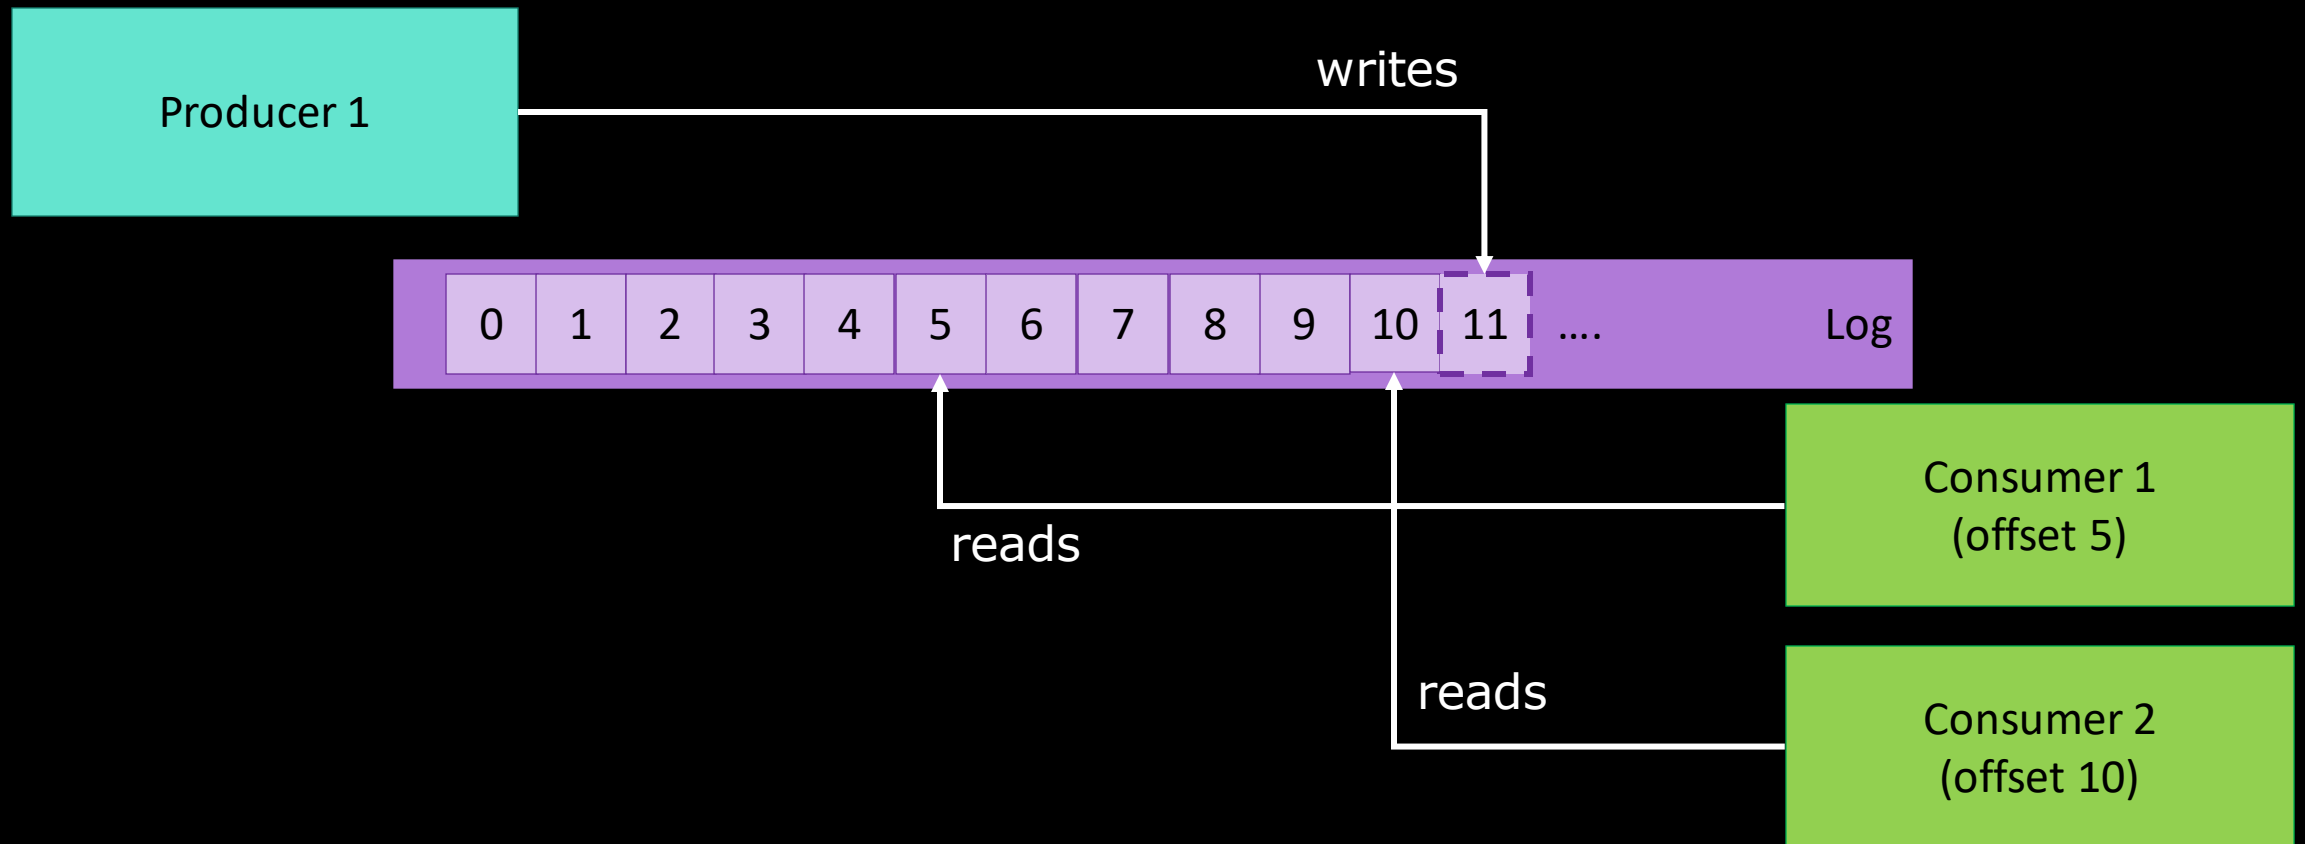
Kafka Commit Log

Abstraction to understand Streaming

- Immutable, append-only data structure (record,event)
- Offset: the position of the record/event in the log

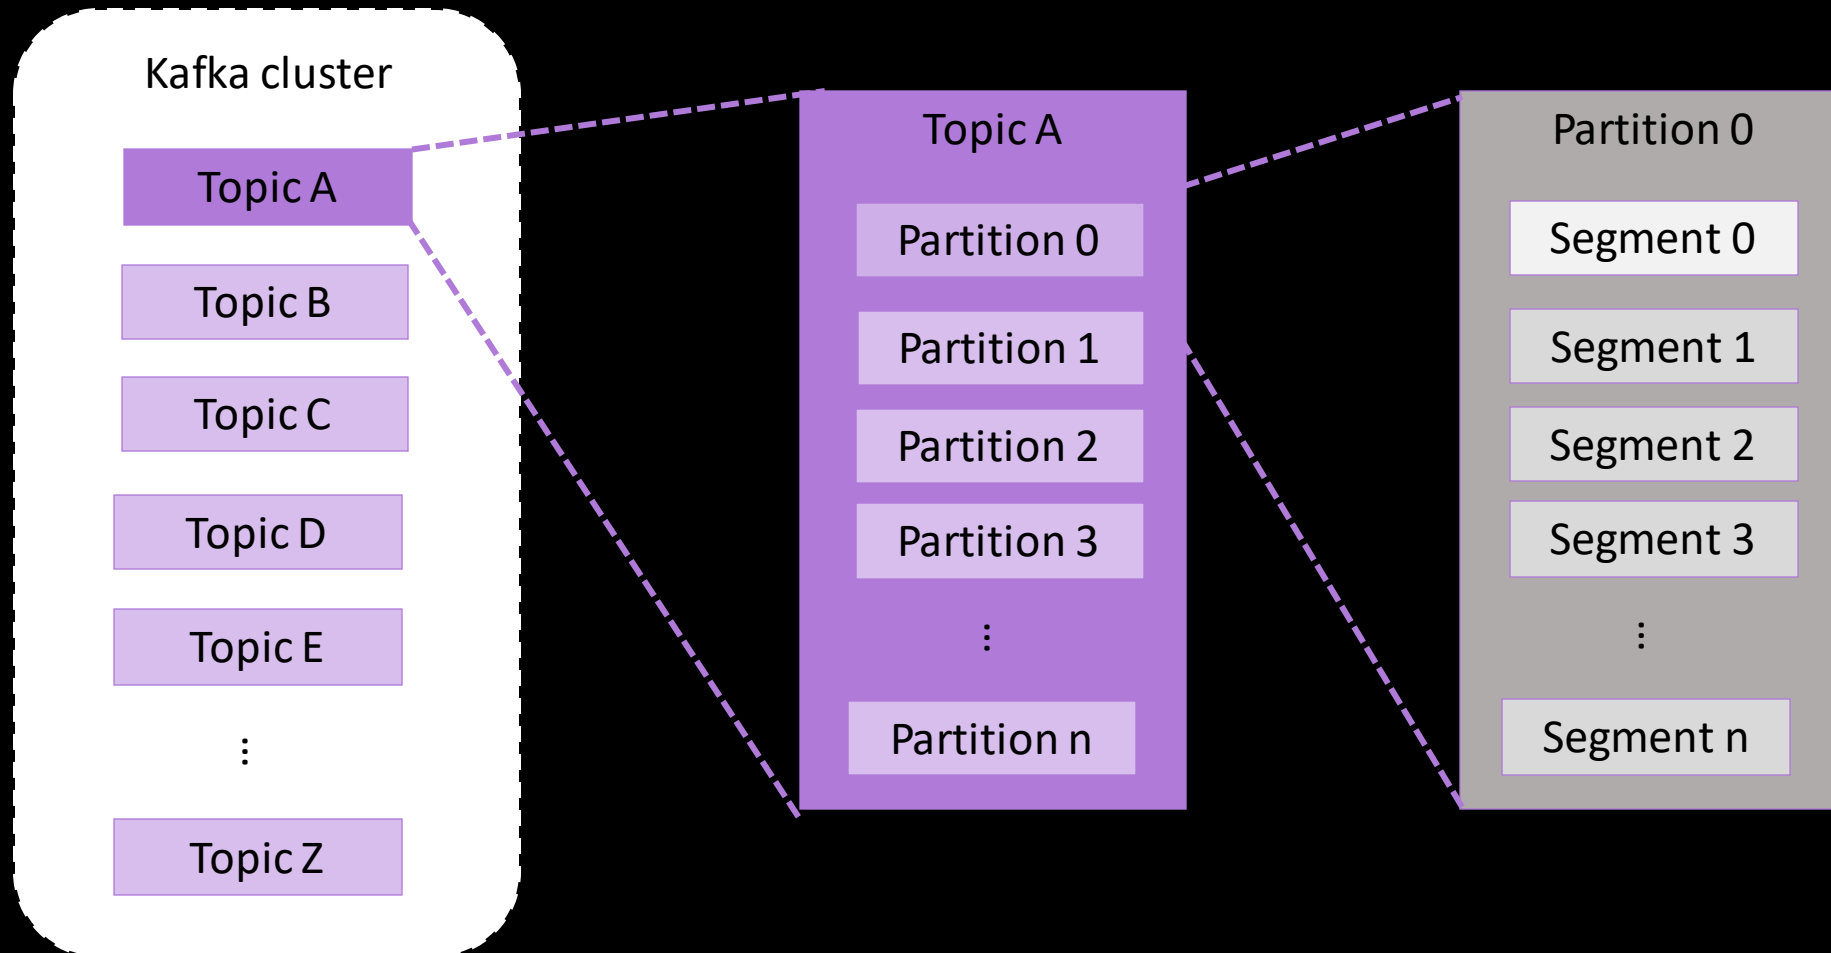


Decoupling Data Producers from Data Consumers

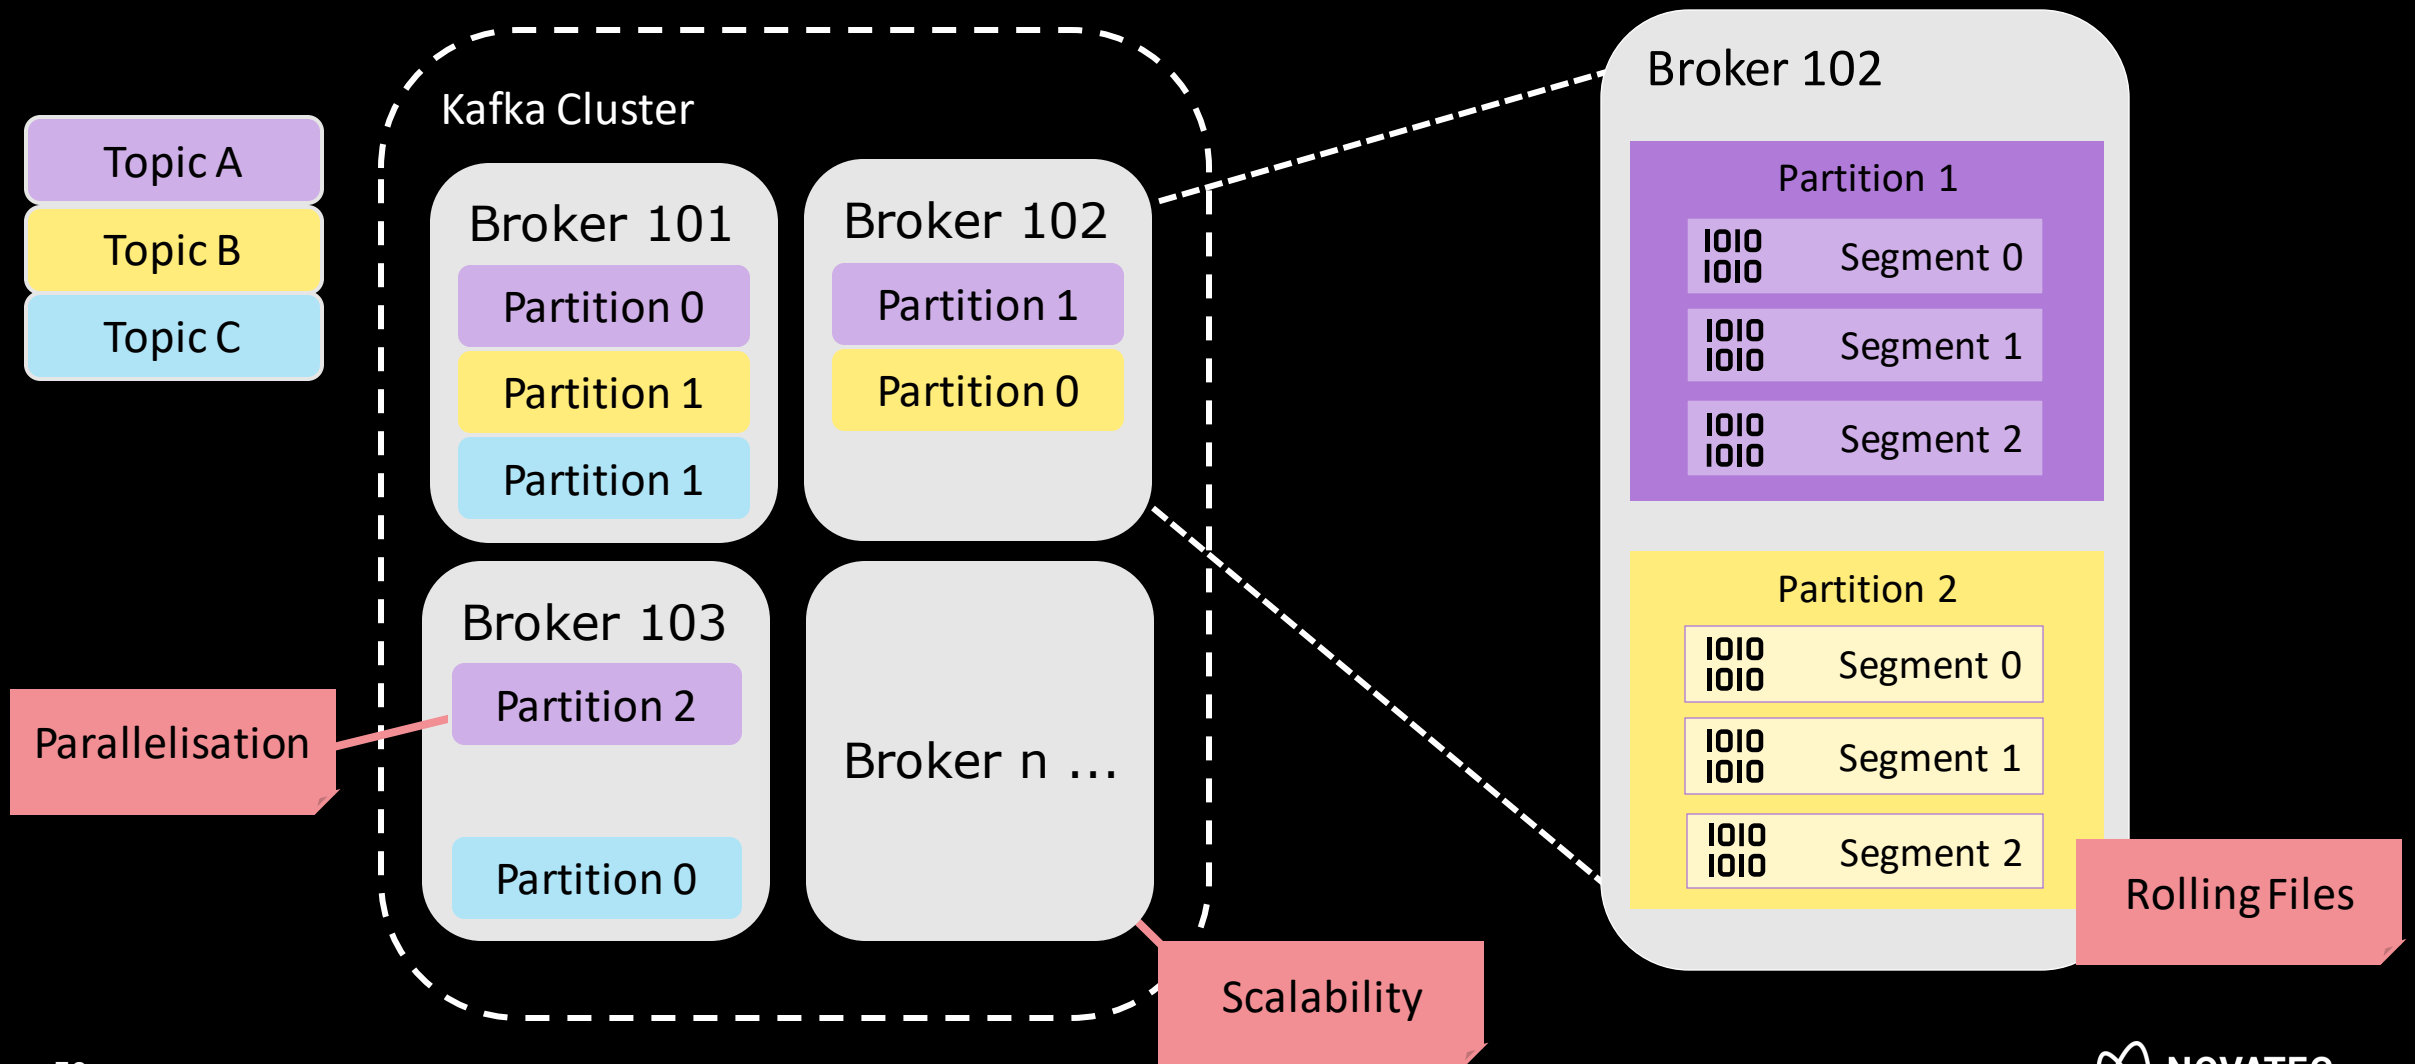


Logical View of Topics, Partitions & Segments

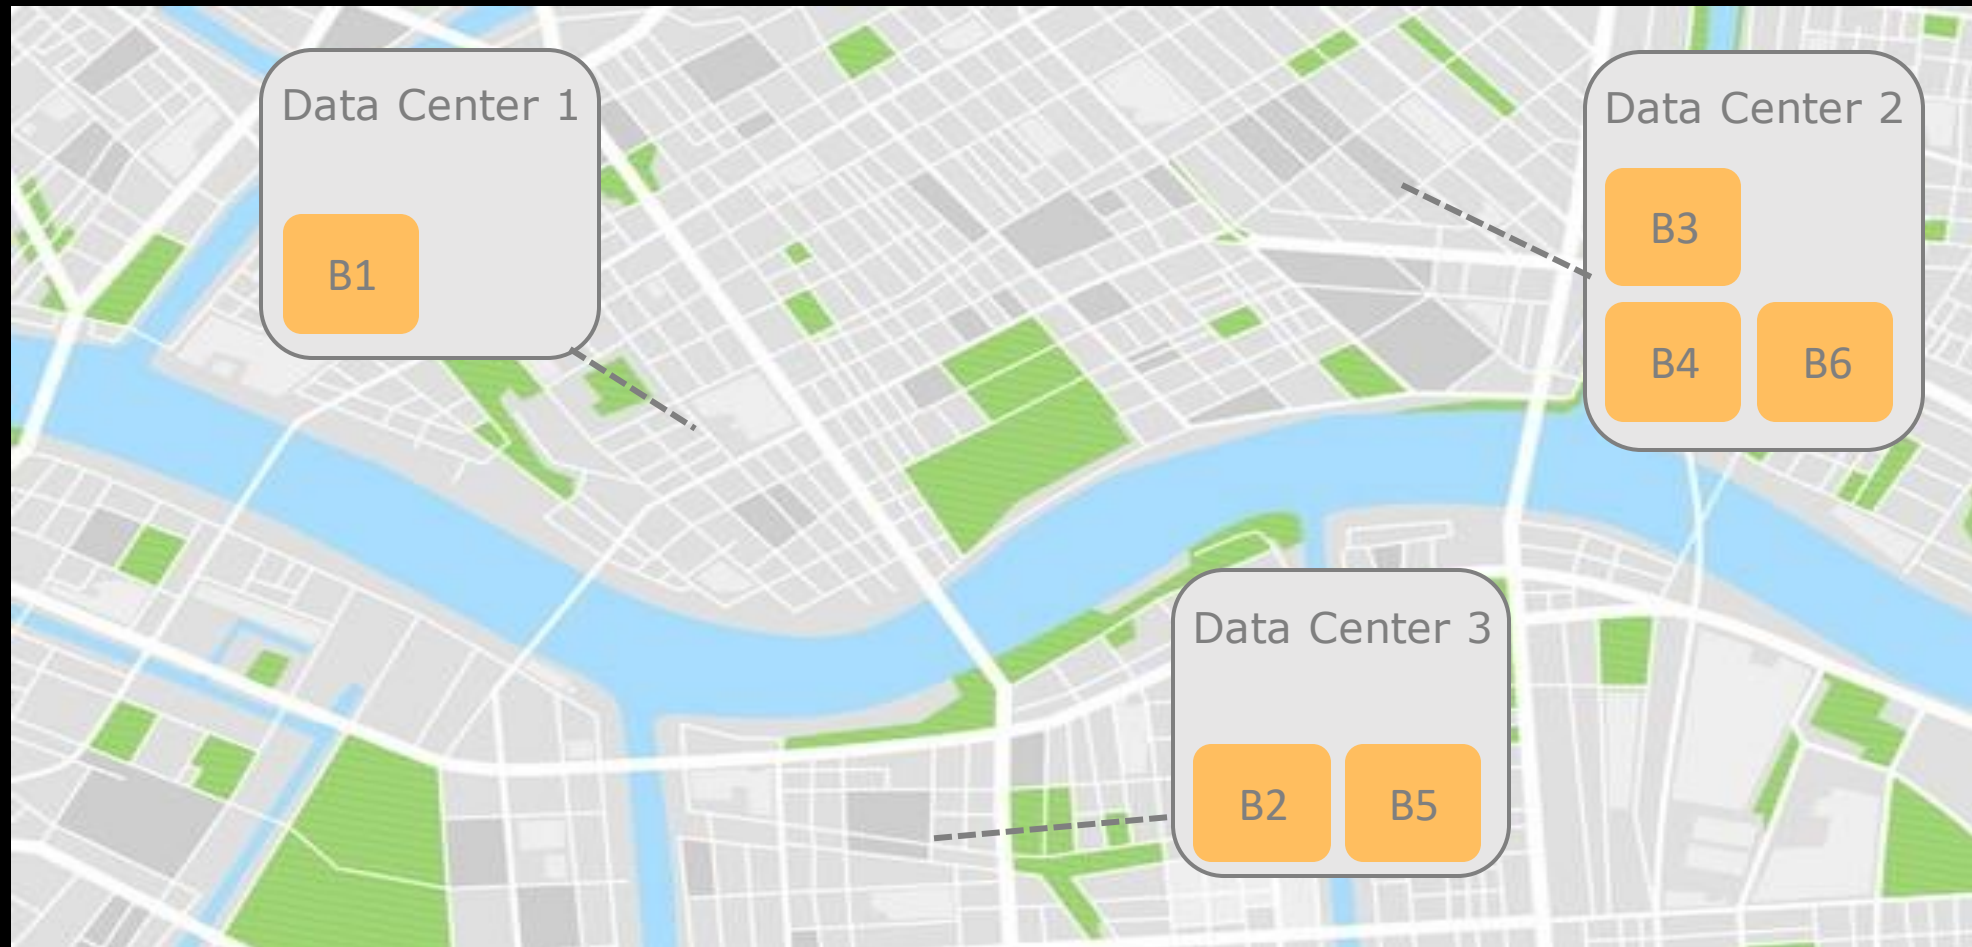
Each Partition is a Commit Log



Physical View of Topics, Partitions & Segments



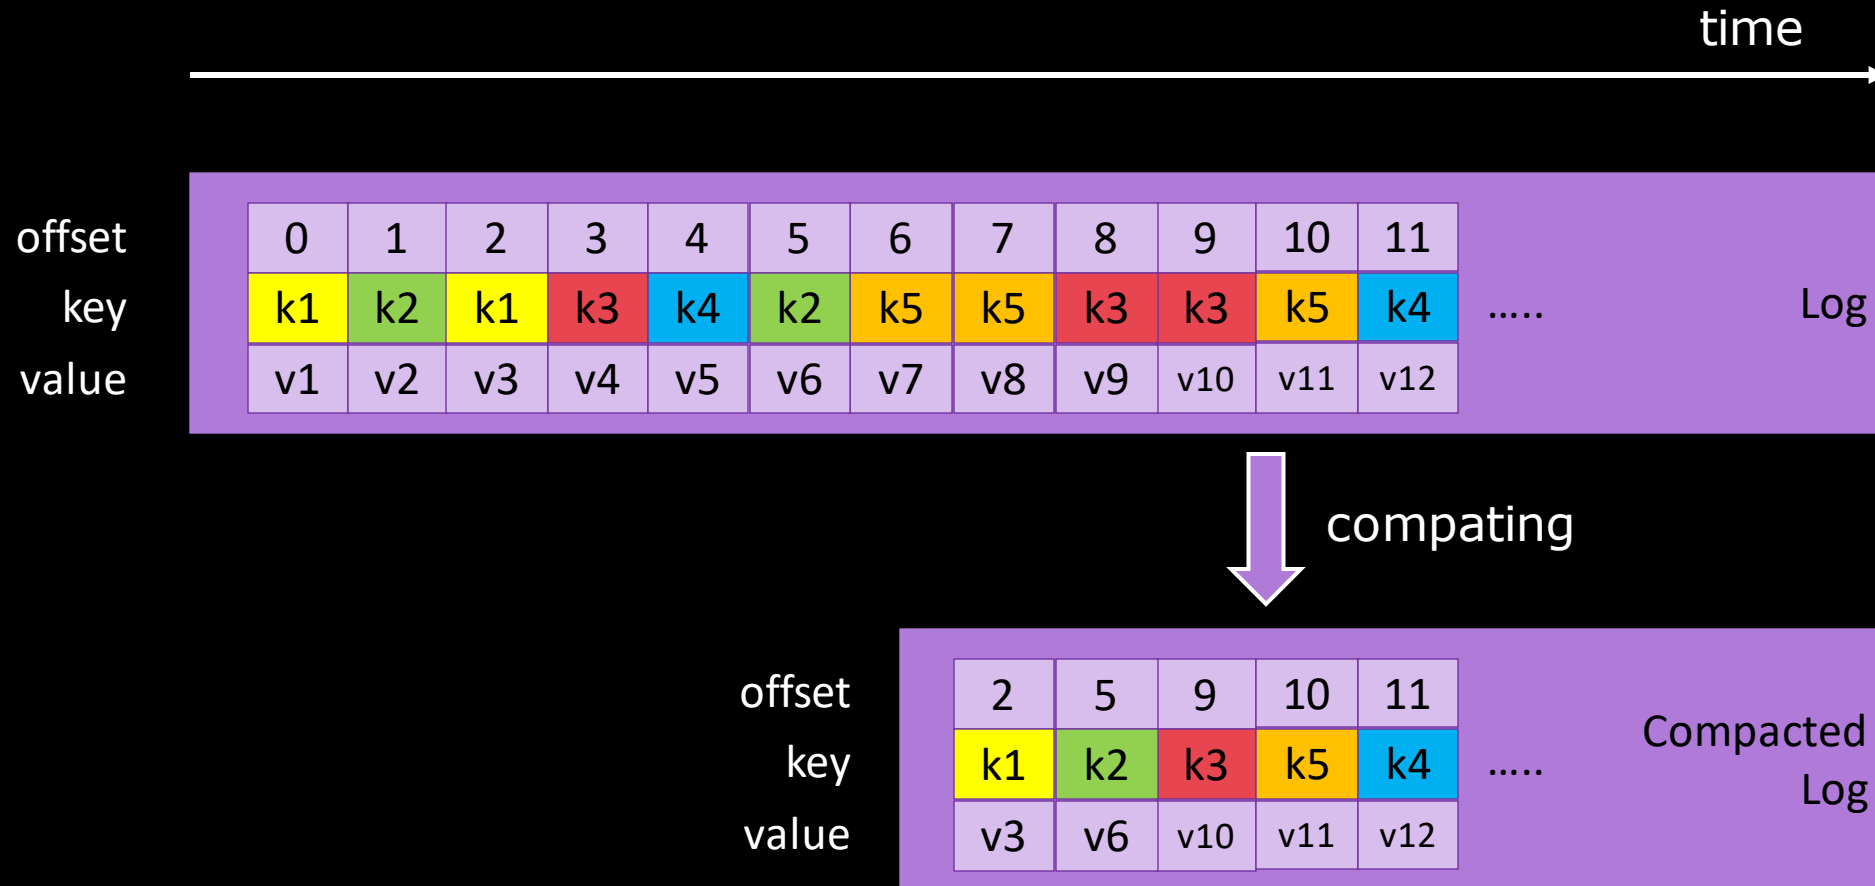
Brokers in several Data Centers



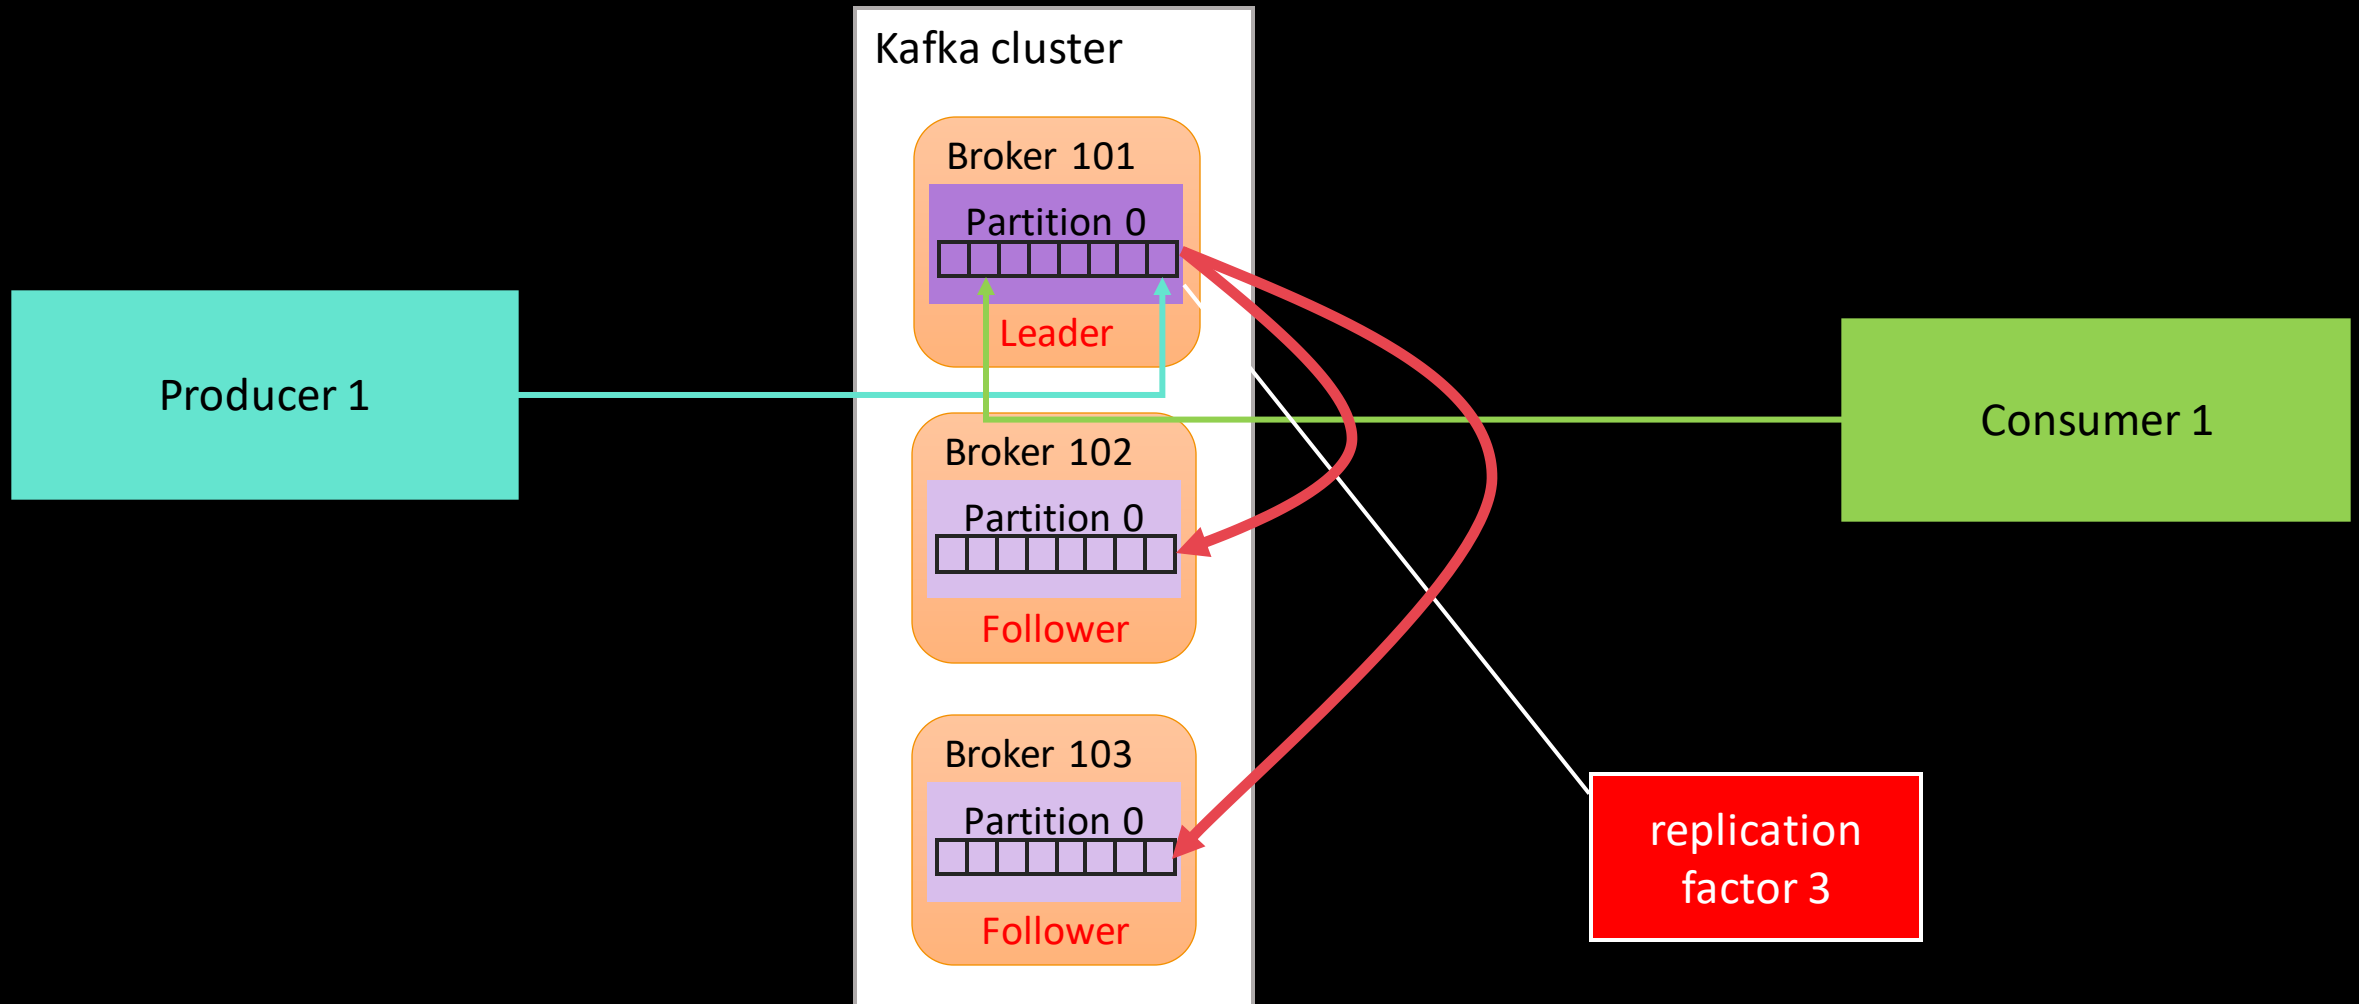
Managing Log File Growth Retention-Policies

- Cleanup.policy
 - delete (default)
 - Segments too old: retention.ms (default 7 days)
 - Partitions too large: retention.bytes (default: -1 unlimited)
 - compact (keep only the freshest value)
 - delete and compact (Example: Order – Management)
- Cleanup applies to Segment-Files.
 - messages are guaranteed to live at least as long as retention time
 - only non-active segment files get deleted upon Cleanup

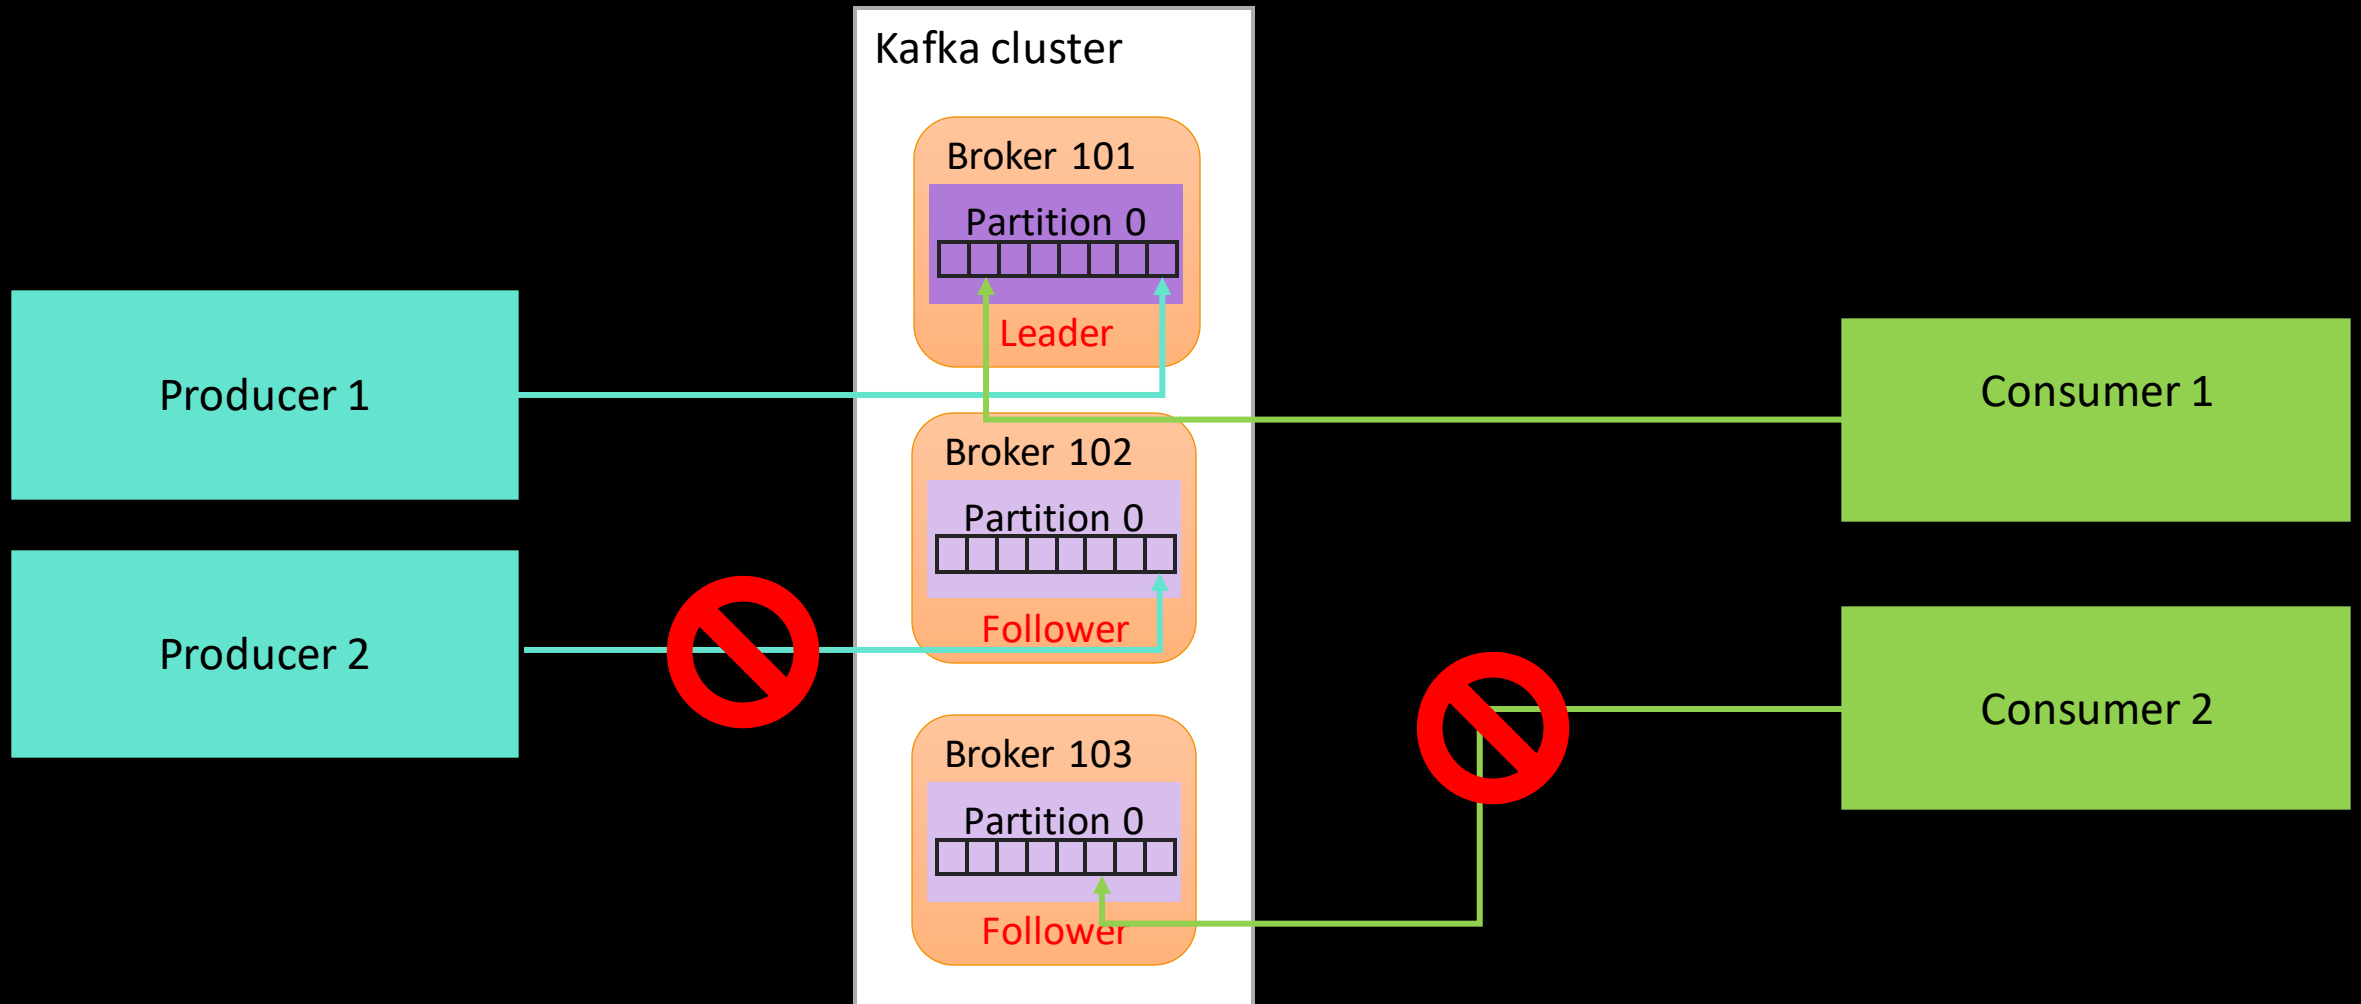
Log Compaction



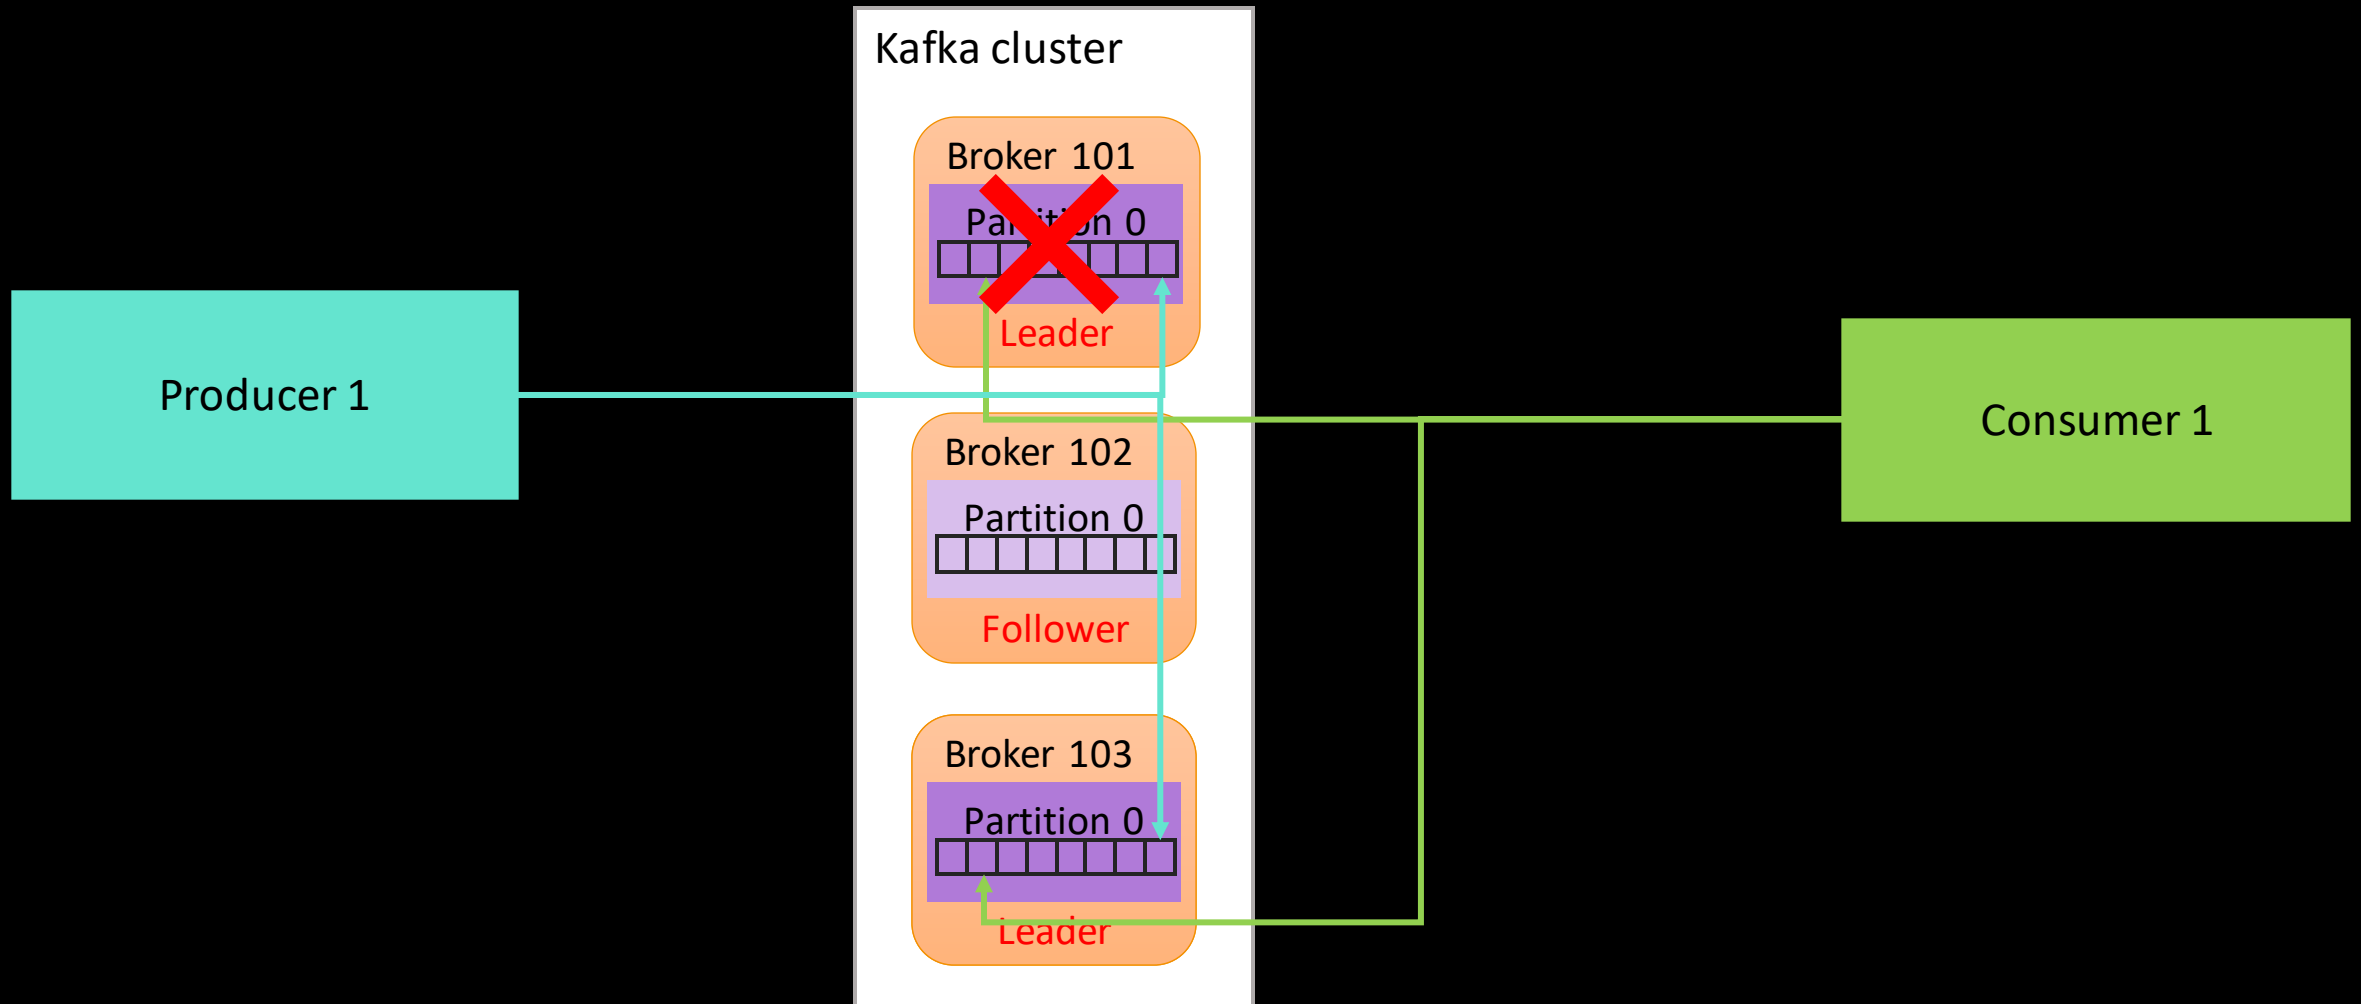
Reliability & Durability: Replication of Partitions



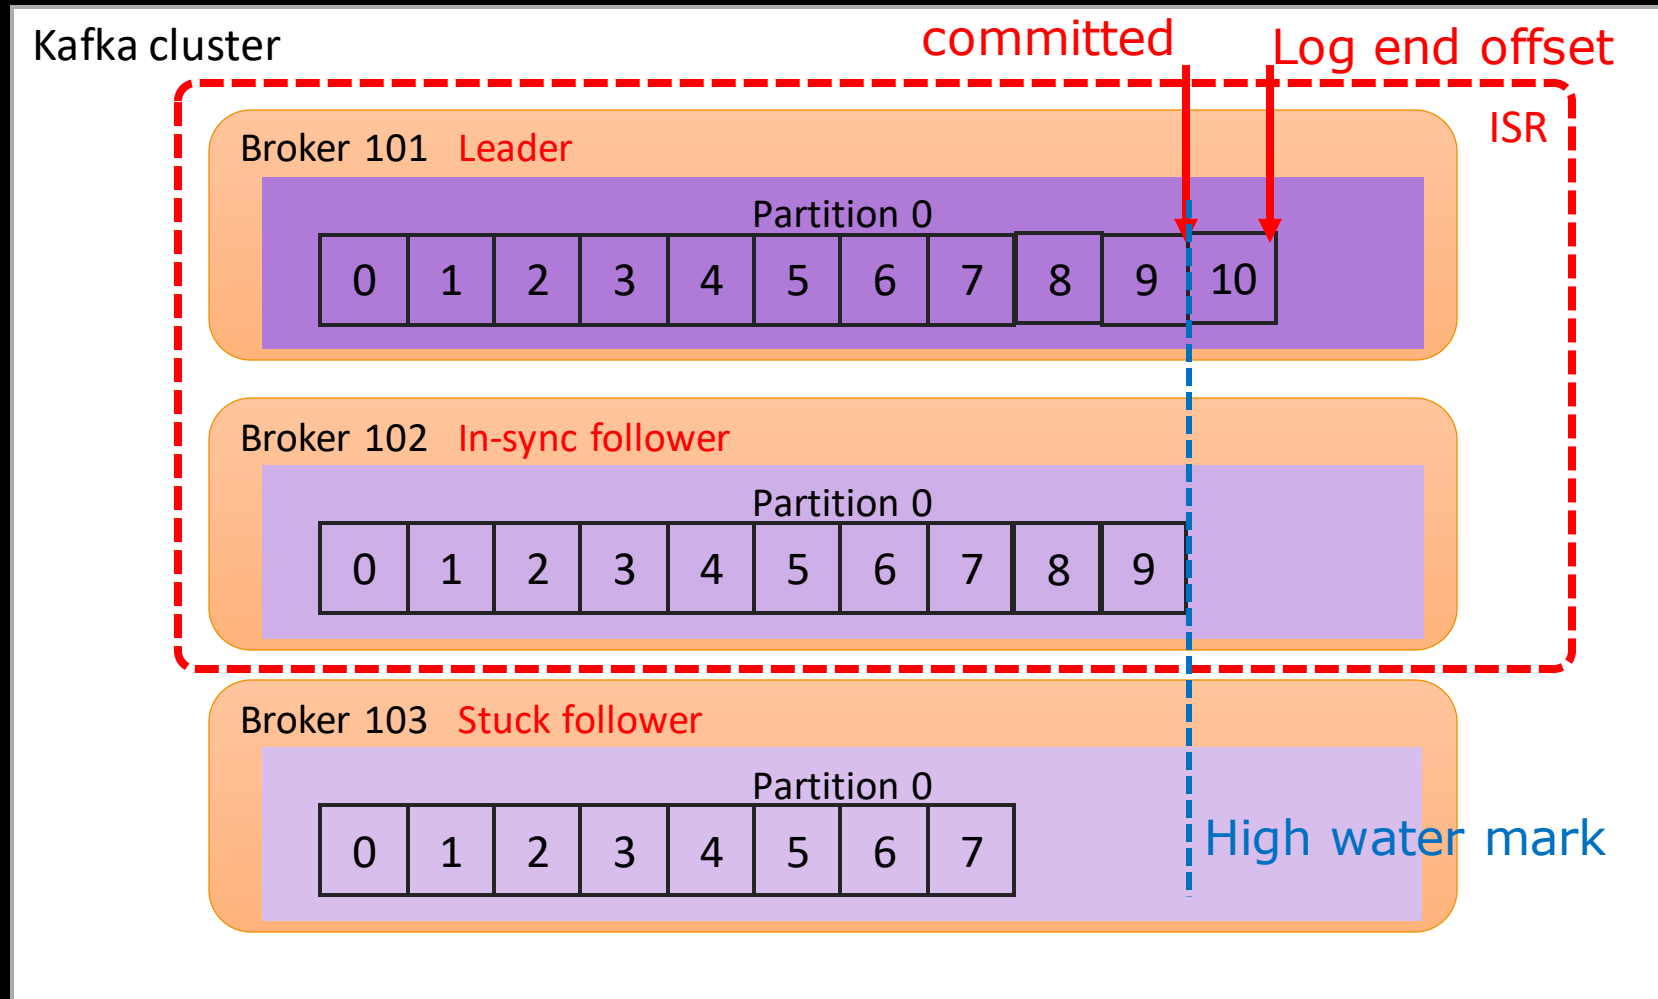
Clients interact with Leaders



Leader Failover



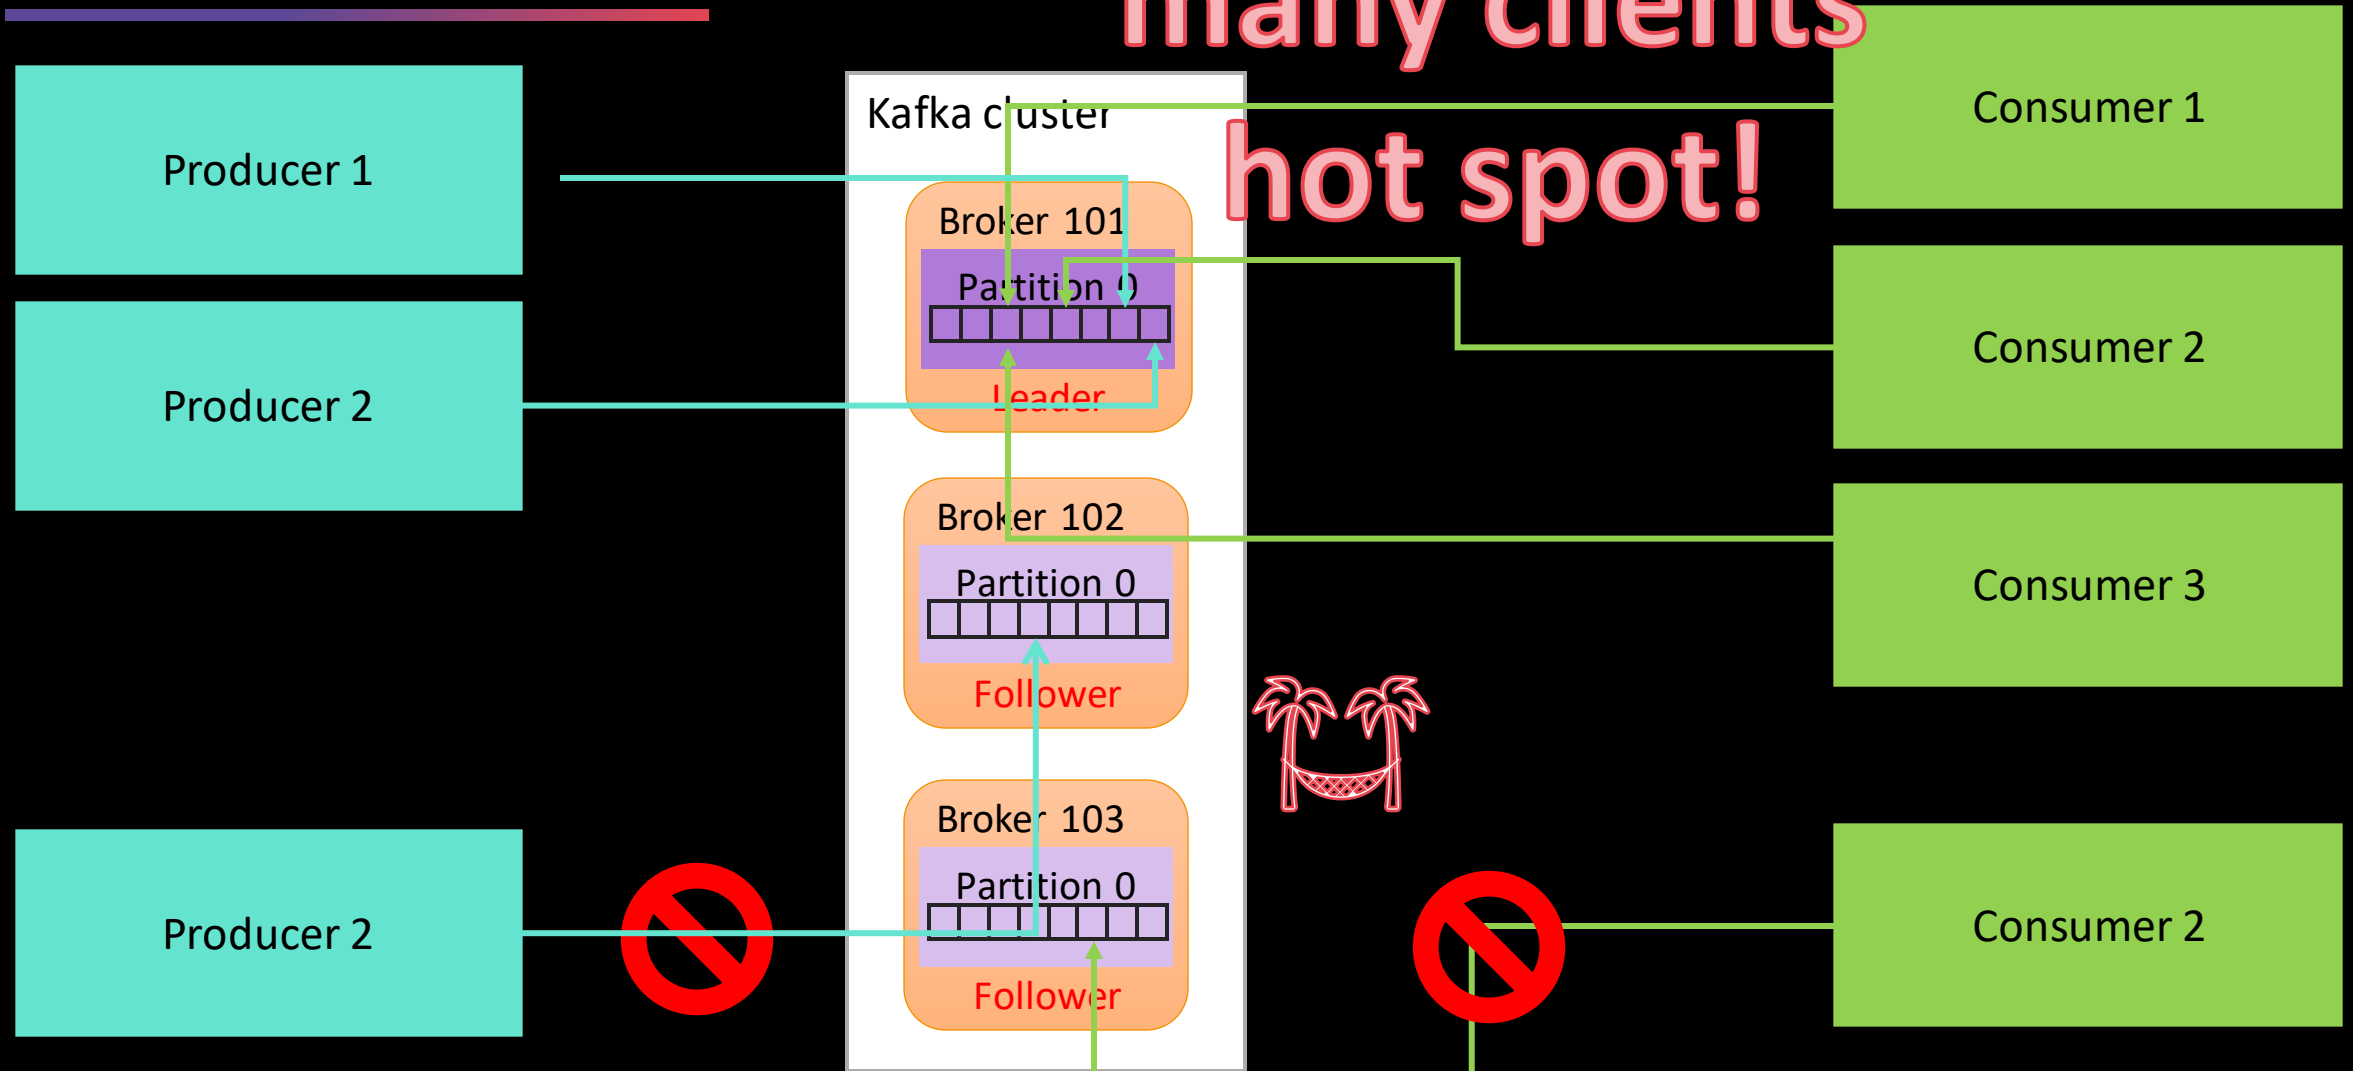
In-Sync Replicas



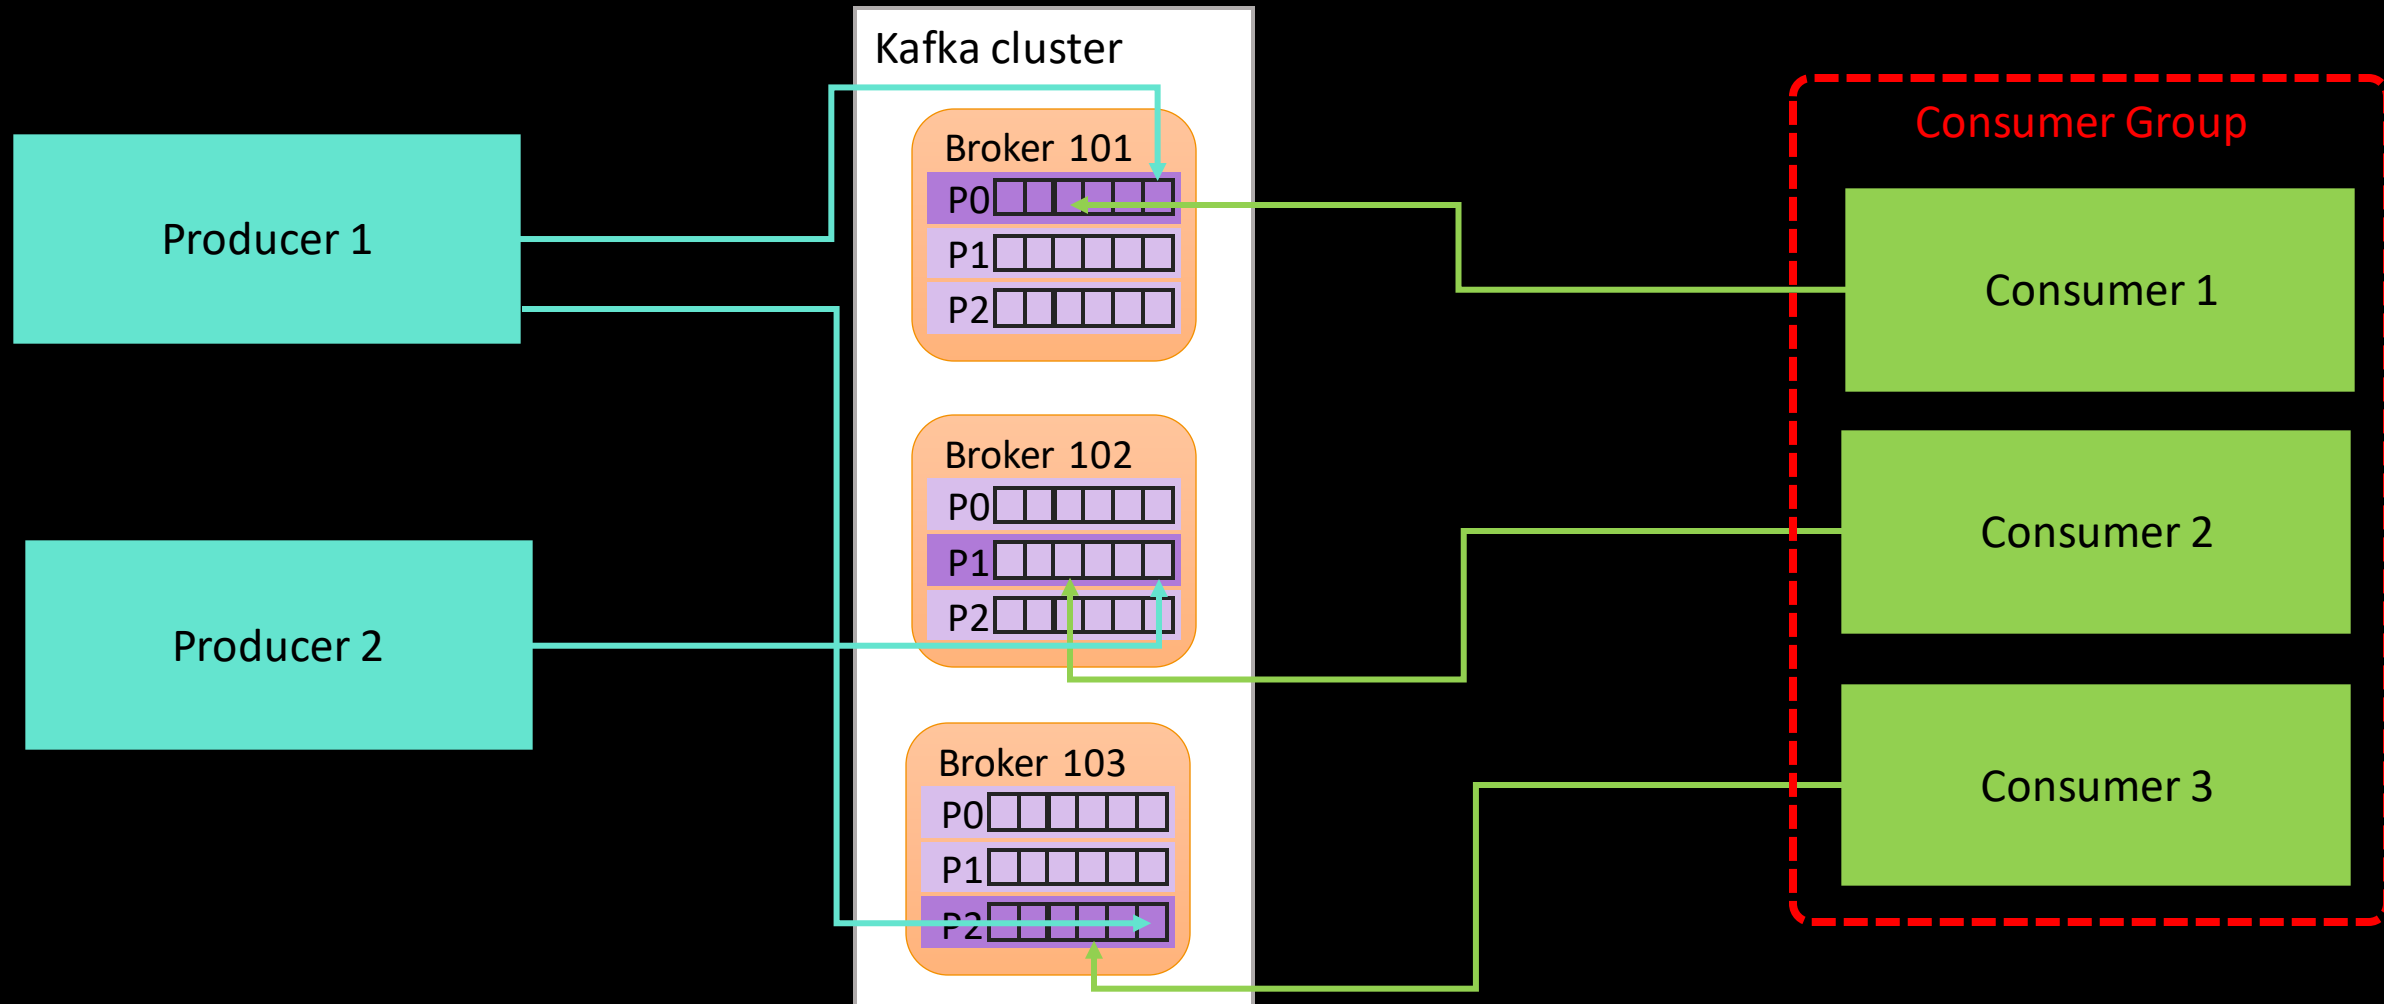
Load Balancing: Partition Leadership

many clients

hot spot!



Load Balancing Partitions Leadership (2)



Zookeeper: What is it good for?

Controller
Election

Cluster
Membership

Topic
Configuration

Quotas

Access
Control Lists

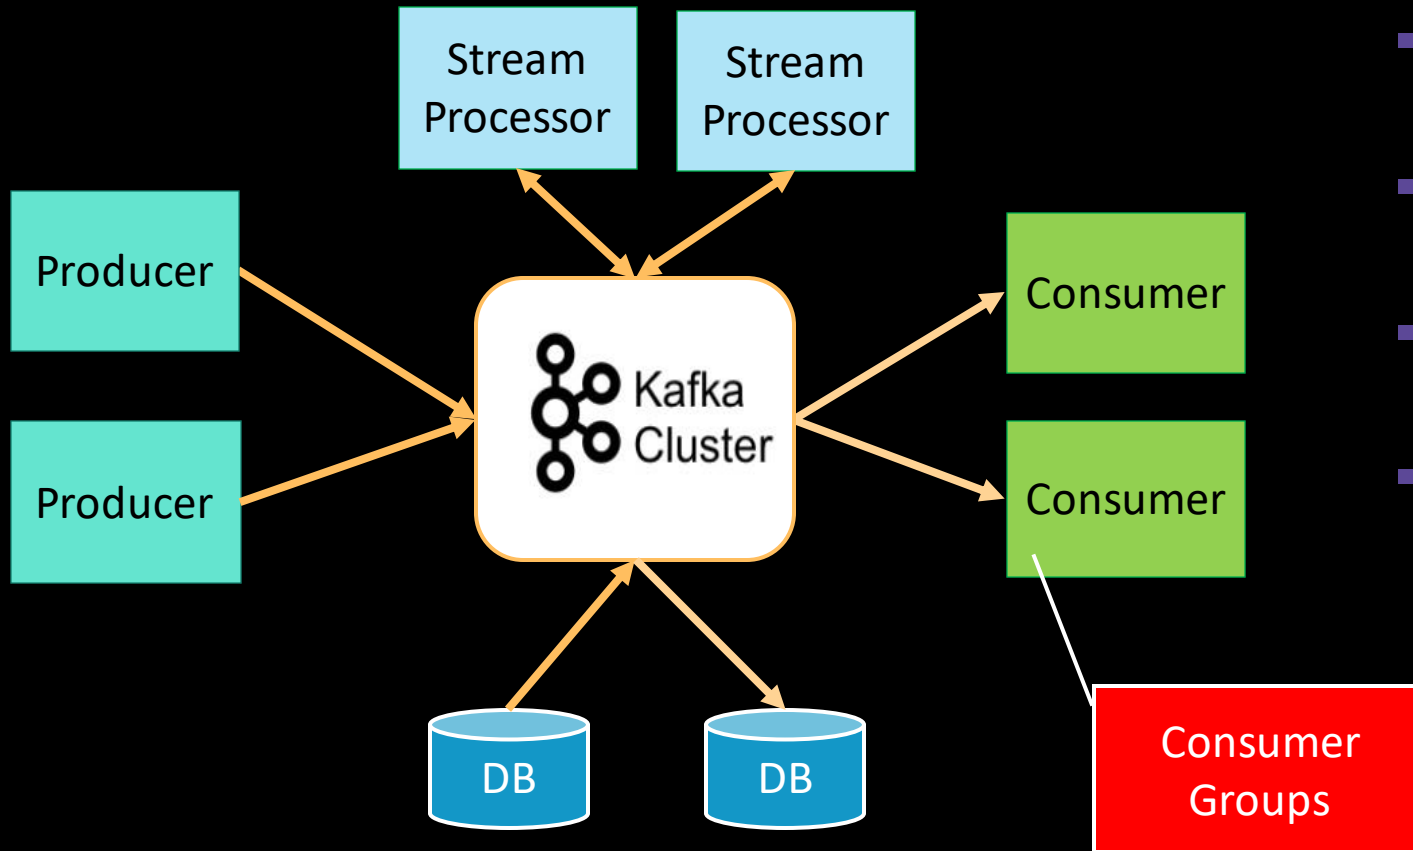
Content

- Integration
- Event-Streaming Platform Kafka
- Asynchronous Communication
- Kafka Basics & Components
- Kafka APIs
- Kafka Examples
- Kafka Exercises



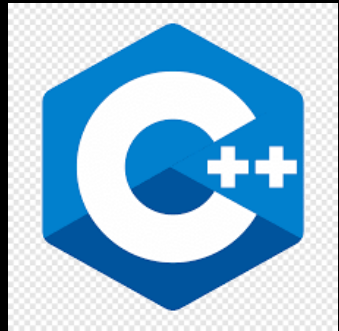
Kafka APIs

Kafka Core Components: 4 Core-APIs



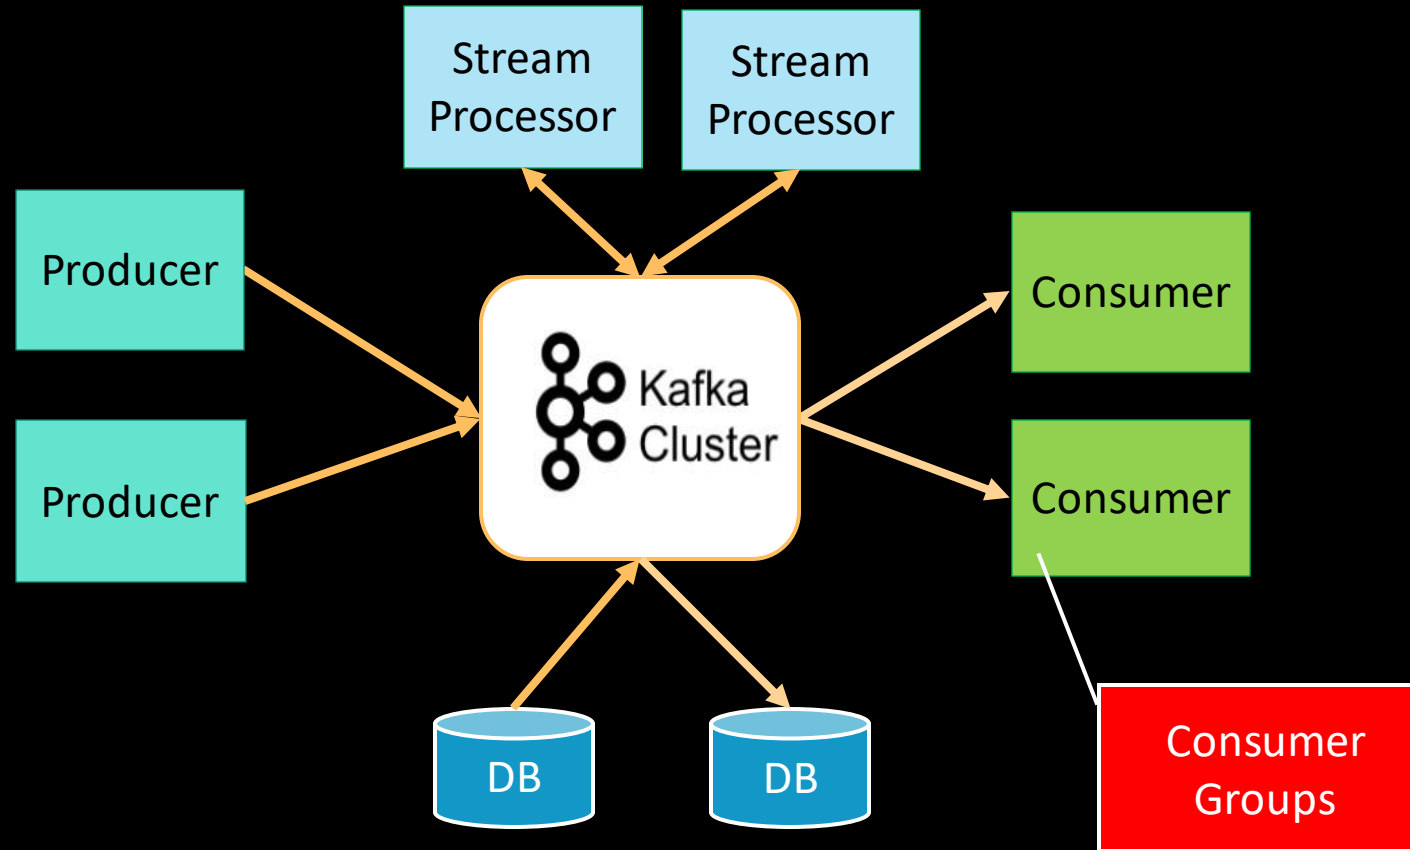
- Producer API: Enables to write messages
- Consumer API: Enables to read messages
- Streams API: Enables to analyze and transform messages
- Connect API: Enables the creation of reusable Clients

Kafka Clients supported by Confluent

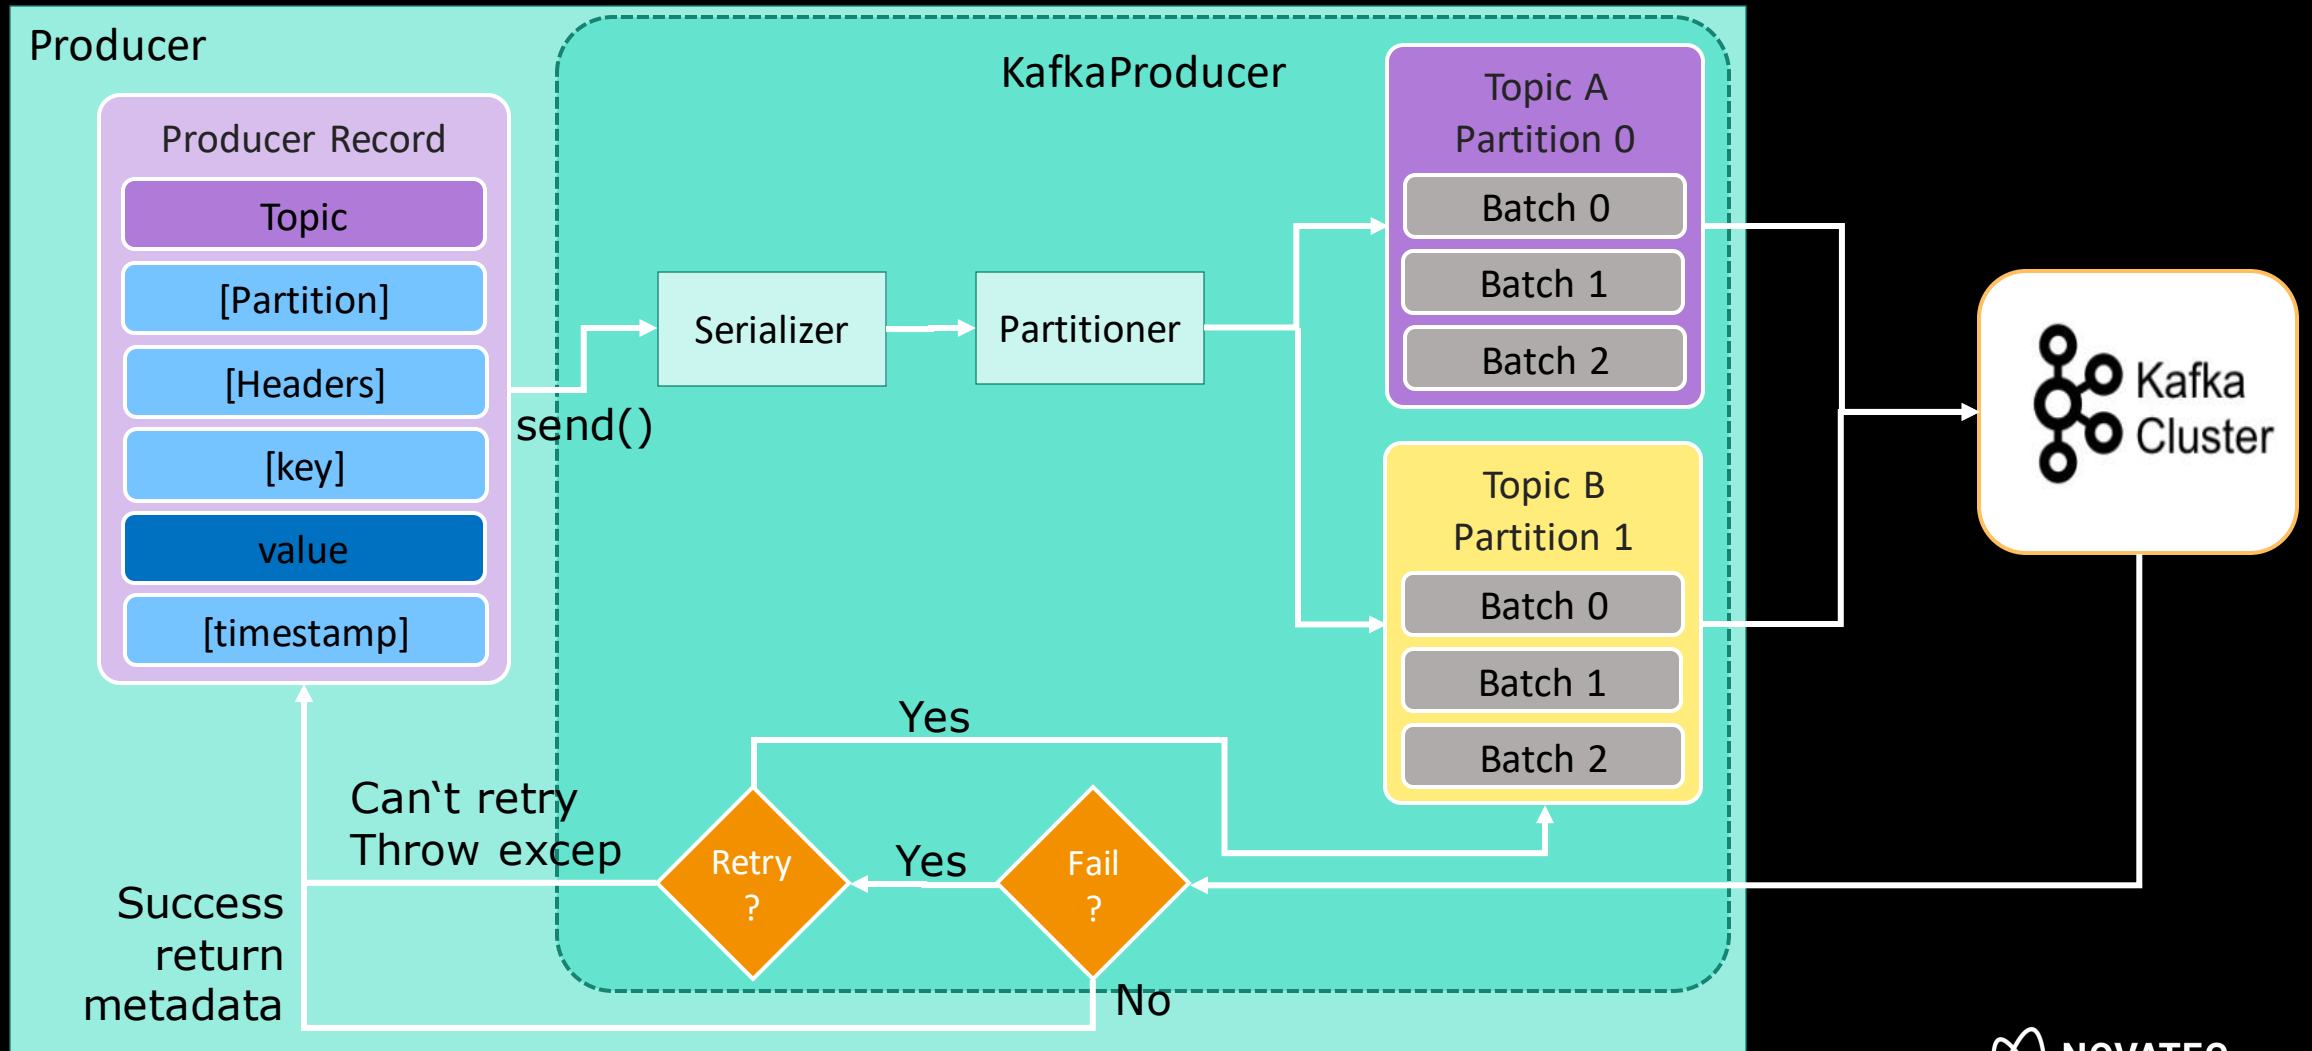


- <https://github.com/confluentinc/examples/tree/5.3.1-post/clients/cloud>
- client code in many languages
- JVM: Java, Groovy, Scala, Kotlin, Clojure
- C-Library: C, C#, Go, NodeJS, Python, Ruby,

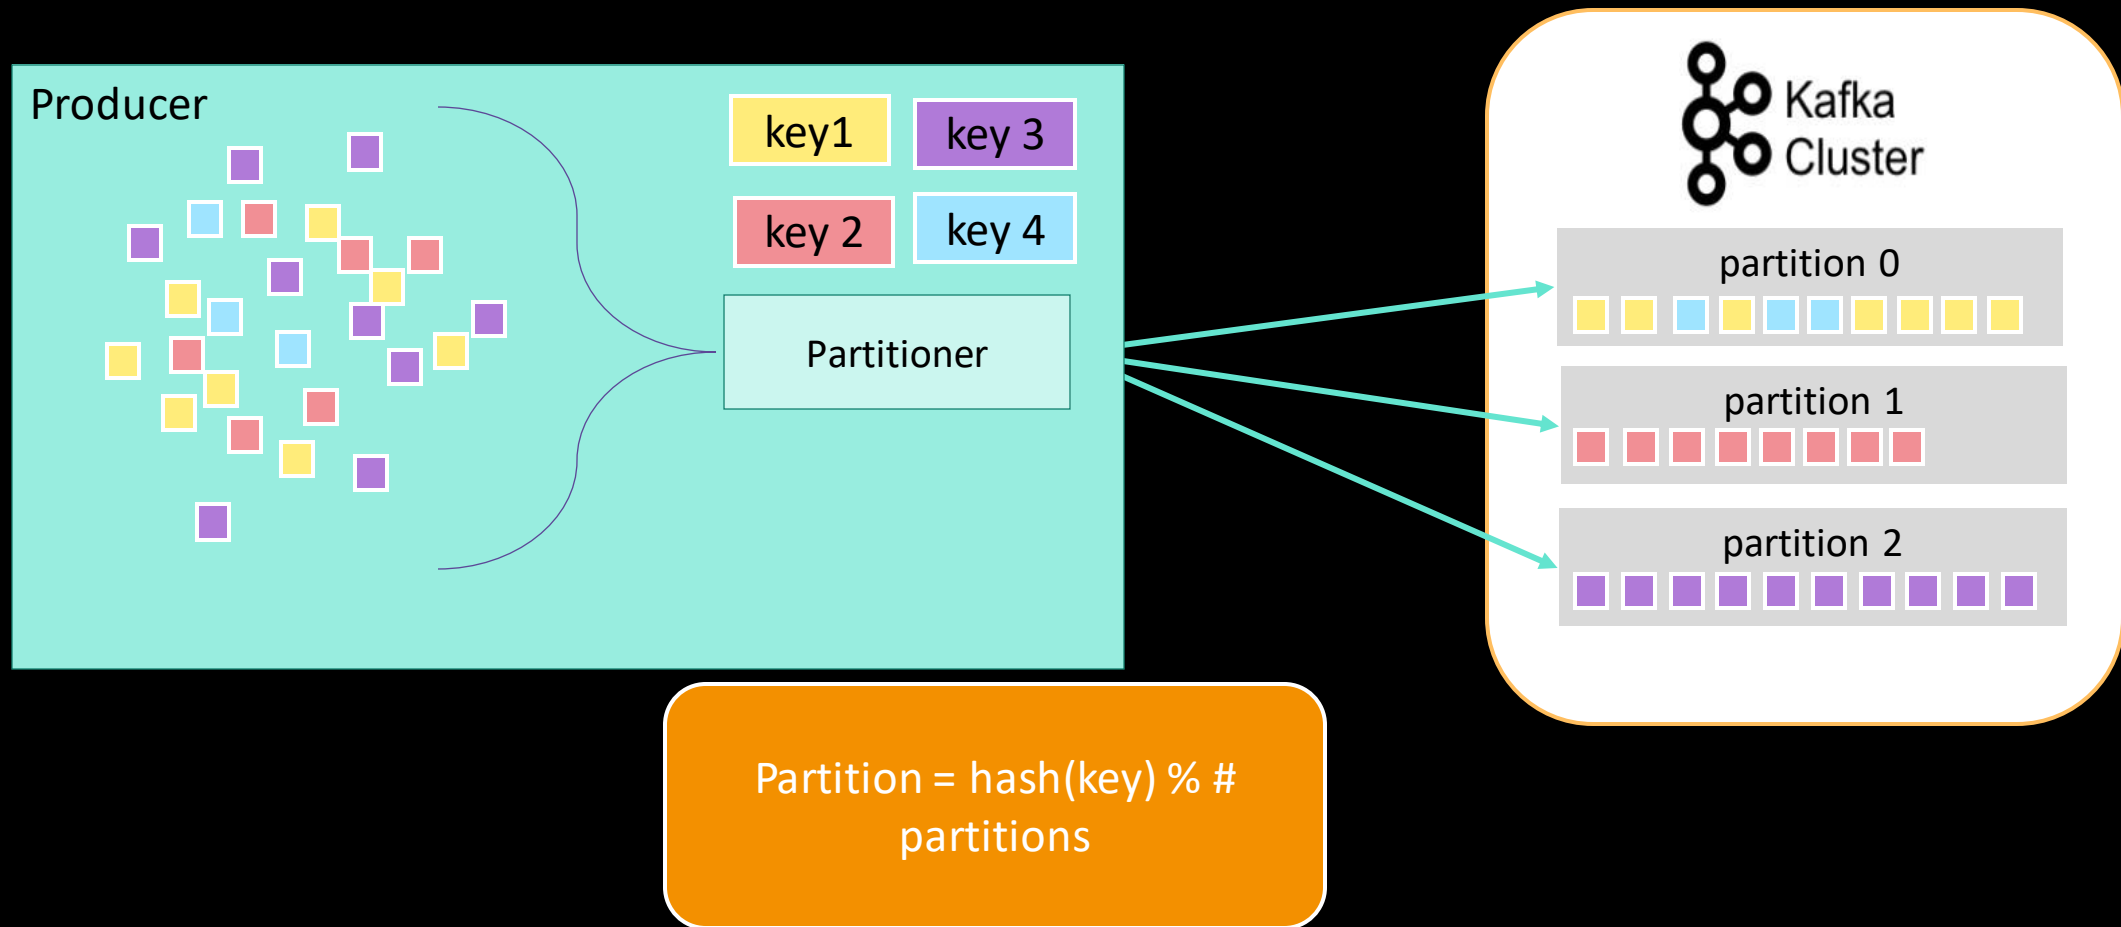
Kafka Core Components: Producer API



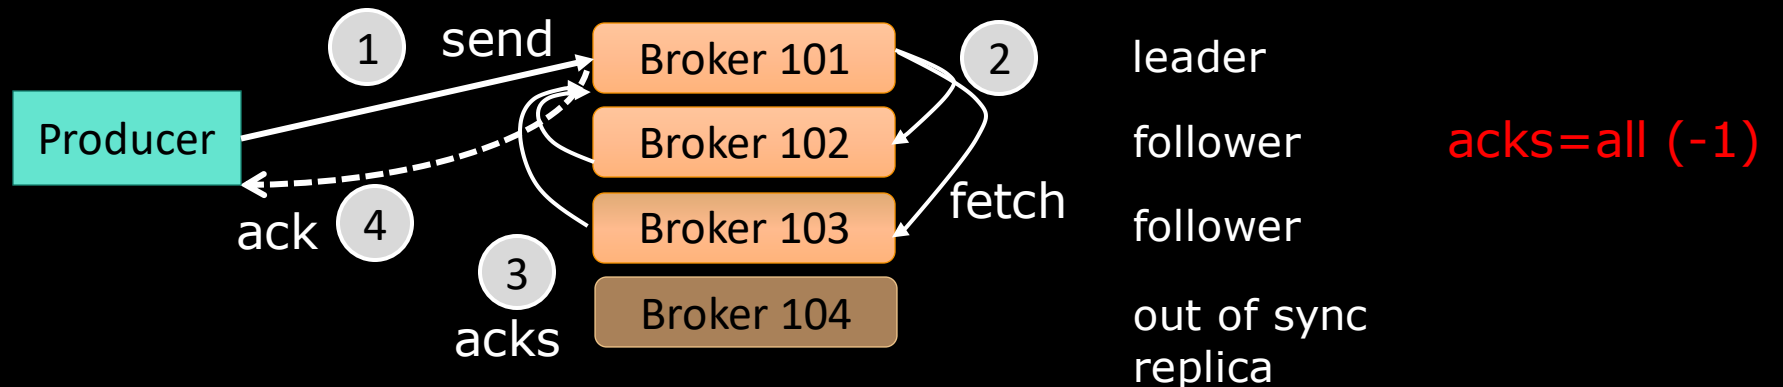
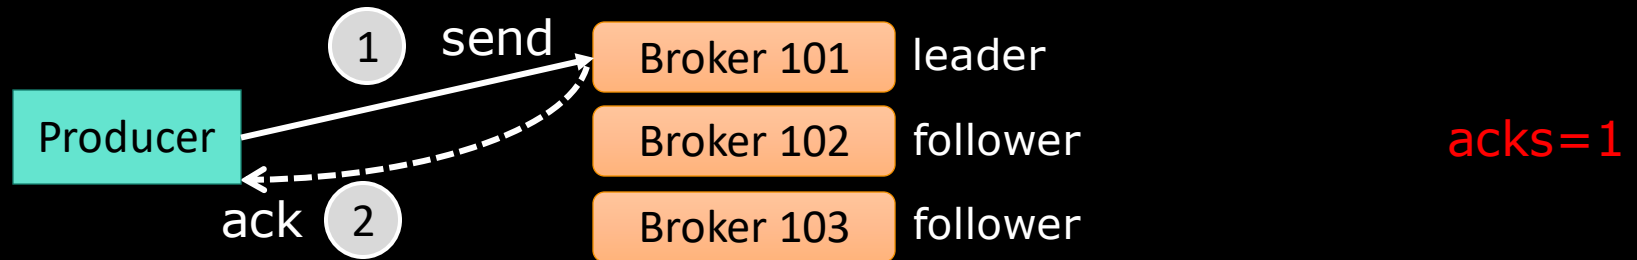
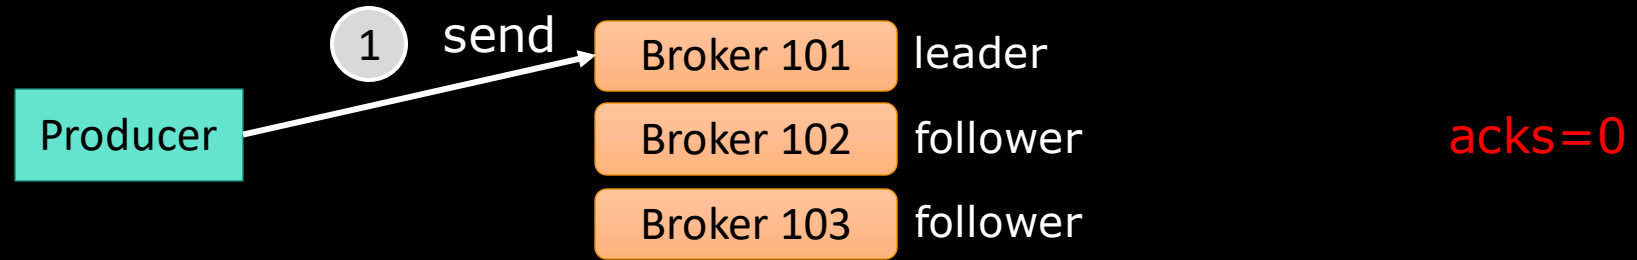
Producer API: High Level Architecture & Design



Default Partitioner



Acknowledgement



Development: A Basic Producer in Java

```
BasicProducer.java x
1 package clients;
2
3 import java.util.Properties;
4 import org.apache.kafka.clients.producer.KafkaProducer;
5 import org.apache.kafka.clients.producer.ProducerRecord;
6
7 public class BasicProducer {
8     public static void main(String[] args) {
9         System.out.println("*** Starting Basic Producer ***");
10
11         Properties settings = new Properties();
12         settings.put("client.id", "basic-producer-v0.1.0");
13         settings.put("bootstrap.servers", "kafka-1:9092,kafka-2:9092");
14         settings.put("key.serializer", "org.apache.kafka.common.serialization.StringSerializer");
15         settings.put("value.serializer", "org.apache.kafka.common.serialization.StringSerializer");
16
17         final KafkaProducer<String, String> producer = new KafkaProducer<>(settings);
18
19         Runtime.getRuntime().addShutdownHook(new Thread(() -> {
20             System.out.println("### Stopping Basic Producer ###");
21             producer.close();
22         }));
23
24         final String topic = "hello_world_topic";
25         for(int i=1; i<=5; i++){
26             final String key = "key-" + i;
27             final String value = "value-" + i;
28             final ProducerRecord<String, String> record = new ProducerRecord<>(topic, key, value);
29             producer.send(record);
30         }
31     }
32 }
```

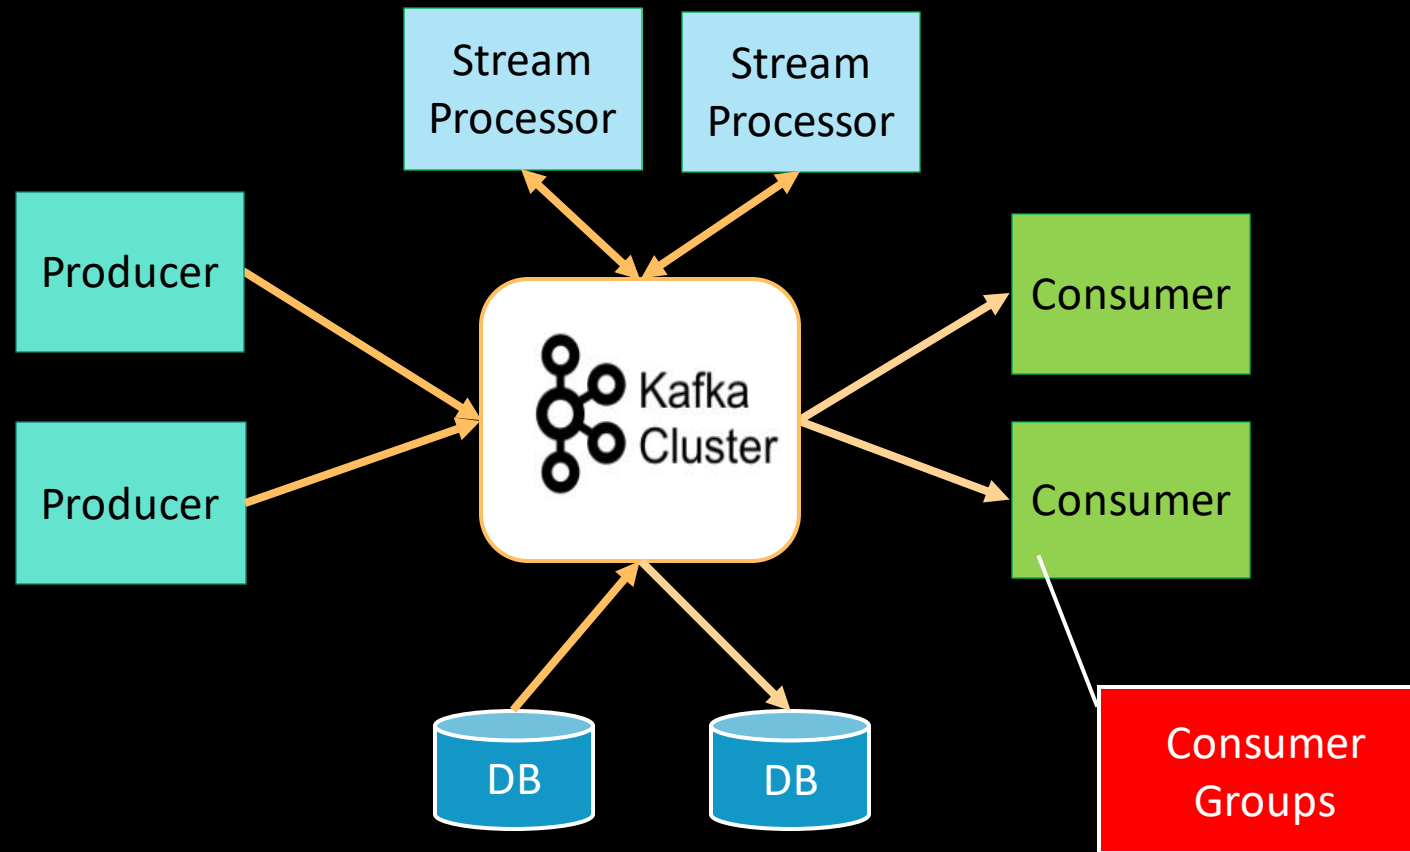
configuration

create producer

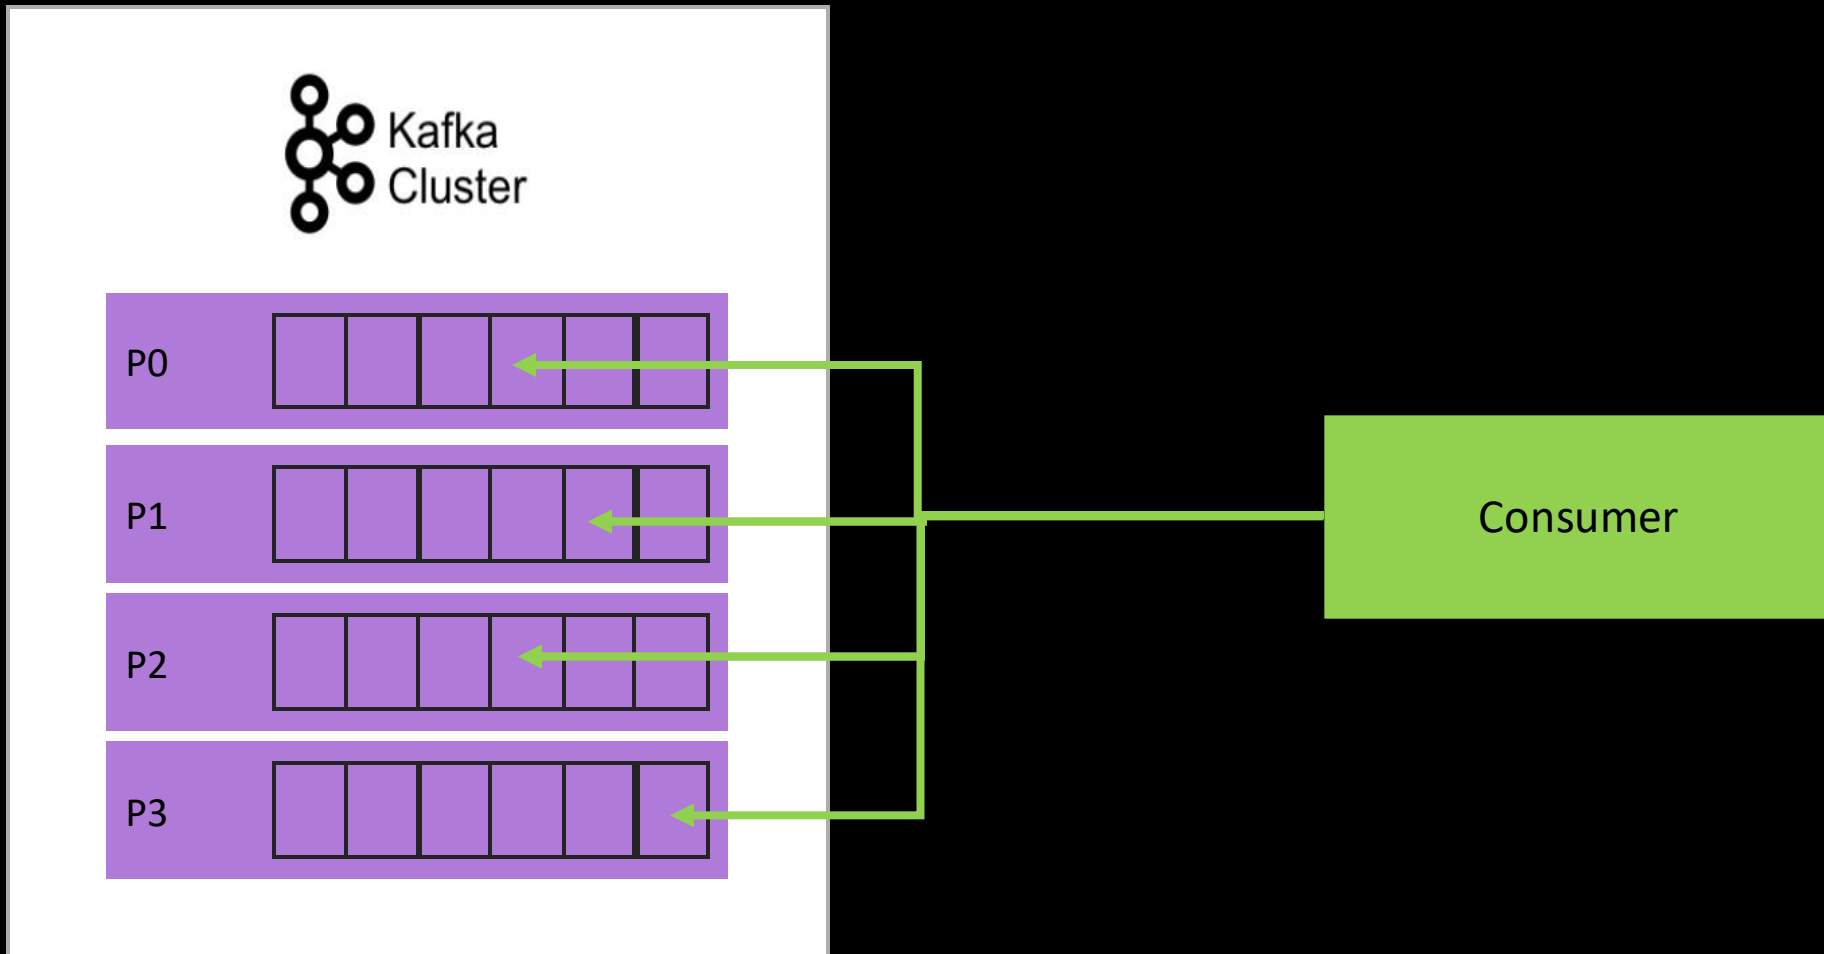
shutdown behaviour

sending data

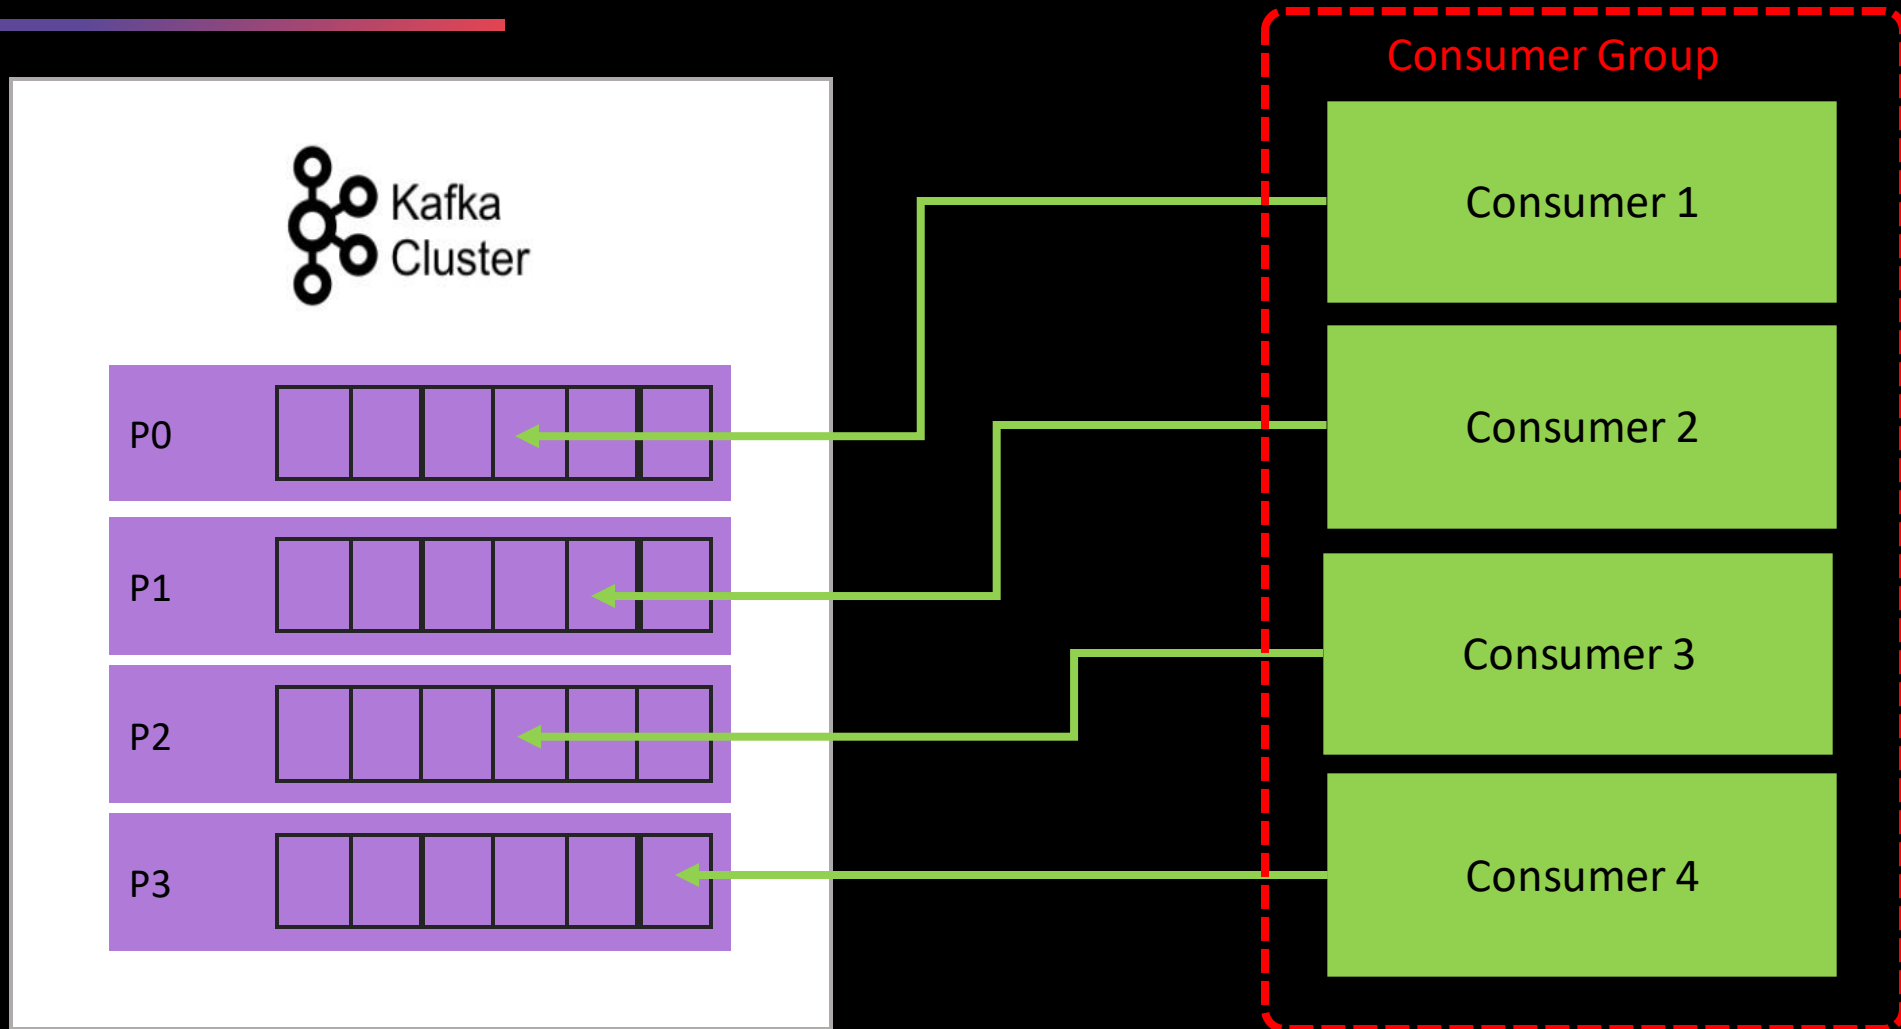
Kafka Core Components: Consumer API



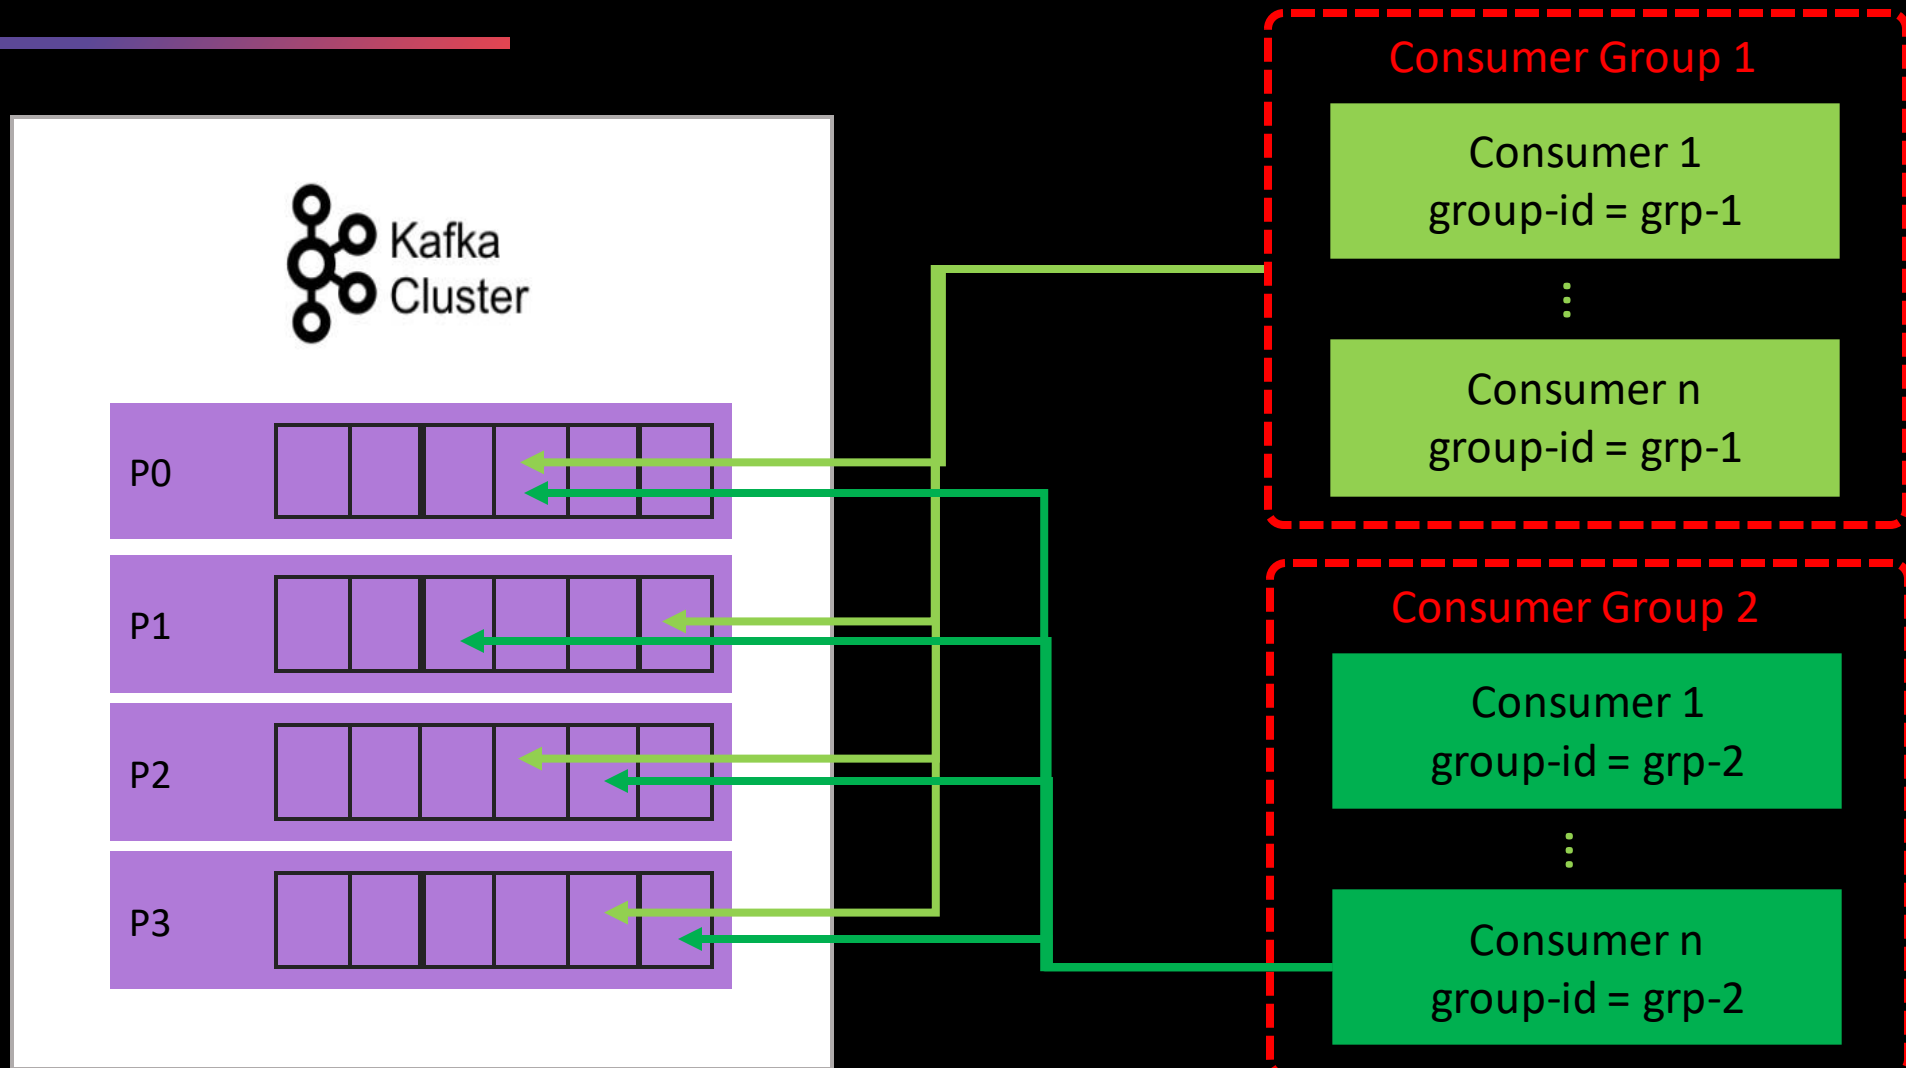
Consuming from Kafka: Single Consumer



Consuming from Kafka: Consumer Group

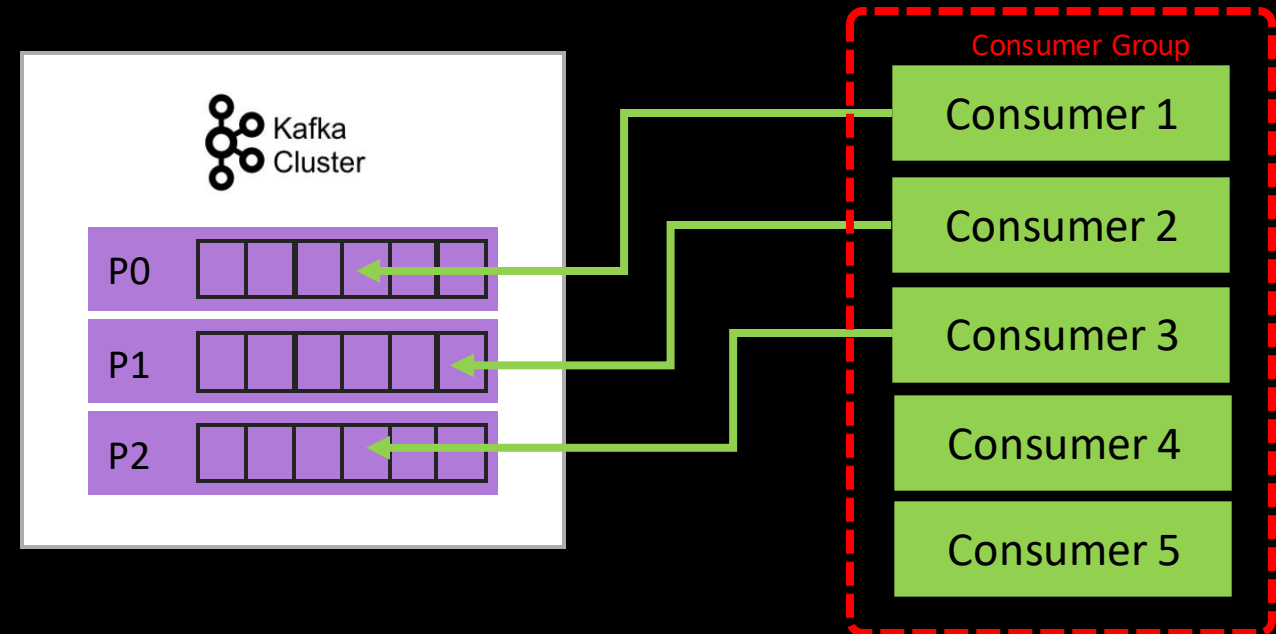
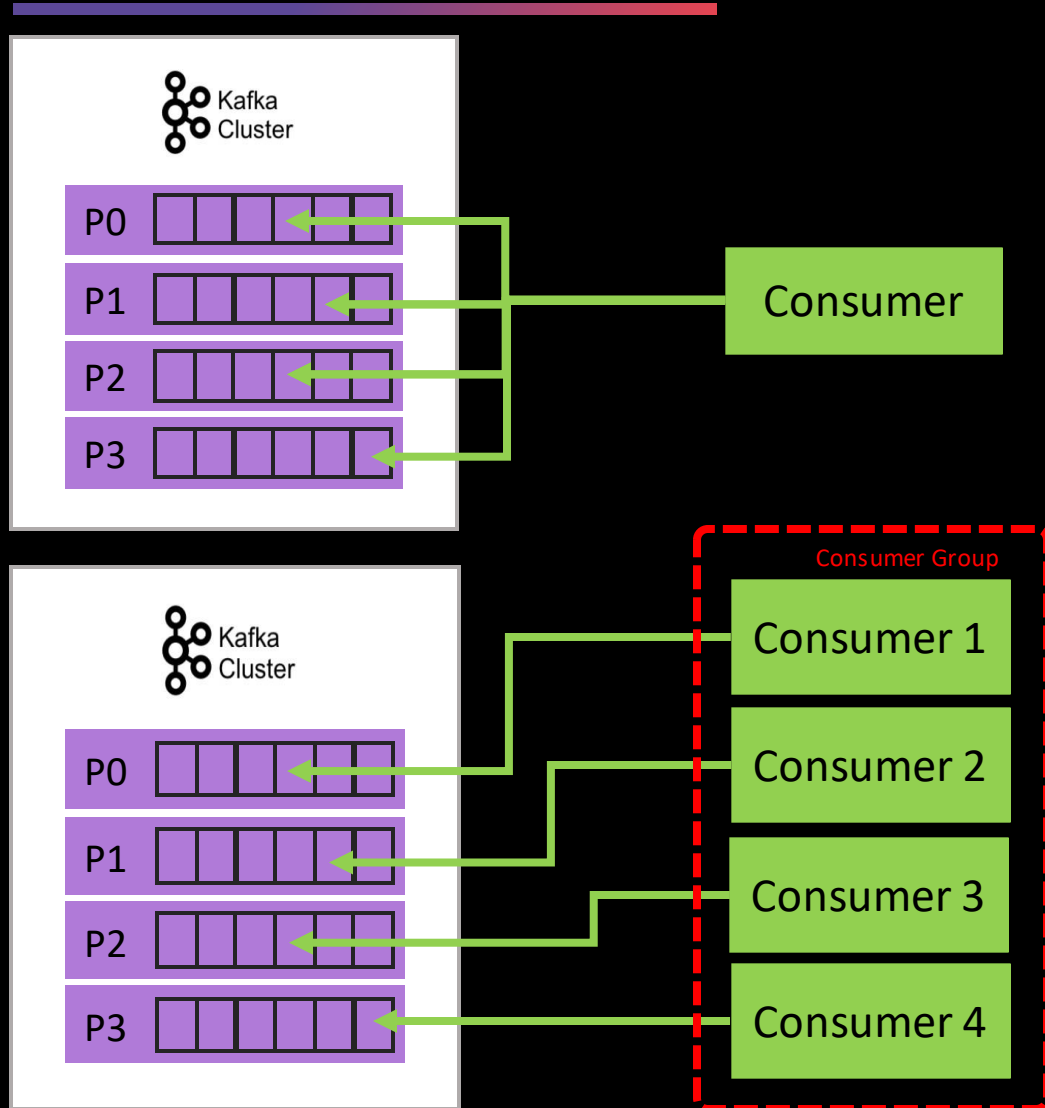


Consuming from Kafka: Multiple Groups

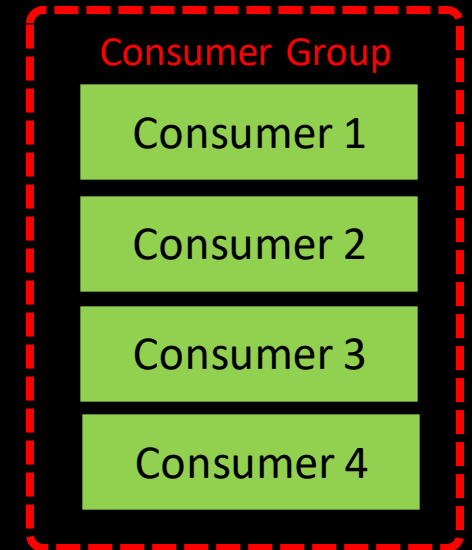
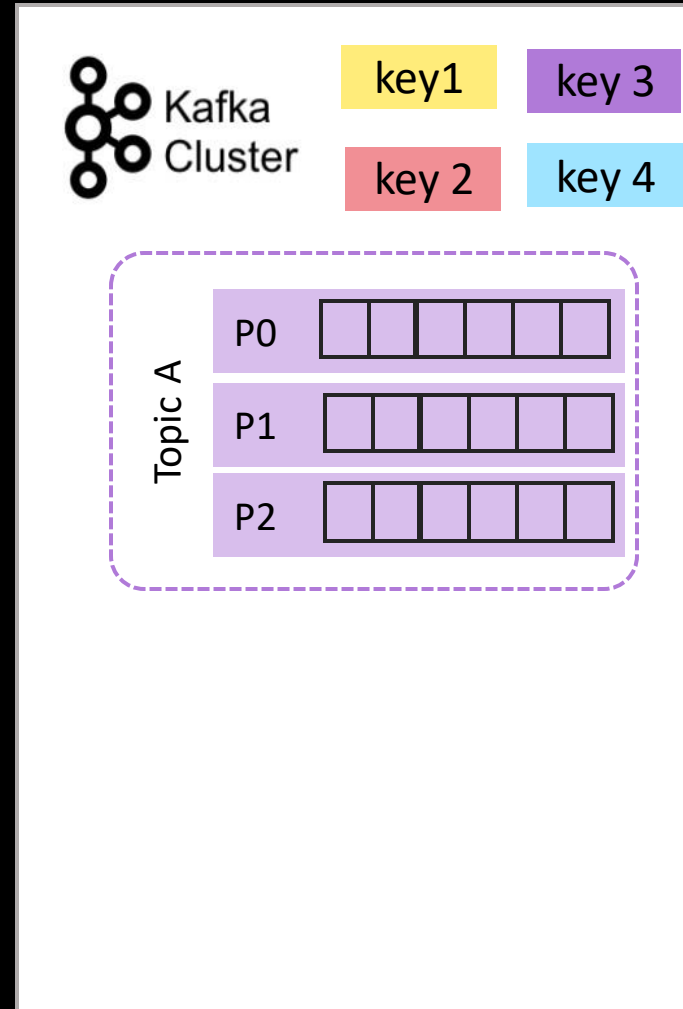
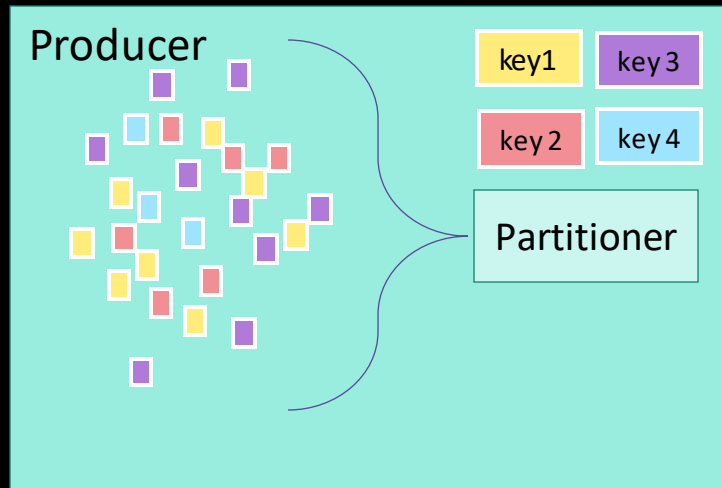


Scalability is limited by Number of Partitions

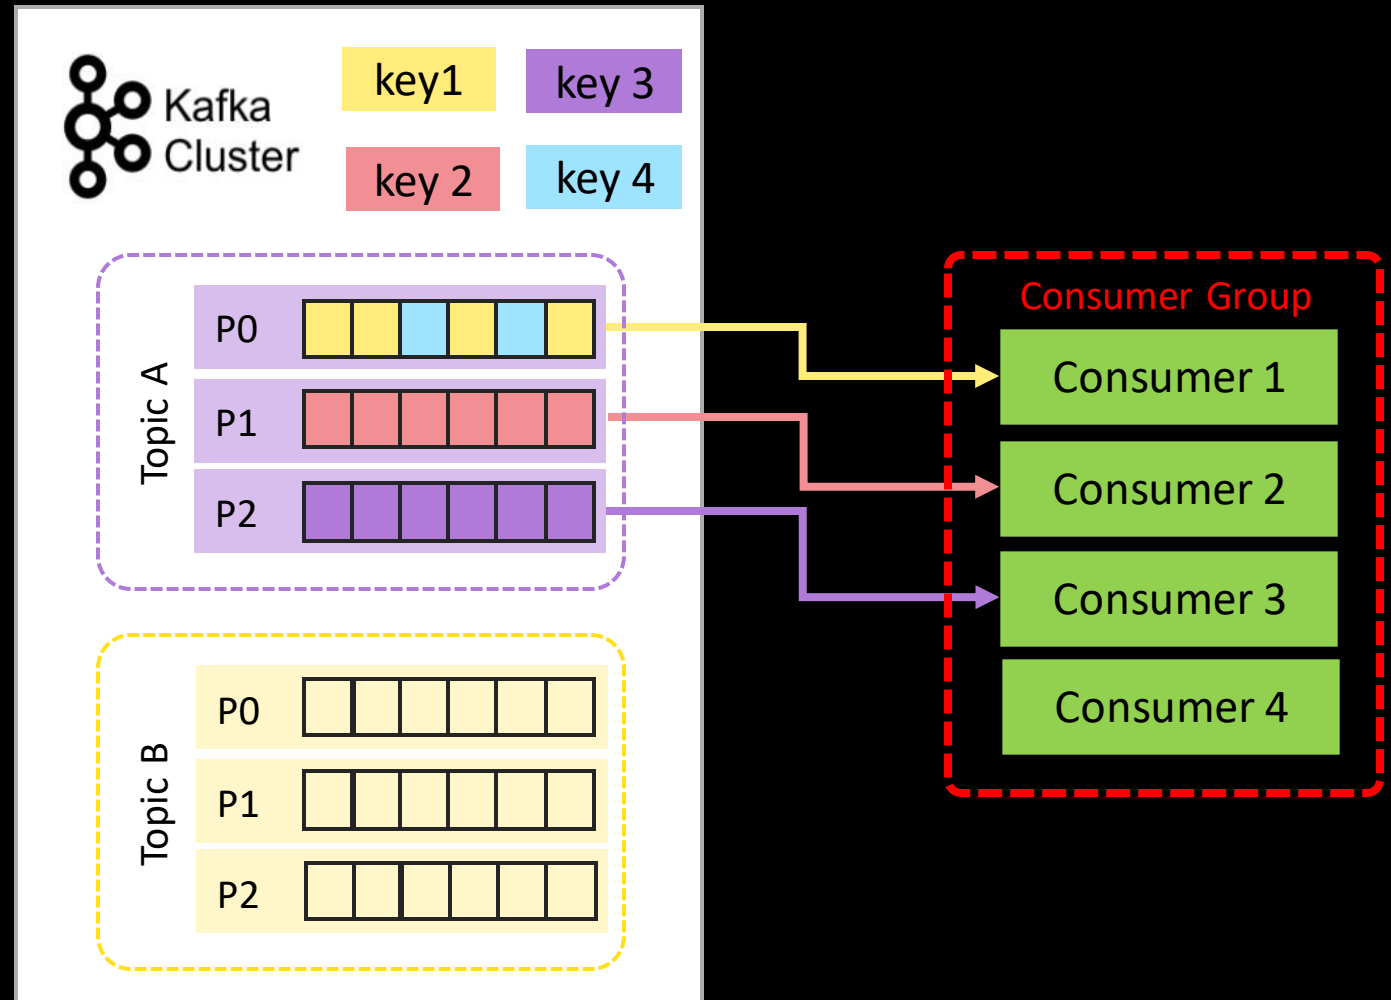
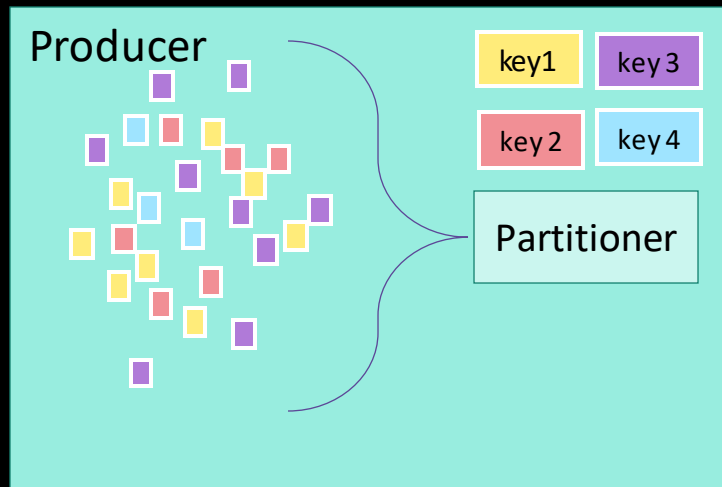
- number of useful consumers in a group is constrained by the number of partitions



How are Partitions assigned to Consumers 1

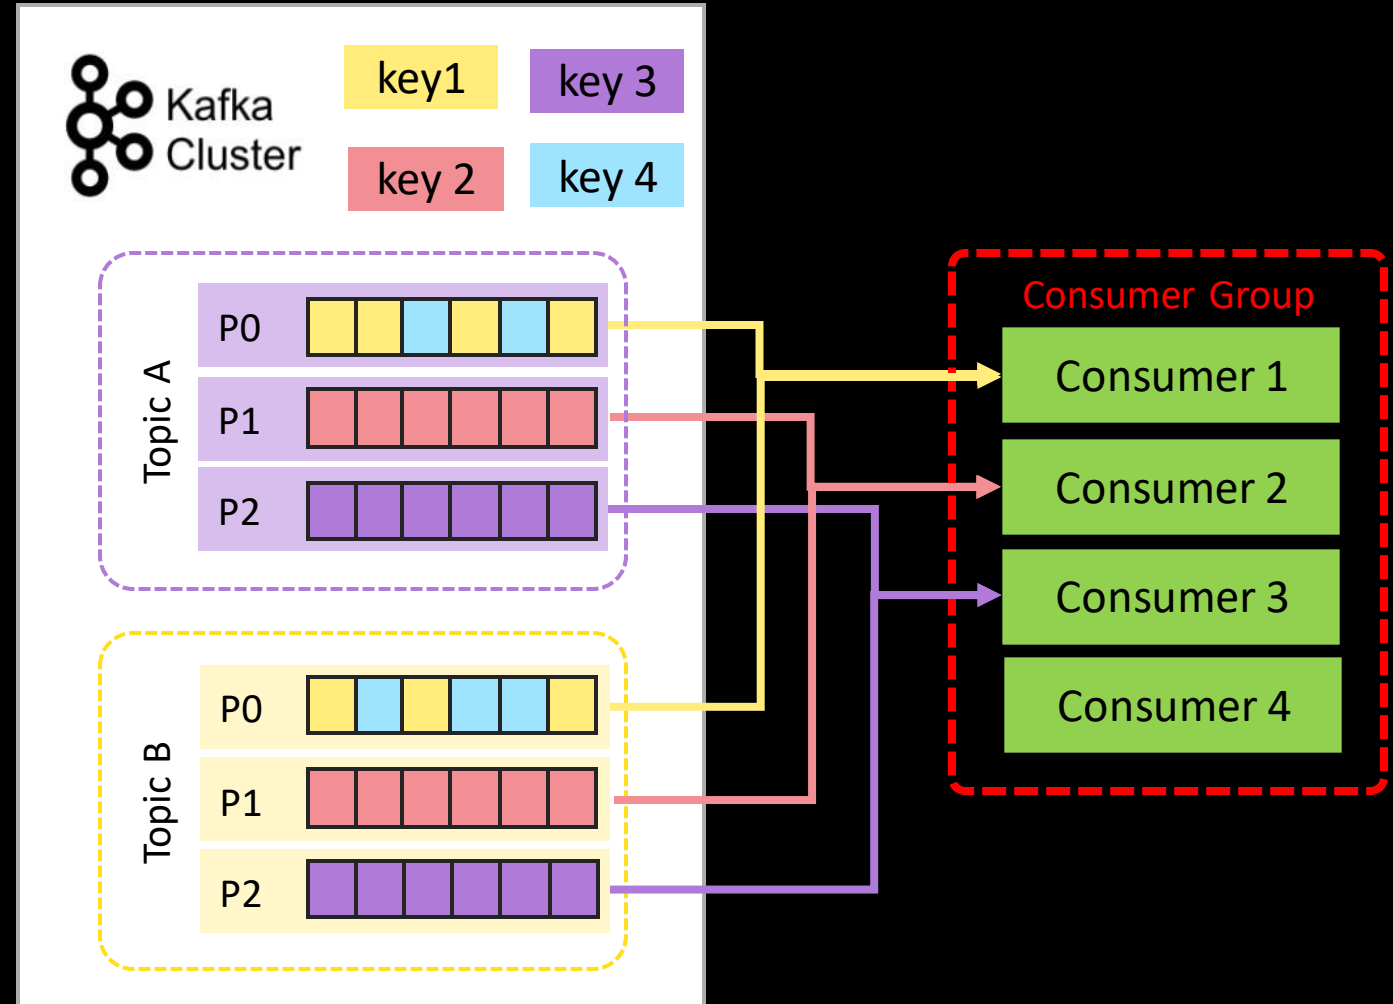


How are Partitions assigned to Consumers 2

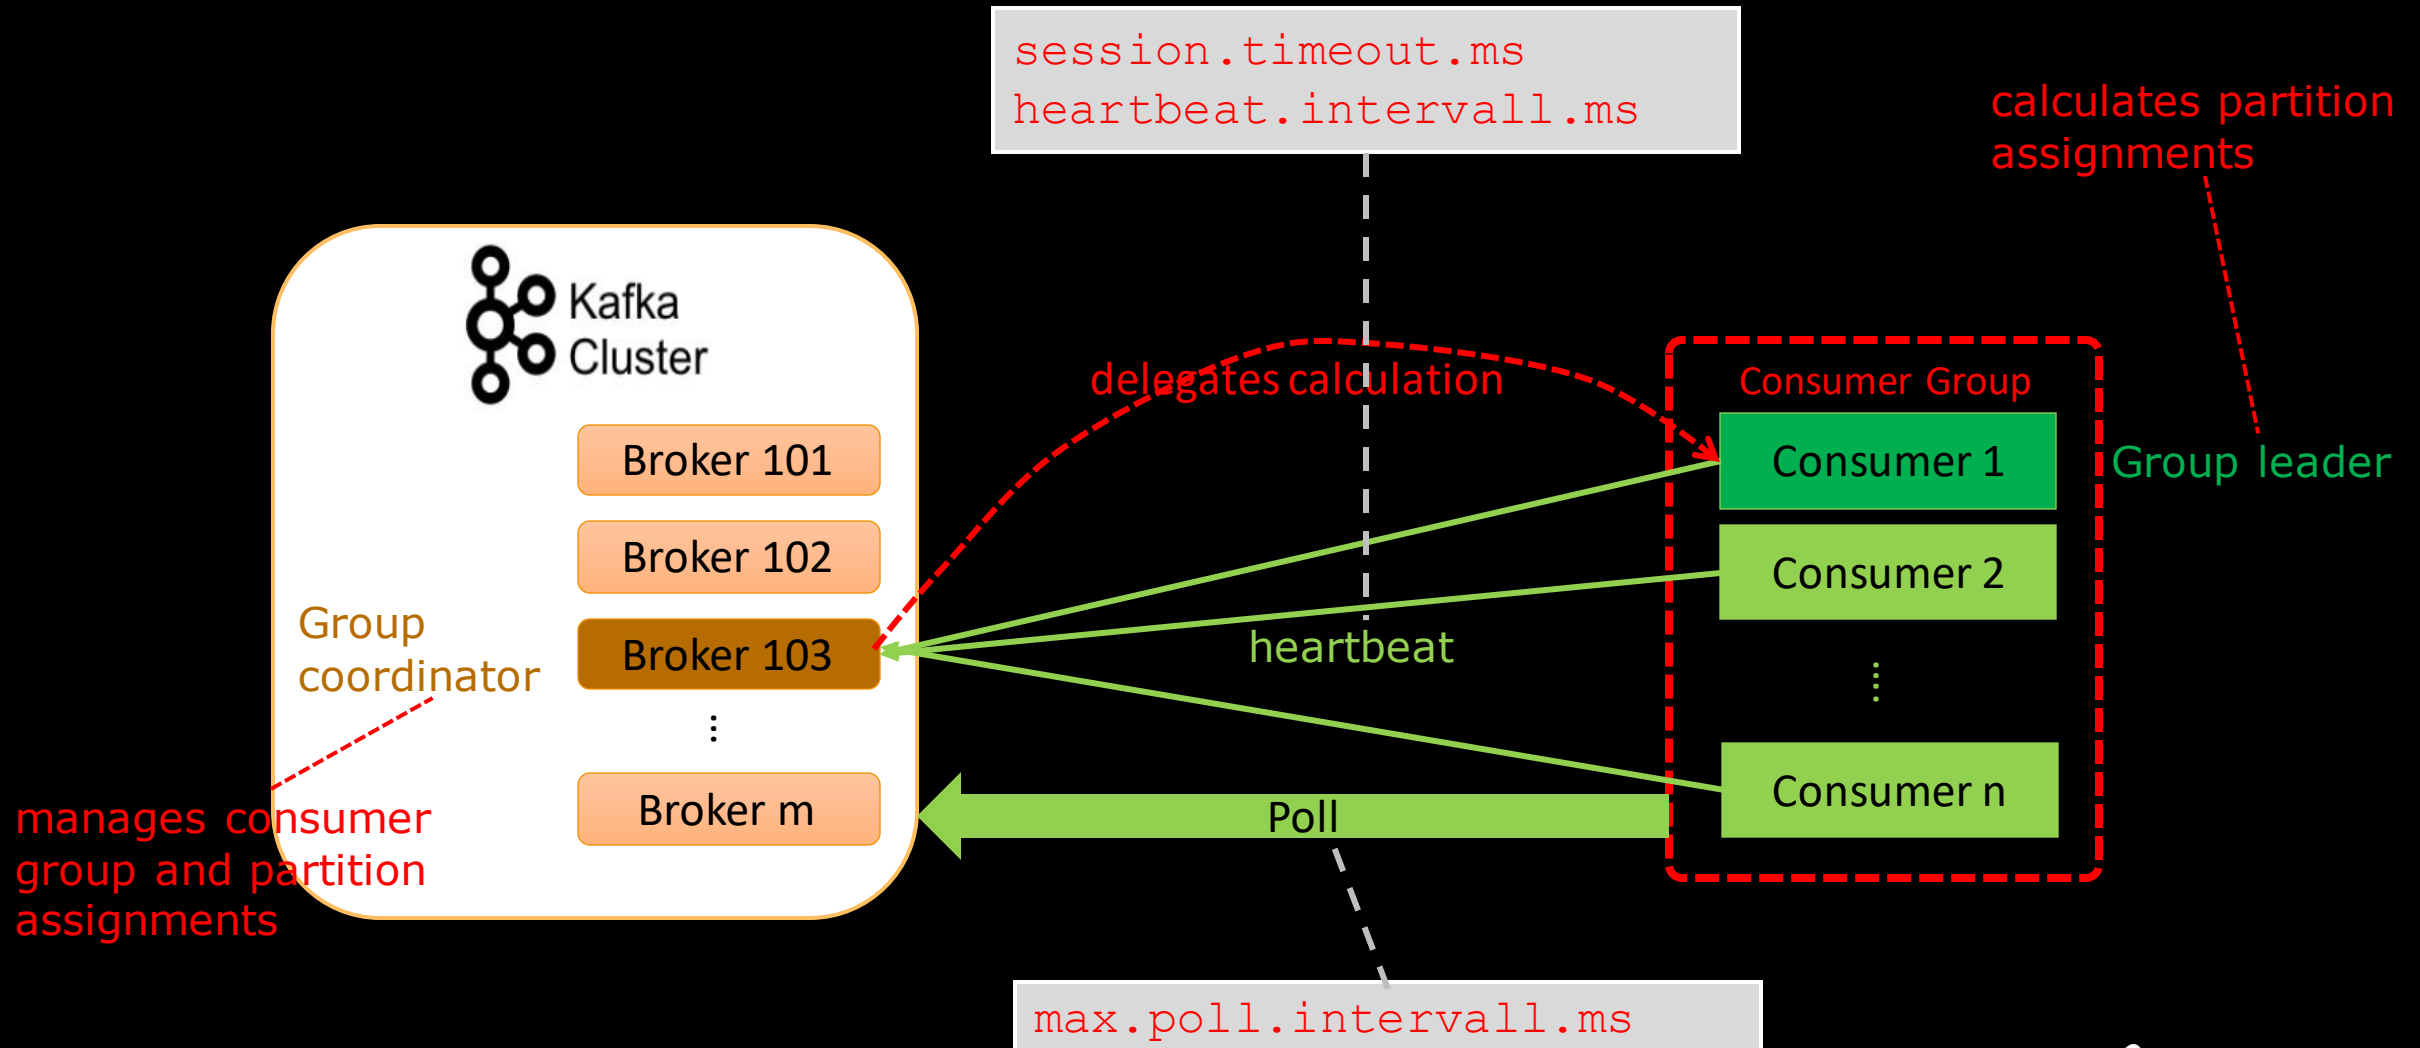


How are Partitions assigned to Consumers 3

- partition.assignment.strategy consumer property:
 - RangeAssignor (used in stream-processing for co-partitioned topics)
 - RoundRobinAssignor
 - StickyAssignor

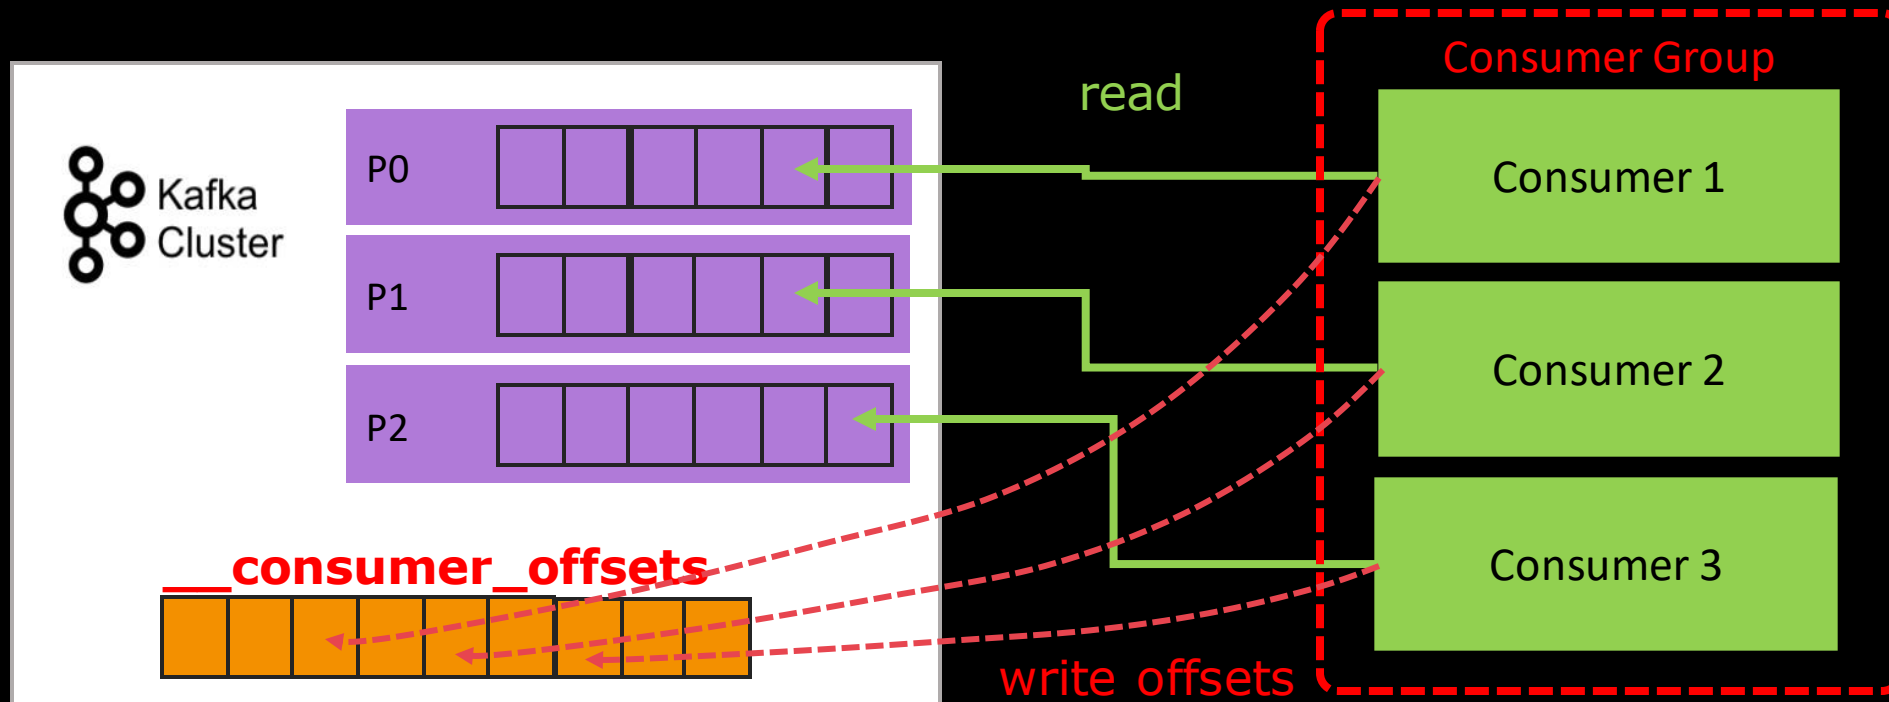


Consumer Liveliness 1-3



Consumers and Offsets

- Offset: Position of a record in the partition
- Group_id, topic, partition is tracked in topic: `__consumer_offsets`
- Consumer Offset Topic tracks which message should be read next



Development: Basic Consumer in .NET/C#

```
8 namespace consumer_net {
  0 references
9 class Program {
  0 references
10 static void Main (string[] args) {
11     Console.WriteLine ("Starting Consumer!");
12     var config = new Dictionary<string, object> {
13         { "group.id", "dotnet-consumer-group" },
14         { "bootstrap.servers", "kafka-1:9092" },
15         { "auto.commit.interval.ms", 5000 },
16         { "auto.offset.reset", "earliest" }
17     };
18
19     var deserializer = new StringDeserializer (Encoding.UTF8);
20     using (var consumer = new Consumer<string, string> (config, deserializer, deserializer)) {
21         consumer.OnMessage += (_, msg) =>
22             Console.WriteLine ($"Read ('{msg.Key}', '{msg.Value}') from: {msg.TopicPartitionOffset}");
23
24         consumer.OnError += (_, error) =>
25             Console.WriteLine ($"Error: {error}");
26
27         consumer.OnConsumeError += (_, msg) =>
28             Console.WriteLine ($"Consume error ({msg.TopicPartitionOffset}): {msg.Error}");
29
30         consumer.Subscribe ("hello_world_topic");
31
32         while (true) {
33             consumer.Poll (TimeSpan.FromMilliseconds (100));
34         }
35     }
36 }
37 }
38 }
```

configuration

message handling

error handling

polling data

subscribing to topic

Content

- Integration
- Event-Streaming Platform Kafka
- Asynchronous Communication
- Kafka Basics & Components
- Kafka APIs
- Kafka Examples
- Kafka Exercises



















Kafka Examples

Bosch Powertools & Deutsche Bahn Passenger Information

Today's Ecosystems are pretty big

- 2.2×10^{12} messages per day (6×10^{15} Byte Petabyte)
- up to 400 Microservices per cluster
- 20 – 200 Broker per cluster



Task Title	Location	Start	Due	Craft	Company	State	News
4.OG Technikzentrale, Stahlbau	-	07/07/20	07/14/20	Rohbauarbeiten	 Company B1 Anne Esch		
Abbau Kran klein	-	09/12/20	09/12/20	Rohbauarbeiten	 Company B1 Anne Esch		
Abbruch Kanal in Baufeld	-	09/11/19	11/09/19	Erdbauarbeiten	 Company B1 Anne Esch		
Abbruch Verbindungsgang	-	08/06/19	08/17/19	Erdbauarbeiten	 Company B1 Anne Esch		
Abdichtung	-	01/11/21	01/16/21	Fliesen	 Assign now		
Abdichtung	-	01/25/21	01/30/21	Fliesen	 Assign now		
Abdichtung	-	12/28/20	01/02/21	Fliesen	 Assign now		
Abdichtung + Dämmung Achse P-O...	-	05/13/20	05/14/20	Rohbauarbeiten	 Company B1 Anne Esch		

RefinemySite: Lean Construction SaaS

Kafka as Single Source of Truth

Confluent Cloud

Kafka as Event-Sourcing Backbone

Replication of Data between Microservices via Topics

Simple Bootstrapping of new Services by reprocessing the Event Stream

Real-Time & Near Realtime Notifications of Users

19 independently deployable Units

29 People across Europe, started in 2016

Topics with different Avro-Schemas

Passenger Information of Deutsche Bahn

- Talk at Confluent Streaming event 11.11.2019



DB Reisendeninformation – RI-Plattform als Single Point of Truth



Wir führen jeden Tag mehrere Changes unterbrechungsfrei an der Produktion durch. Die Anpassungen beinhalten im wesentlichen neue Features, Fehlerbehebungen und Konfigurationsänderungen.



Die RI-Plattform wird auf Basis von Echtzeiten entwickelt und kontinuierlich getestet. Dabei verarbeiten wir aktuell im Livestream ca. 120 Millionen Nachrichten pro Tag.



Die Produktionsumgebung besteht aktuell aus ca. 110 virtuellen Servern in der Cloud, um das jetzige Mengengerüst in Echtzeit zu verarbeiten.



Die RI-Plattform wird aktuell von 11 Scrumteams und dem Betriebsteam mit ca. 90 Mitarbeitern entwickelt.



Technologien der RI-Plattform

CLOUD COMPUTING

Hochverfügbar, skalierbar, virtualisiert, unterbrechungsfrei.

BIG DATA

Große Datenmengen, effiziente Verarbeitung und Analyse.

SPOT* IN ECHTZEIT

Kontinuierliche Datenverarbeitung, unterbrechungsfreier Datenstrom.

*Single Point of Truth

CD/CI

Continuous Deployment und Integration, vollautomatisierter Test, kurze Release-Zyklen, dynamische Umgebungsübersicht.

MICROSERVICE

Loose Kopplung, orchestrierte Instanzen, dynamisch skalierbar.

MESSAGING/STREAMING

Performant, skalierbar, Verarbeitung in Echtzeit.

EVENT SOURCING/MATCHING

Robust durch „full context“, hohe Qualität der Datenverarbeitung durch „geometrische“ Datenströme.

MODERNE OPEN SOURCE PRODUKTE

Apache Kafka + Streams, Kubernetes, Hbase, Hazelcast, Cassandra, Flink, Jupyter ...

Kafka-Project: DB Passenger Information

Facts and Figures

- ~100 persons in 12 Scrum Teams
- 24/7 running (DevOps)
- multiple daily deployments in production
- ~100 virtual servers
- ~100 Microservices

Kafka

Broker

**6 in 3
AZs**

Msg In

**300M
/day**

Volume In

**1TB
/day**

Topic/
Partition

**320
/3,3k**

Msg In (Avg)

**4k
/sec**

Volume In

**13MB
/sec**

Content

- Integration
- Event-Streaming Platform Kafka
- Asynchronous Communication
- Kafka Basics & Components
- Kafka APIs
- Kafka Examples
- Kafka Exercises



Kafka Exercises

Introduction and preparation of the next unit

Sources

Sources

1. <https://www.enterpriseintegrationpatterns.com/patterns/messaging/>
2. Enterprise Integration Patterns, Gregor Hohpe and Bobby Woolf: ISBN 0321200683
3. <https://kafka.apache.org/>
4. <https://www.confluent.io/what-is-apache-kafka/>
5. <https://www.confluent.io/resources/>
6. <https://www.informatik-aktuell.de/betrieb/verfuegbarkeit/apache-kafka-eine-schluesselplattform-fuer-hochskalierbare-systeme.html>
7. <https://www.slideshare.net/KaiWaehner/apache-kafka-vs-integration-middleware-mq-etl-esb?ref=https://www.kai-waehner.de/blog/2019/03/07/apache-kafka-middleware-mq-etl-esb-comparison/>
8. <https://www.confluent.io/blog/apache-kafka-vs-enterprise-service-bus-esb-friends-enemies-or-frenemies/>
9. <https://microservices.io>



Christopher Uldack

Senior Consultant

Christopher.Uldack@novatec-gmbh.de

Oliver Berger

Senior Managing Consultant

Oliver.Berger@novatec-gmbh.de

Novatec Consulting GmbH

Dieselstraße 18/1

D-70771 Leinfelden-Echterdingen

T. +49 711 22040-700

info@novatec-gmbh.de

www.novatec-gmbh.de