


Using Kafka to Optimize Data Movement and System Integration

Alex Holmes @  JavaOne™



https://www.flickr.com/photos/tom_bennett/7095600611

THIS SUCKS



E

T

L

(CIRCA 2560 B.C.E.)

A FEW YEARS LATER ...

2,014 C.E.

I NEED YOU TO COPY
DATA FROM ORACLE
TO HADOOP ...



3 WEEKS LATER ...


THIS SUCKS!



Agenda

- Look at challenges with data and system integration
- Provide an overview of Kafka, including its performance
- Examine Kafka data and system integration patterns

whoami



- Alex Holmes
- Software engineer
- Working on distributed systems for many years
- @grep_alex
- grepalex.com

Challenges with data integration

I NEED YOU TO COPY DATA
FROM ORACLE INTO HADOOP
FOR OUR DATA SCIENTISTS







Problems you'll need to solve:

- Reliability/fault tolerance
- Scalability/throughput
- Handle large tables
- Throttle network IO
- Idempotent writes
- Scheduling

OH, AND WE NEED TO USE THE
SAME DATA TO CALCULATE
AGGREGATES IN REAL-TIME





Sqoop

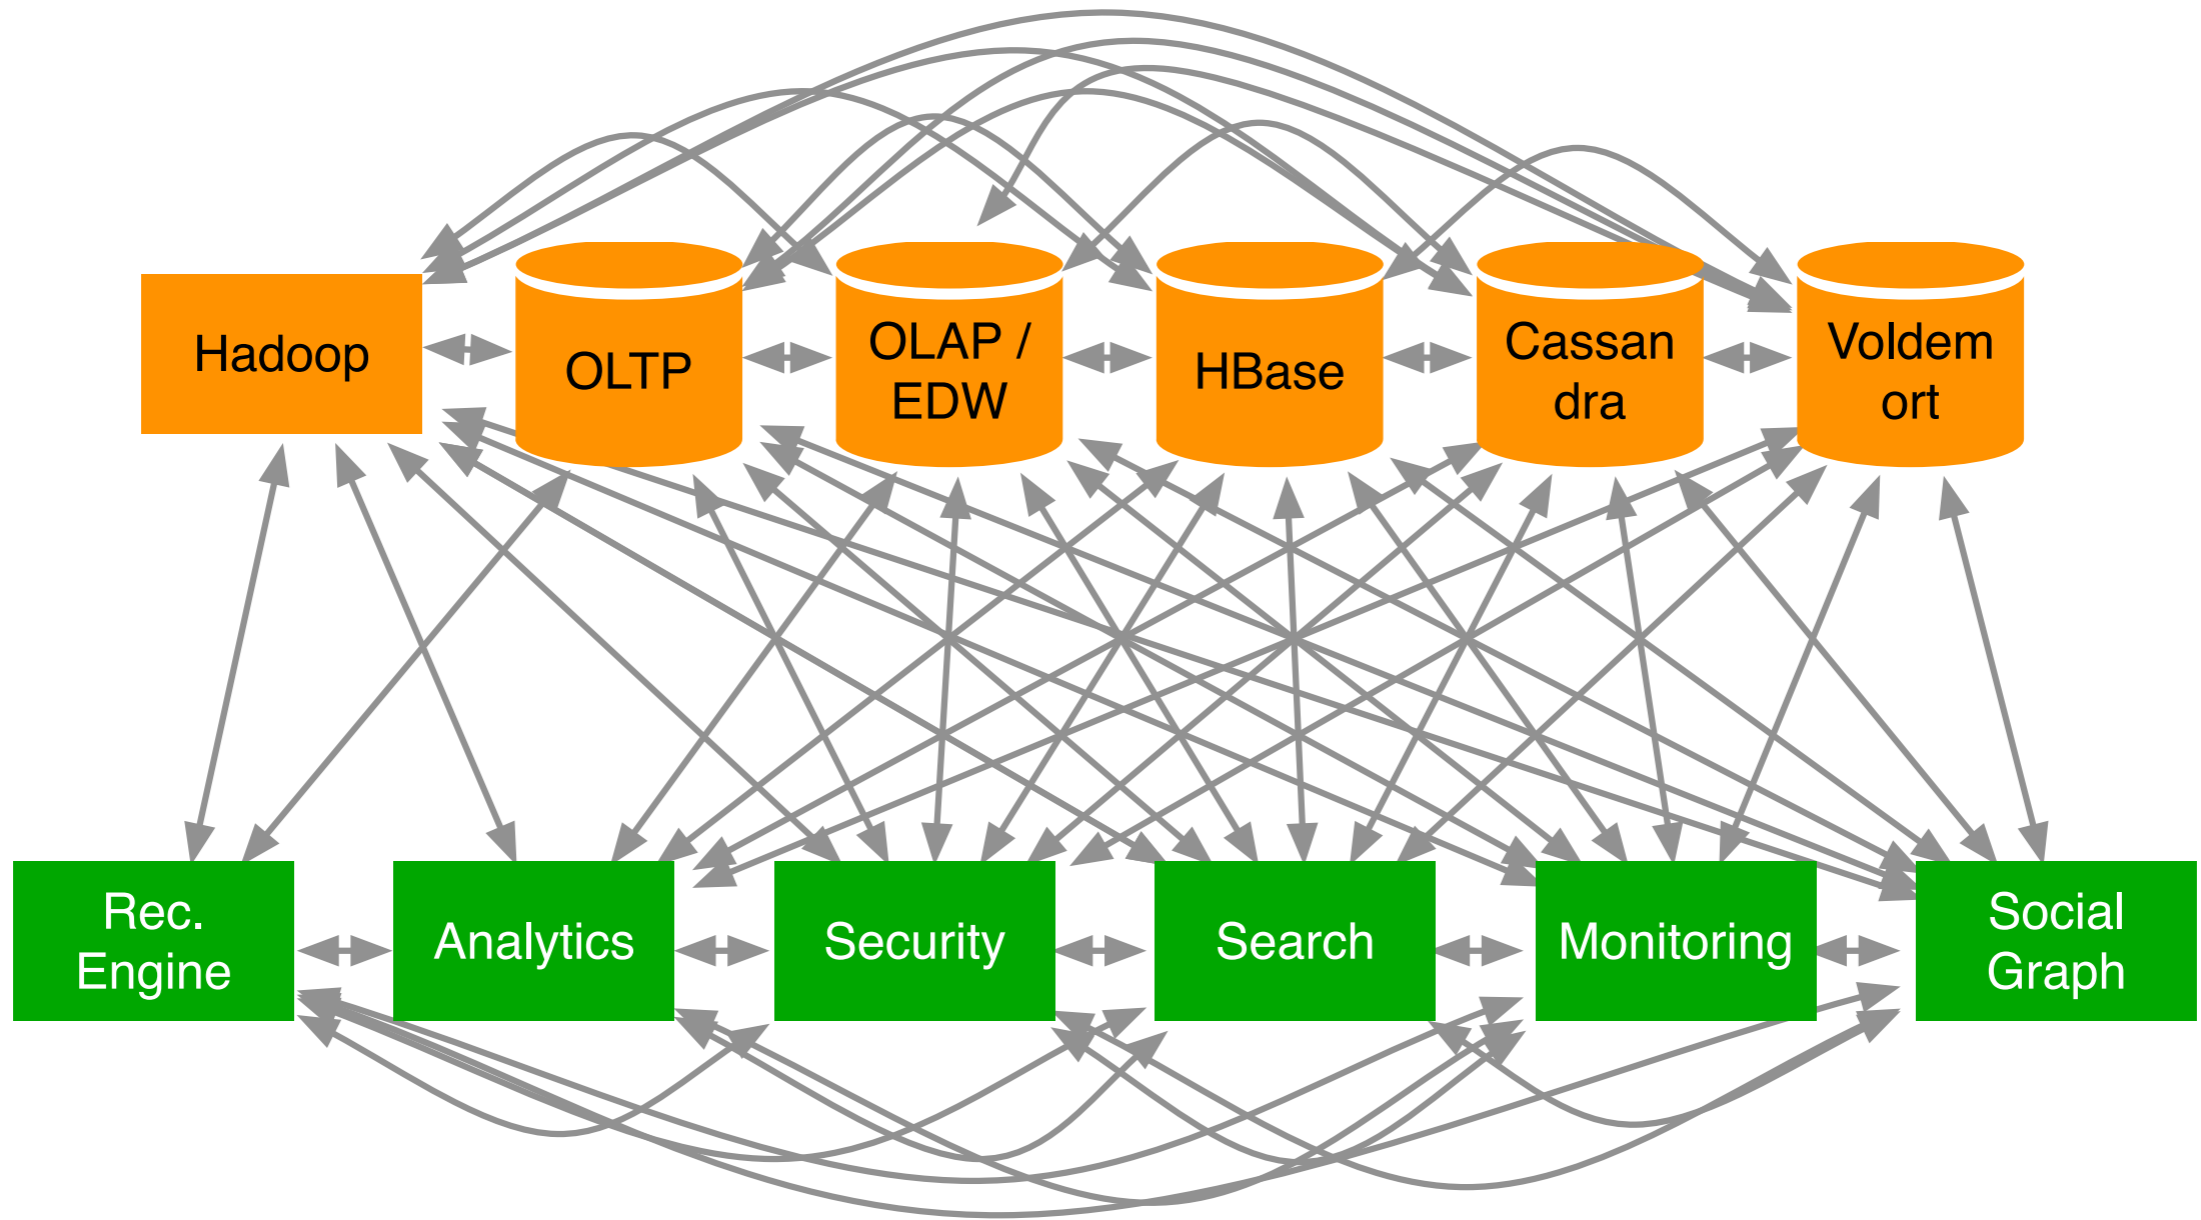
?

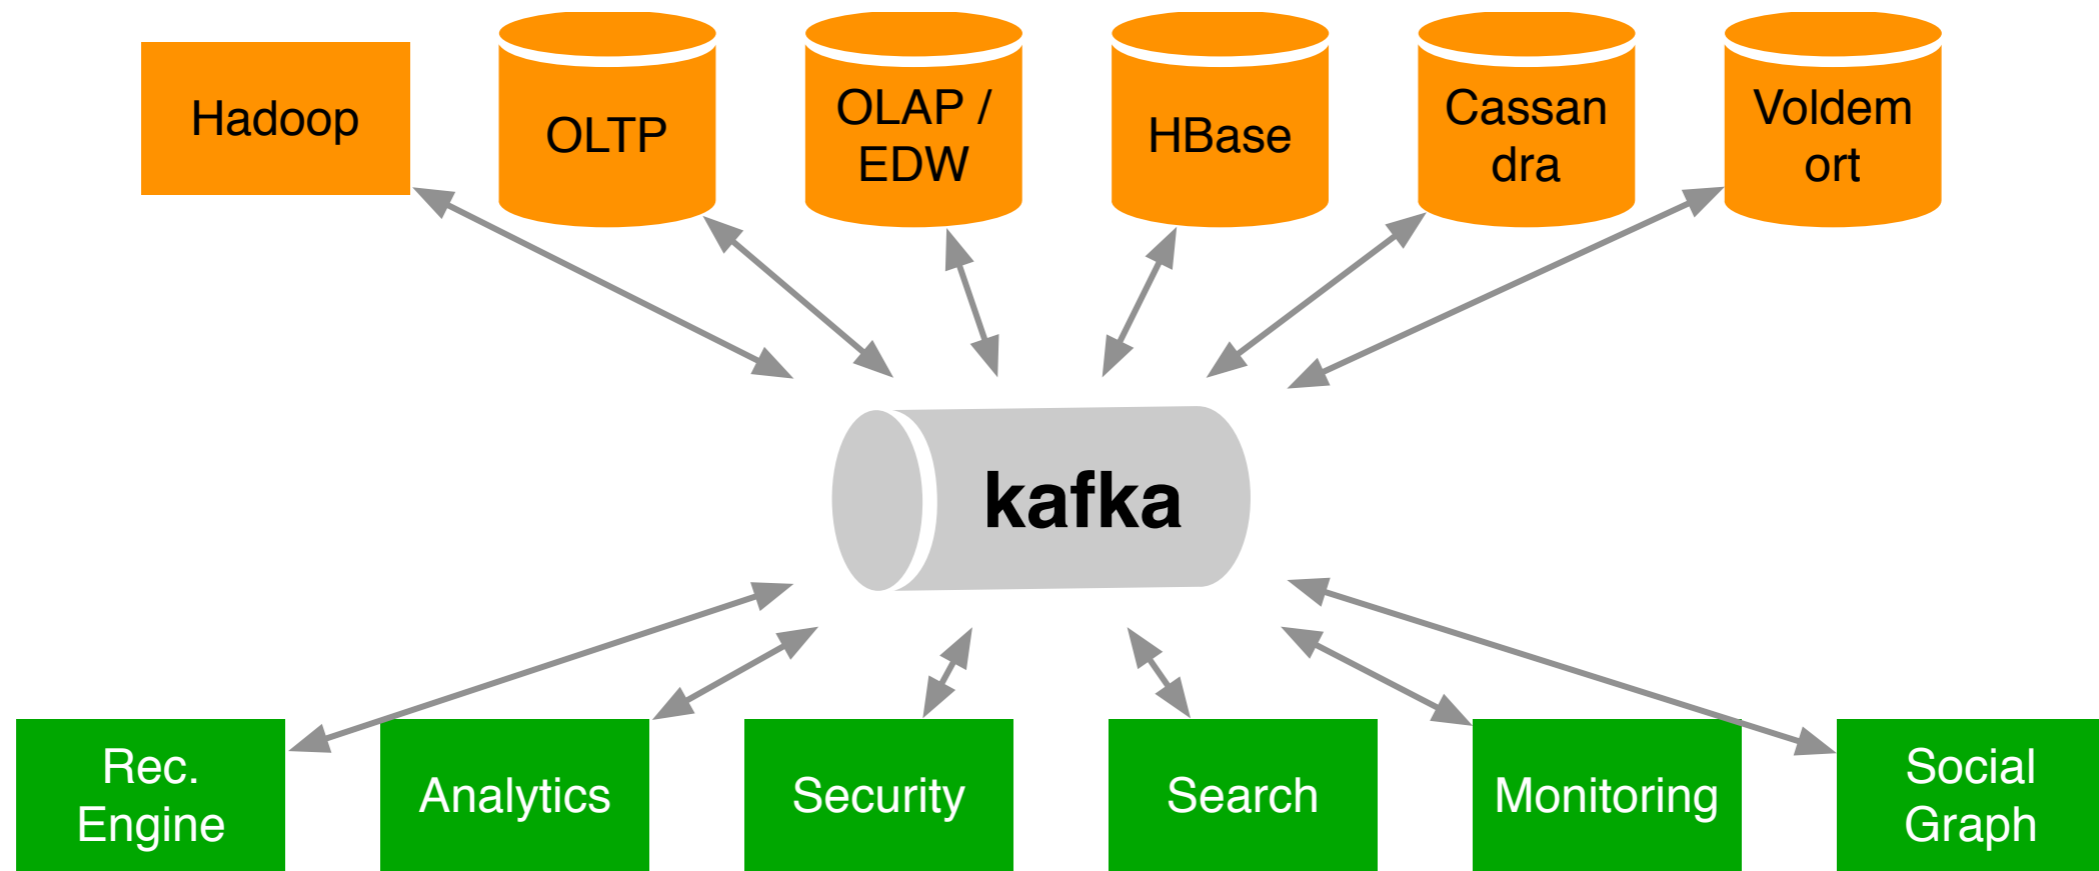


STORM



NoSQL





Kafka

Background

- Apache project
- Originated from LinkedIn
- Open-sourced in 2011
- Written in Scala and Java
- Borrows concepts in messaging systems and logs

Powered by

- [LinkedIn](#) - Apache Kafka is used at LinkedIn for activity stream data and operational metrics. This powers various products like LinkedIn Newsfeed, LinkedIn Today in ad
- [Twitter](#) - As part of their Storm stream processing infrastructure, e.g. [this](#).
- [Netflix](#) - Real-time monitoring and event-processing [pipeline](#).
- [Square](#) - We use Kafka as a bus to move all systems events through our various datacenters. This includes metrics, logs, custom events etc. On the consumer side, we
- [Spotify](#) - Kafka is used at Spotify as part of their [log delivery system](#).
- [Pinterest](#) - Kafka is used with [Secor](#) as part of their [log collection pipeline](#).
- [Uber](#)
- [Tumblr](#) - [See this](#)
- [Box](#) - At Box, Kafka is used for the production analytics pipeline & real time monitoring infrastructure. We are planning to use Kafka for some of the new products & featur
- [Airbnb](#) - Used in our event pipeline, exception tracking & more to come.
- [Mozilla](#) - Kafka will soon be replacing part of our current production system to collect performance and usage data from the end-users browser for projects like Telemetry
- [Tagged](#) - Apache Kafka drives our new pub sub system which delivers real-time events for users in our latest game - Deckadence. It will soon be used in a host of new u
- [Foursquare](#) - Kafka powers online to online messaging, and online to offline messaging at Foursquare. We integrate with monitoring, production systems, and our offline
- [StumbleUpon](#) - Data collection platform for analytics.
- [Coursera](#) - At Coursera, Kafka powers education at scale, serving as the data pipeline for realtime learning analytics/dashboards.
- [Shopify](#) - Access logs, A/B testing events, domain events ("a checkout happened", etc.), metrics, delivery to HDFS, and customer reporting. We are now focusing on con
- [Mate1.com Inc.](#) - Apache kafka is used at Mate1 as our main event bus that powers our news and activity feeds, automated review systems, and will soon power real tim
- [Boundary](#) - Apache Kafka aggregates high-flow message streams into a unified distributed pubsub service, brokering the data for other internal systems as part of Bounc
- [Ancestry.com](#) - Kafka is used as the [event log processing pipeline](#) for delivering better personalized product and service to our customers.
- [DataSift](#) - Apache Kafka is used at DataSift as a collector of monitoring events and to track user's consumption of data streams in real time. <http://highscalability.com/blog>
- [Spongecell](#) - We use Kafka to run our entire analytics and monitoring pipeline driving both real-time and ETL applications for our customers.
- [Wooga](#) - We use Kafka to aggregate and process tracking data from all our facebook games (which are hosted at various providers) in a central location.
- [AddThis](#) - Apache Kafka is used at AddThis to collect events generated by our data network and broker that data to our analytics clusters and real-time web analytics pla
- [Urban Airship](#) - At Urban Airship we use Kafka to buffer incoming data points from mobile devices for processing by our analytics infrastructure.
- [Metamarkets](#) - We use Kafka to ingest real-time event data, stream it to Storm and Hadoop, and then serve it from our Druid cluster to feed our interactive analytics dast Druid.
- [Simple](#) - We use Kafka at Simple for log aggregation and to power our analytics infrastructure.
- [Gnip](#) - Kafka is used in their [twitter ingestion and processing pipeline](#).
- [Loggly](#) - Loggly is the world's most popular cloud-based log management. Our cloud-based log management service helps DevOps and technical teams make sense of t [processing infrastructure](#).
- [RichRelevance](#) - Real-time tracking event pipeline.
- [SocialTwist](#) - We use Kafka internally as part of our reliable email queueing system.
- [Countandra](#) - We use a hierarchical distributed counting engine, uses Kafka as a primary speedy interface as well as routing events for cascading counting
- [FlyHaji.com](#) - We use Kafka to collect all metrics and events generated by the users of the website.
- [uSwitch](#) - [See this blog](#).
- [InfoChimps](#) - Kafka is part of the [InfoChimps real-time data platform](#).
- [Visual Revenue](#) - We use Kafka as a distributed queue in front of our web traffic stream processing infrastructure (Storm).
- [Oolya](#) - Kafka is used as the primary high speed message queue to power Storm and our real-time analytics/event ingestion pipelines.
- [Datadog](#) - Kafka brokers data to most systems in our metrics and events ingestion pipeline. Different modules contribute and consume data from it, for streaming CEP (h Cassandra, S3), or batch analysis (Hadoop).
- [VisualDNA](#) We use Kafka 1. as an infrastructure that helps us bring continuously the tracking events from various datacenters into our central hadoop cluster for offline p for future inference and recommendation engines
- [Sematext](#) - in SPM (performance monitoring + alerting), Kafka is used for metrics collection and feeds SPM's in-memory data aggregation (OLAP cube creation) as well e (search analytics) Kafka is used in search and click stream collection before being aggregated and persisted. In [Logsene \(log analytics\)](#) Kafka is used to pass logs and o
- [Wize Commerce](#) - At Wize Commerce (previously, NexTag), Kafka is used as a distributed queue in front of Storm based processing for search index generation. We pla into various data sinks like Hadoop, HBase, etc.
- [Quixey](#) - At Quixey, The Search Engine for Apps, Kafka is an integral part of our eventing, logging and messaging infrastructure.
- [LinkSmart](#) - Kafka is used at LinkSmart as an event stream feeding Hadoop and custom real time systems.
- [LucidWorks Big Data](#) - We use Kafka for syncing LucidWorks Search (Solr) with incoming data from Hadoop and also for sending LucidWorks Search logs back to Hadoc
- [Cloud Physics](#) - Kafka is powering our high-flow event pipeline that aggregates over 1.2 billion metric series from 1000+ data centers for near-to-real time data center op
- [Graylog2](#) - Graylog2 is a free and open source log management and data analysis system. It's using Kafka as default transport for Graylog2 Radio. The use case is desc
- [Yieldbot](#) - Yieldbot uses kafka for real-time events, camus for batch loading, and mirrormakers for x-region replication.
- [LivePerson](#) - Using Kafka as the main data bus for all real time events.
- [Retention Science](#) - Click stream ingestion and processing.
- [Strava](#) - Powers our analytics pipeline, activity feeds denorm and several other production services.
- [Outbrain](#) - We use Kafka in production for real time log collection and processing, and for cross-DC cache propagation.
- [SwiftKey](#) - We use Apache Kafka for analytics event processing.
- [Yeller](#) - Yeller uses Kafka to process large streams of incoming exception data for it's customers. Rate limiting, throttling and batching are all built on top of Kafka.
- [Emerging Threats](#) - Emerging threats uses Kafka in our event pipeline to process billions of malware events for search indices, alerting systems, etc.
- [Hotels.com](#) - Hotels.com uses Kafka as pipeline to collect real time events from multiple sources and for sending data to HDFS.
- [Helpprace](#) - Kafka is used as a distributed high speed message queue in our help desk software as well as our real-time event data aggregation and analytics.

Twitter

LinkedIn

Netflix

Spotify


Mozilla

A log

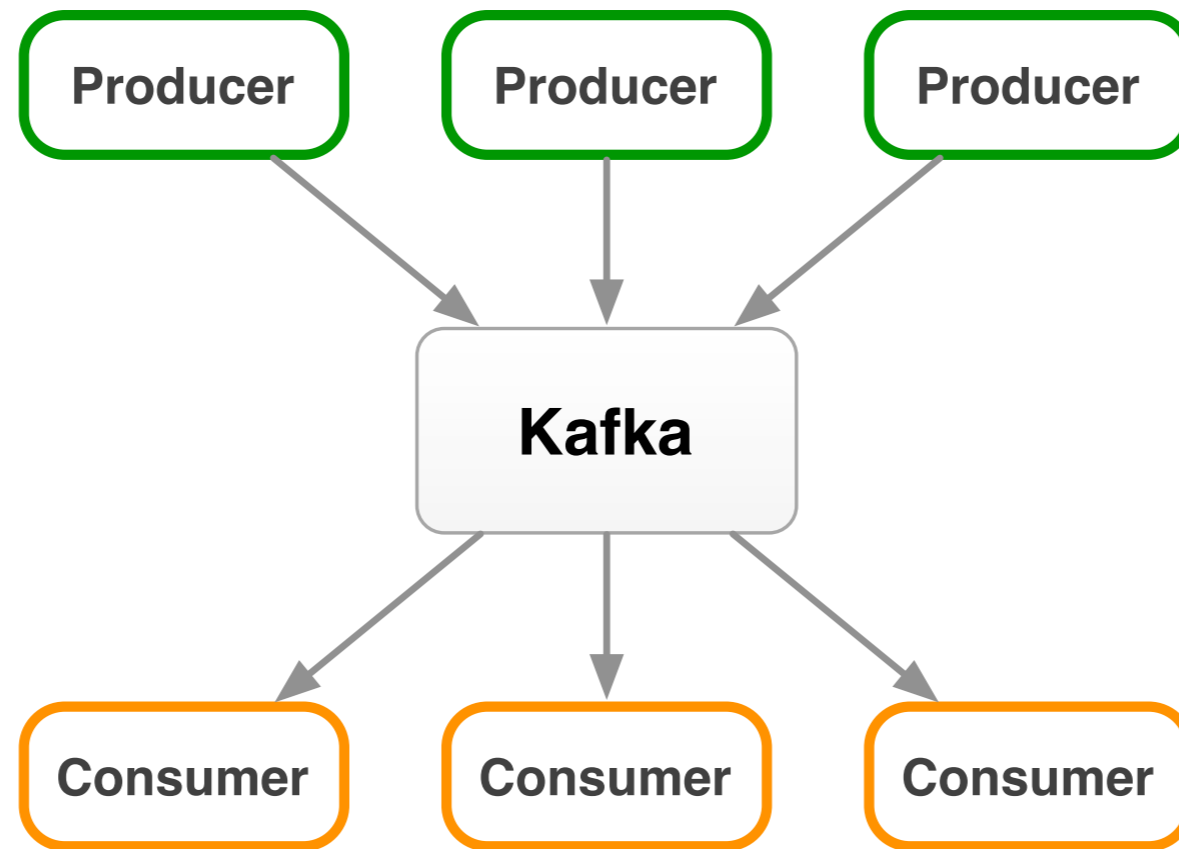
(partition in Kafka parlance)

First record

Most recent record



Producer-consumer system



Kafka != JMS

- Doesn't use the JMS API
- Ordering guarantees
- Mid-to-long term message storage
- Rewind capabilities
- Many-to-many messaging
- Broker doesn't maintain delivery state

Concepts

- A Kafka **cluster** consists of 1..N brokers and manages 1..N topics
- **Topics** are a category or feed name where messages are published
- Topics have one or more **partitions**, which are ordered, immutable sequence of messages that are continually appended to, much like a commit log

Data scaling and parallelism

Partition 0

| | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|----|----|-----|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | ... |
|---|---|---|---|---|---|---|---|---|---|----|----|-----|

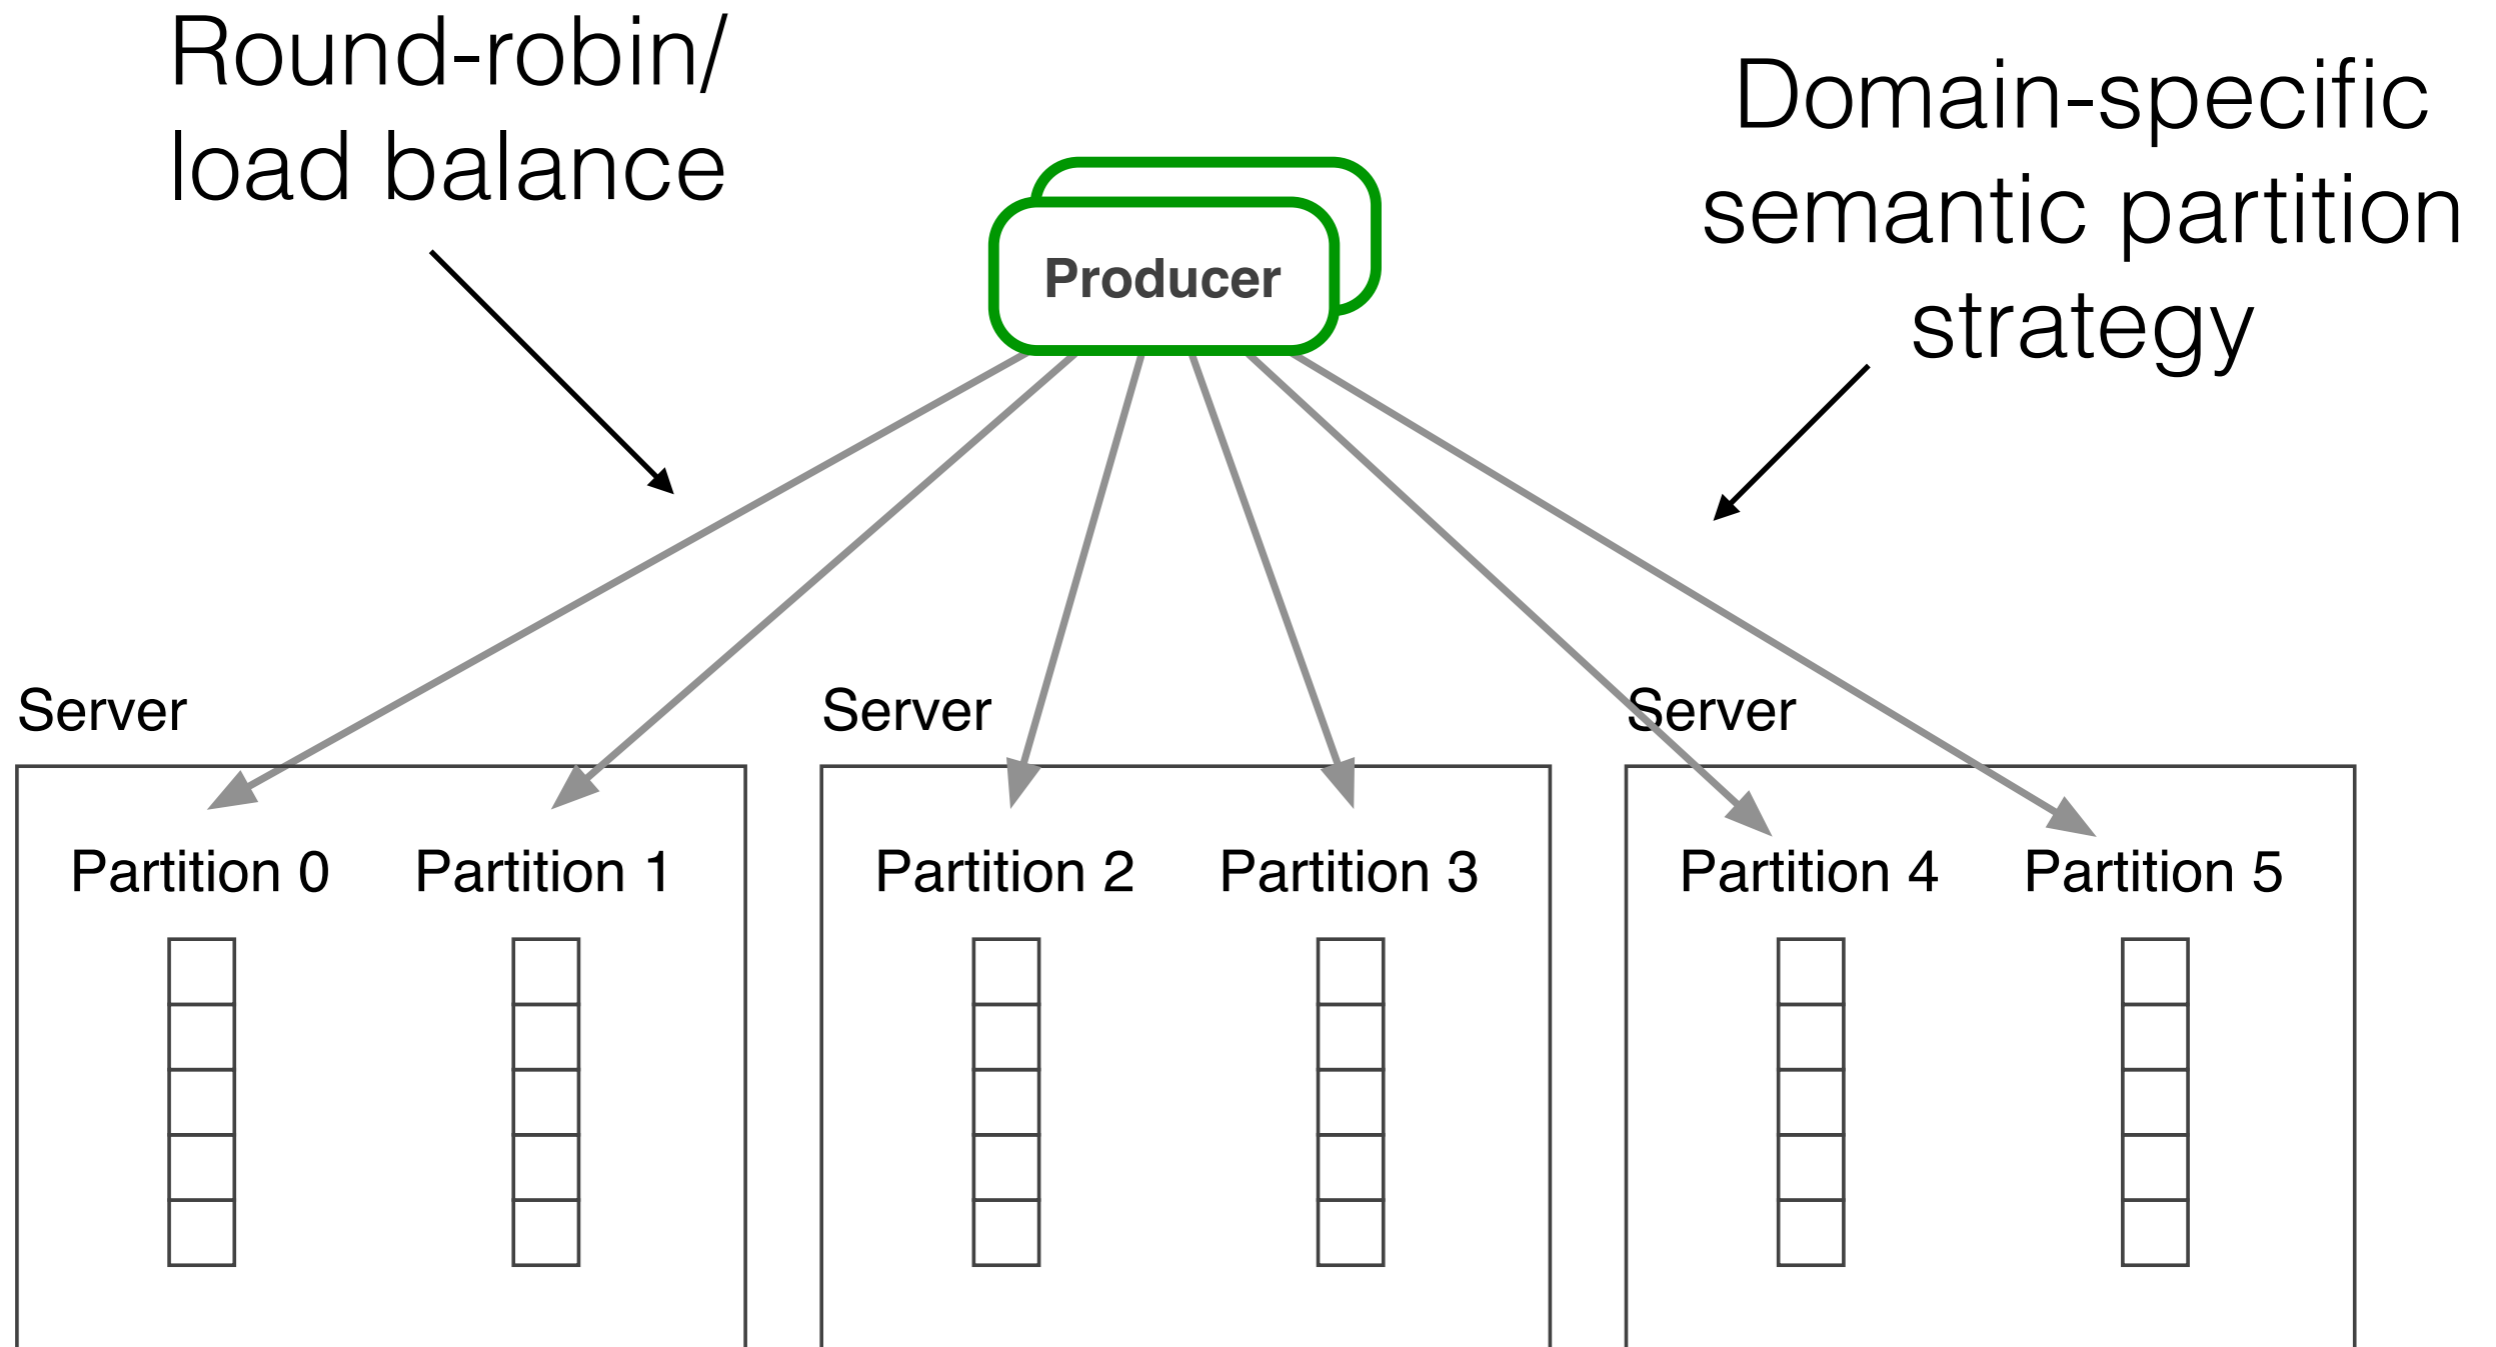
Partition 1

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|-----|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | ... |
|---|---|---|---|---|---|---|---|---|-----|

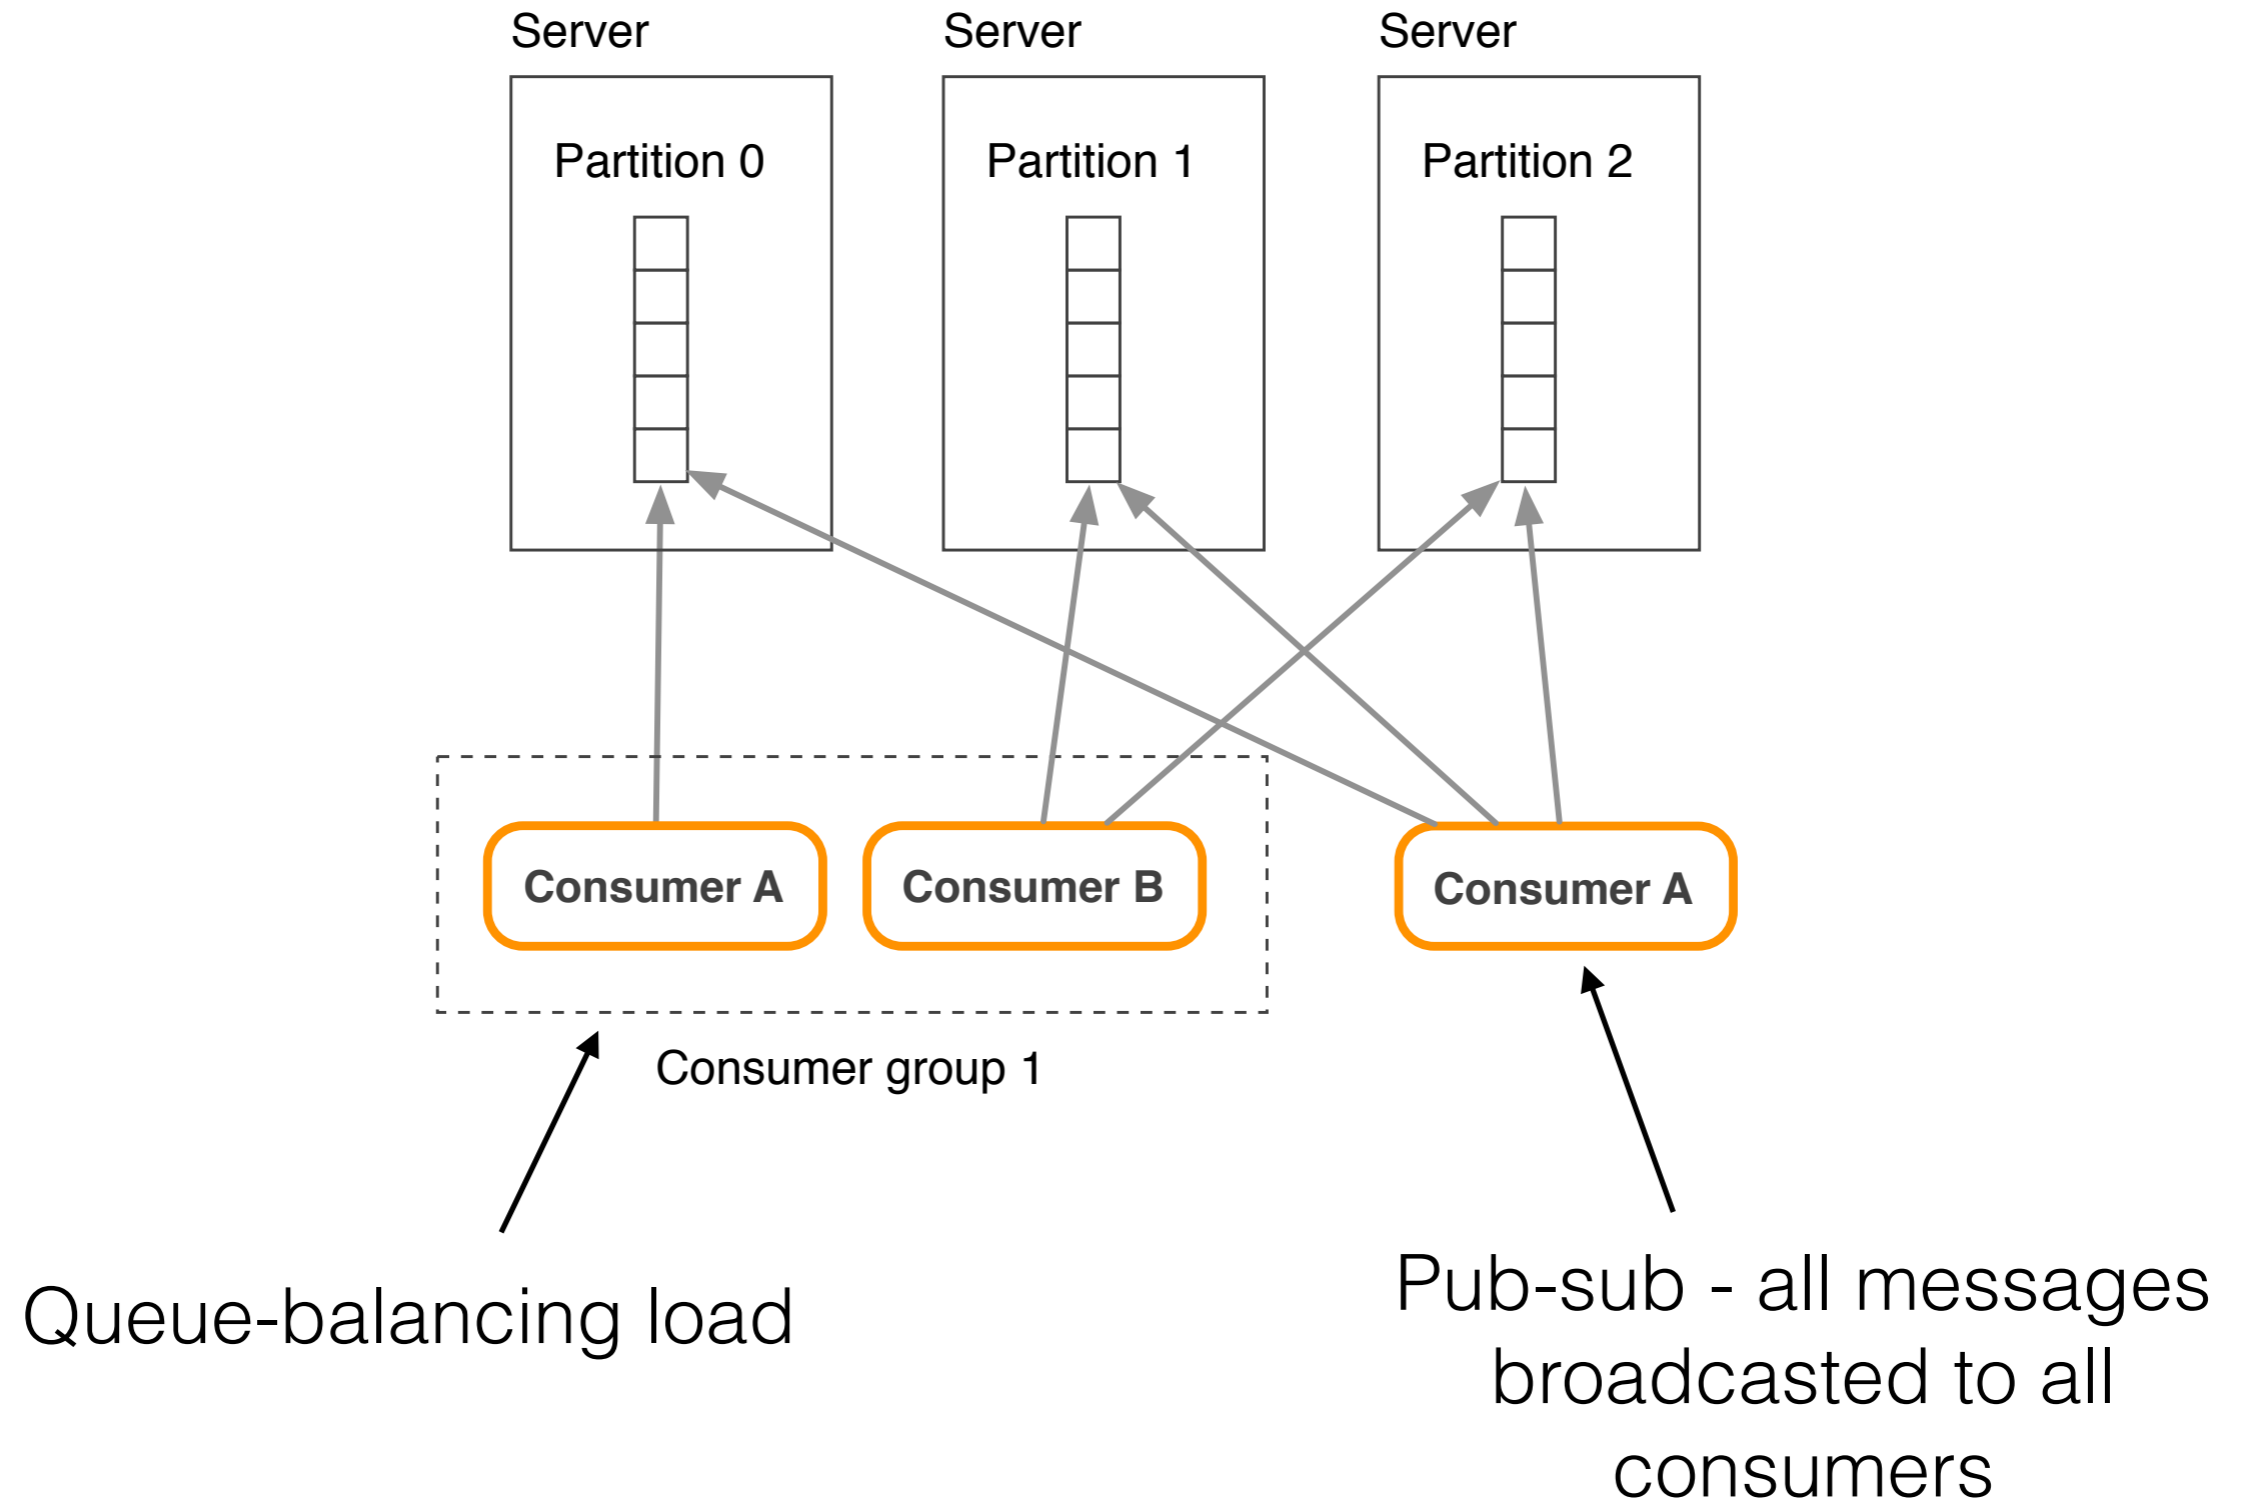
Partition 2

| | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|----|----|----|----|-----|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | ... |
|---|---|---|---|---|---|---|---|---|---|----|----|----|----|-----|

Partition distribution

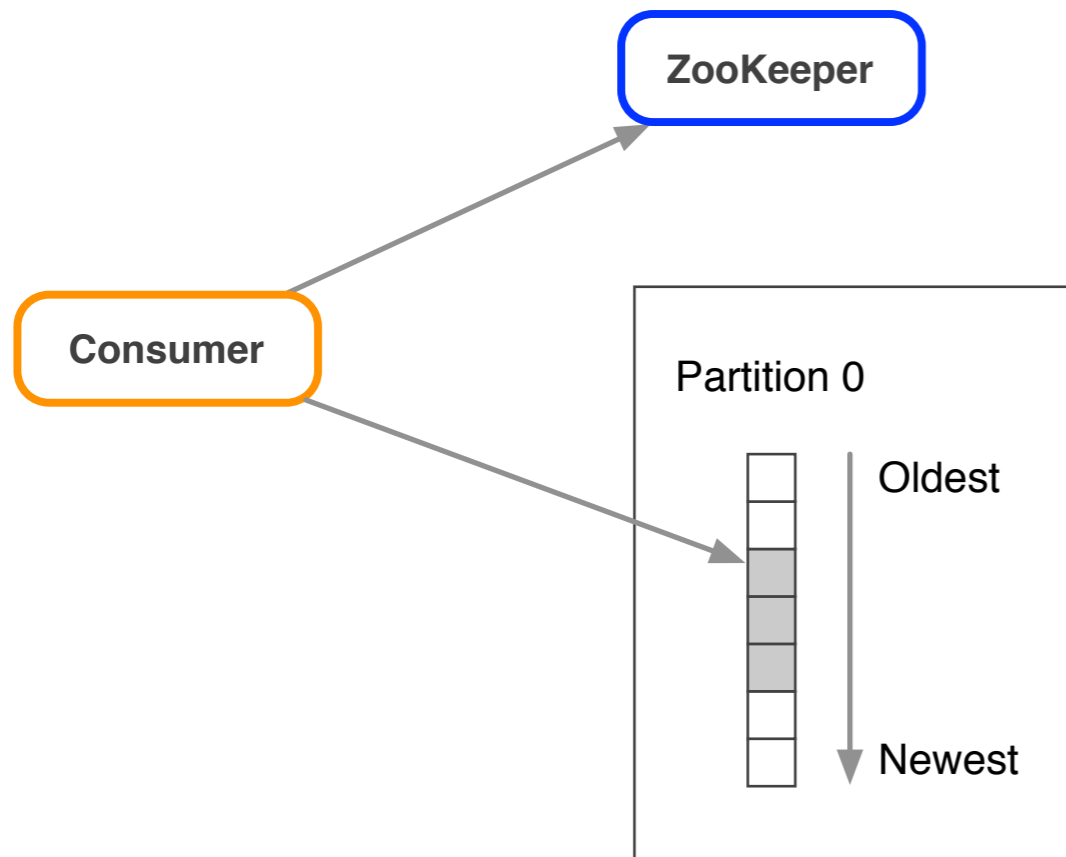


Consumers

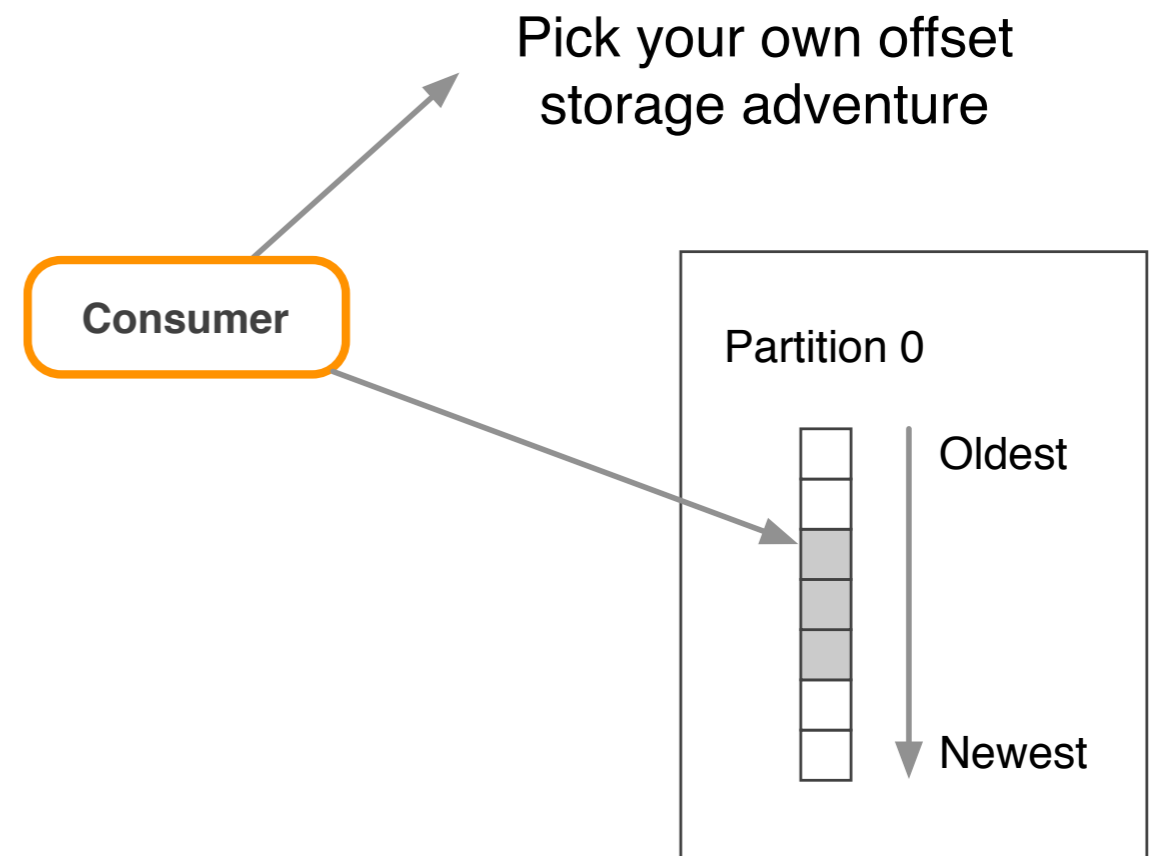


Consumer state

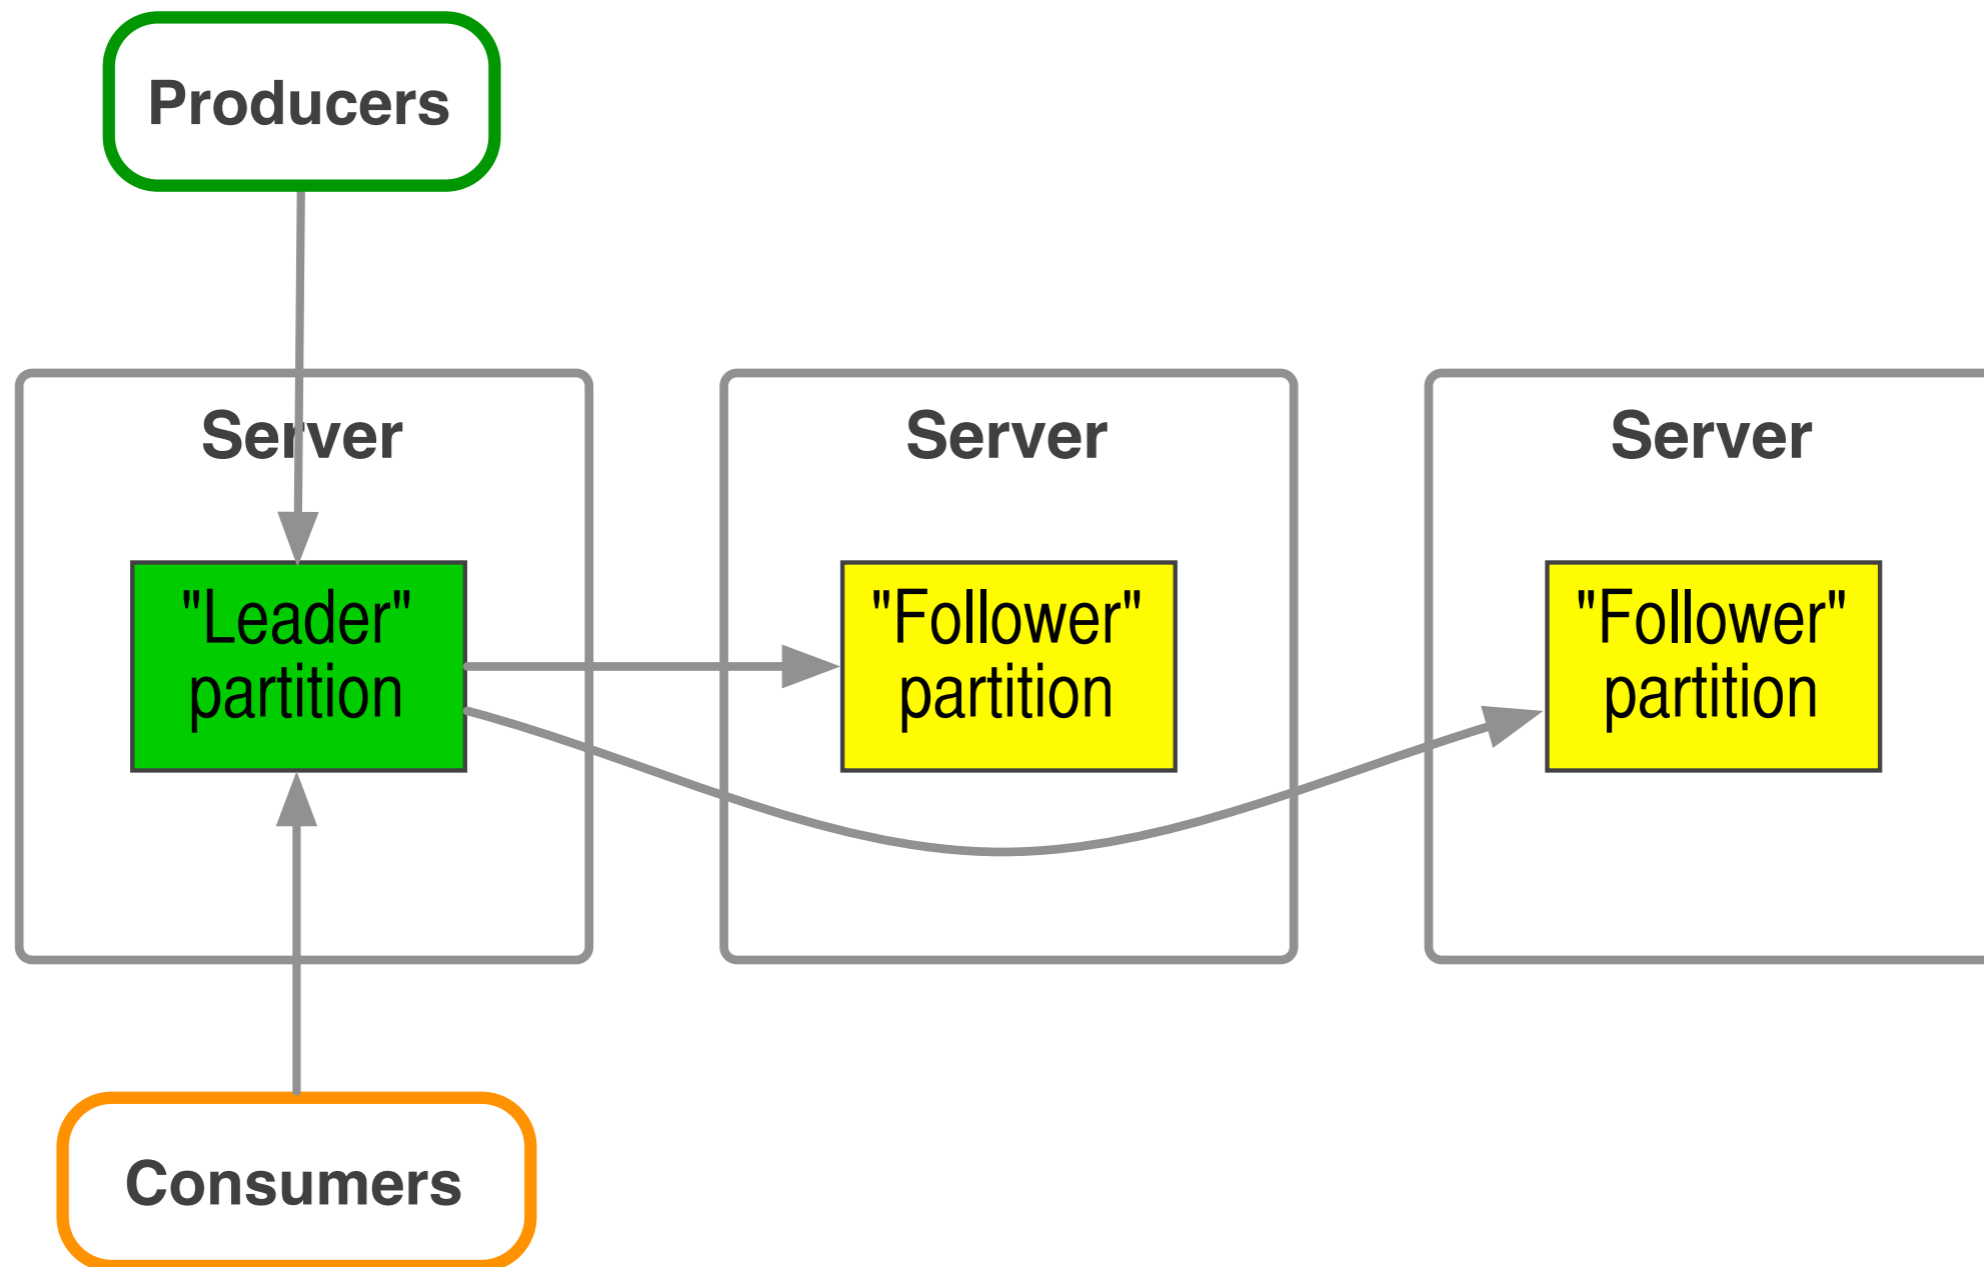
High-level consumer



Low-level consumer
(manage your own state)

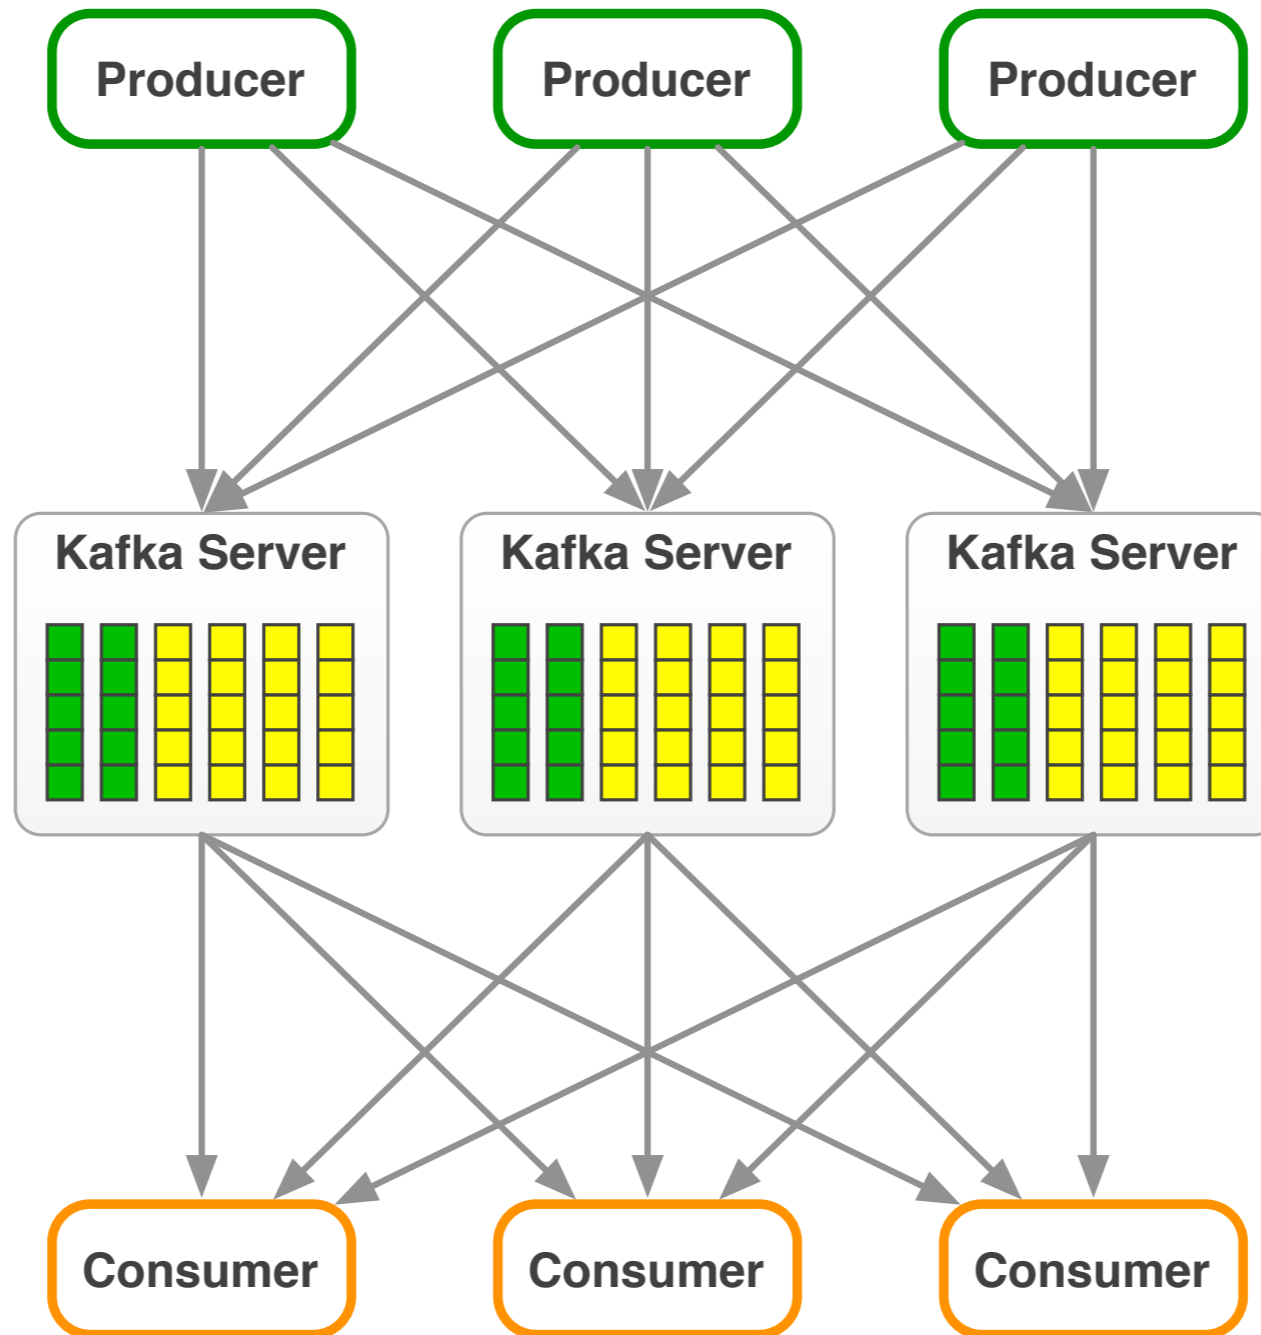


Partition replication

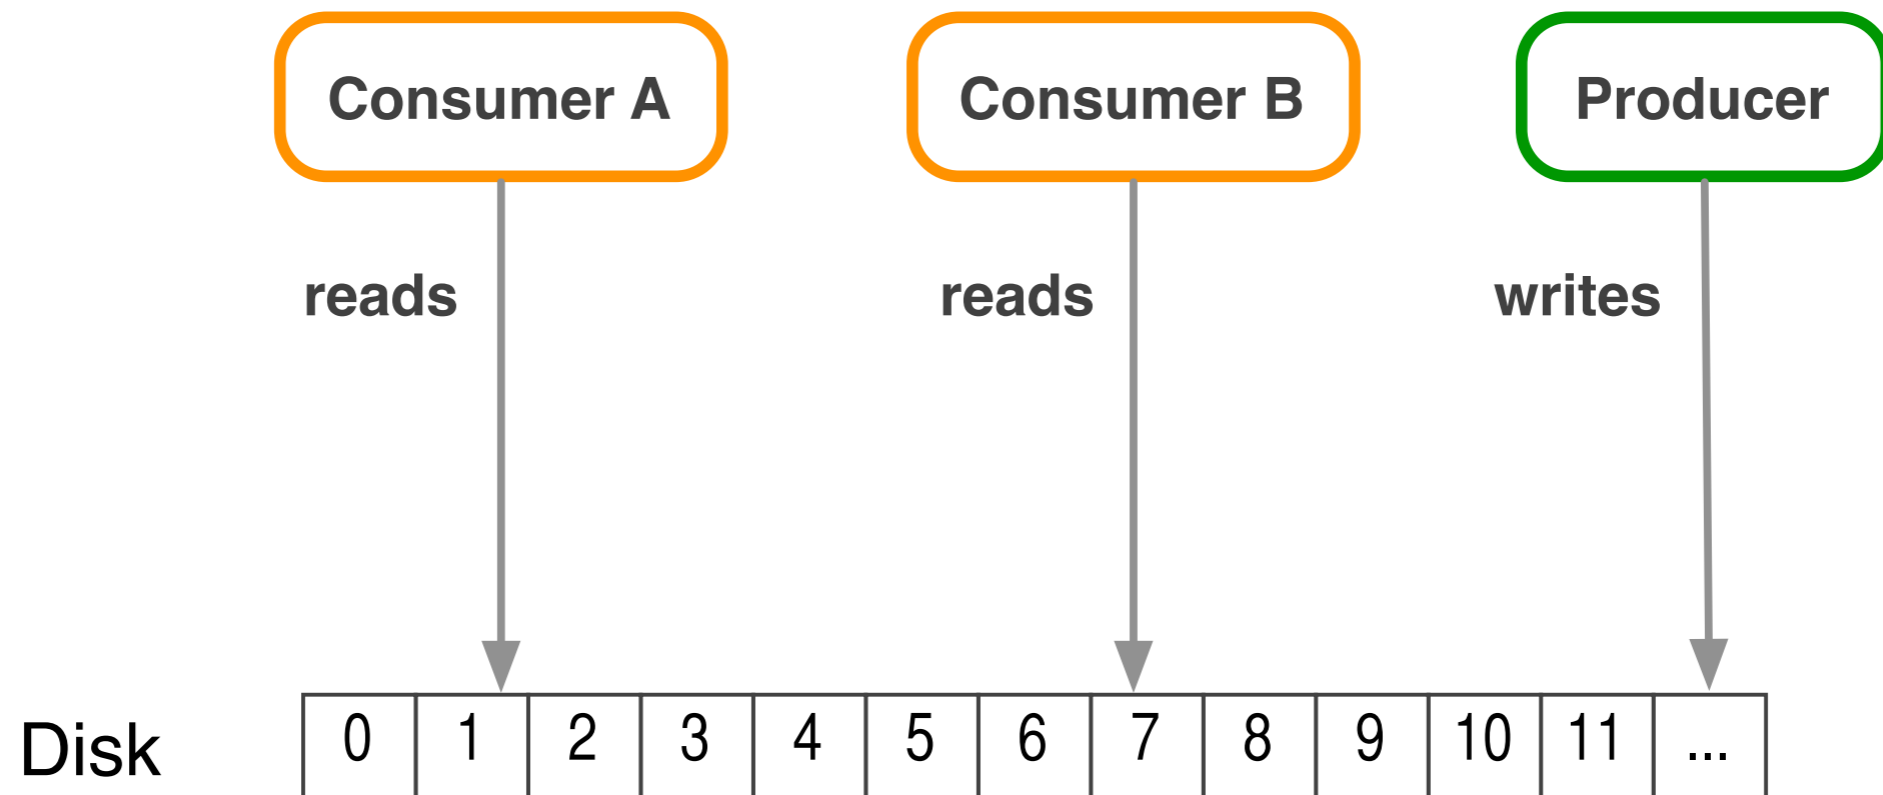


Performance + Scalability

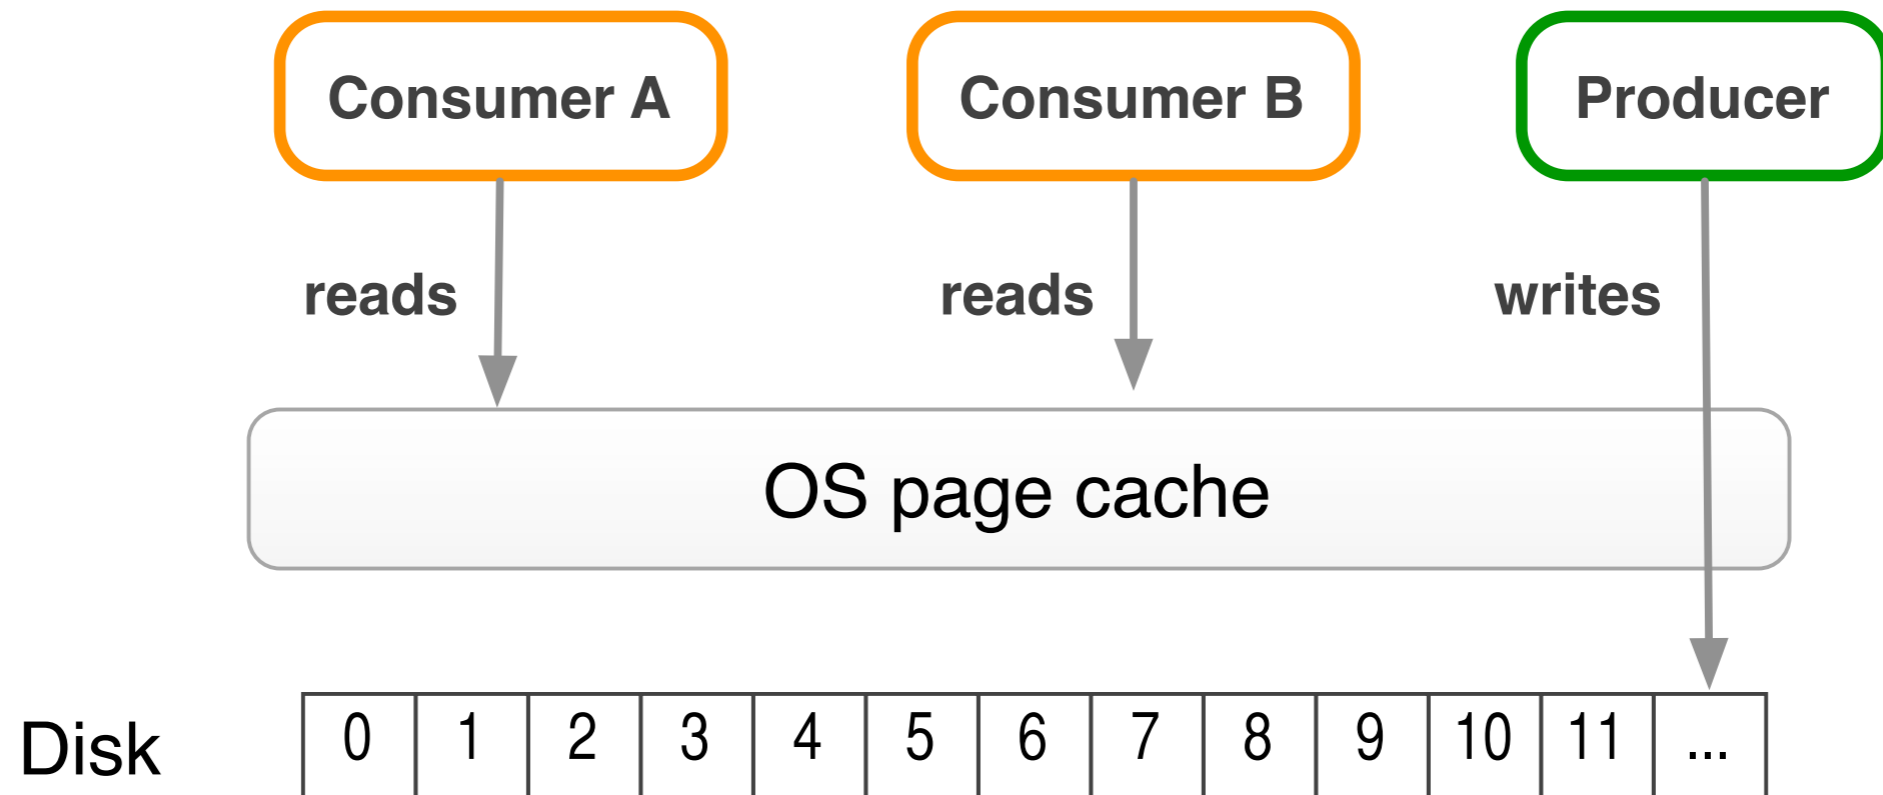
Scalability



Sequential disk IO

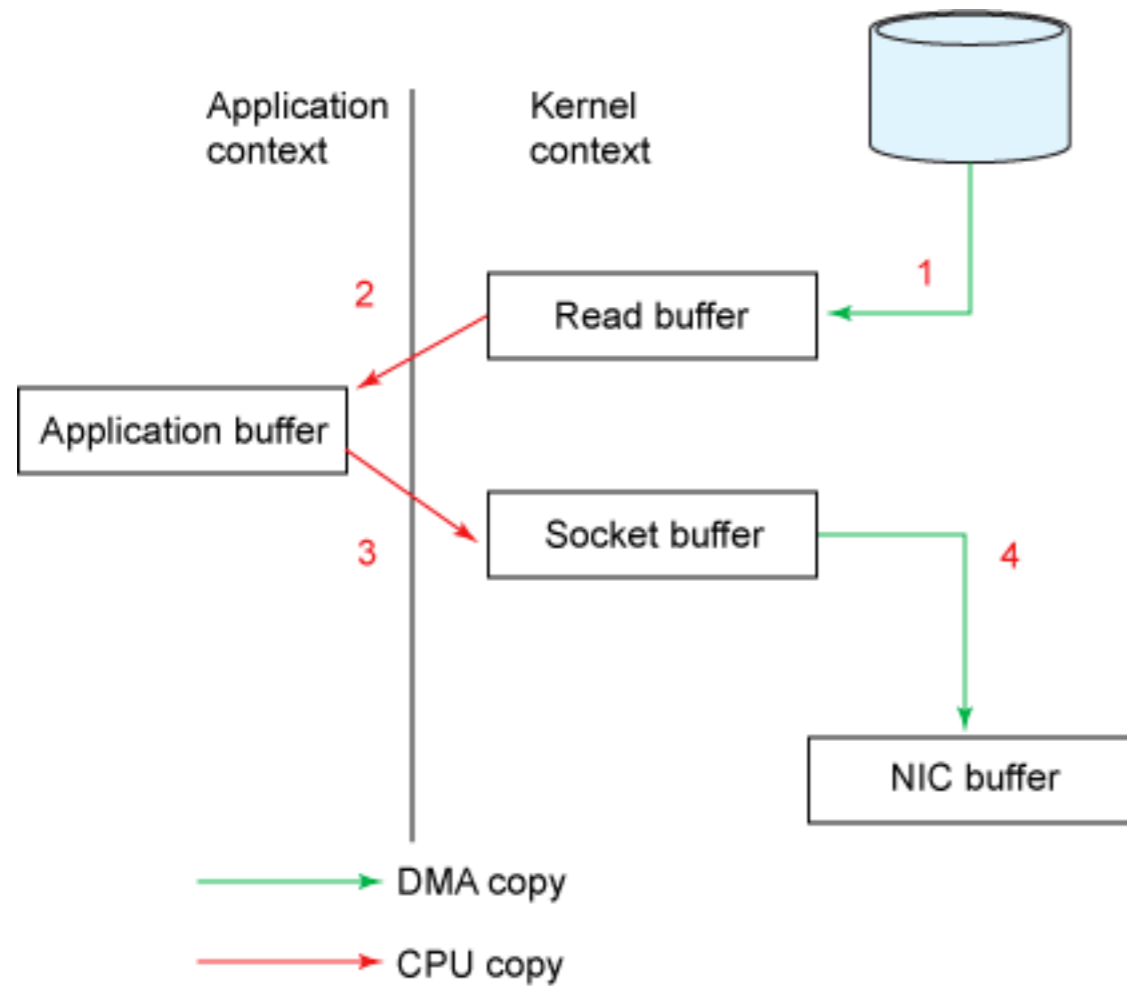


O.S. page cache is leveraged

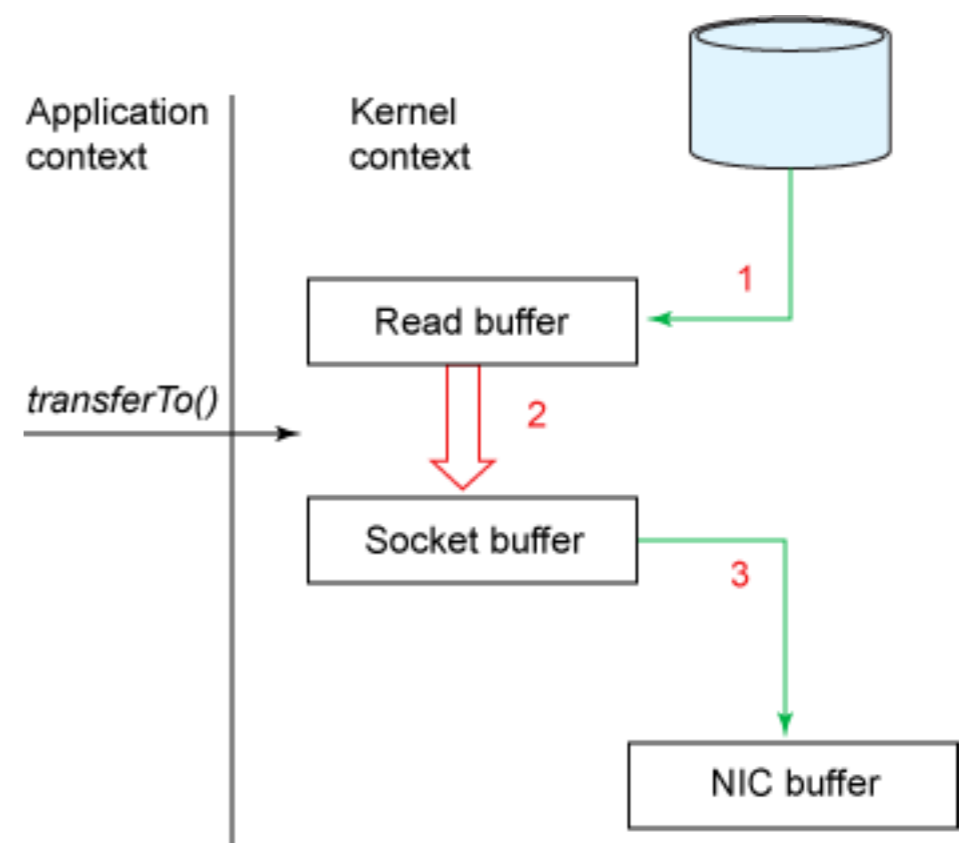


Zero-copy IO

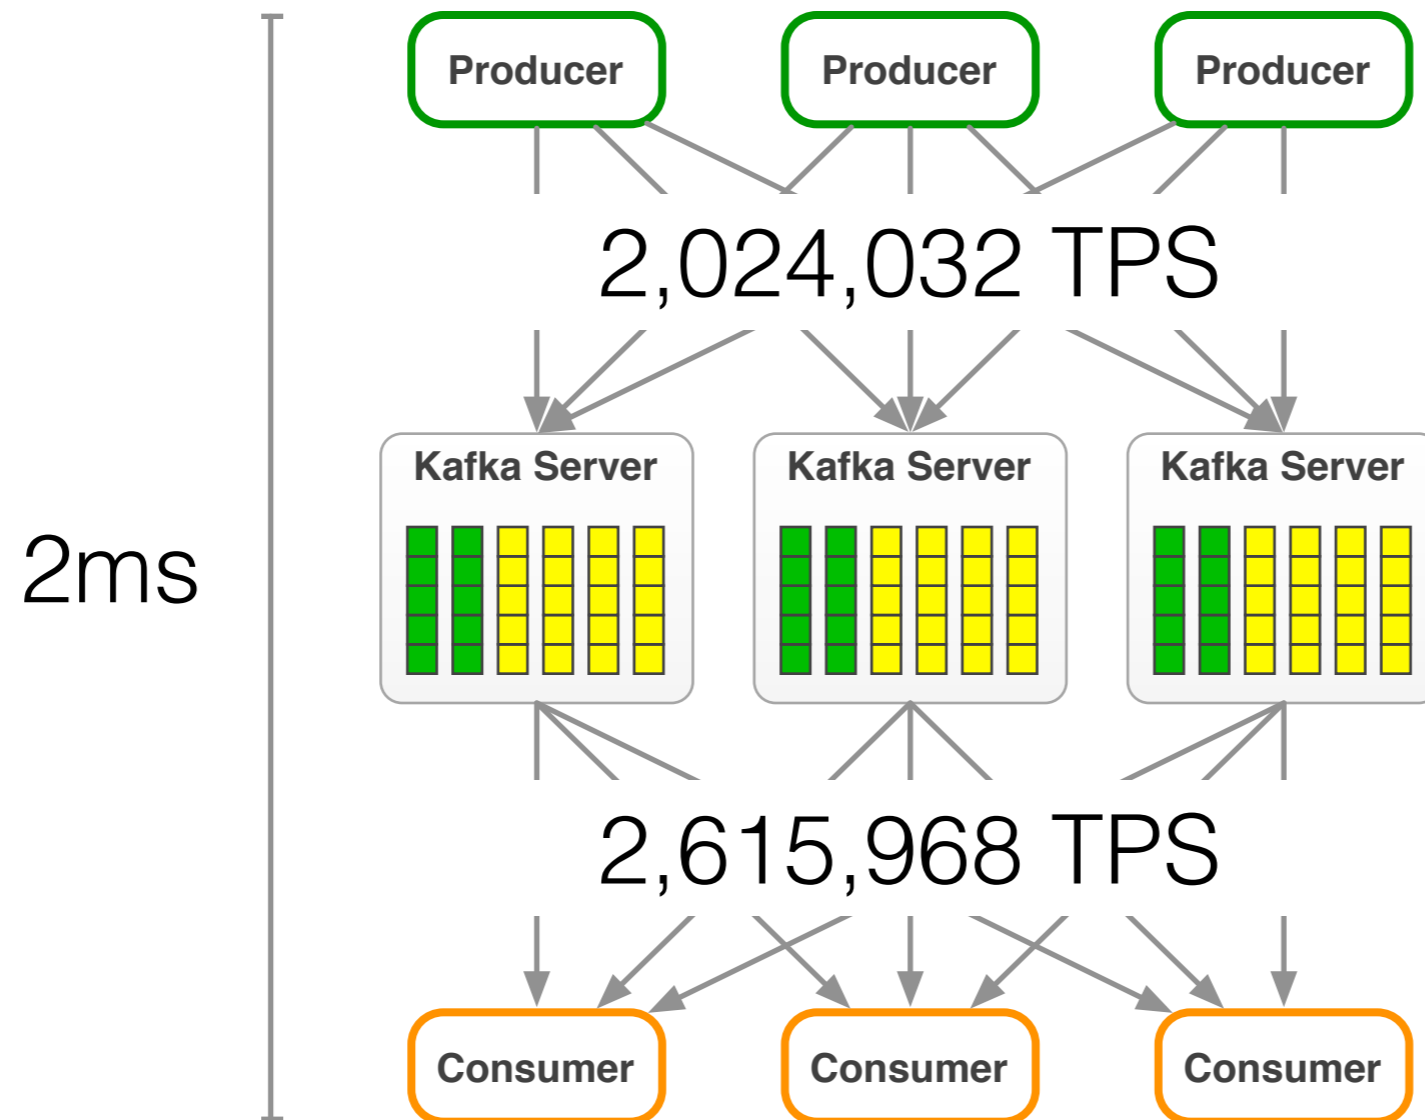
Traditional



Zero-copy

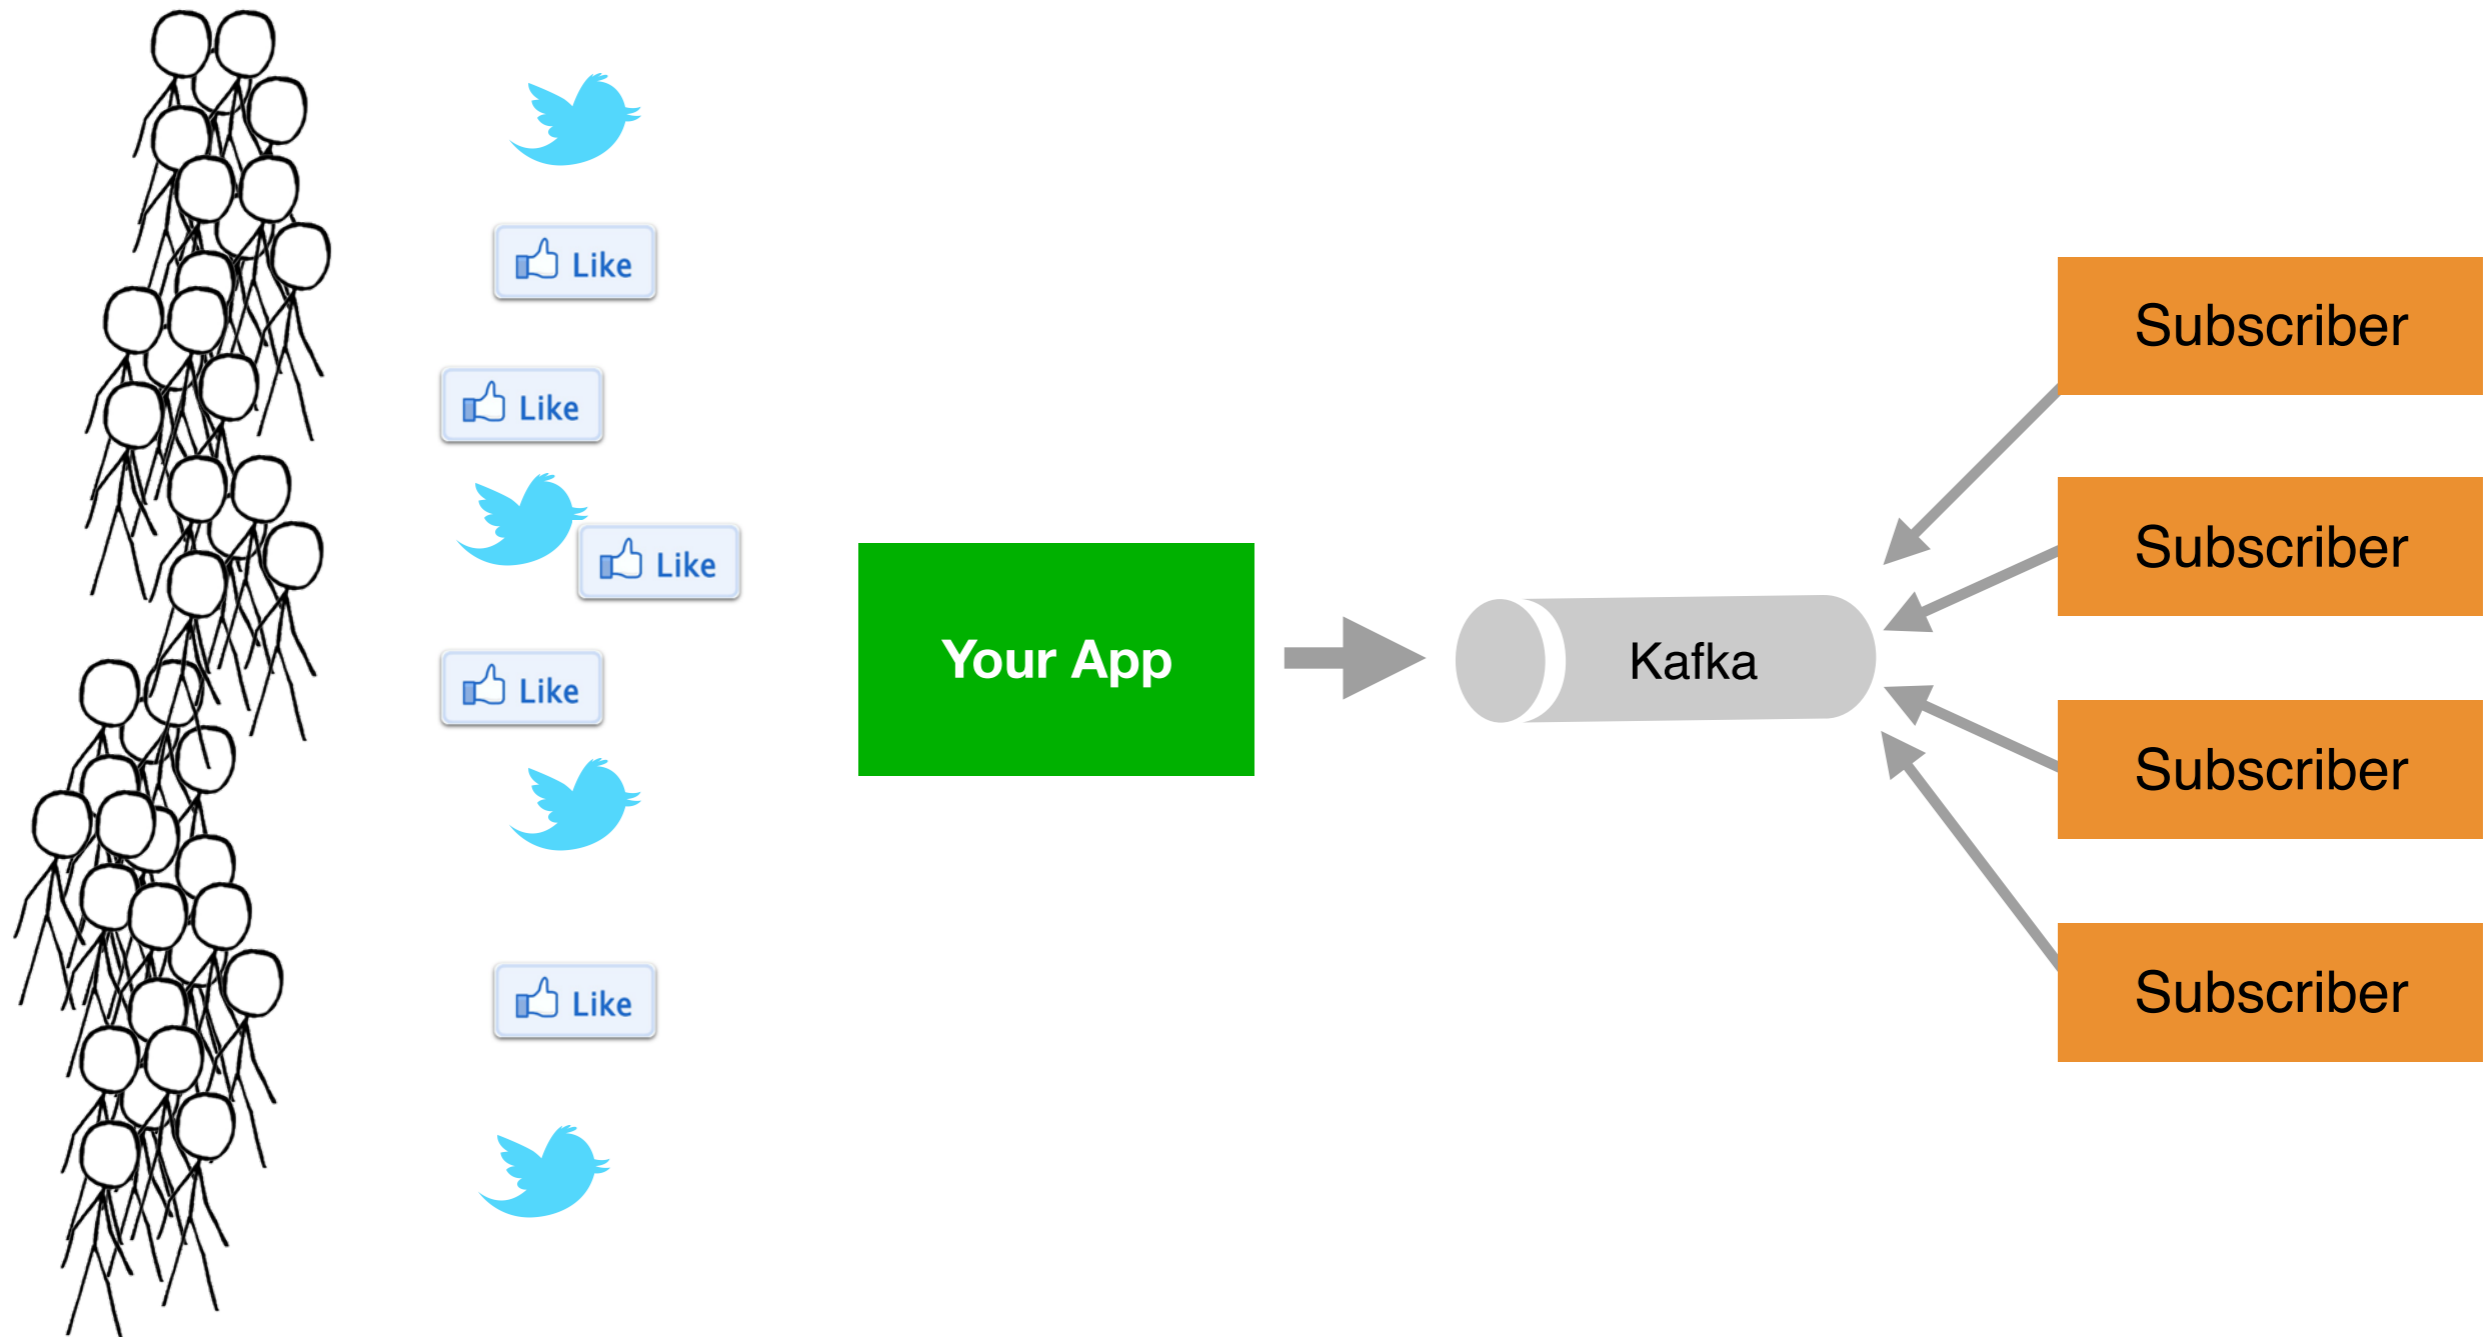


Throughput

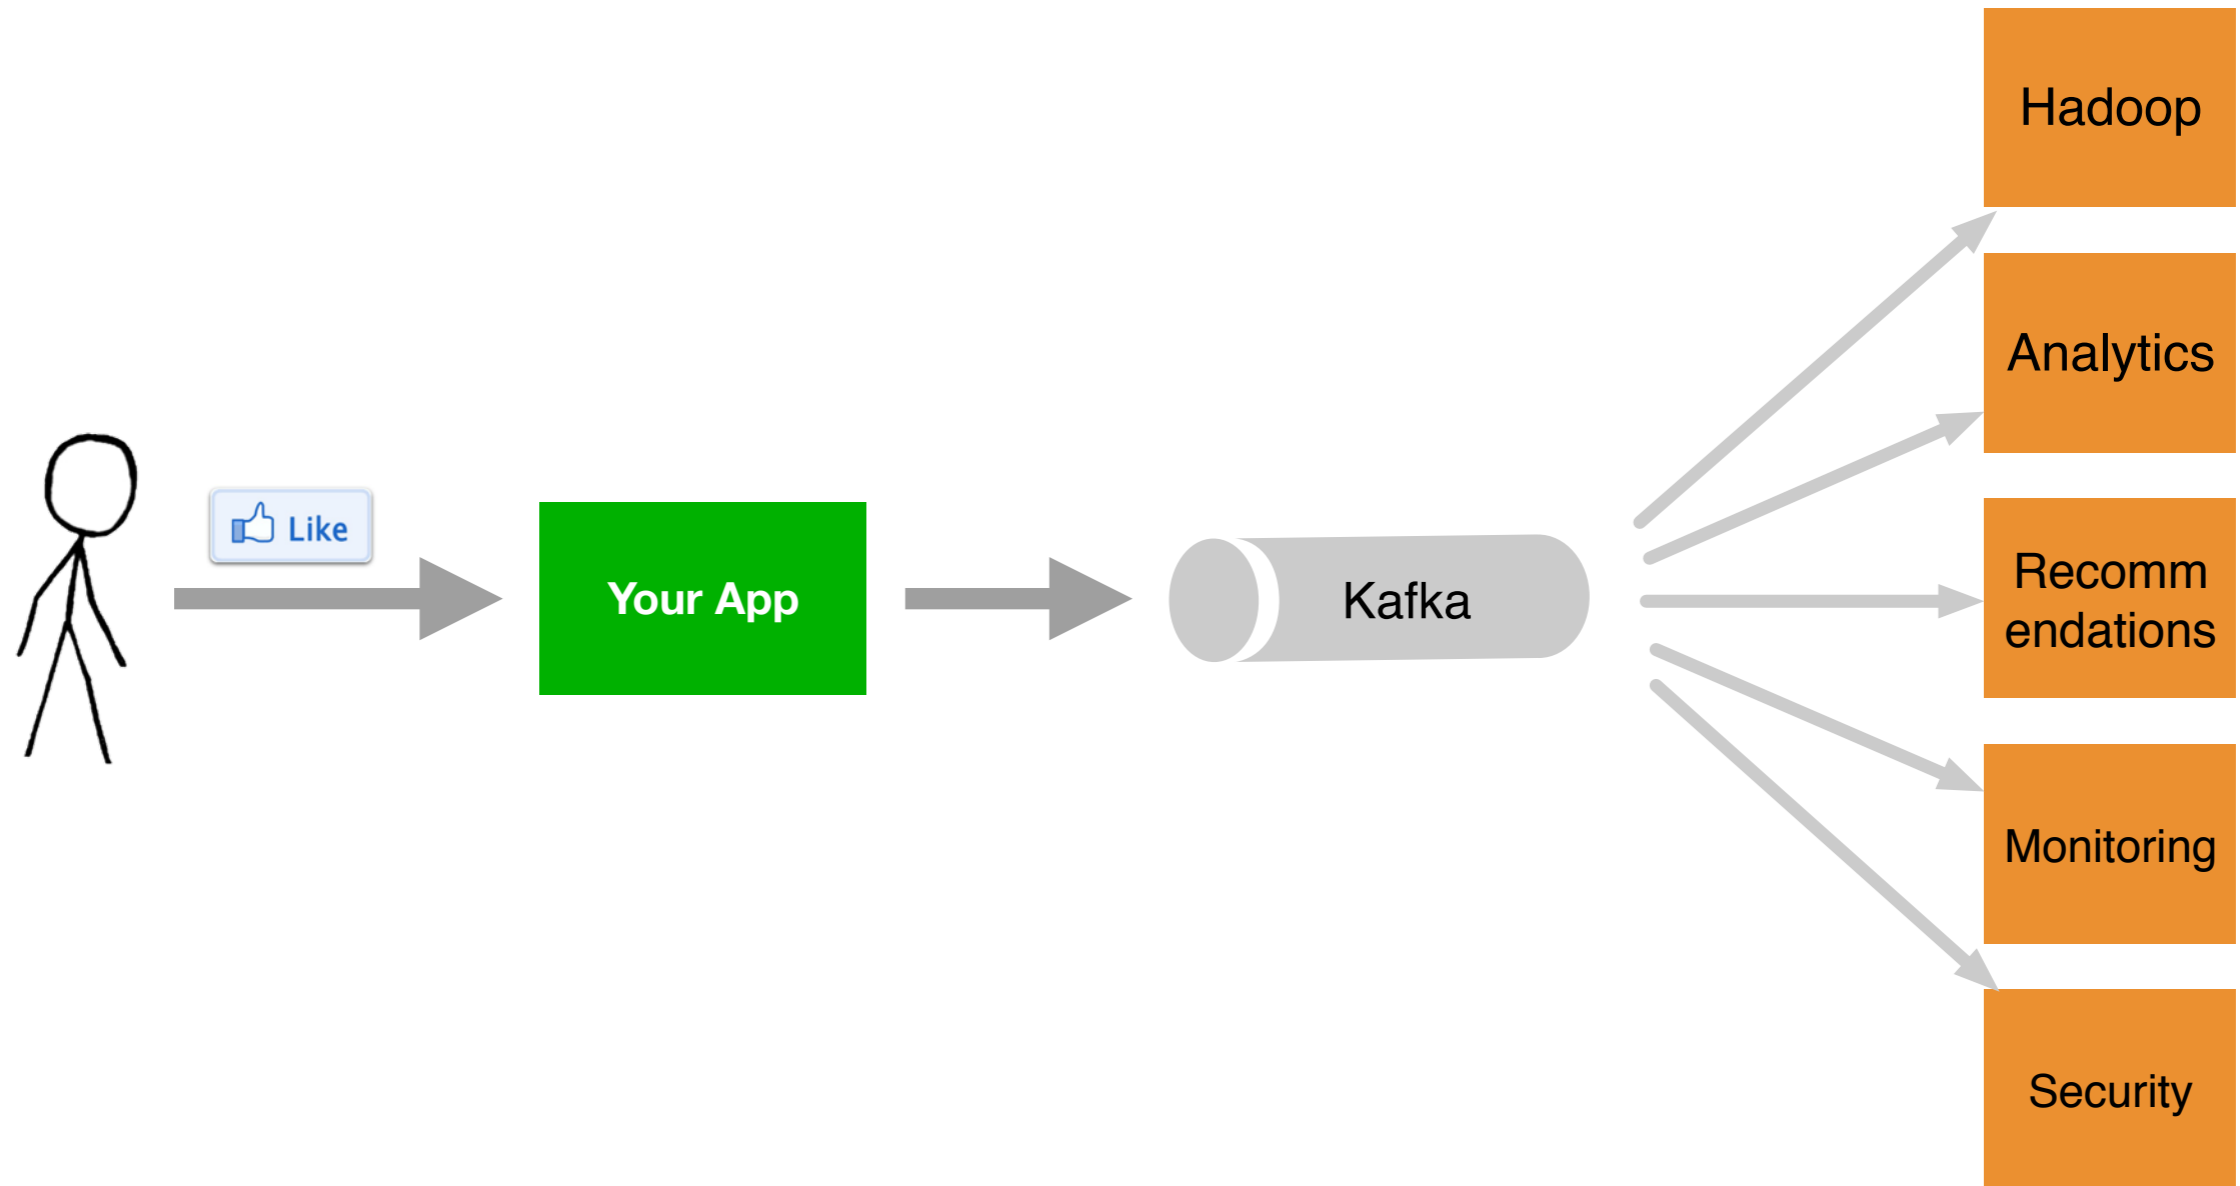


Kafka data integration patterns

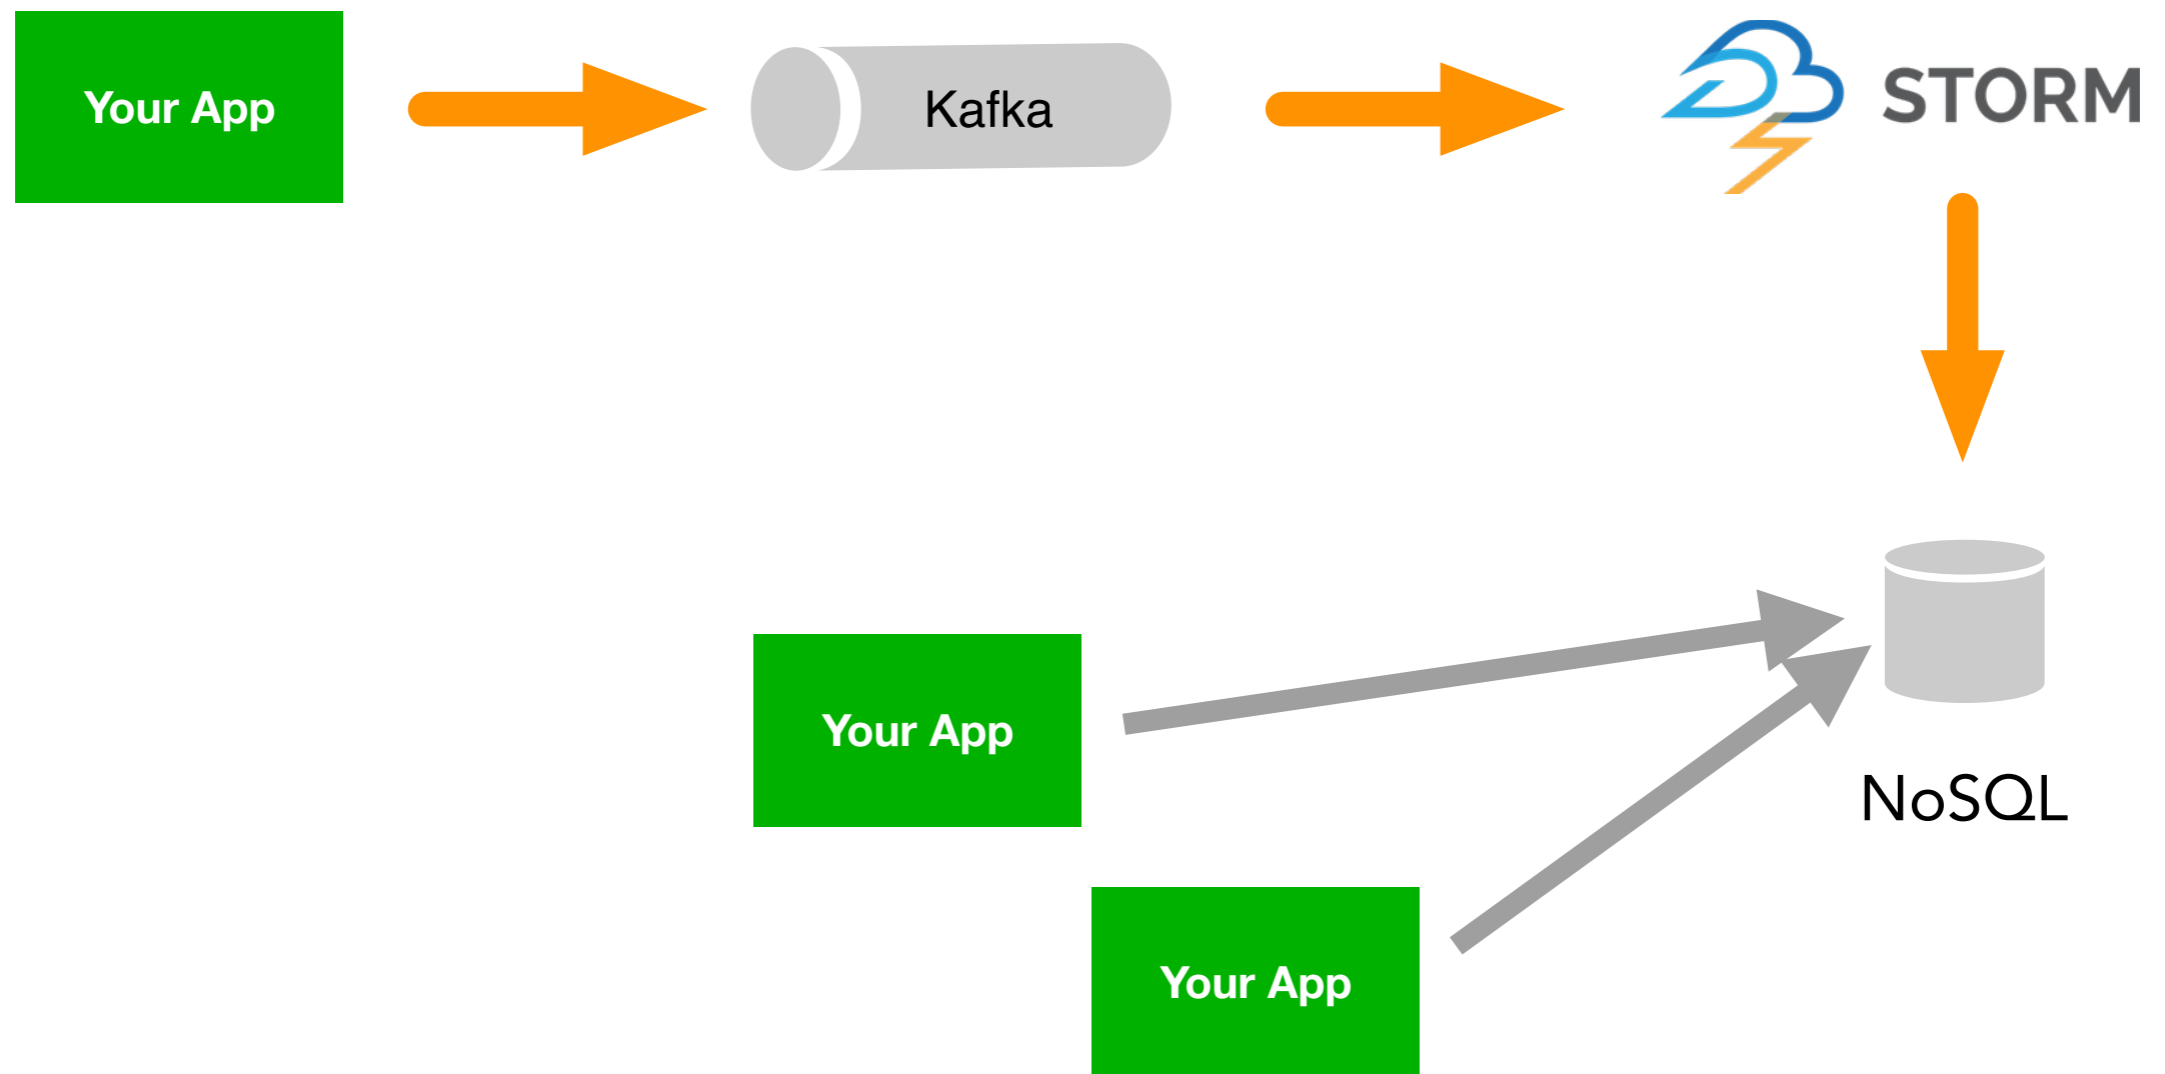
External data subscription



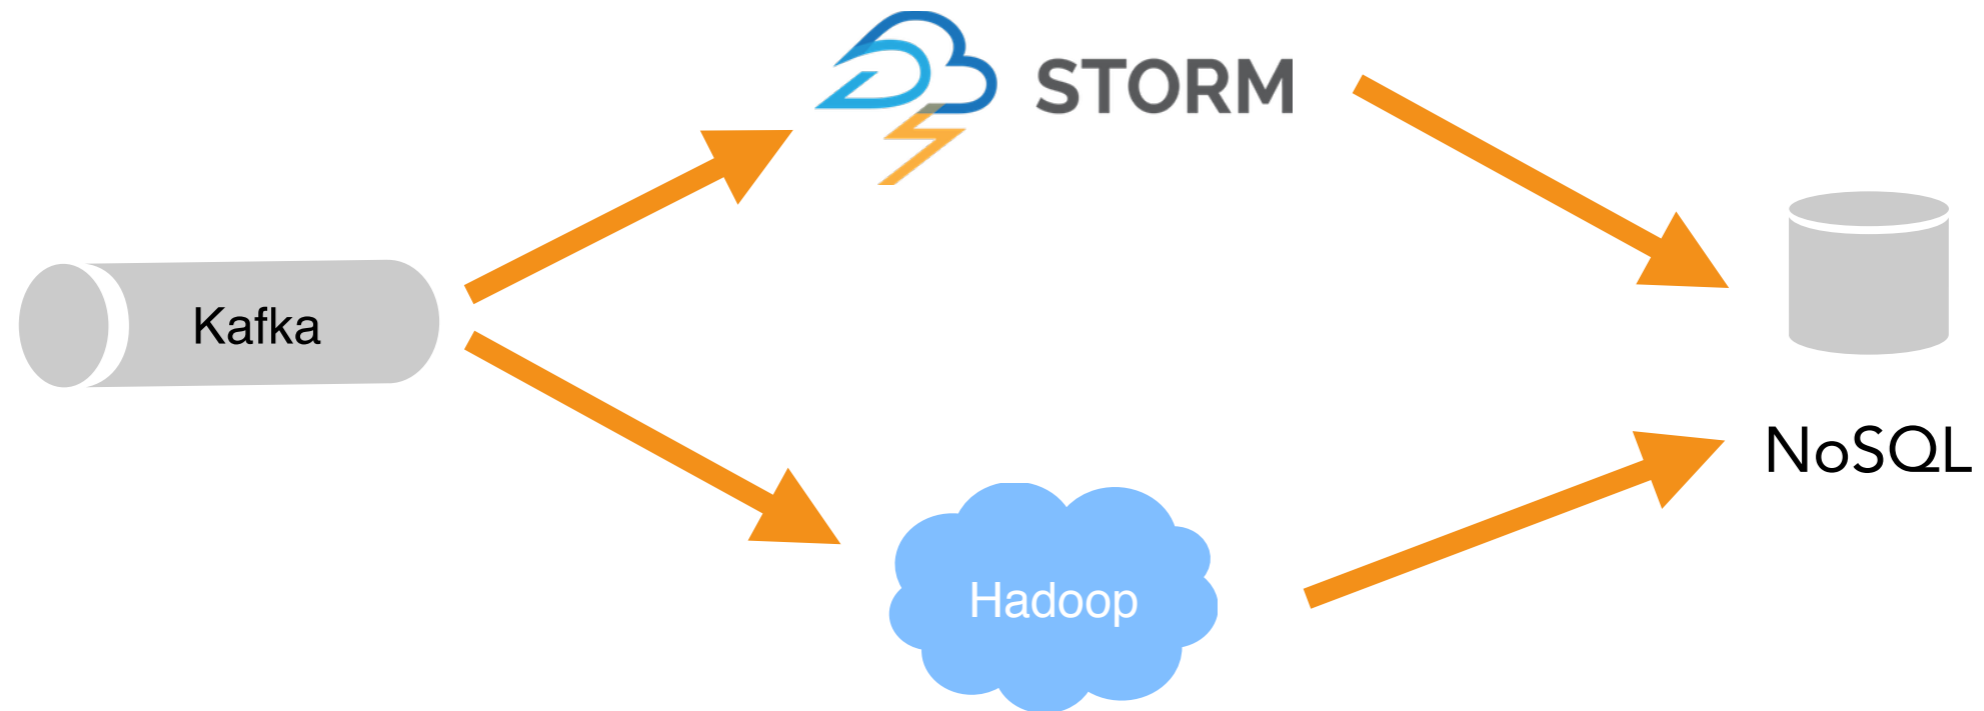
Internal data subscription



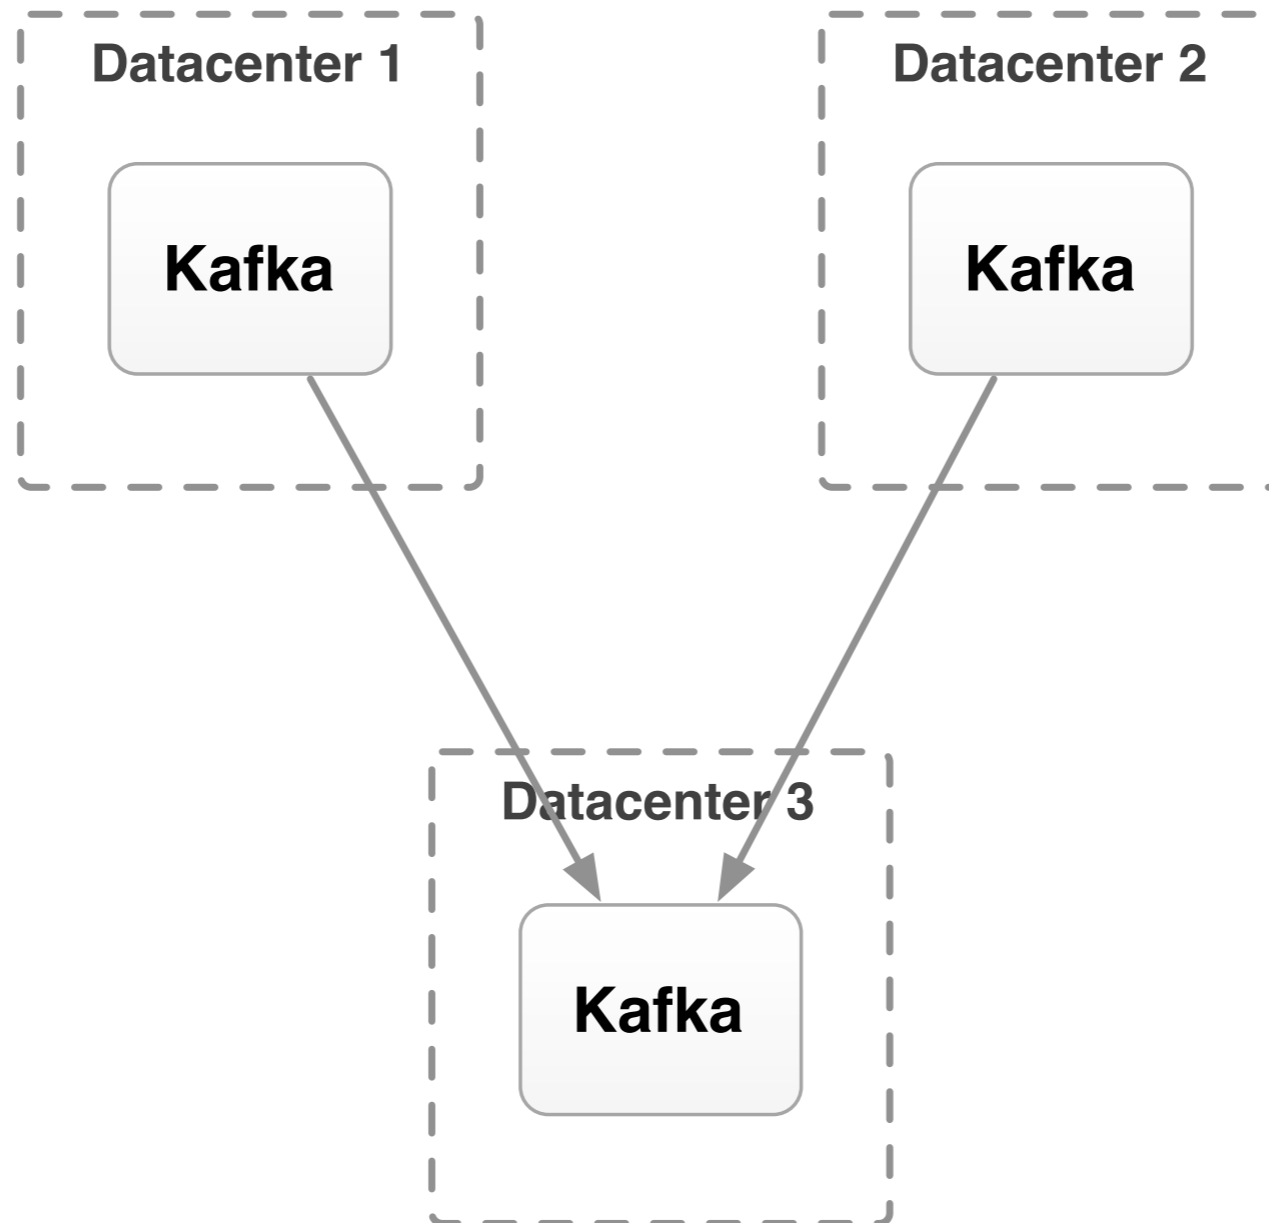
Stream processing



Lambda architecture



Multi-datacenter



Additional patterns

- Partition on user ID so all events for a user land on a single partition
- Publish system and application logs for downstream aggregation and monitoring
- System of record for your data


Data standardization

- Important to pick a lingua franca for your data
- Considerations should include compactness, schema evolution, tooling and third-party support
- Avro is a good fit (coupled with Parquet for on-disk)

Conclusion

- Kafka is a scalable system with useful traits such as ordered messages, durability, availability
- Popular for data movement and feeding multiple downstream systems using the same data pipe
- Excellent documentation available at <http://kafka.apache.org/documentation.html>

Shameless plug



Book signing at the JavaOne bookstore @ 1pm tomorrow

CON3515 - The Top 10 Hadoop Patterns and Anti-patterns @ 11am tomorrow Parc 55 - Embarcadero