

# The Many Faces of Apache Kafka: Leveraging Real-time Data at Scale

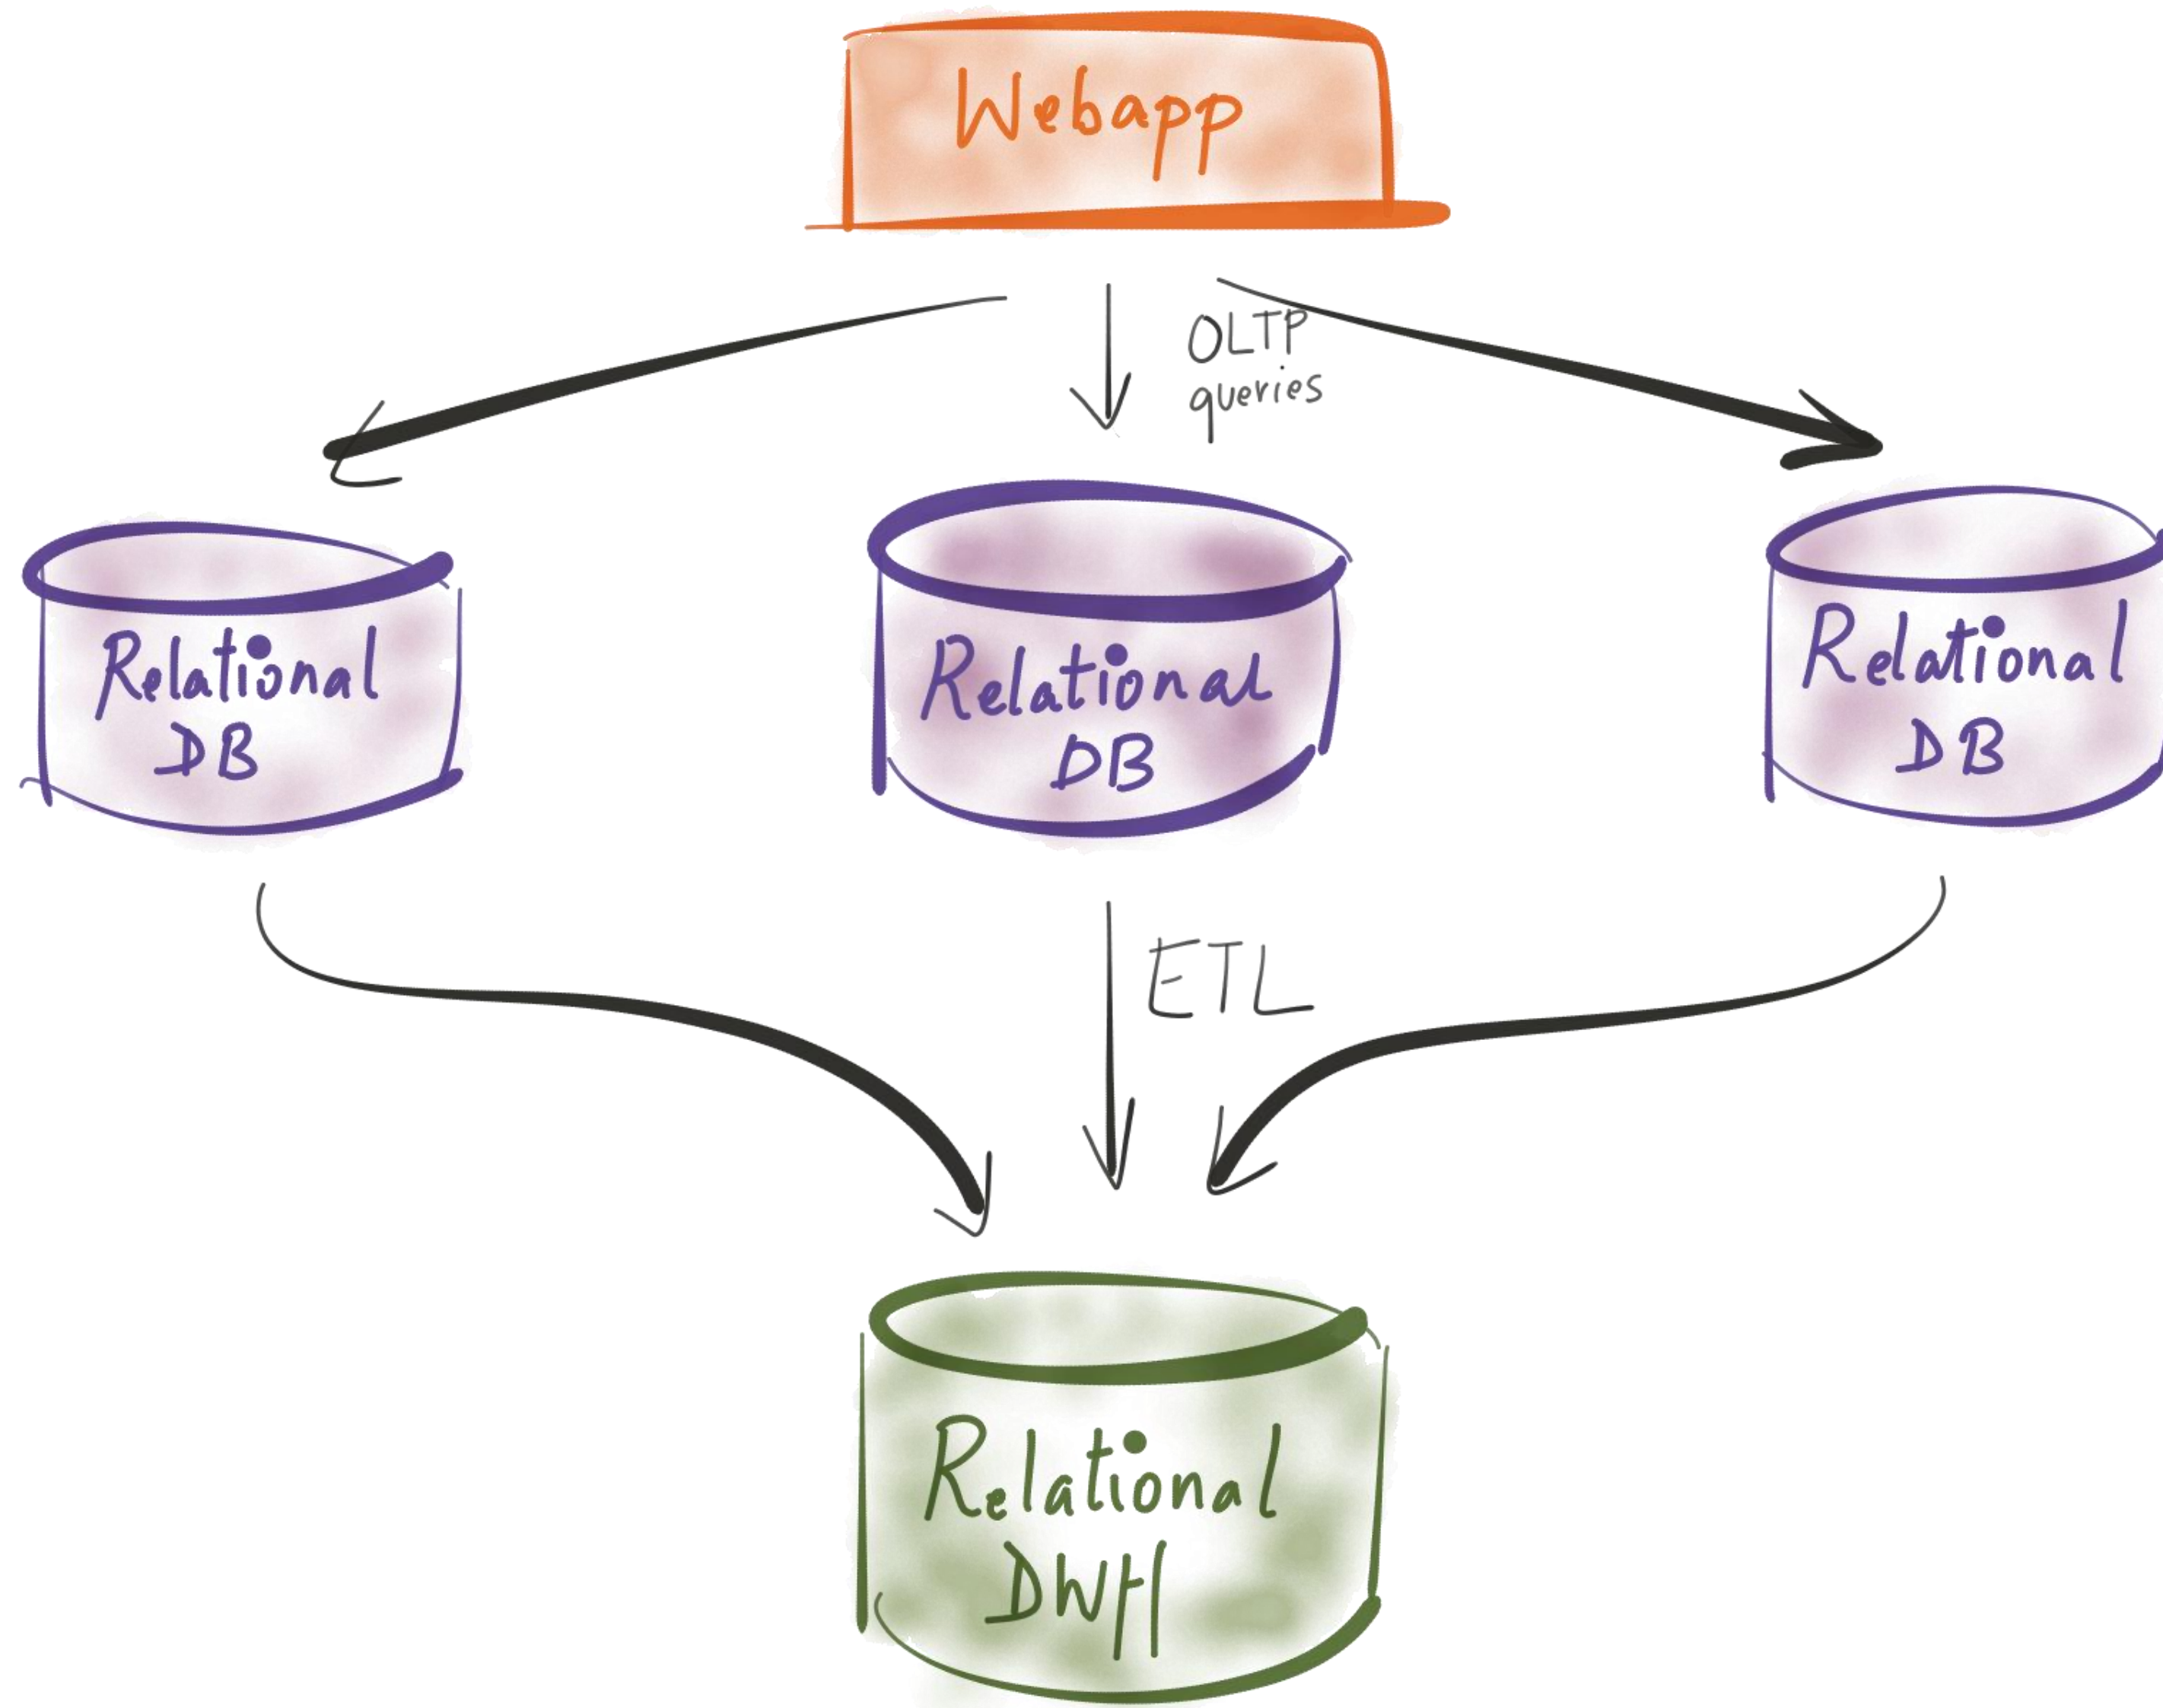
Neha Narkhede, Confluent



circa 2009

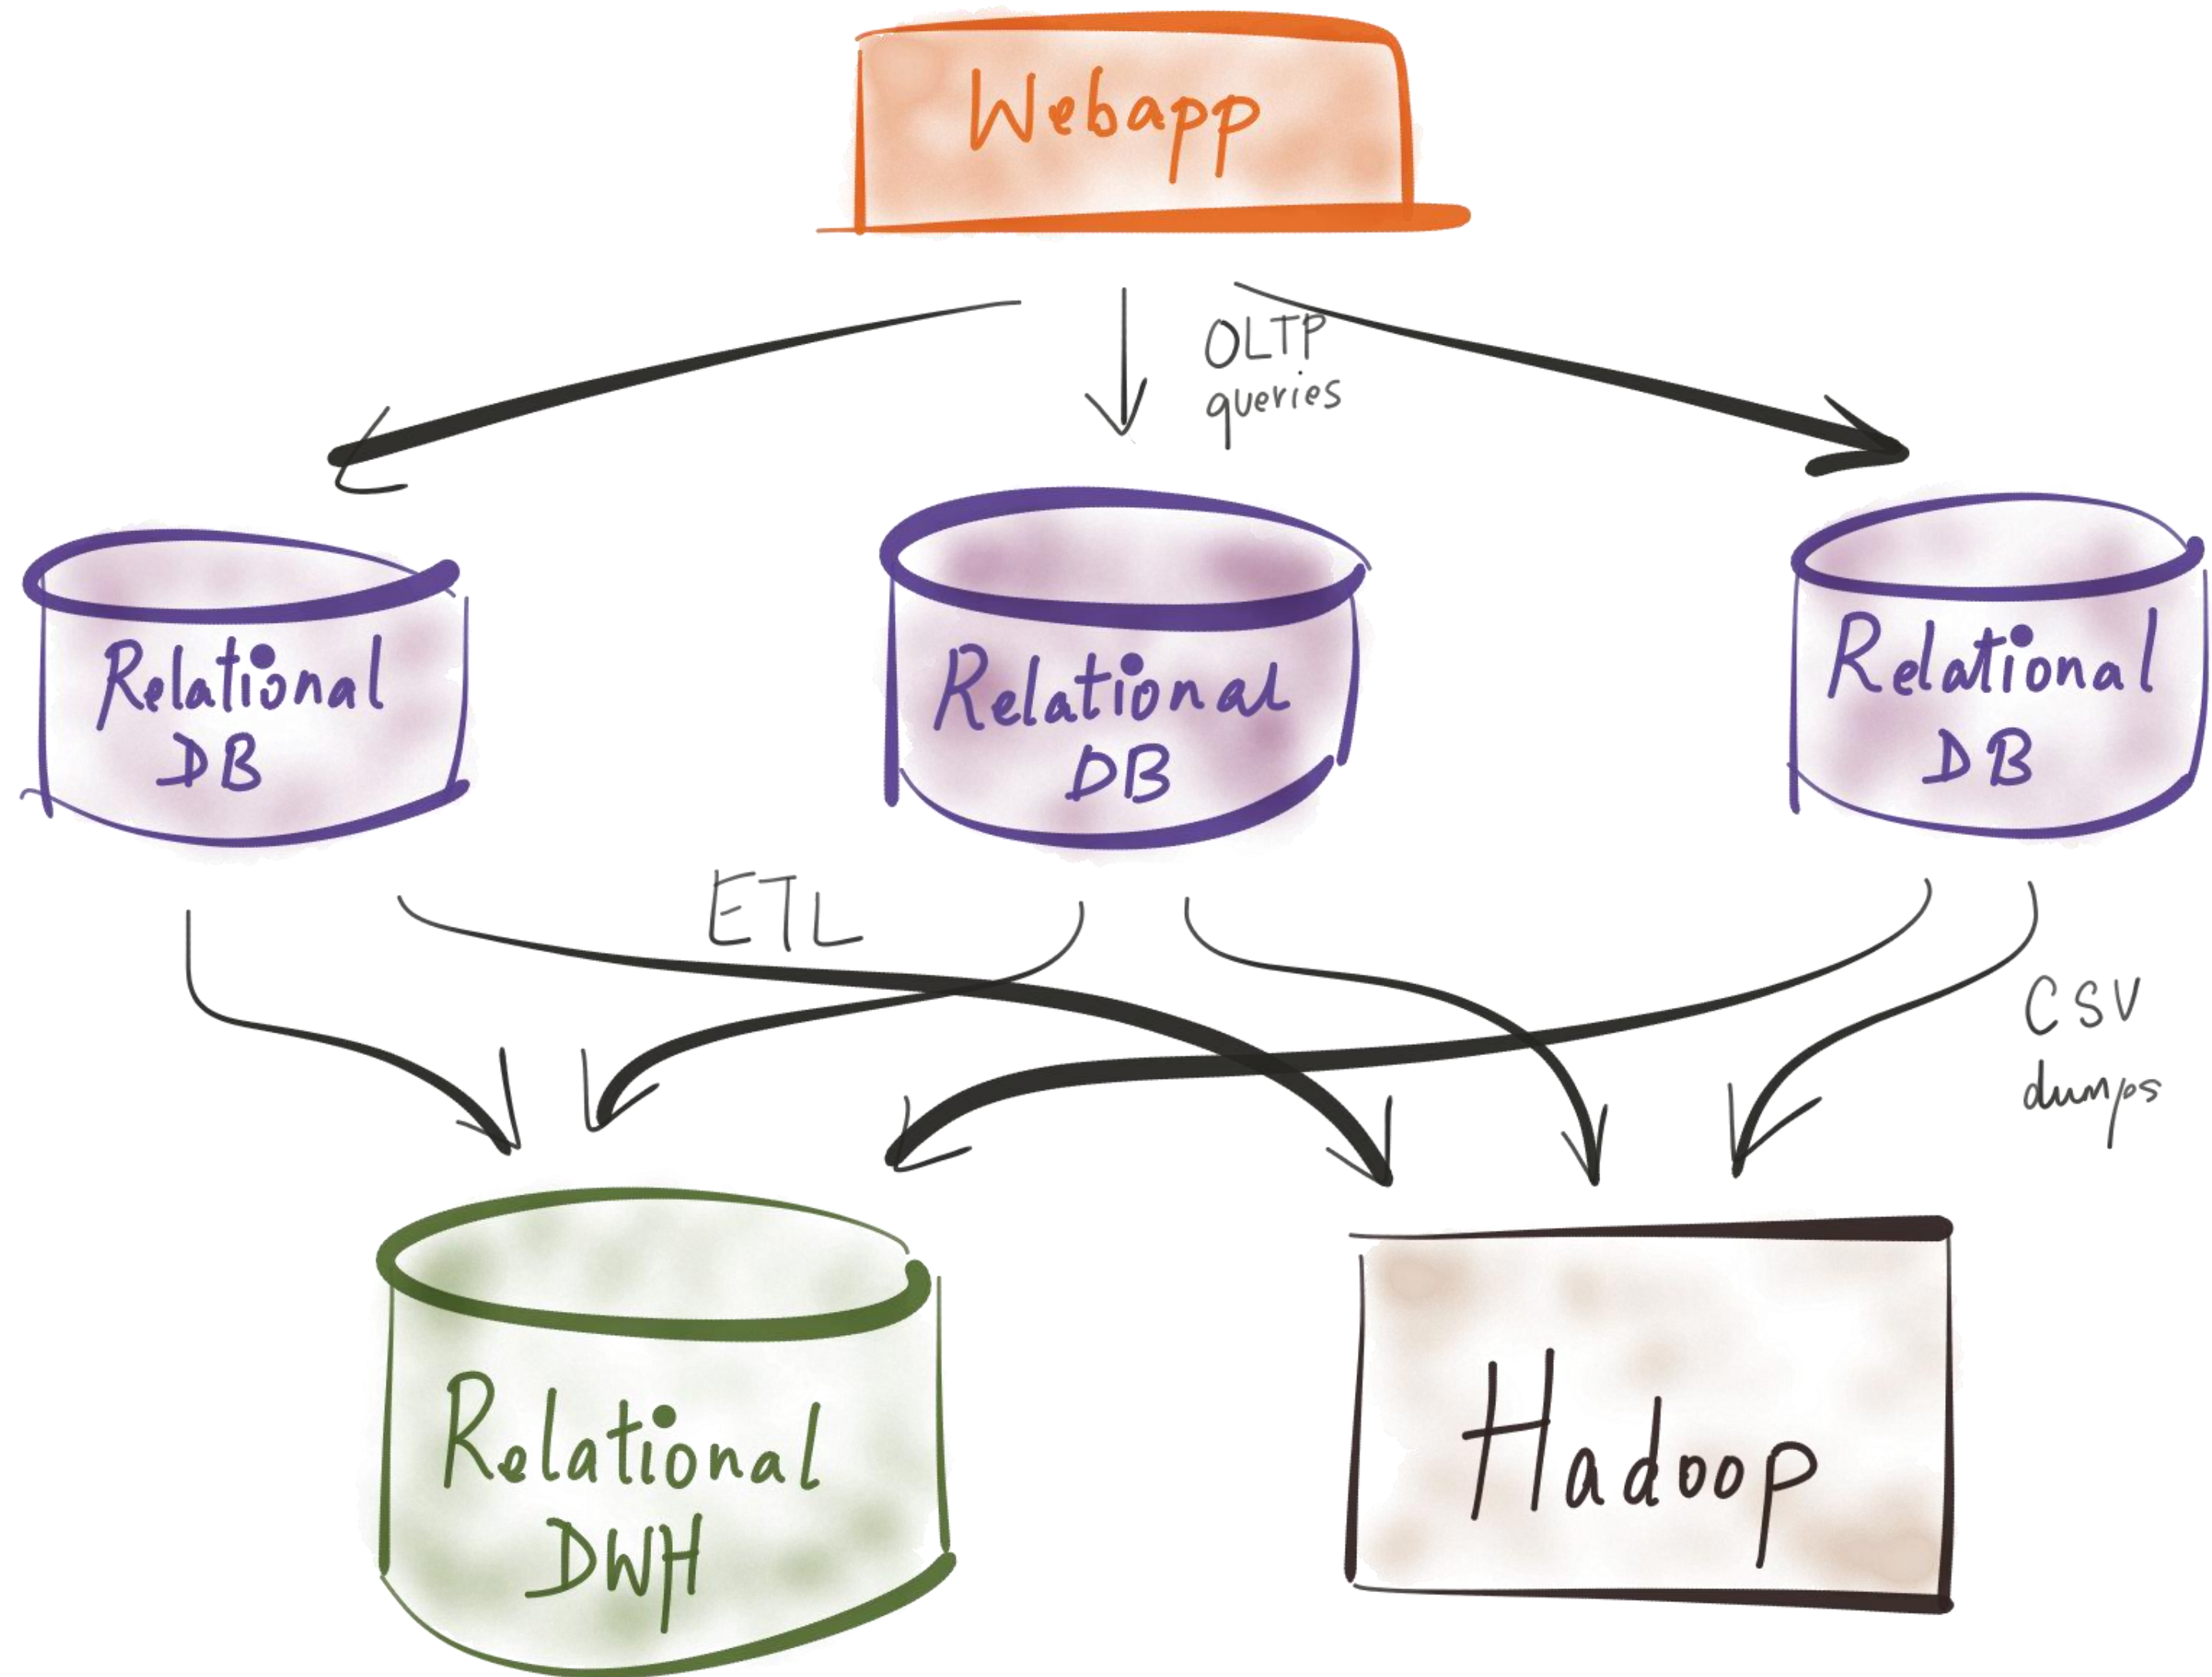


# RELATIONAL DB CHANGES



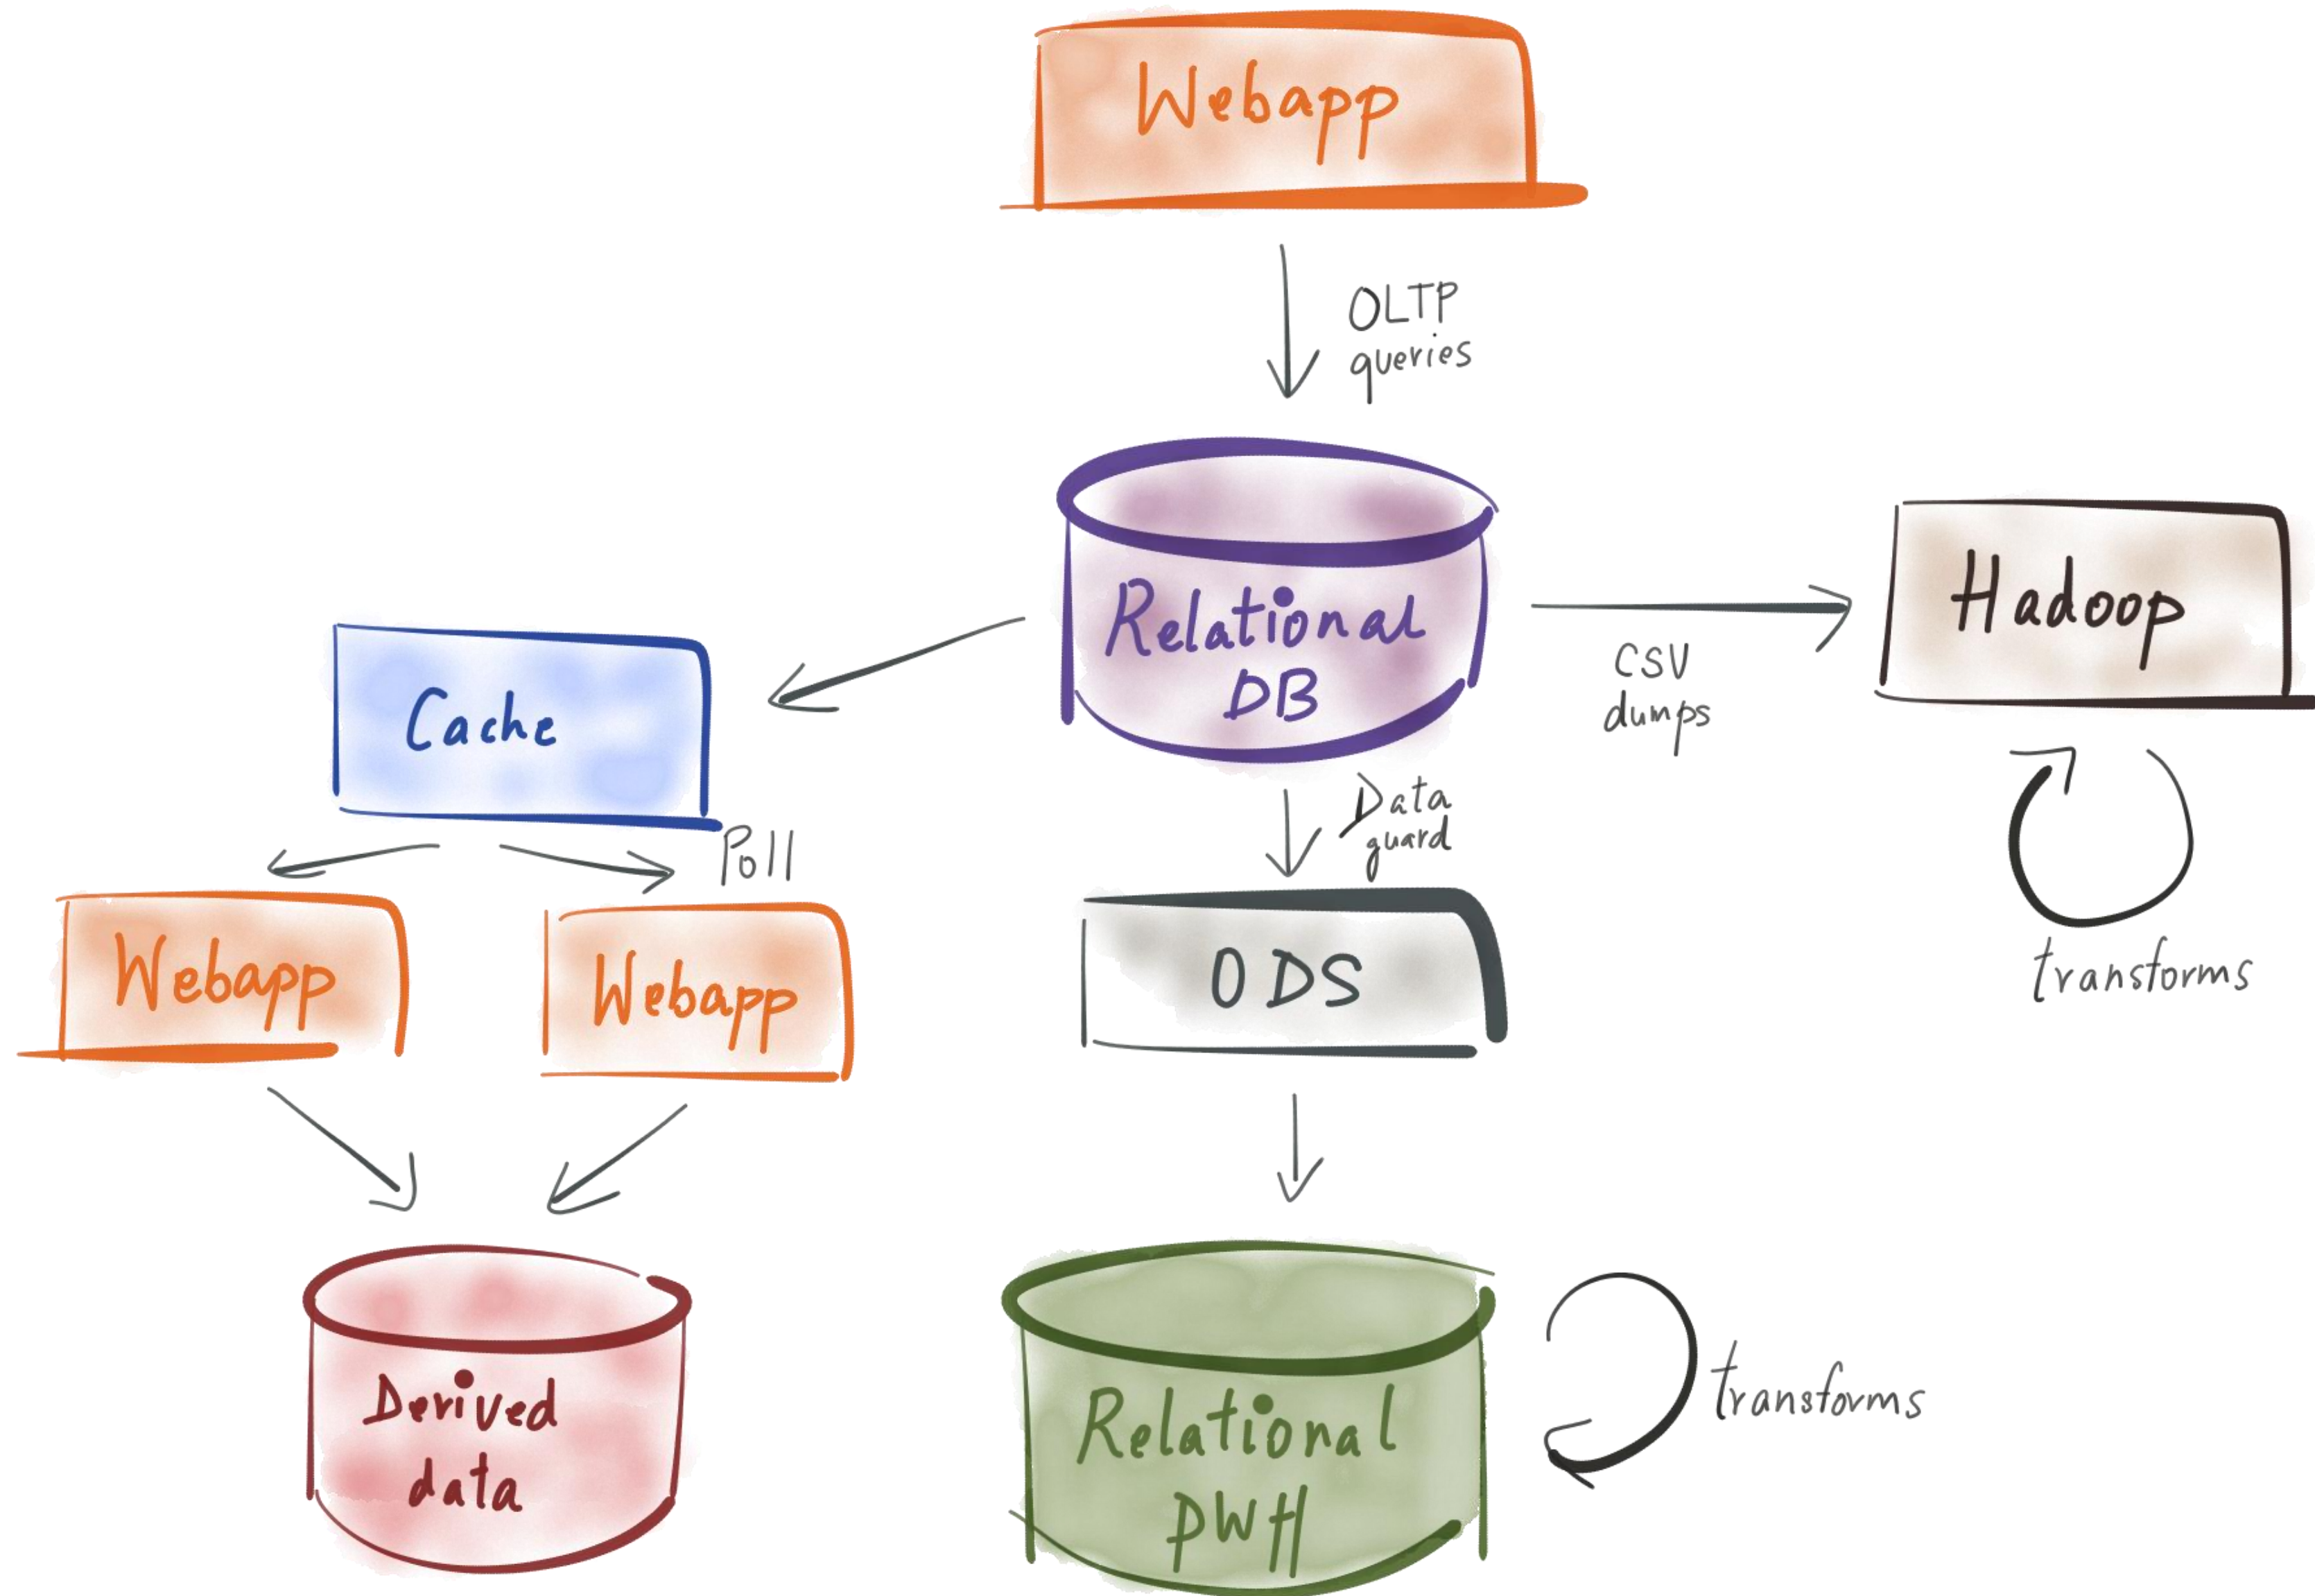


# RELATIONAL DB CHANGES



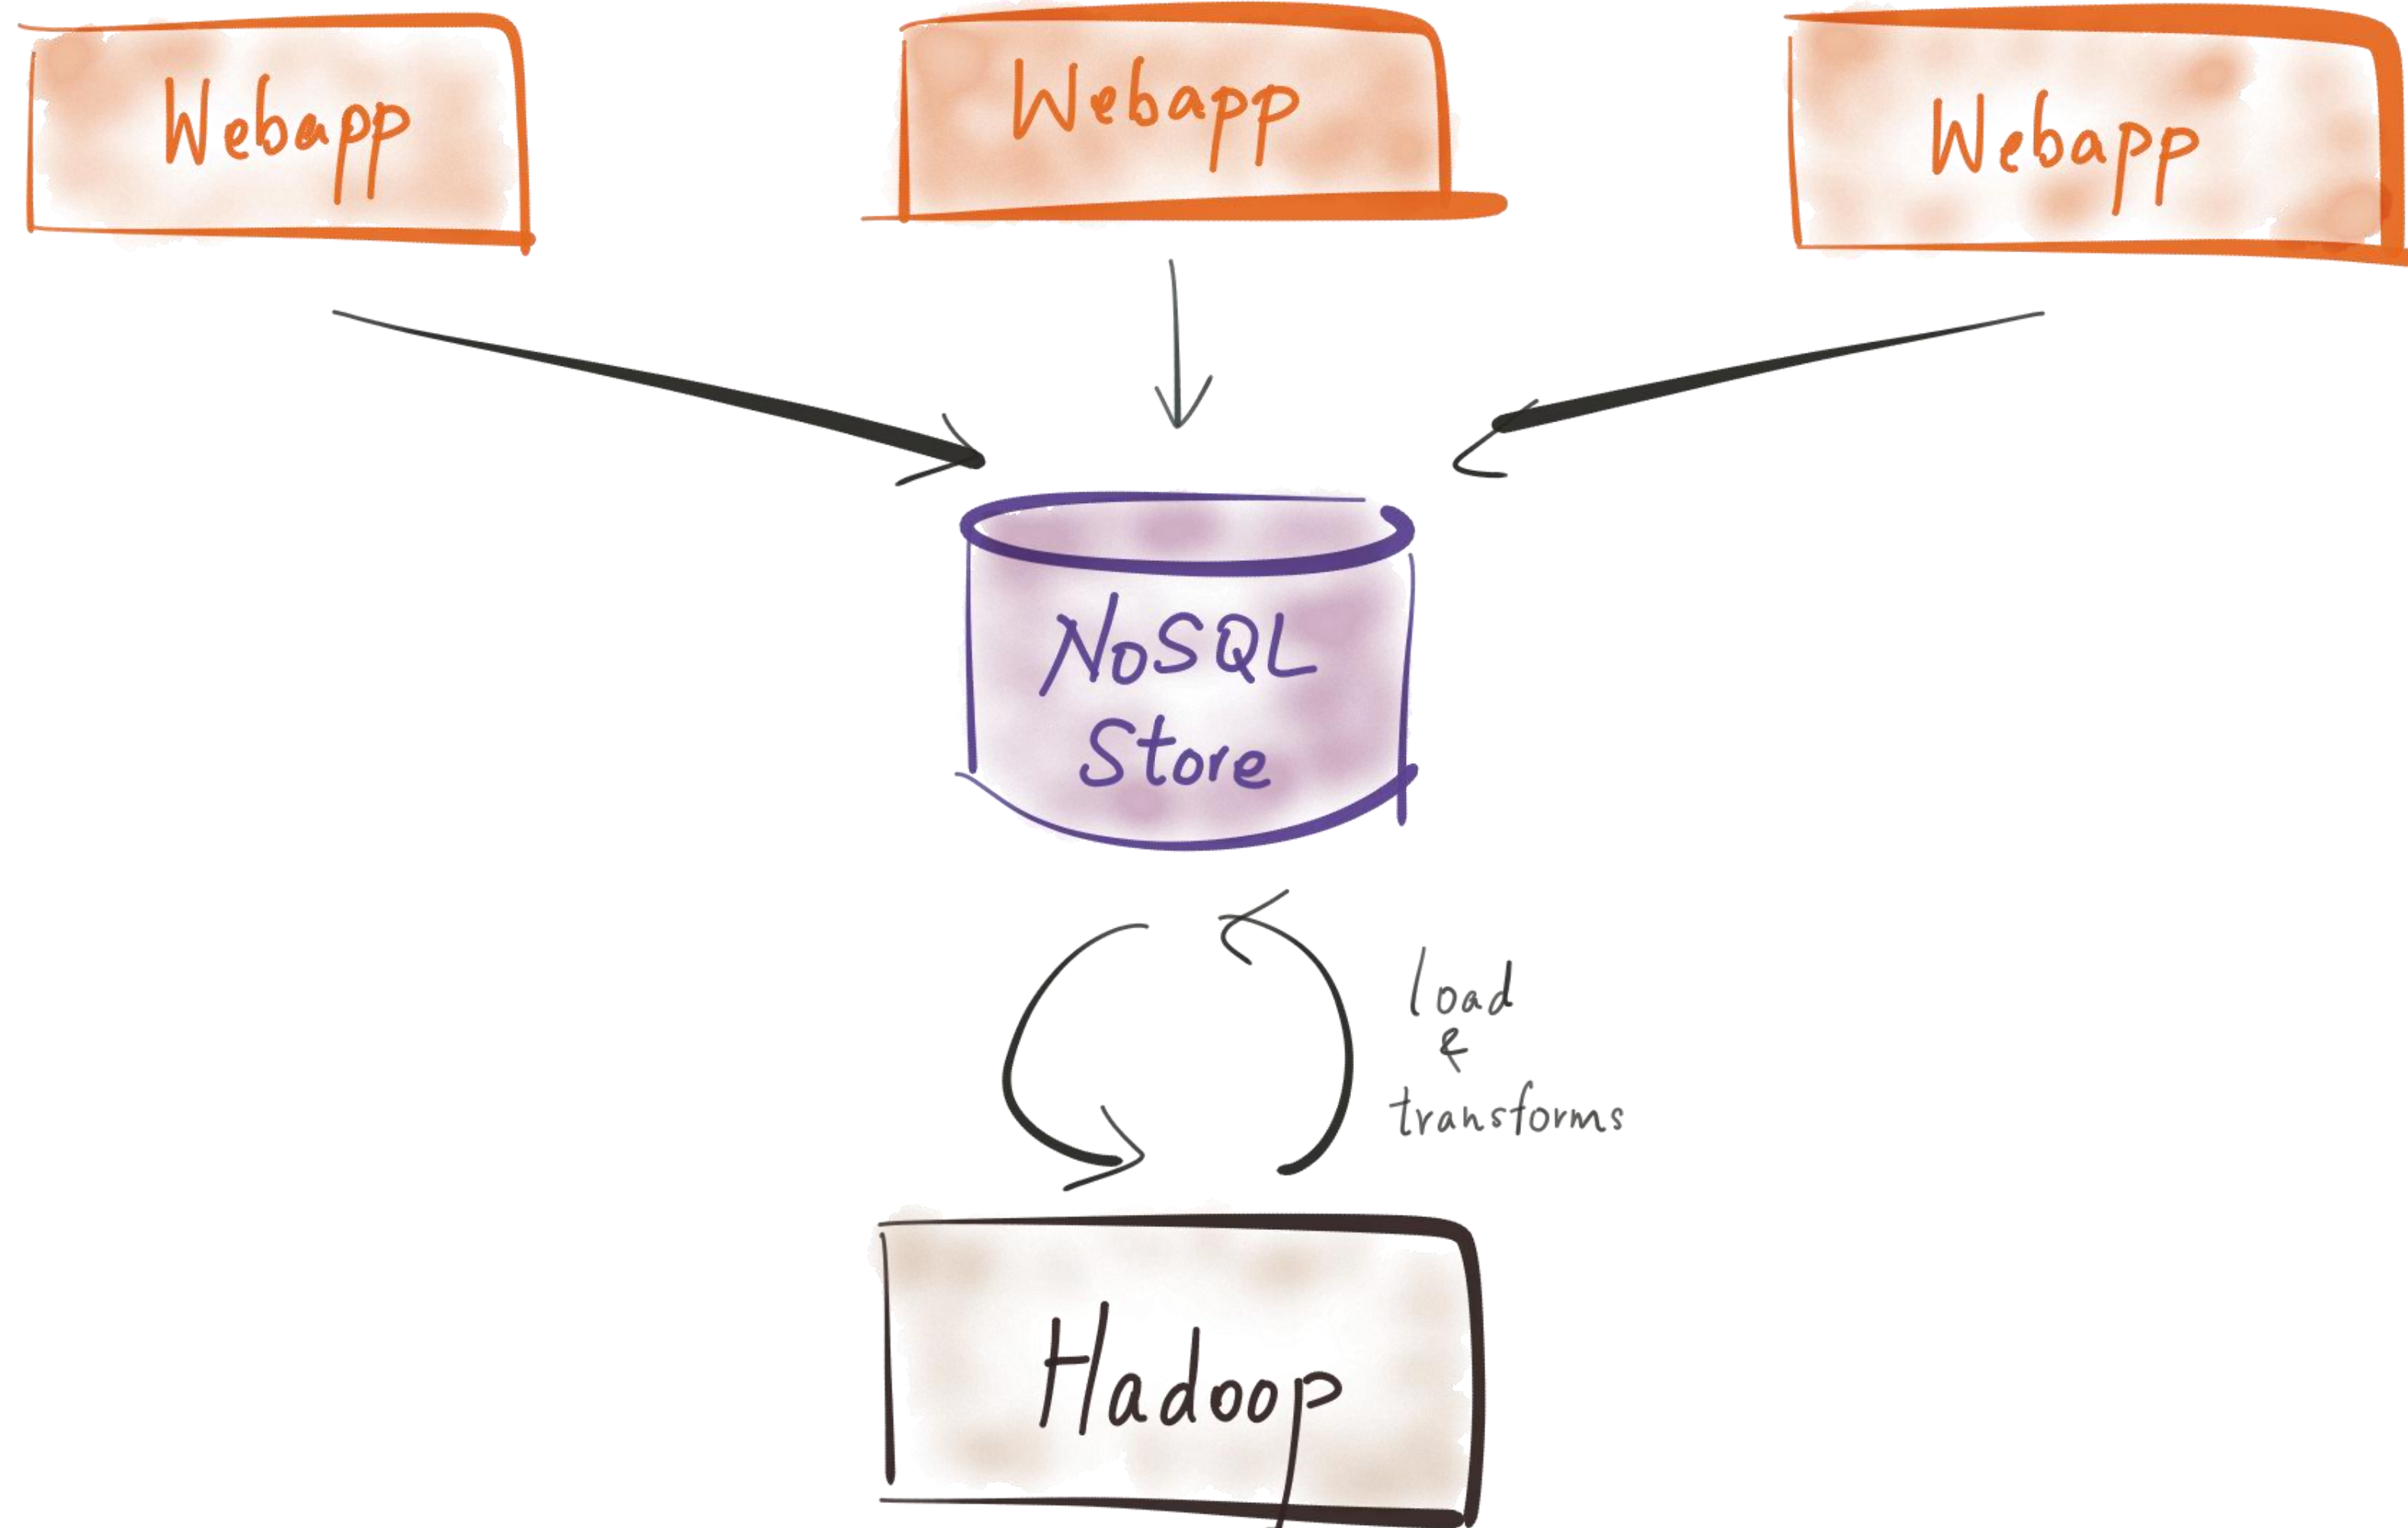


# RELATIONAL DB CHANGES



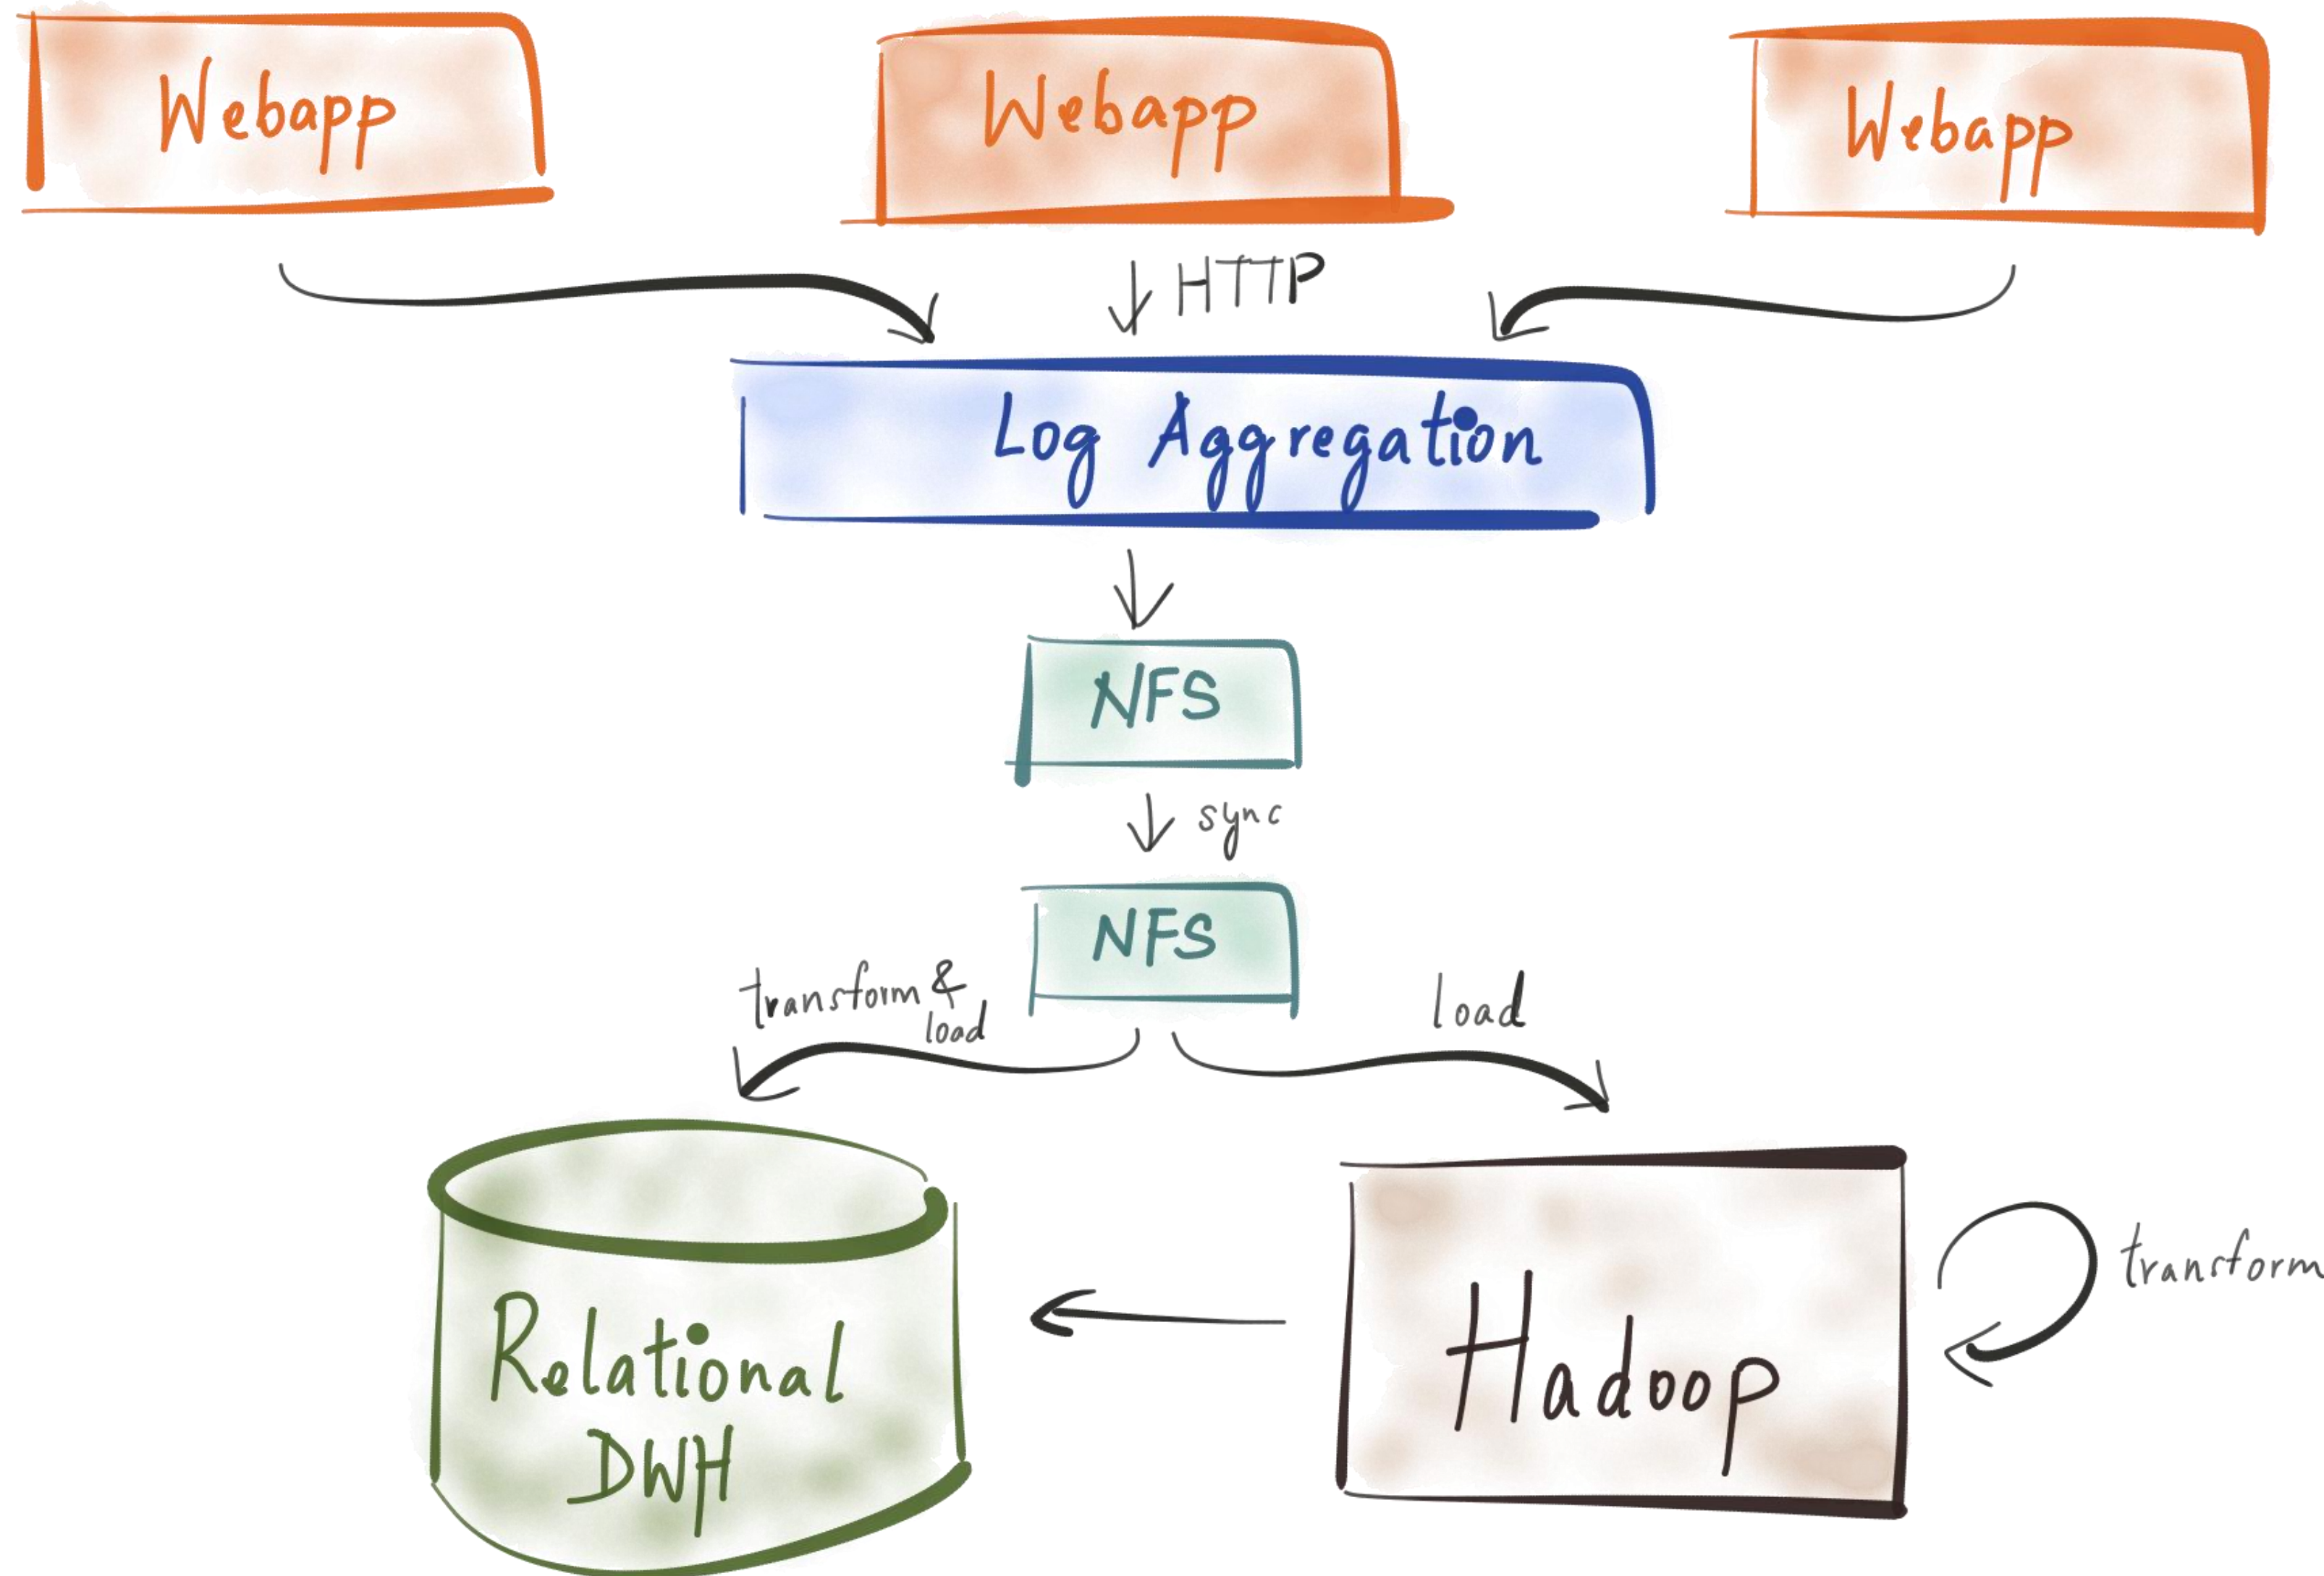


# NoSQL CHANGES



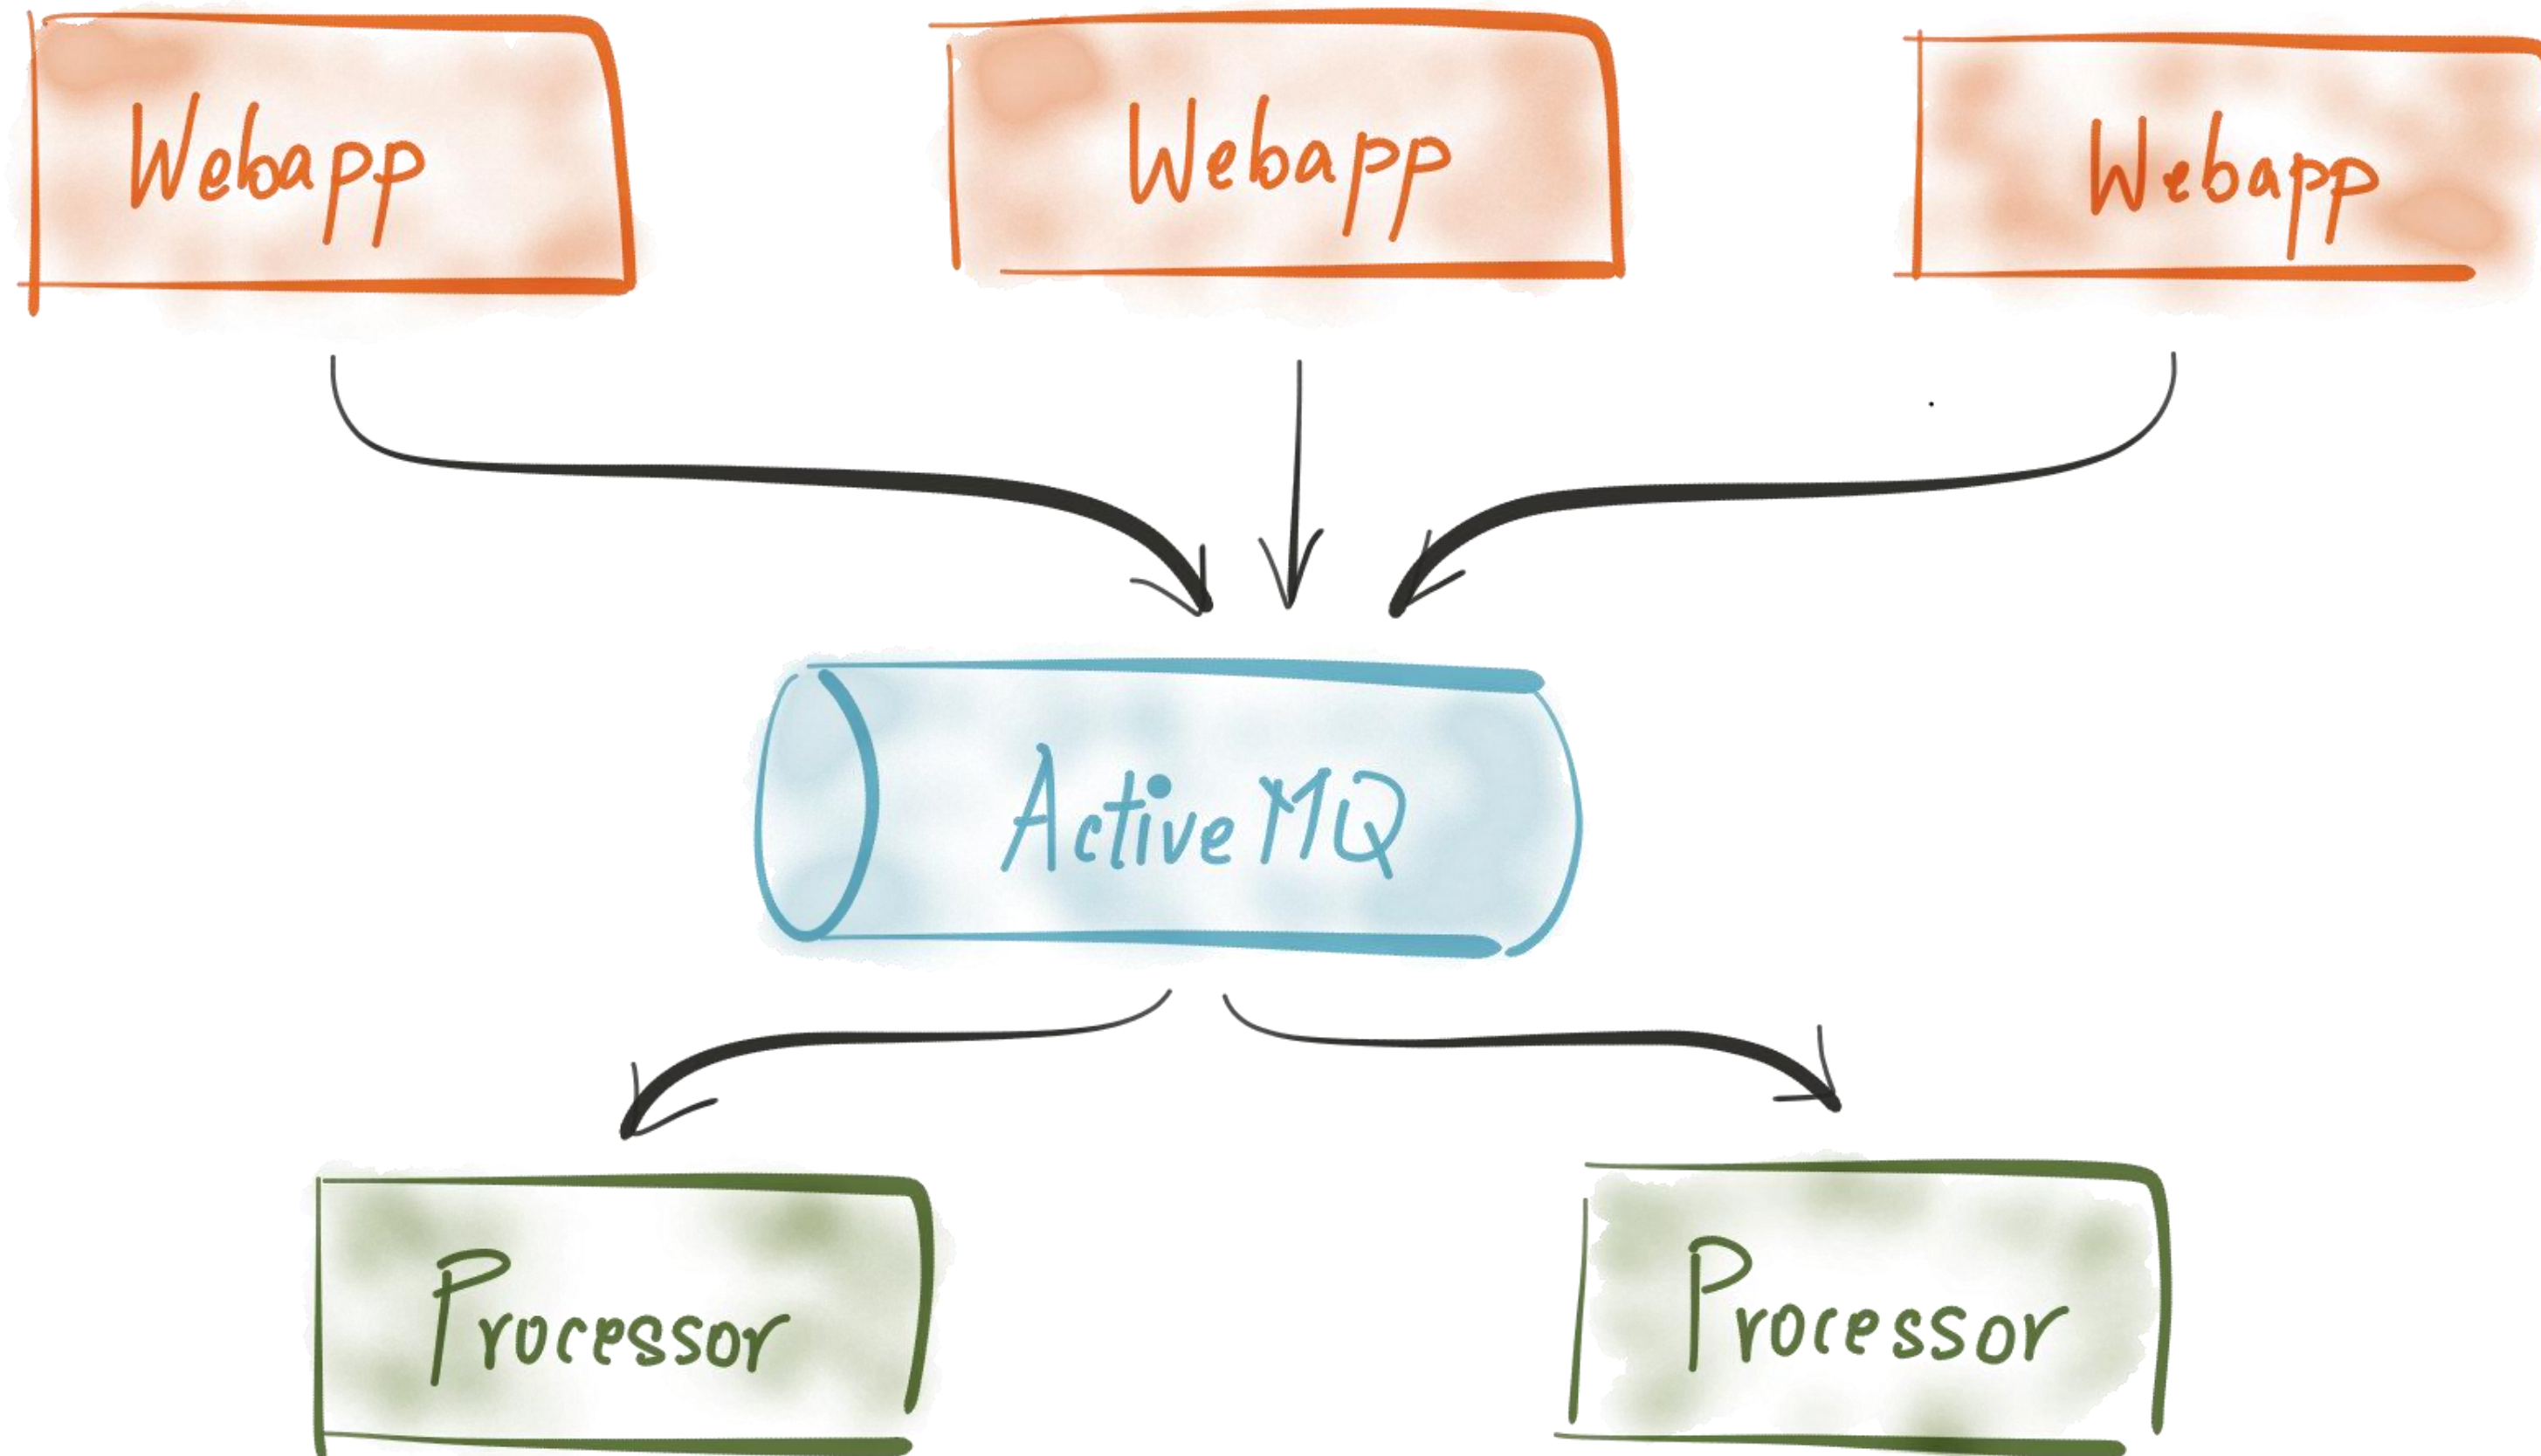


# USER ACTIVITY DATA





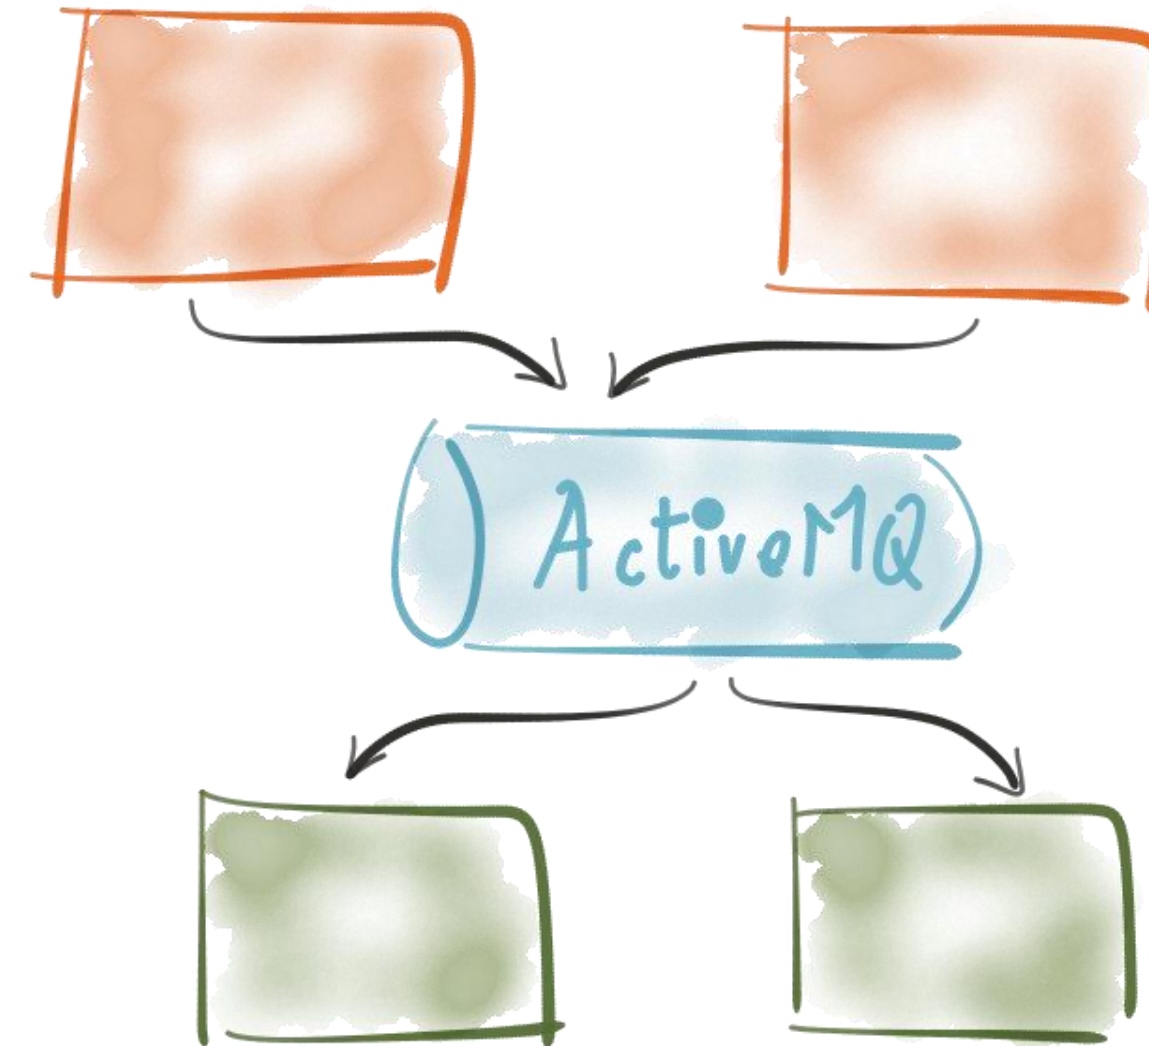
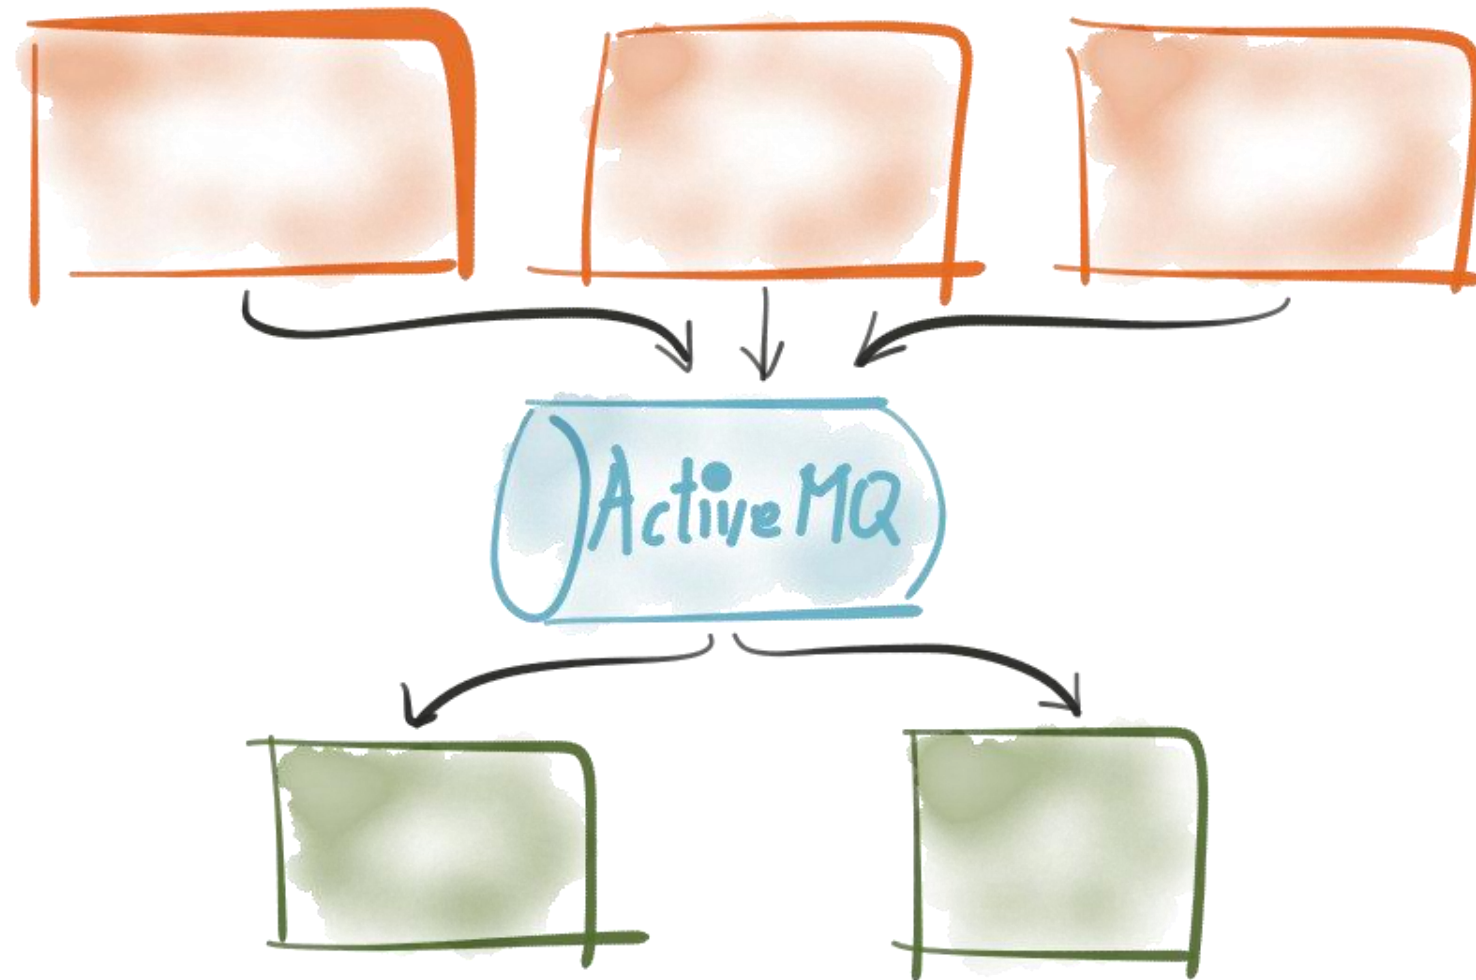
# MESSAGING





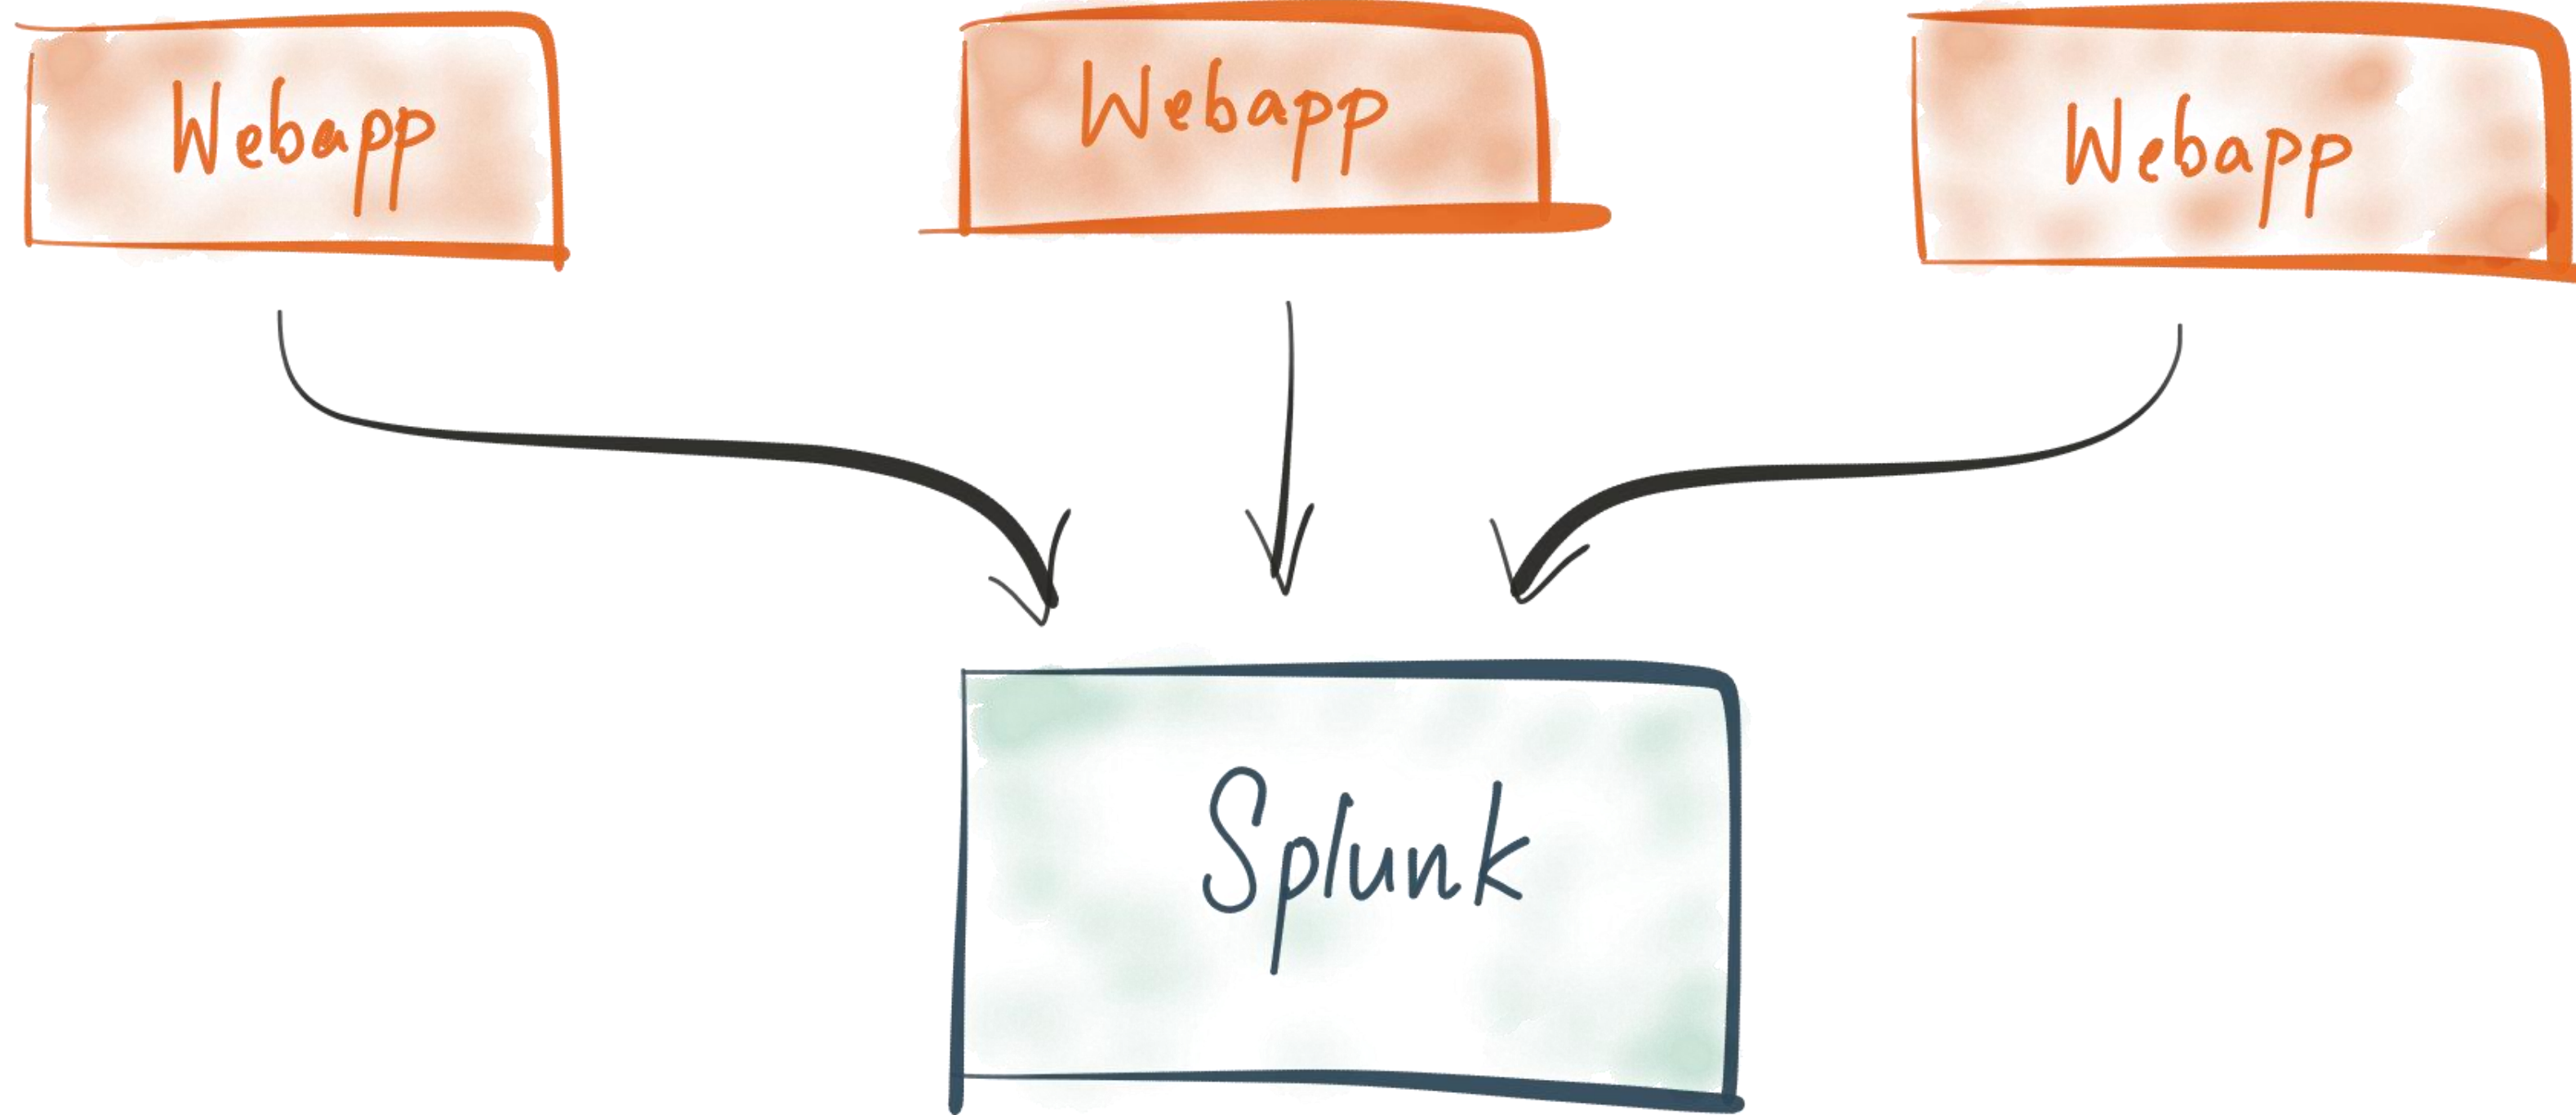


# MESSAGING



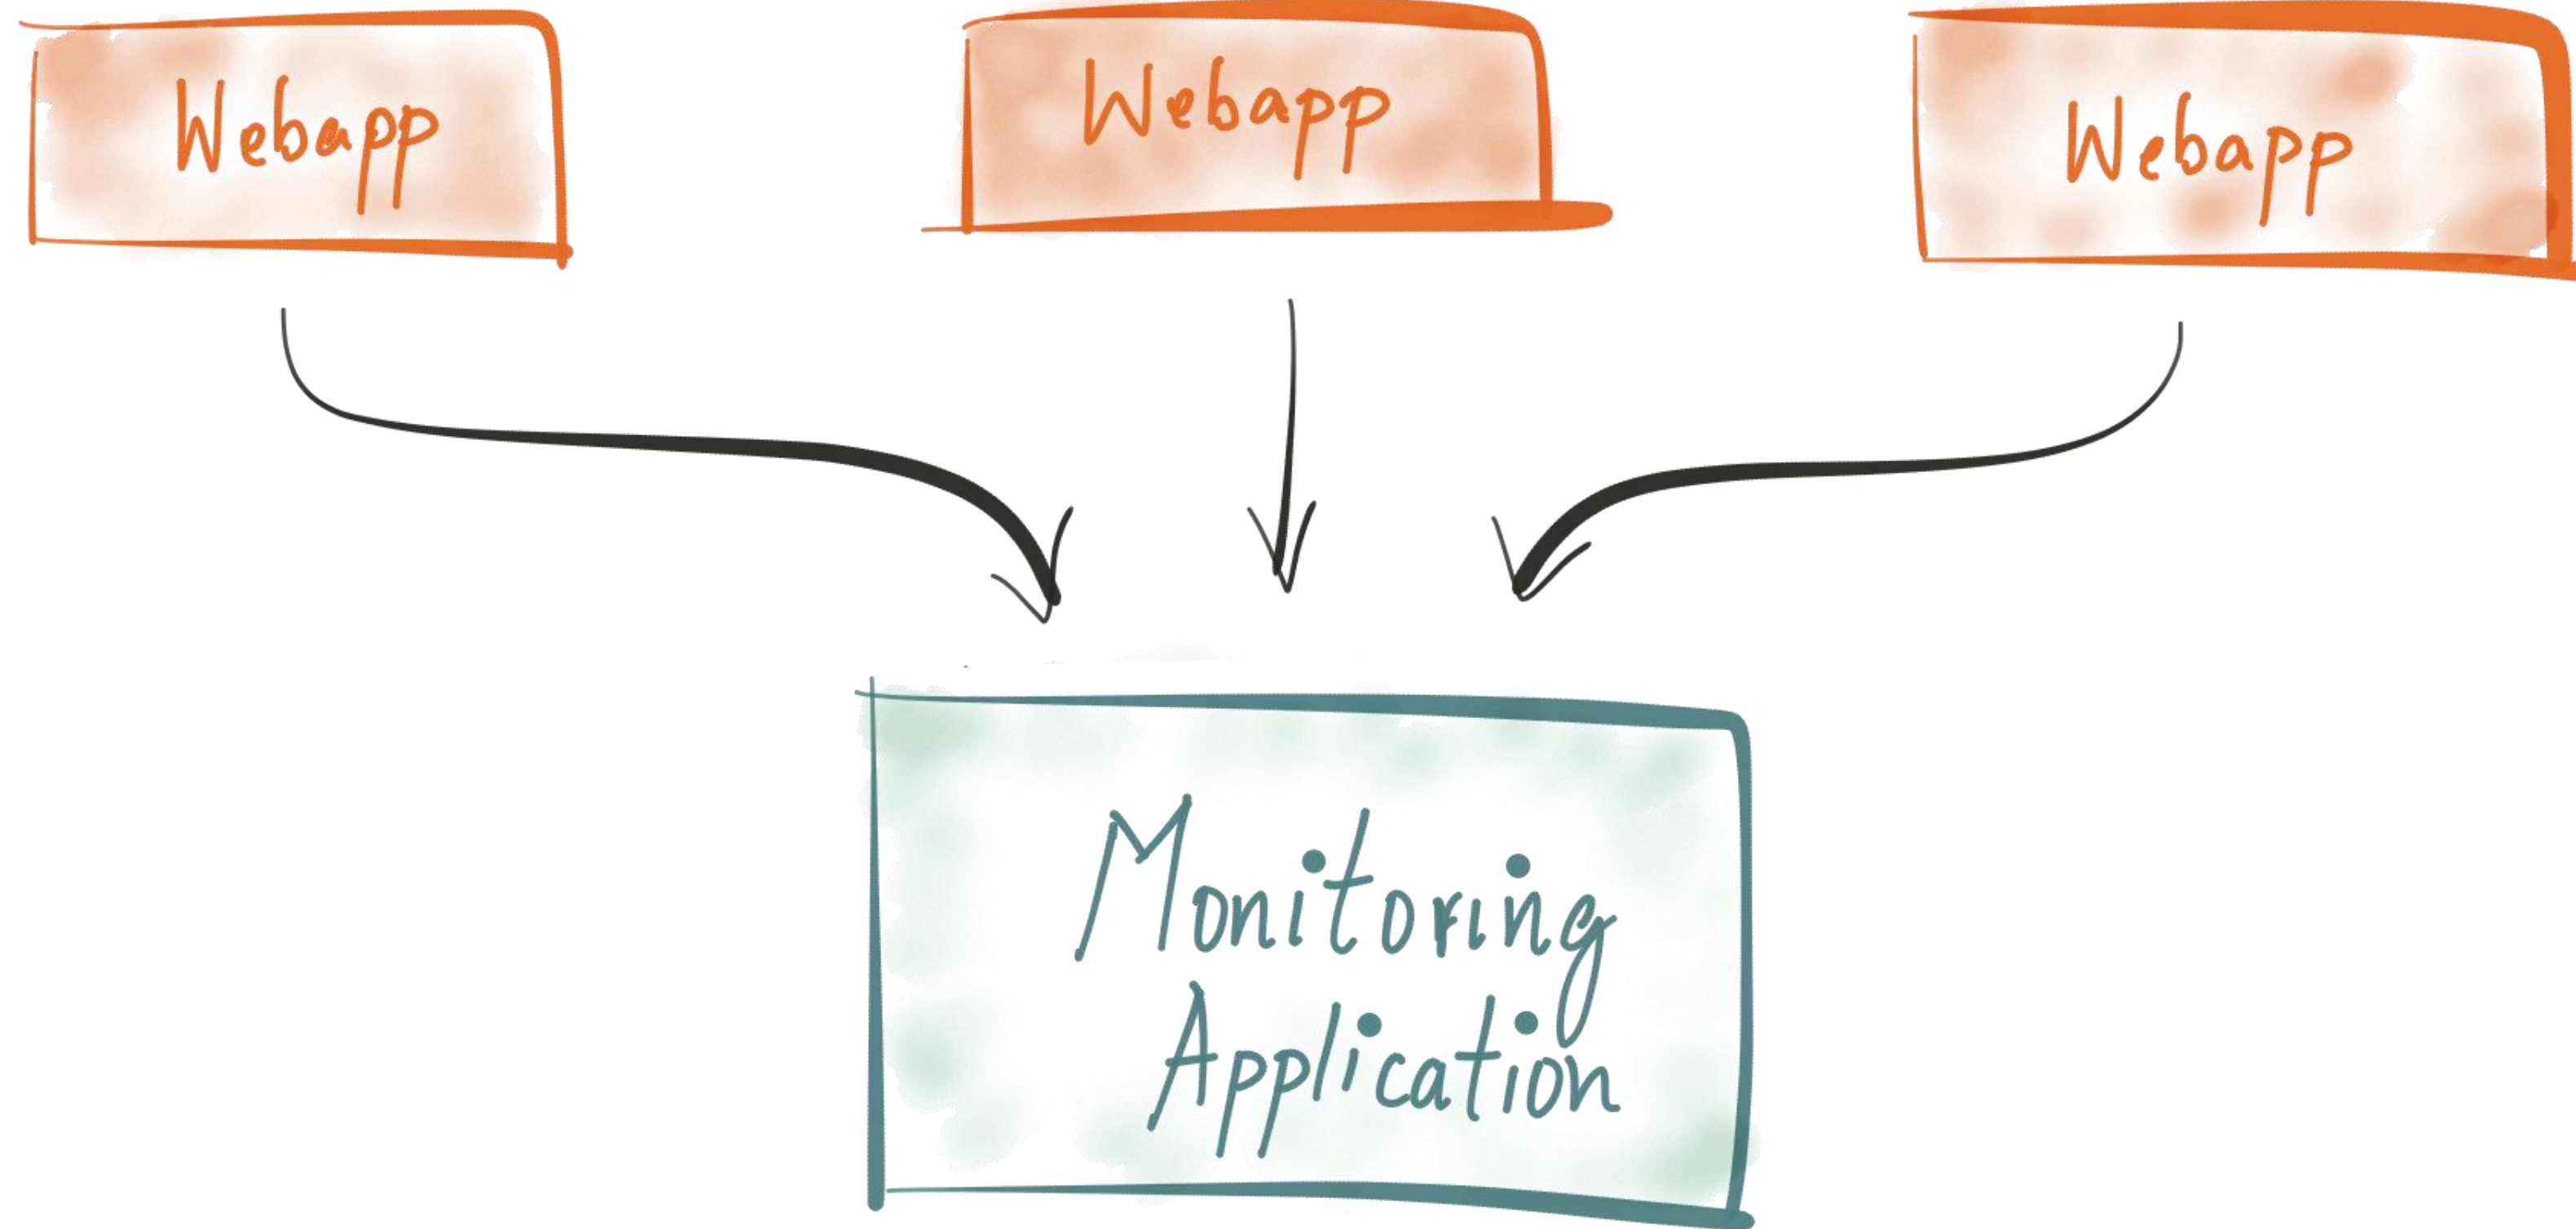


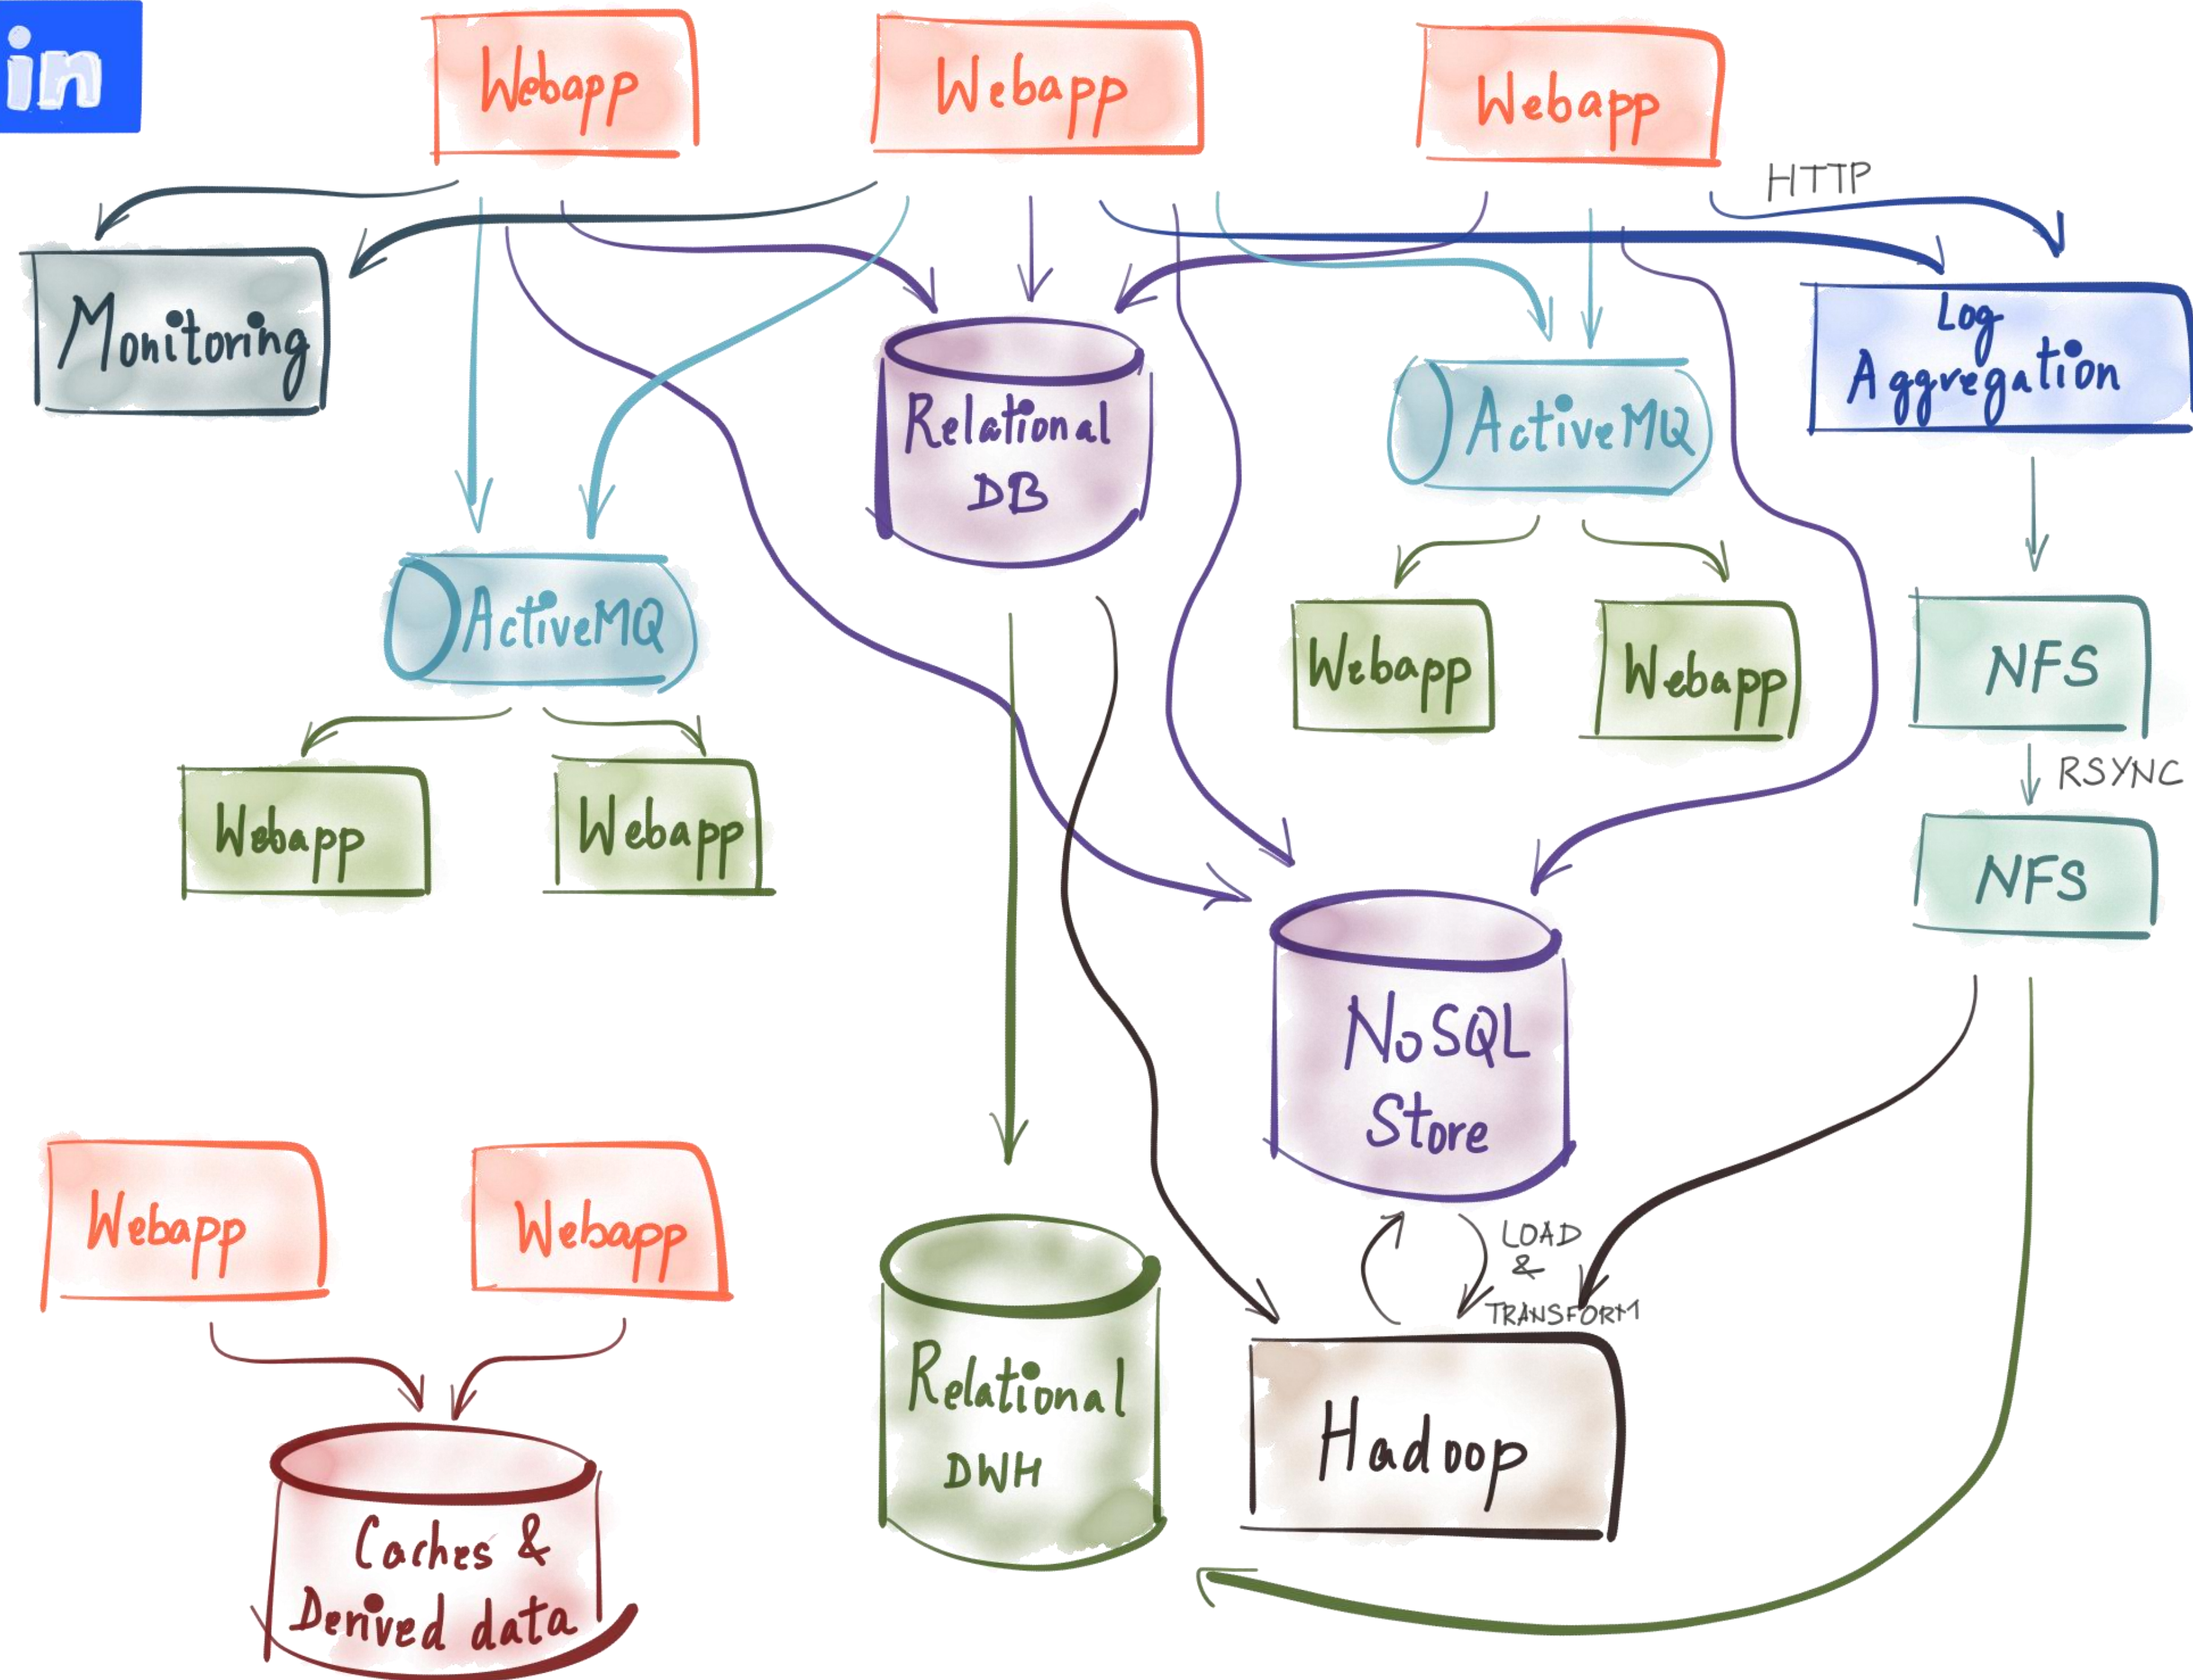
# APPLICATION LOGS





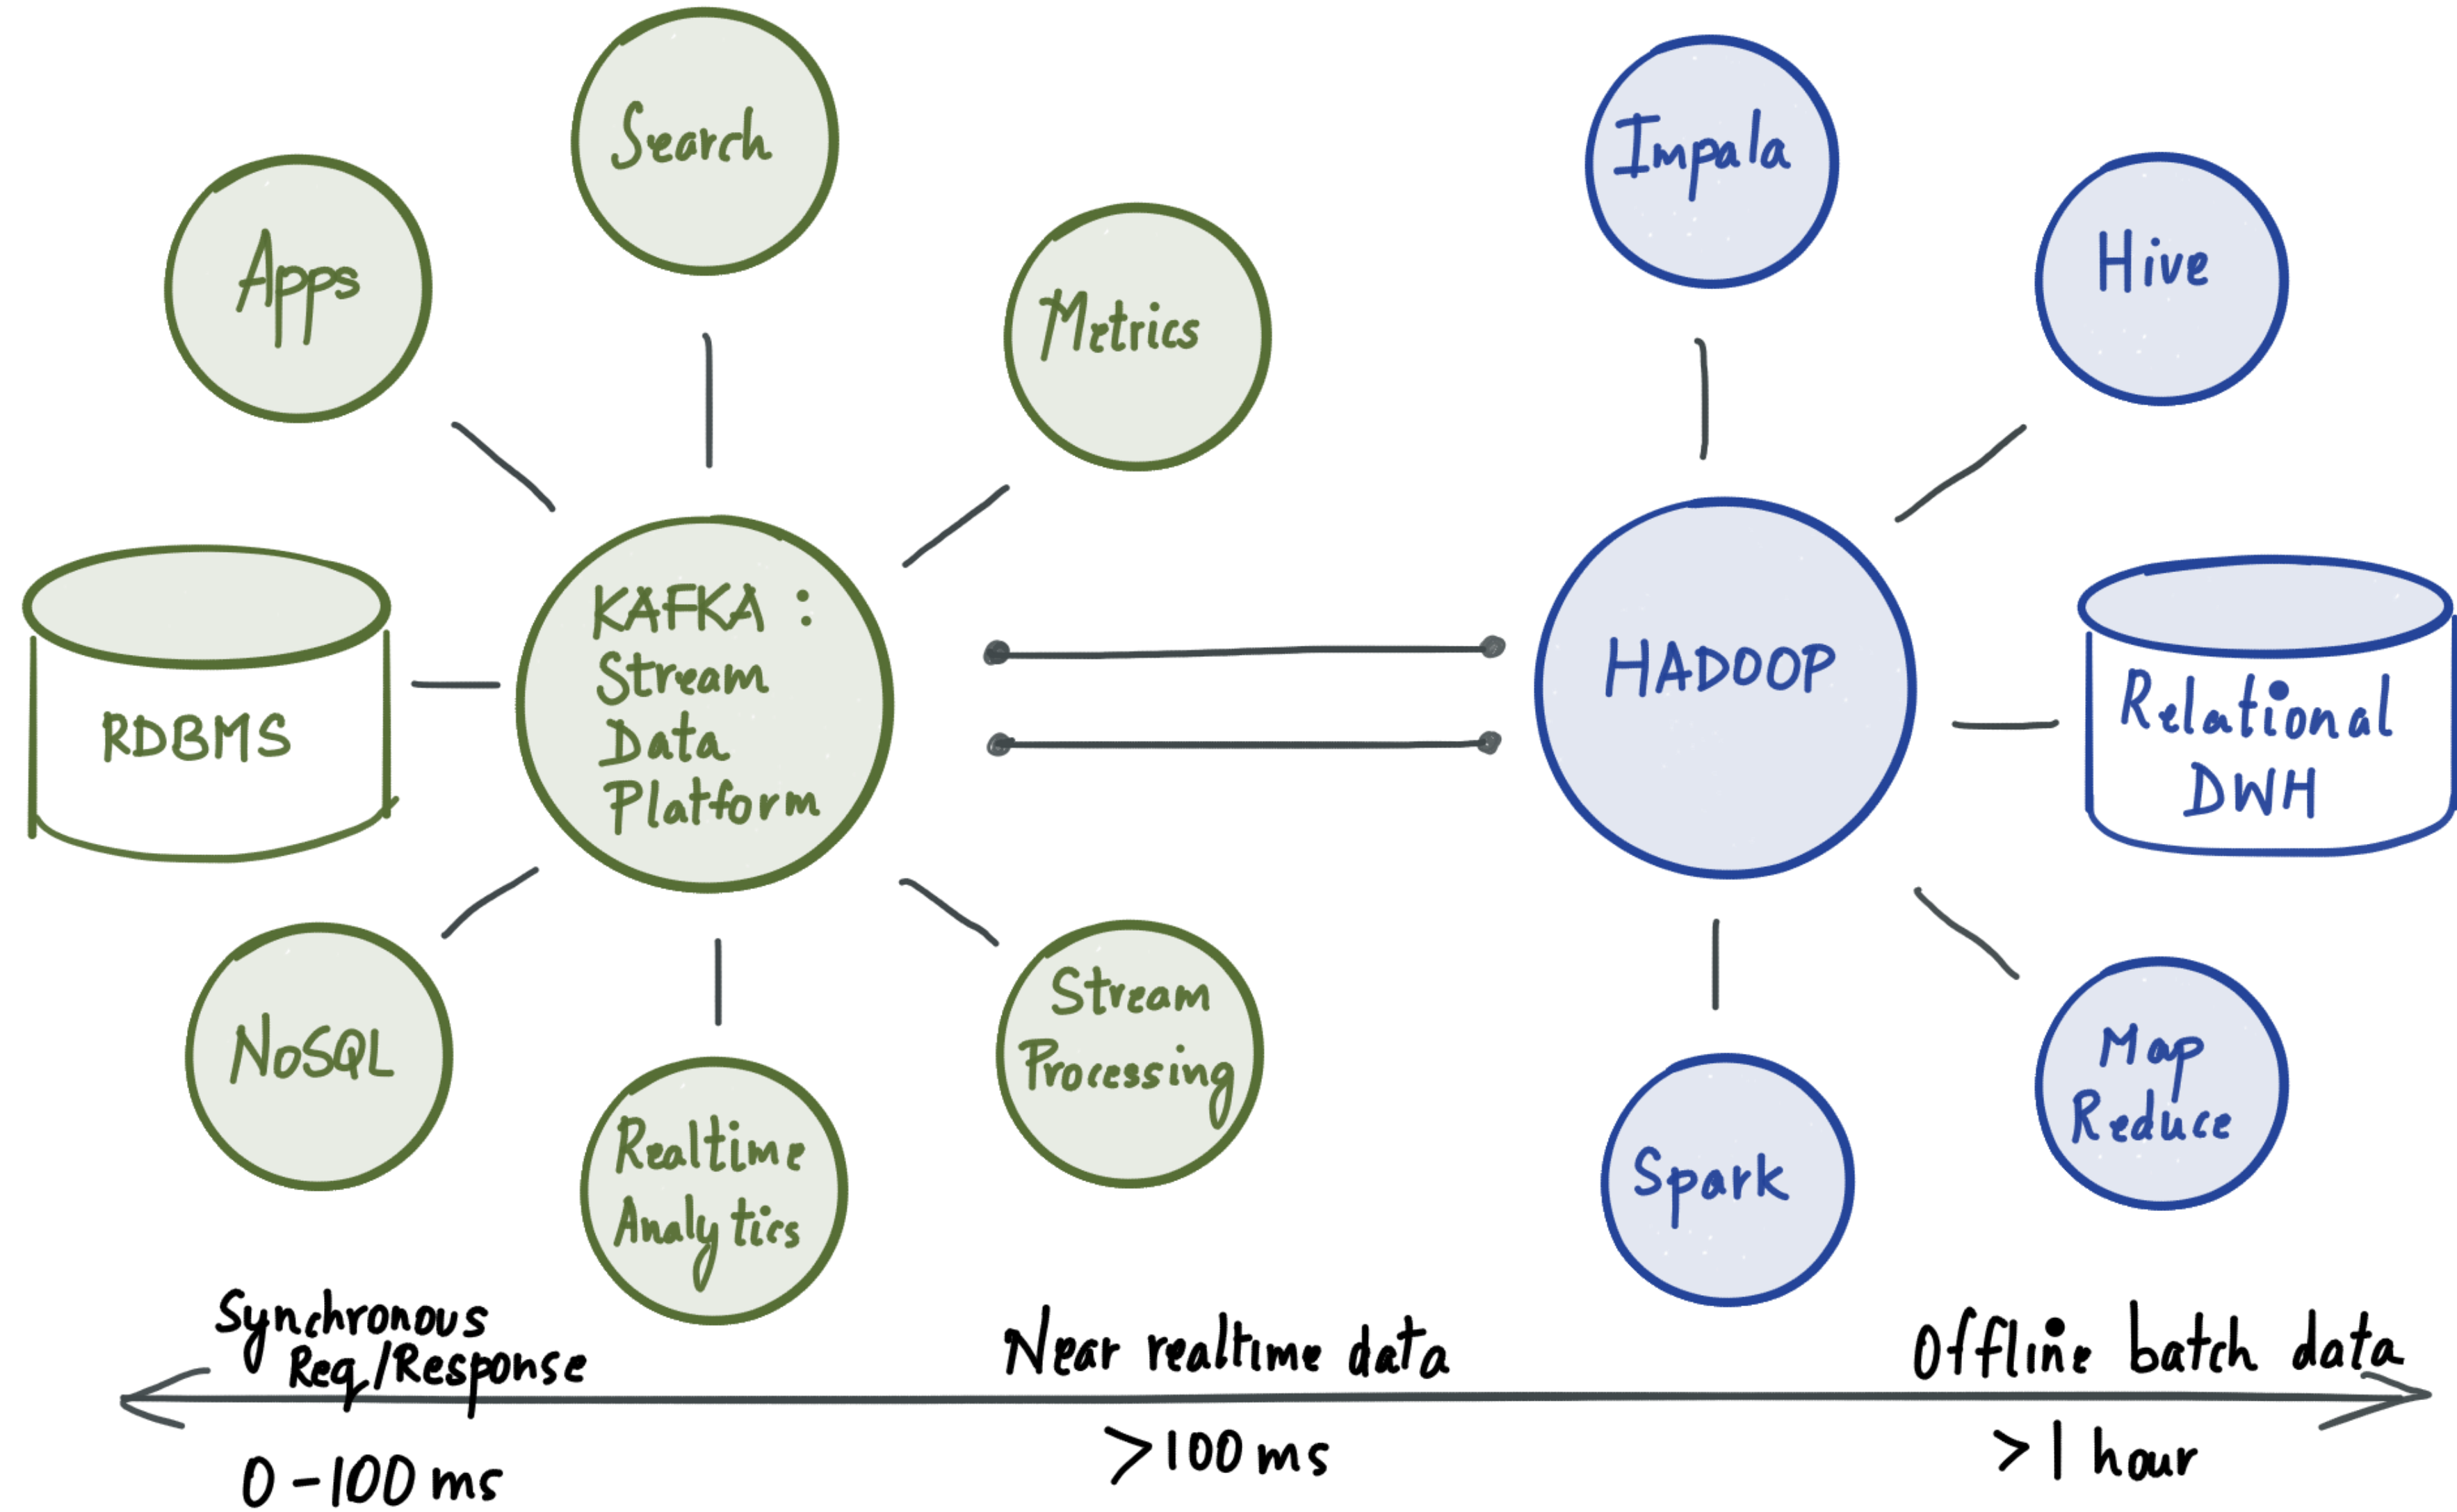
# APPLICATION METRICS







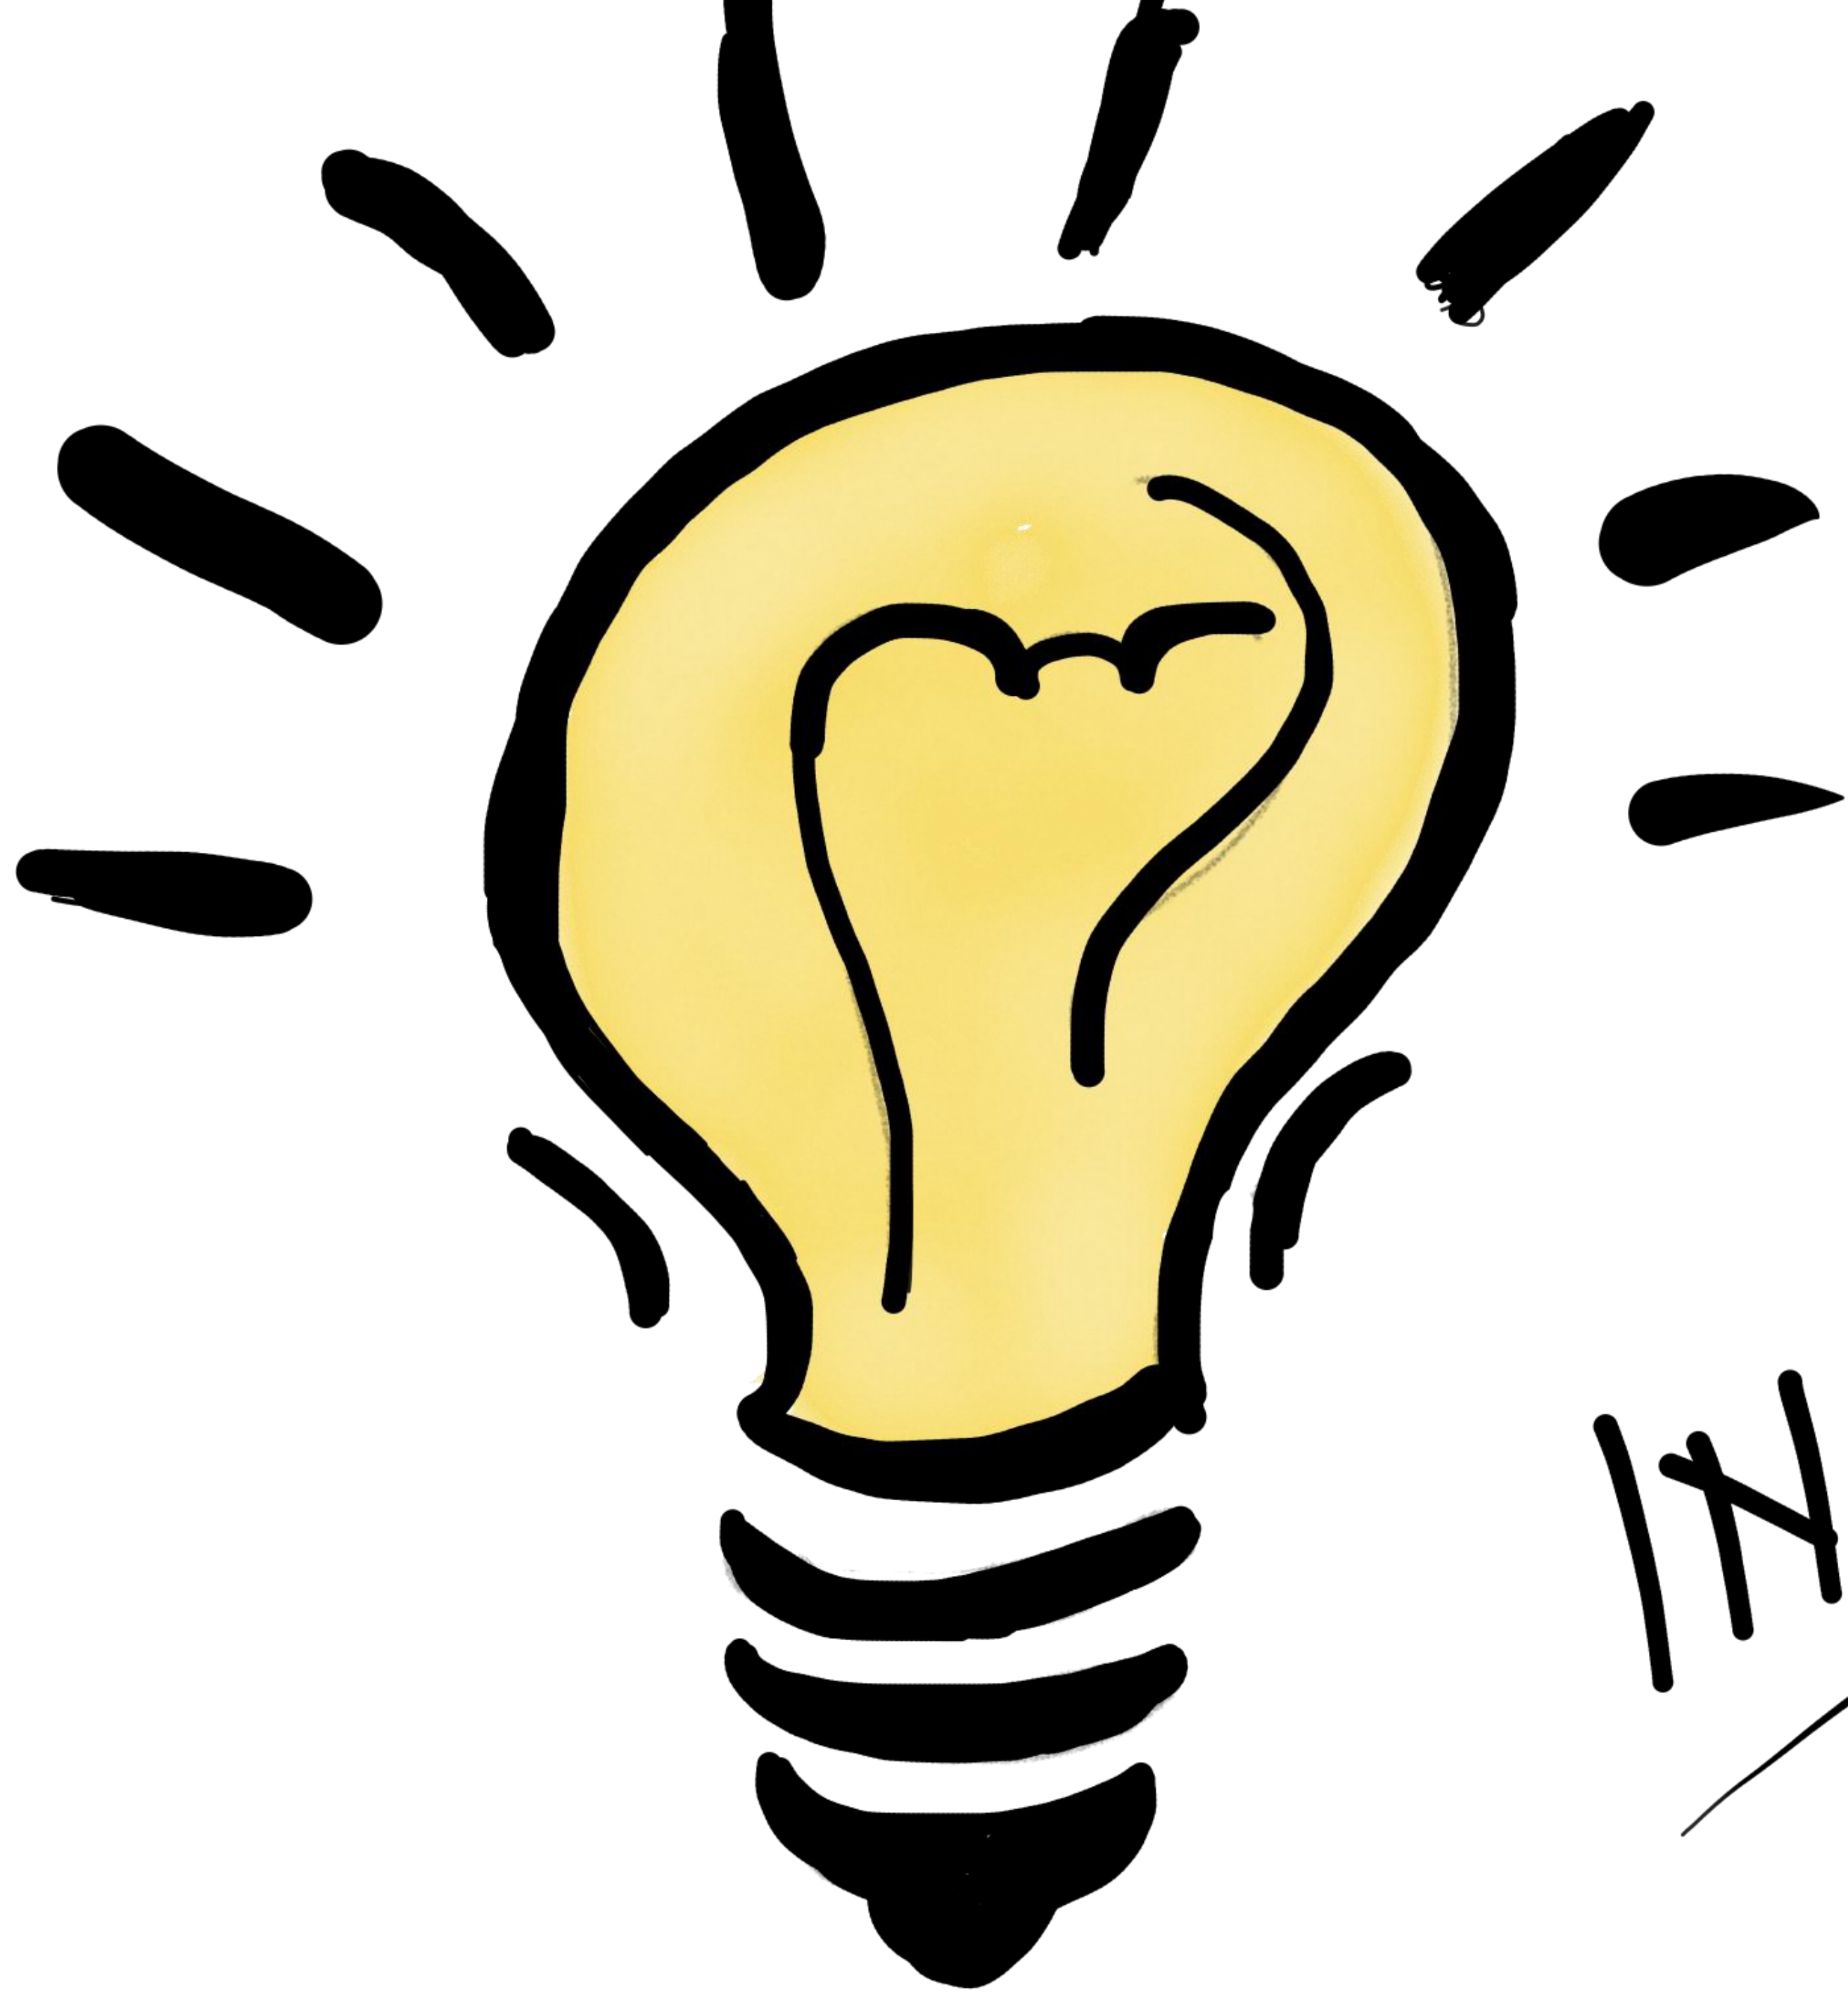
# STREAM DATA PLATFORM



# DATA INTEGRATION

## KEY CHALLENGE

“Making sure the data ends up in all the right places”



INFRA



FIRST ATTEMPT: MESSAGE QS

**ActiveMQ**

 **RabbitMQ**

# MESSAGE QS : PROBLEMS

THROUGHPUT

PERSISTENCE

BATCH  
SYSTEMS

SHARDING

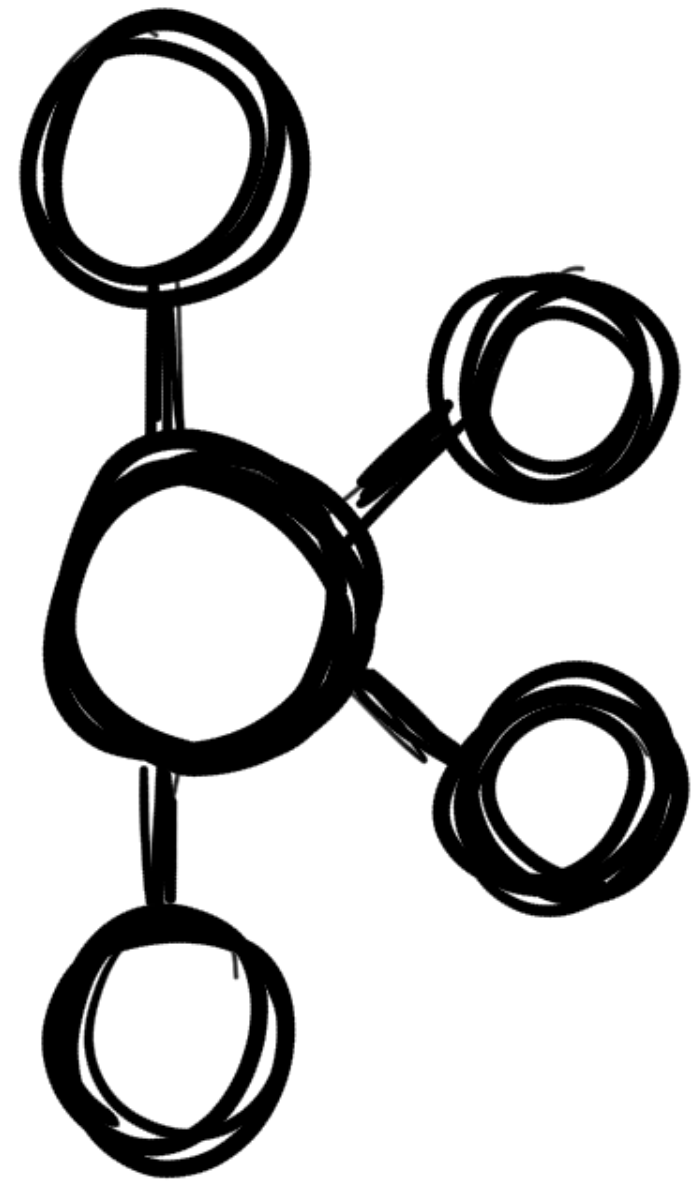
STREAM  
PROCESSING

ORDERING  
GUARANTEES

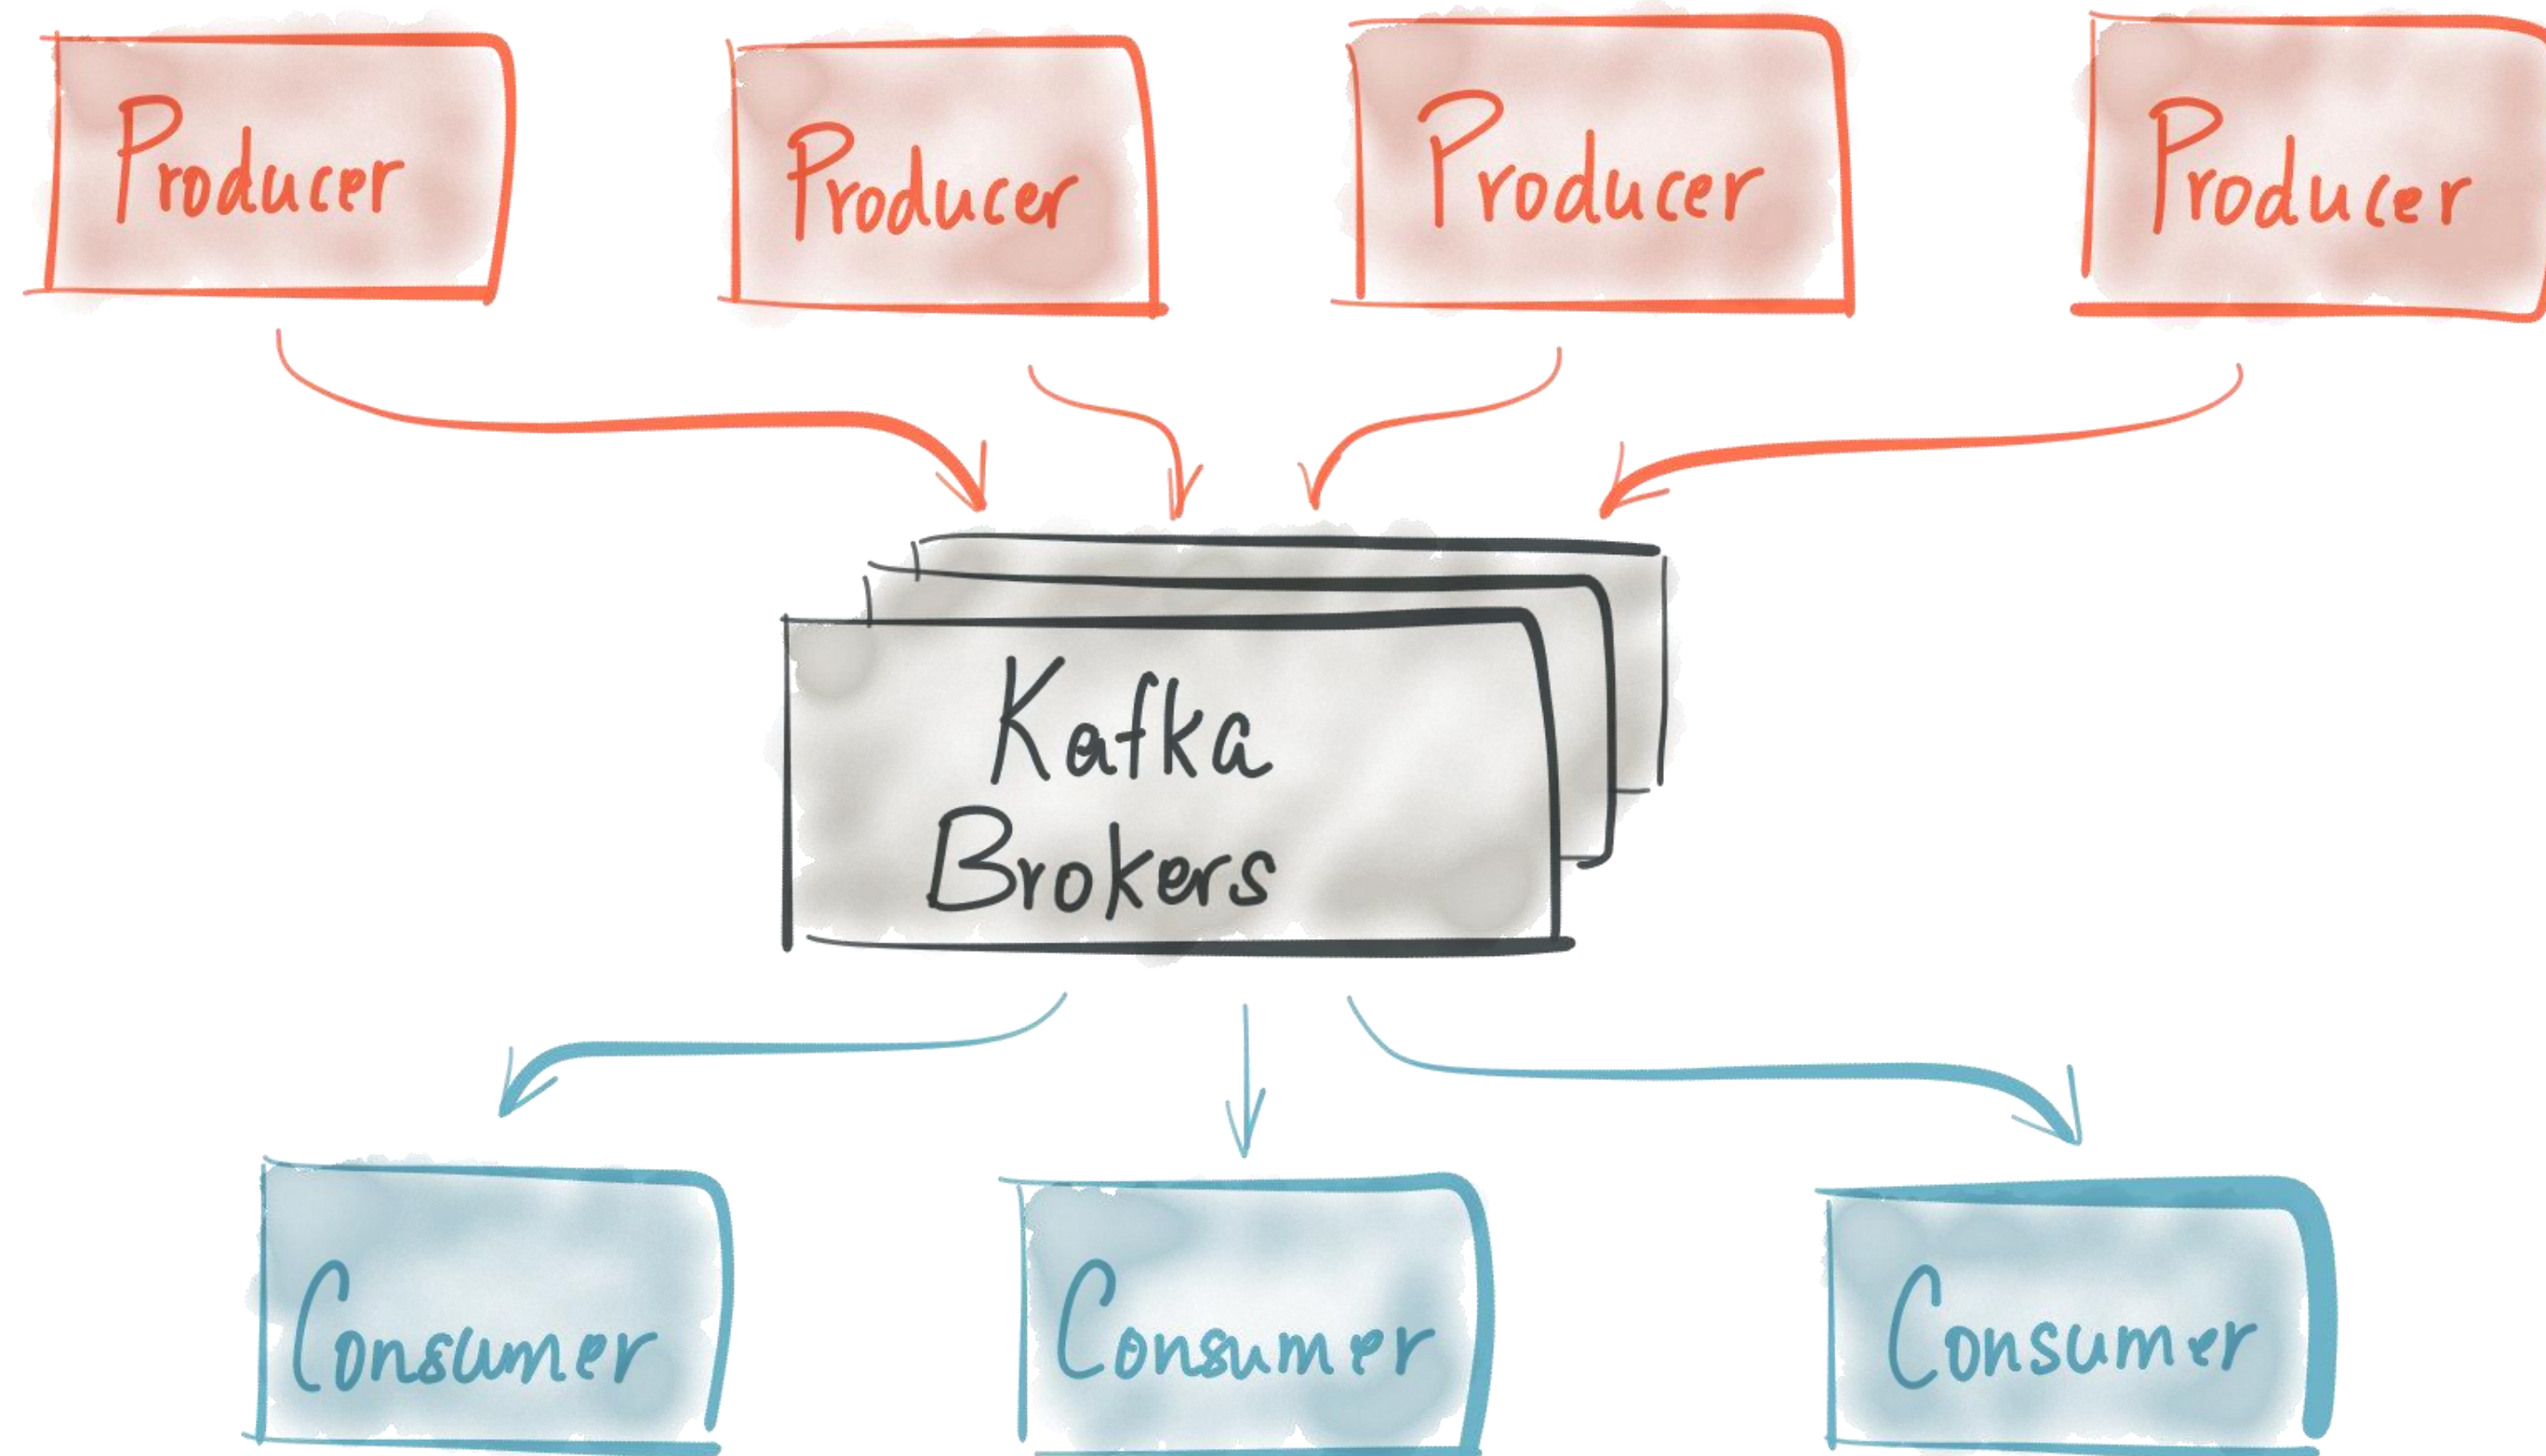
SUCCESSFUL

ATTEMPT:

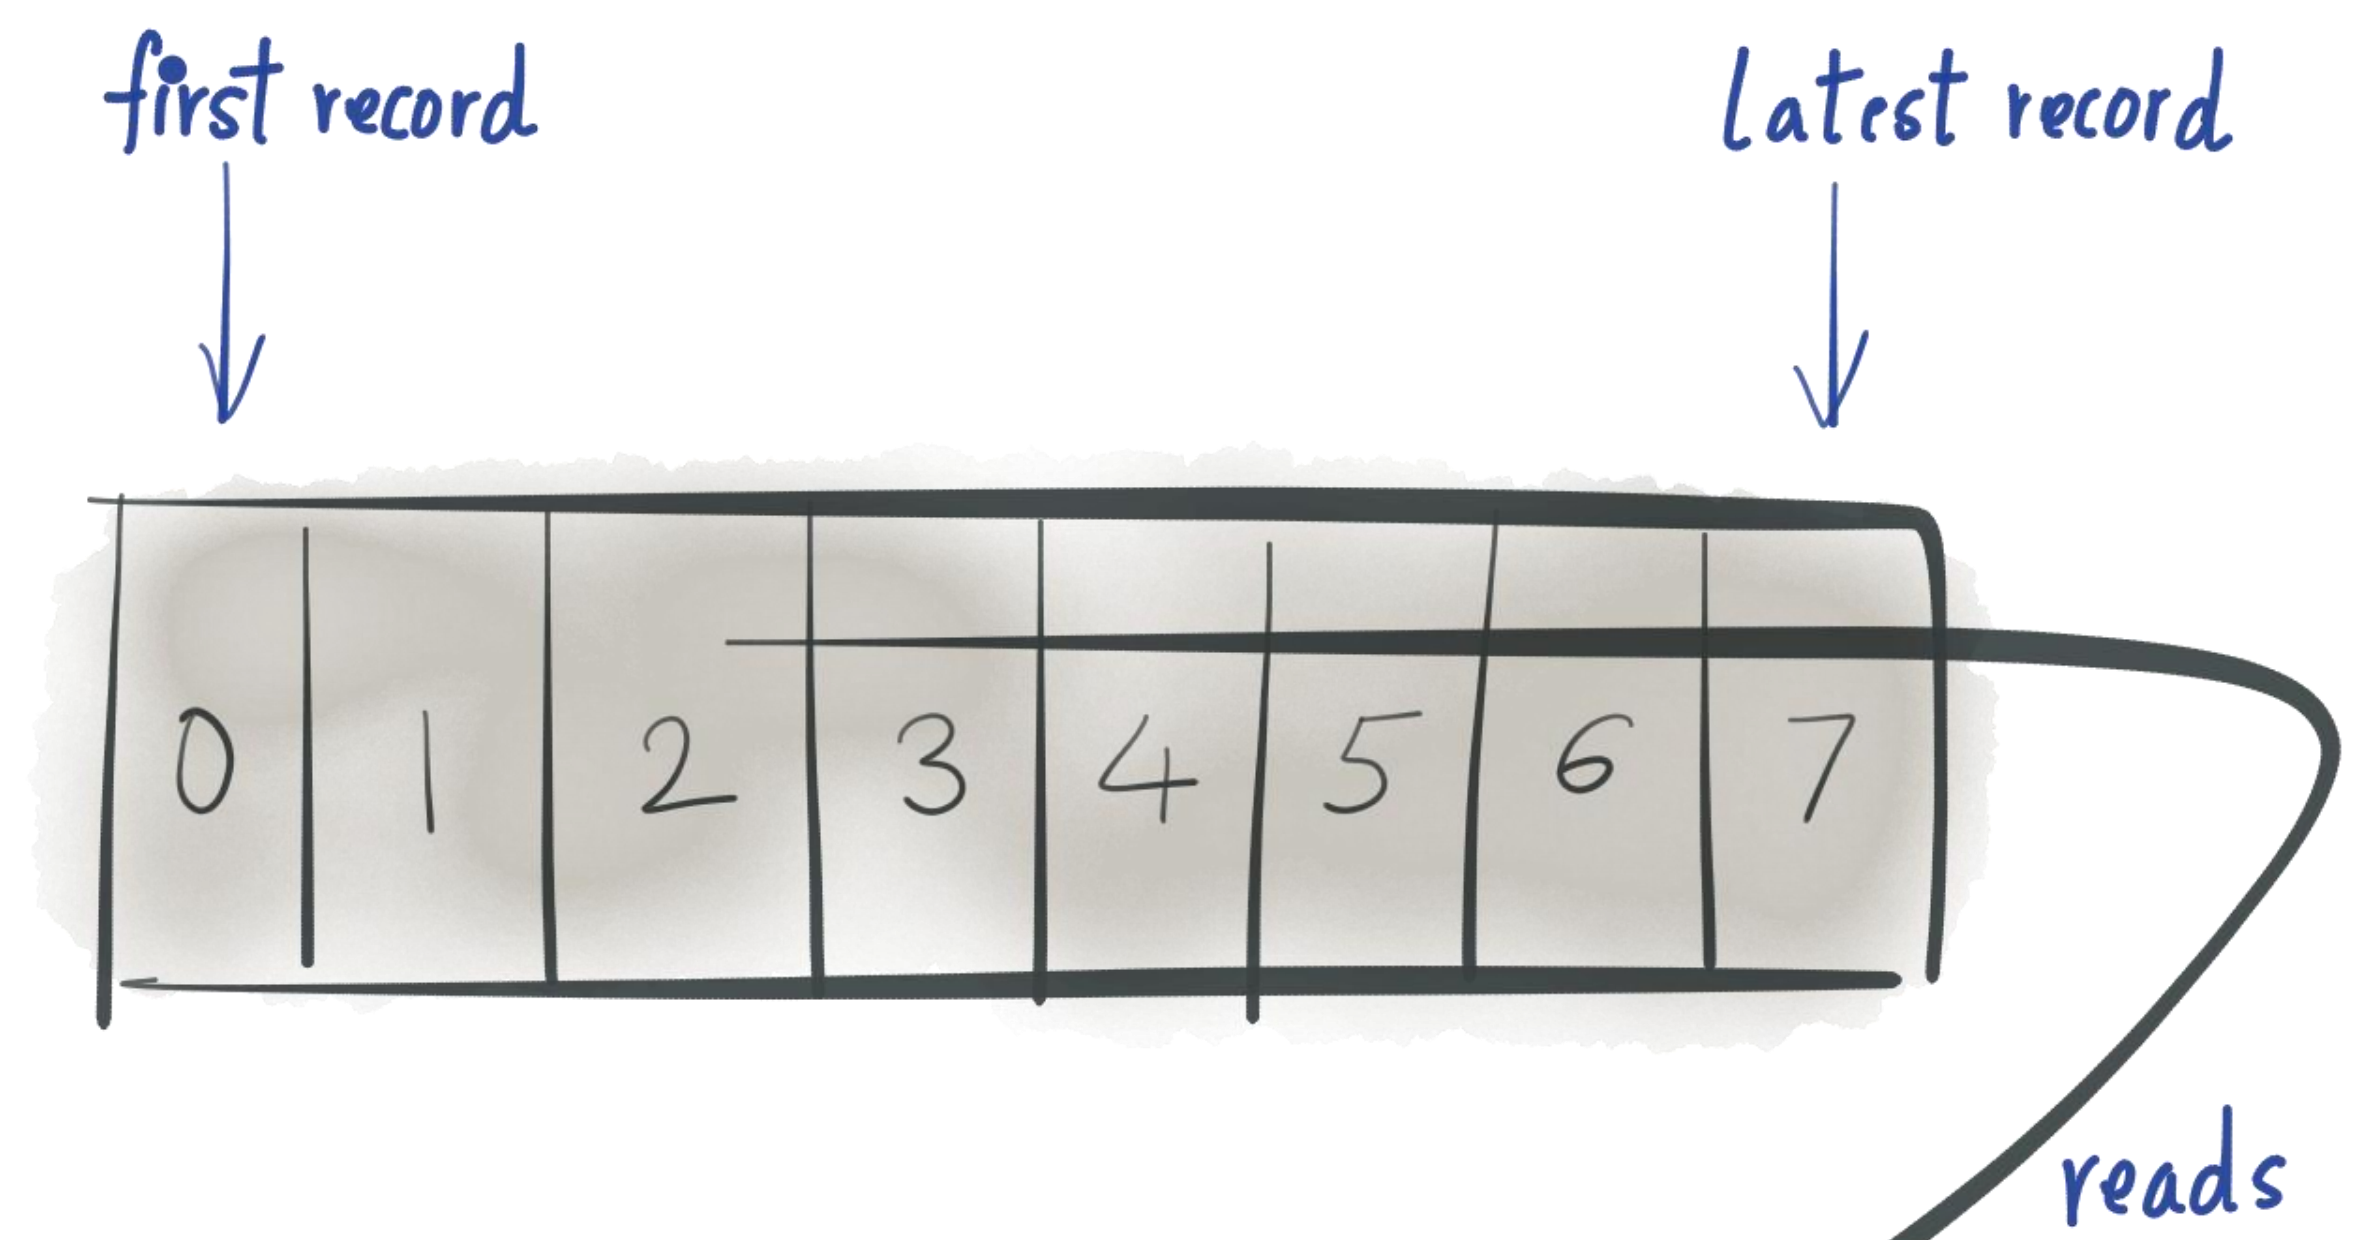
KAFKA



# APACHE KAFKA : 10,000 FT VIEW

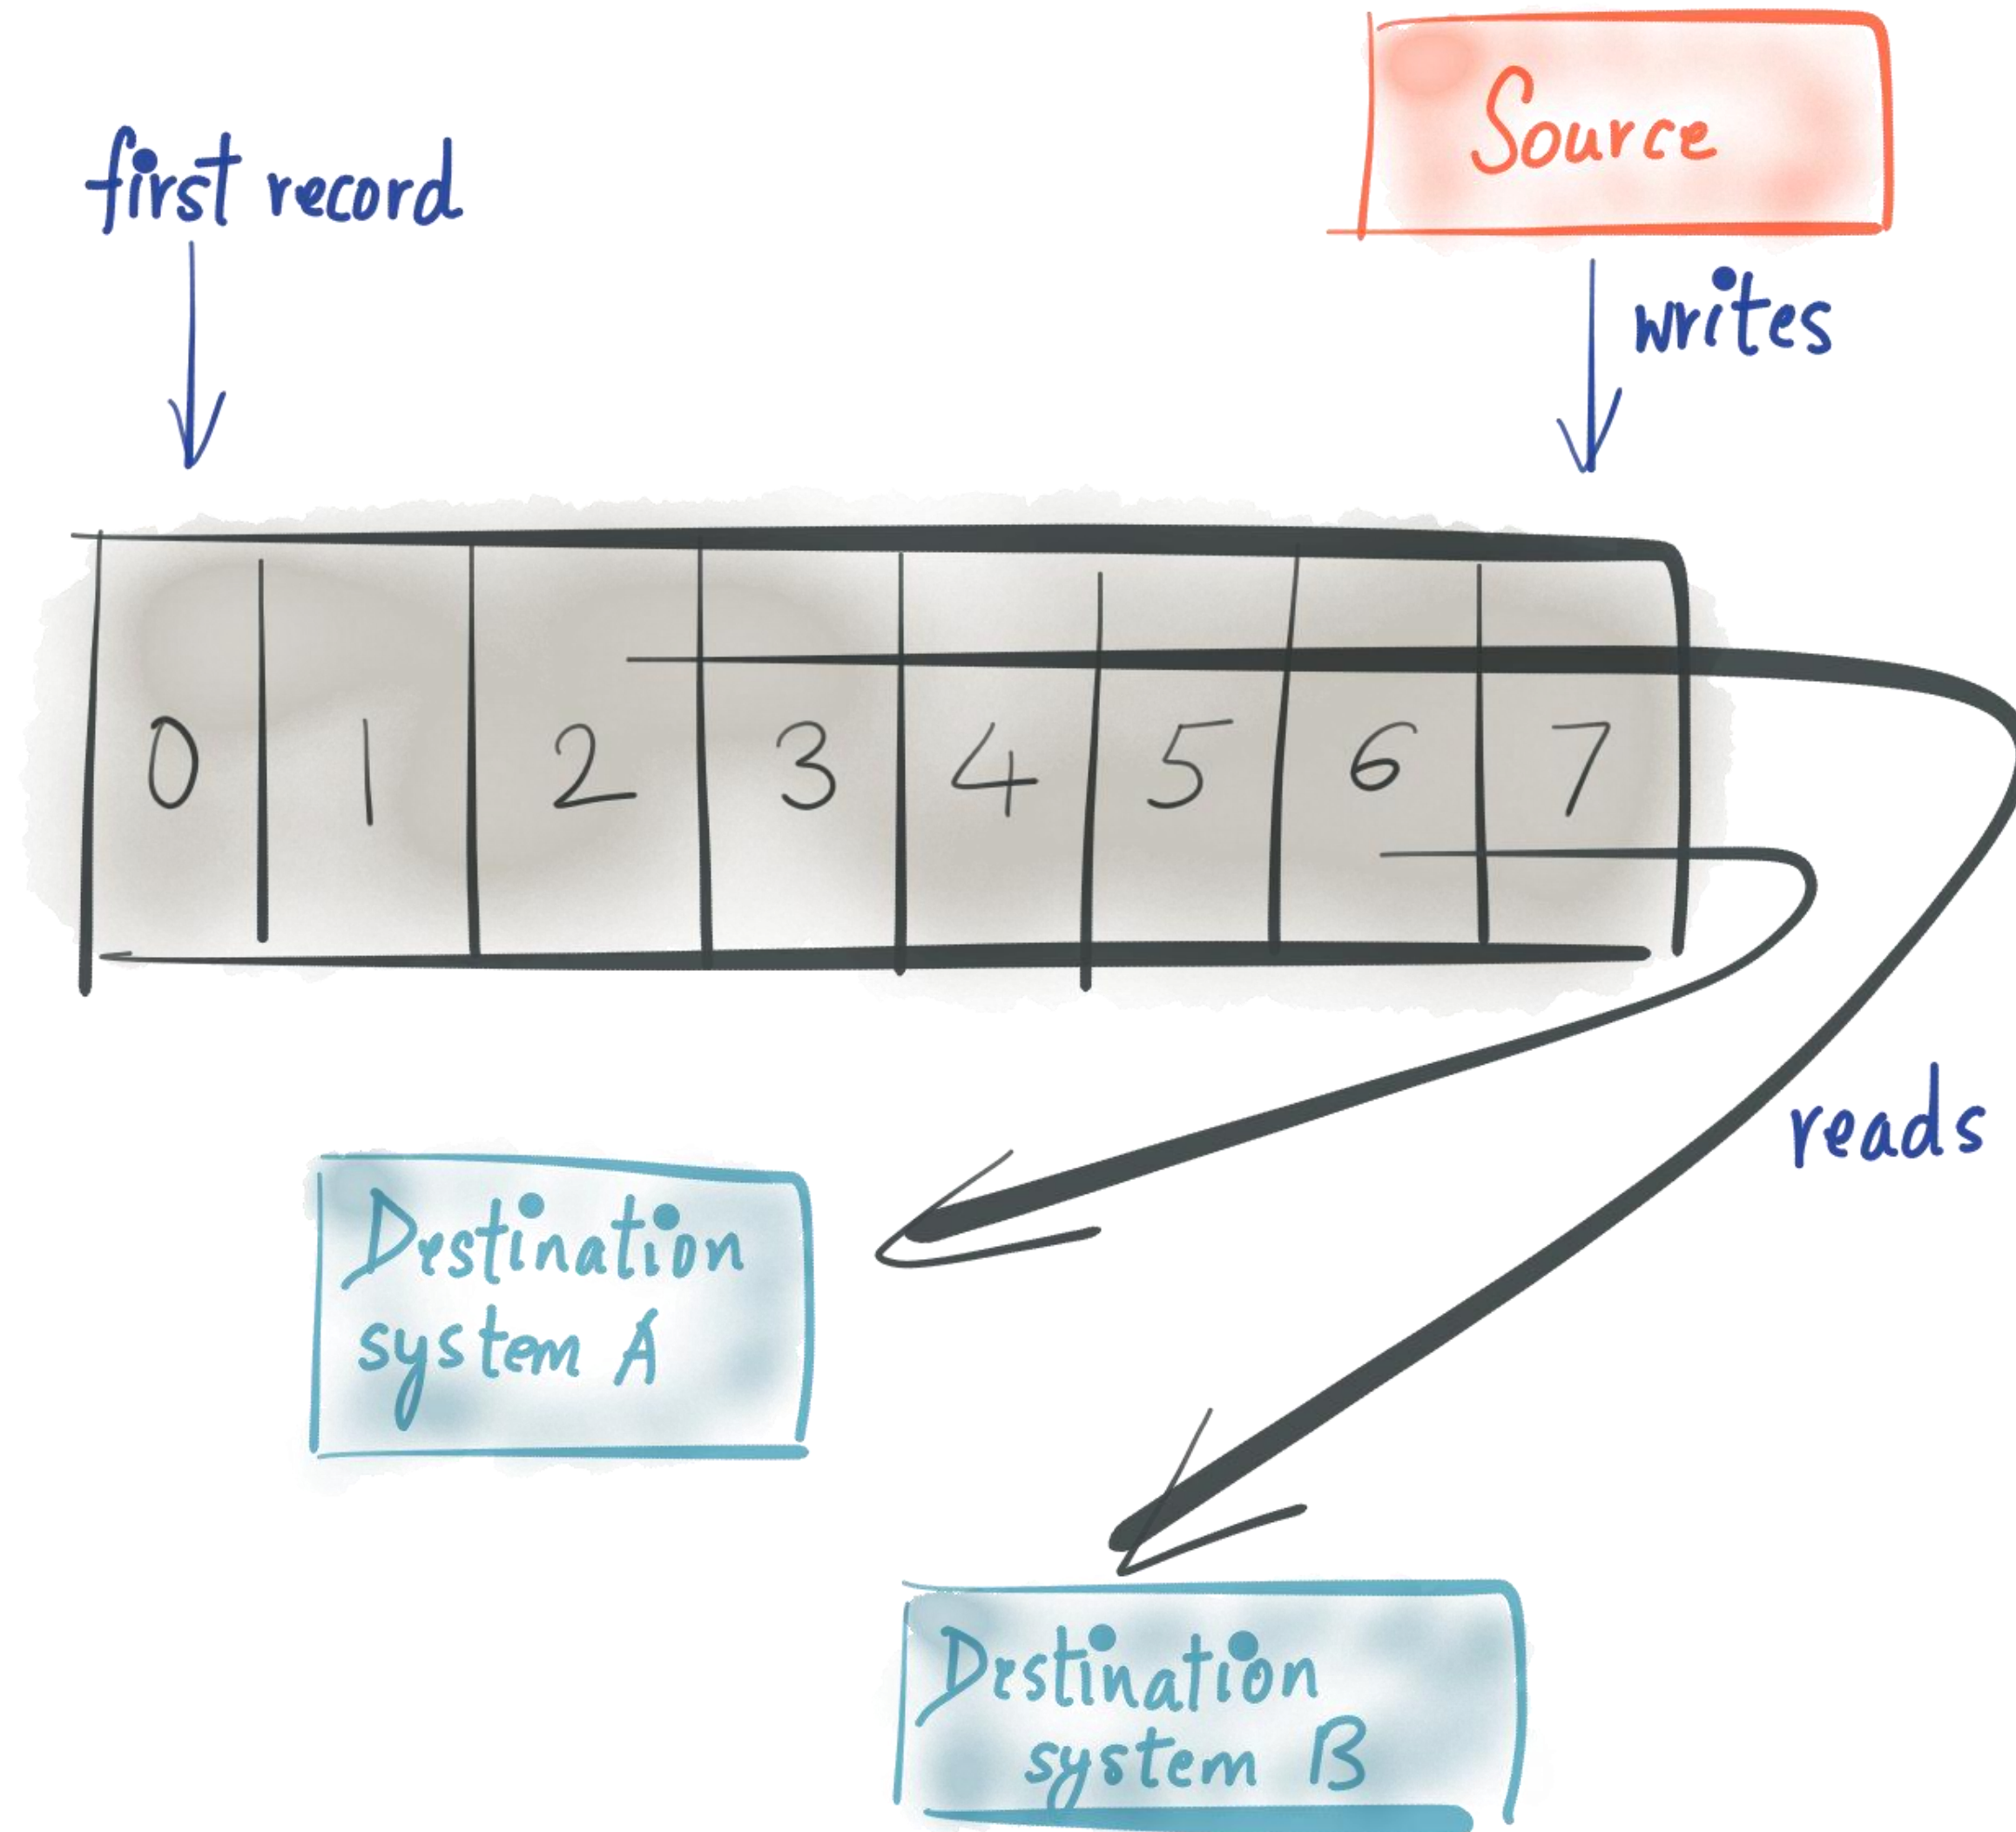


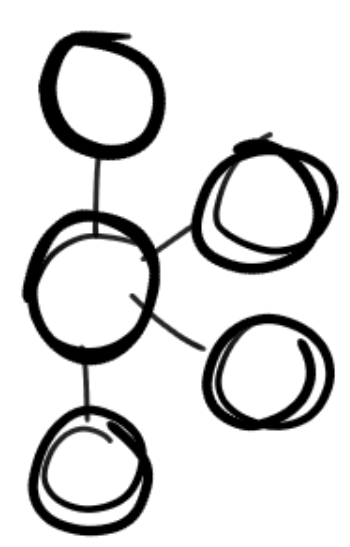
# KAFKA KEY IDEA : LOG



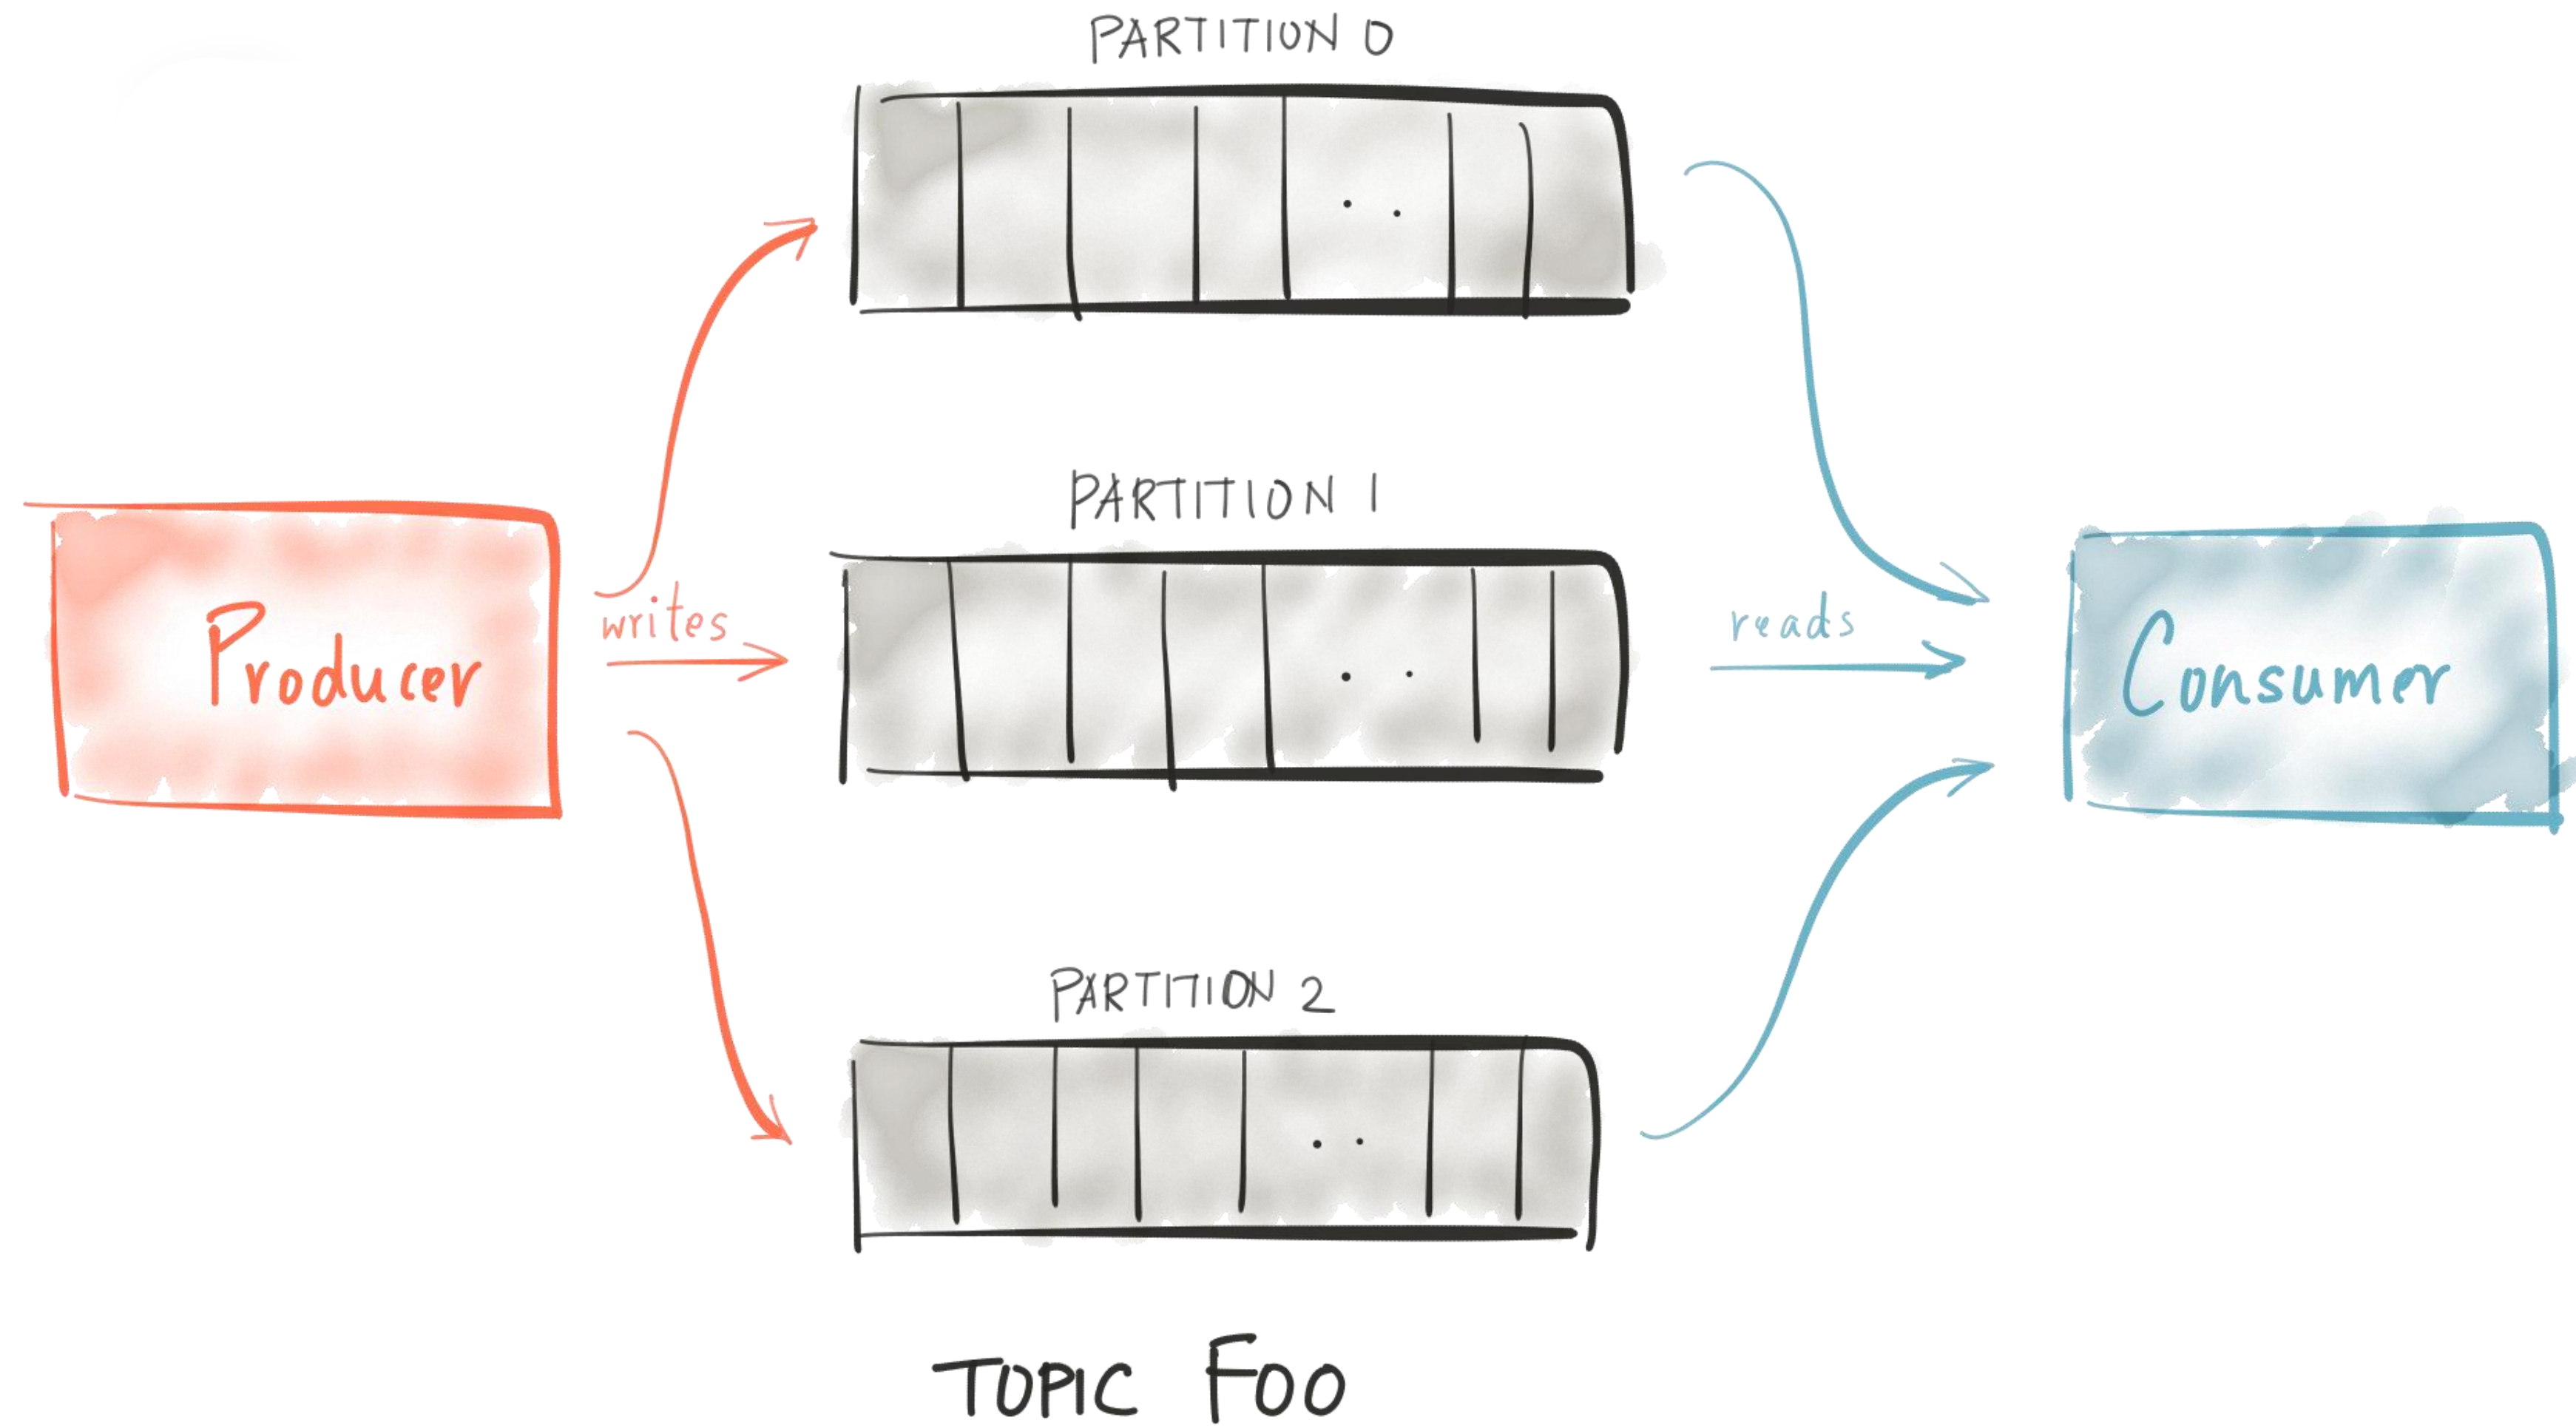
Sequential access  
=  
High performance

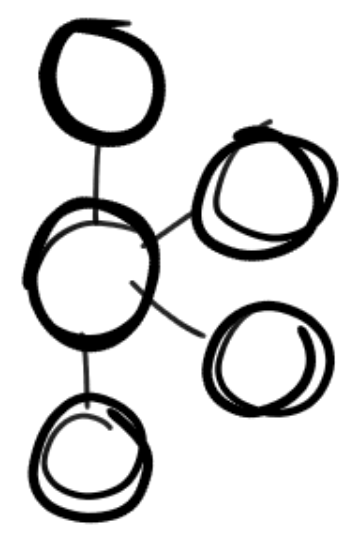
# LOGS & PUB-SUB



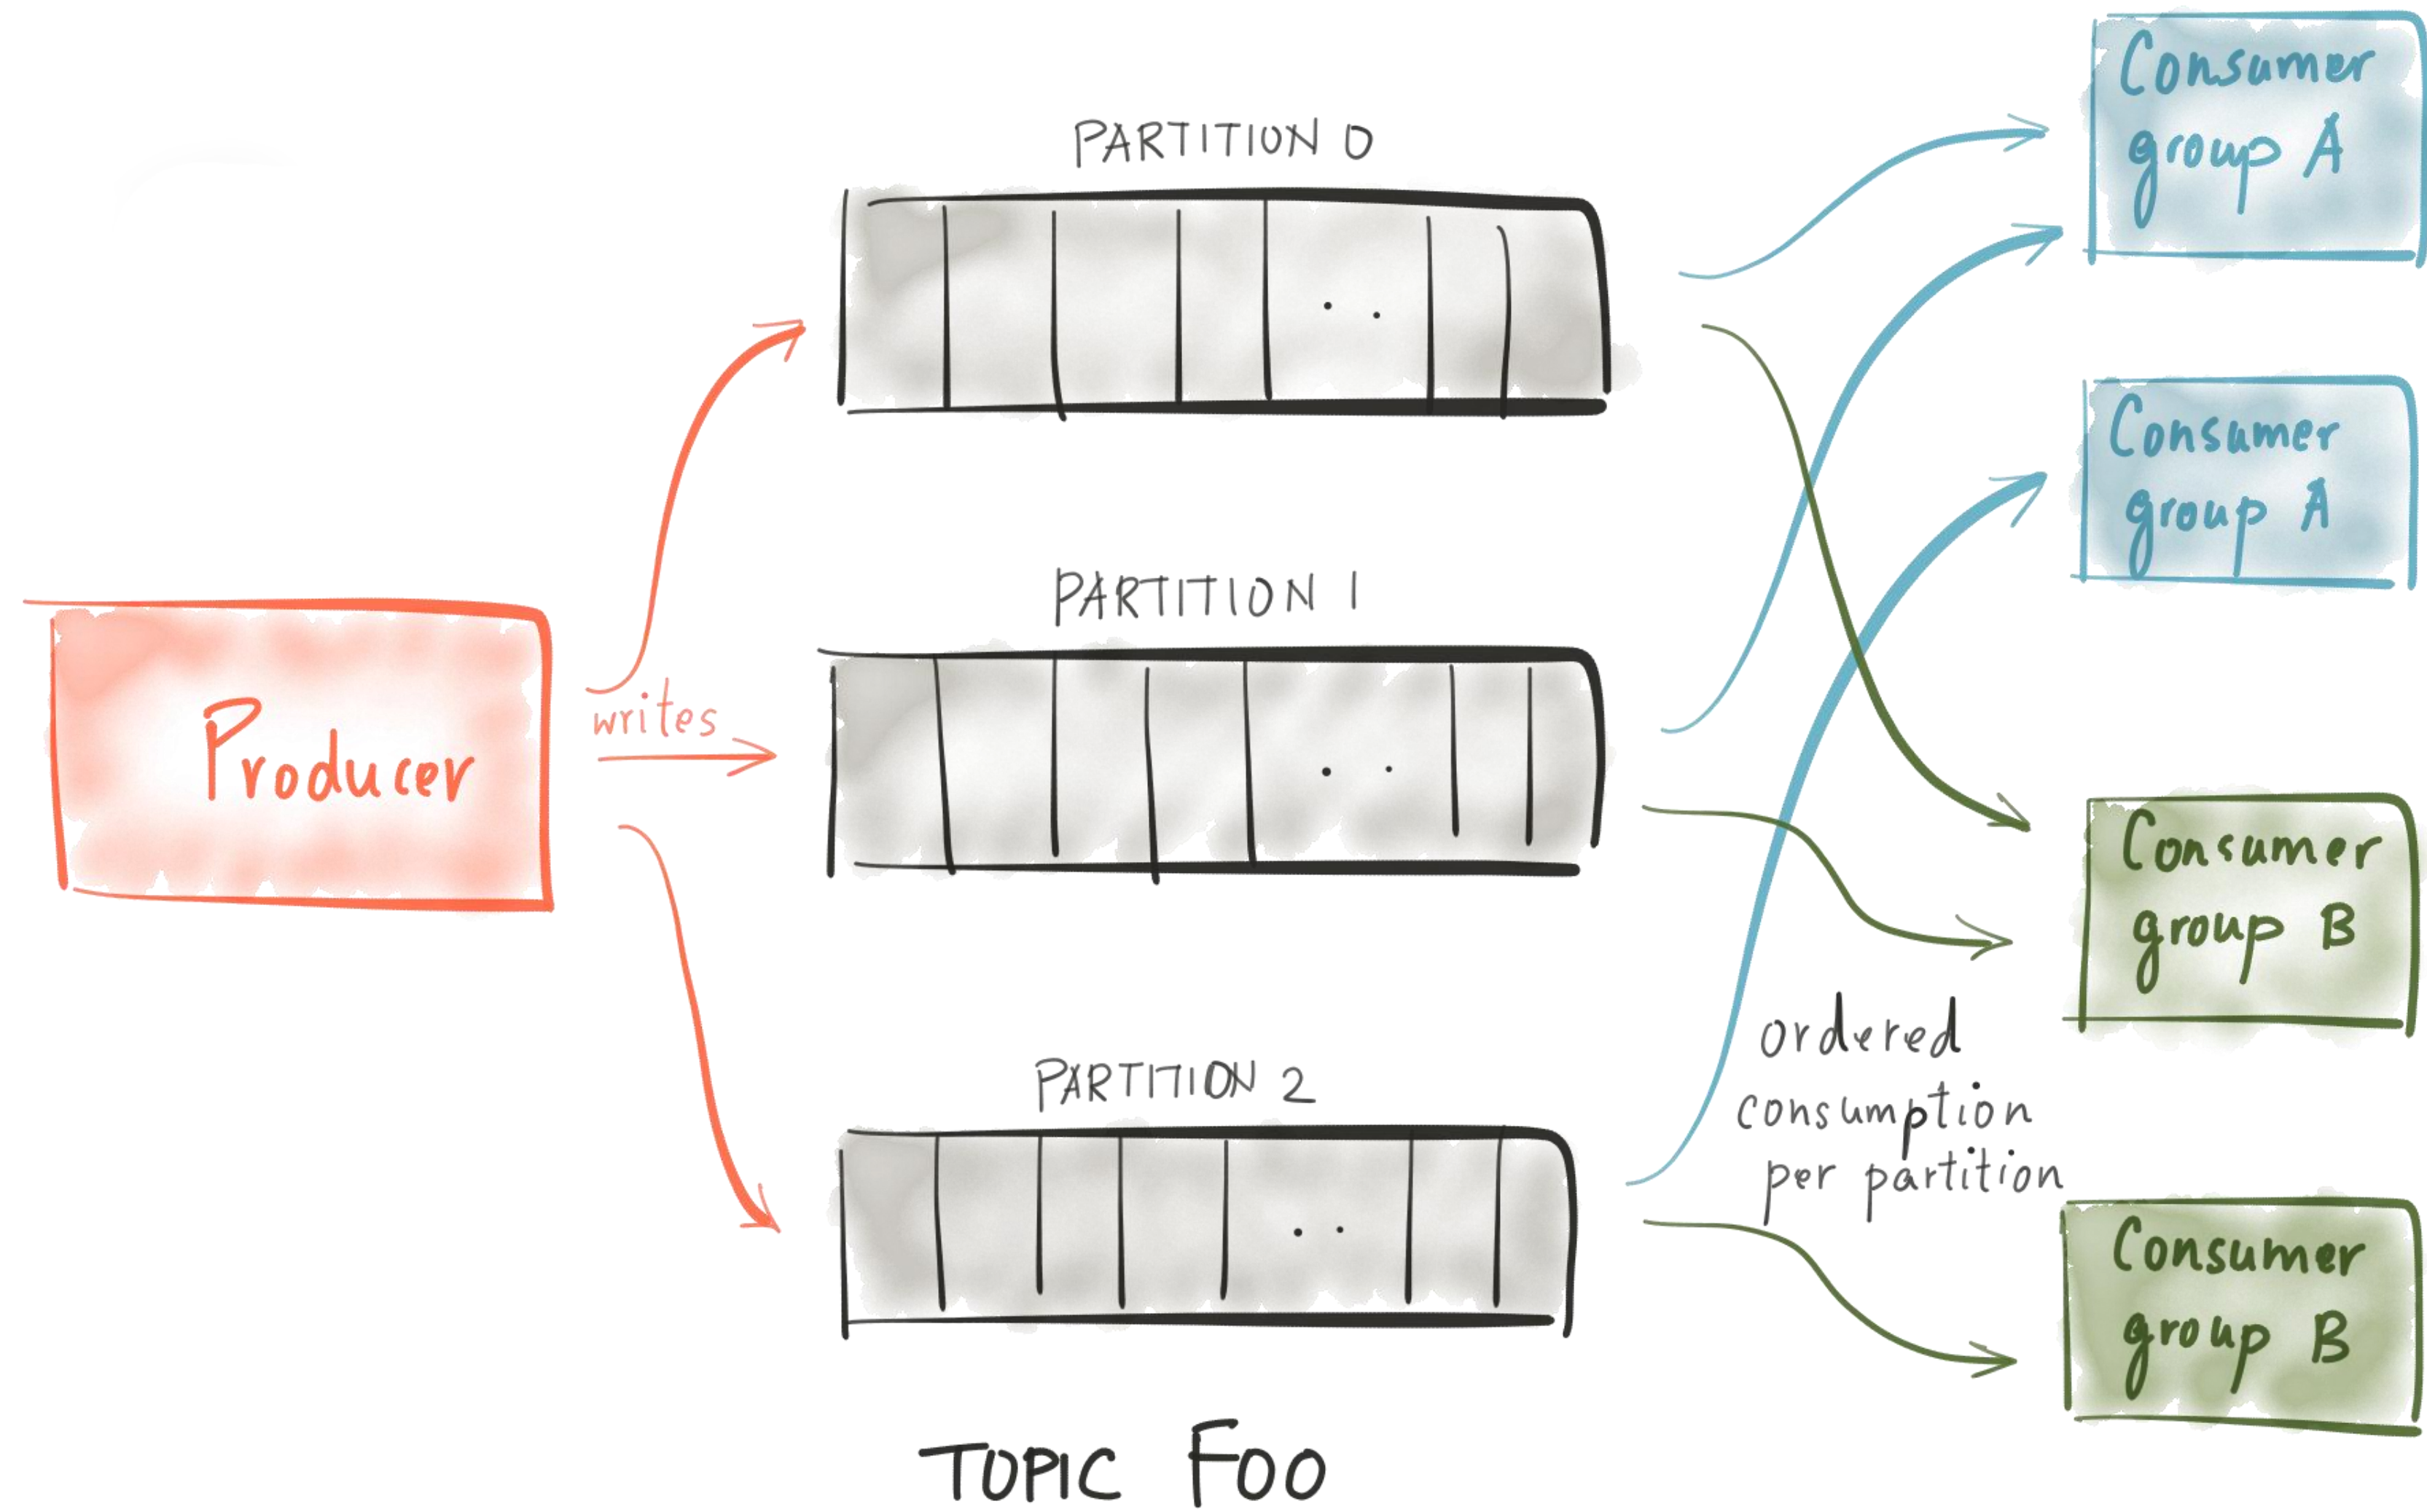


Topic  $\approx$  PARTITIONED LOG

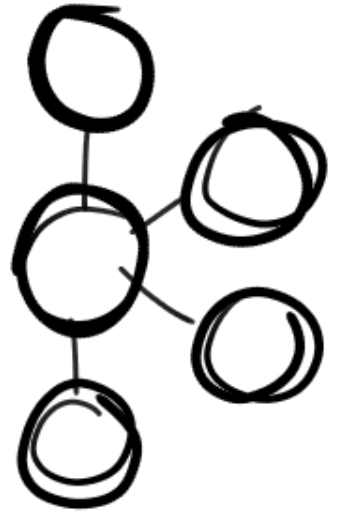




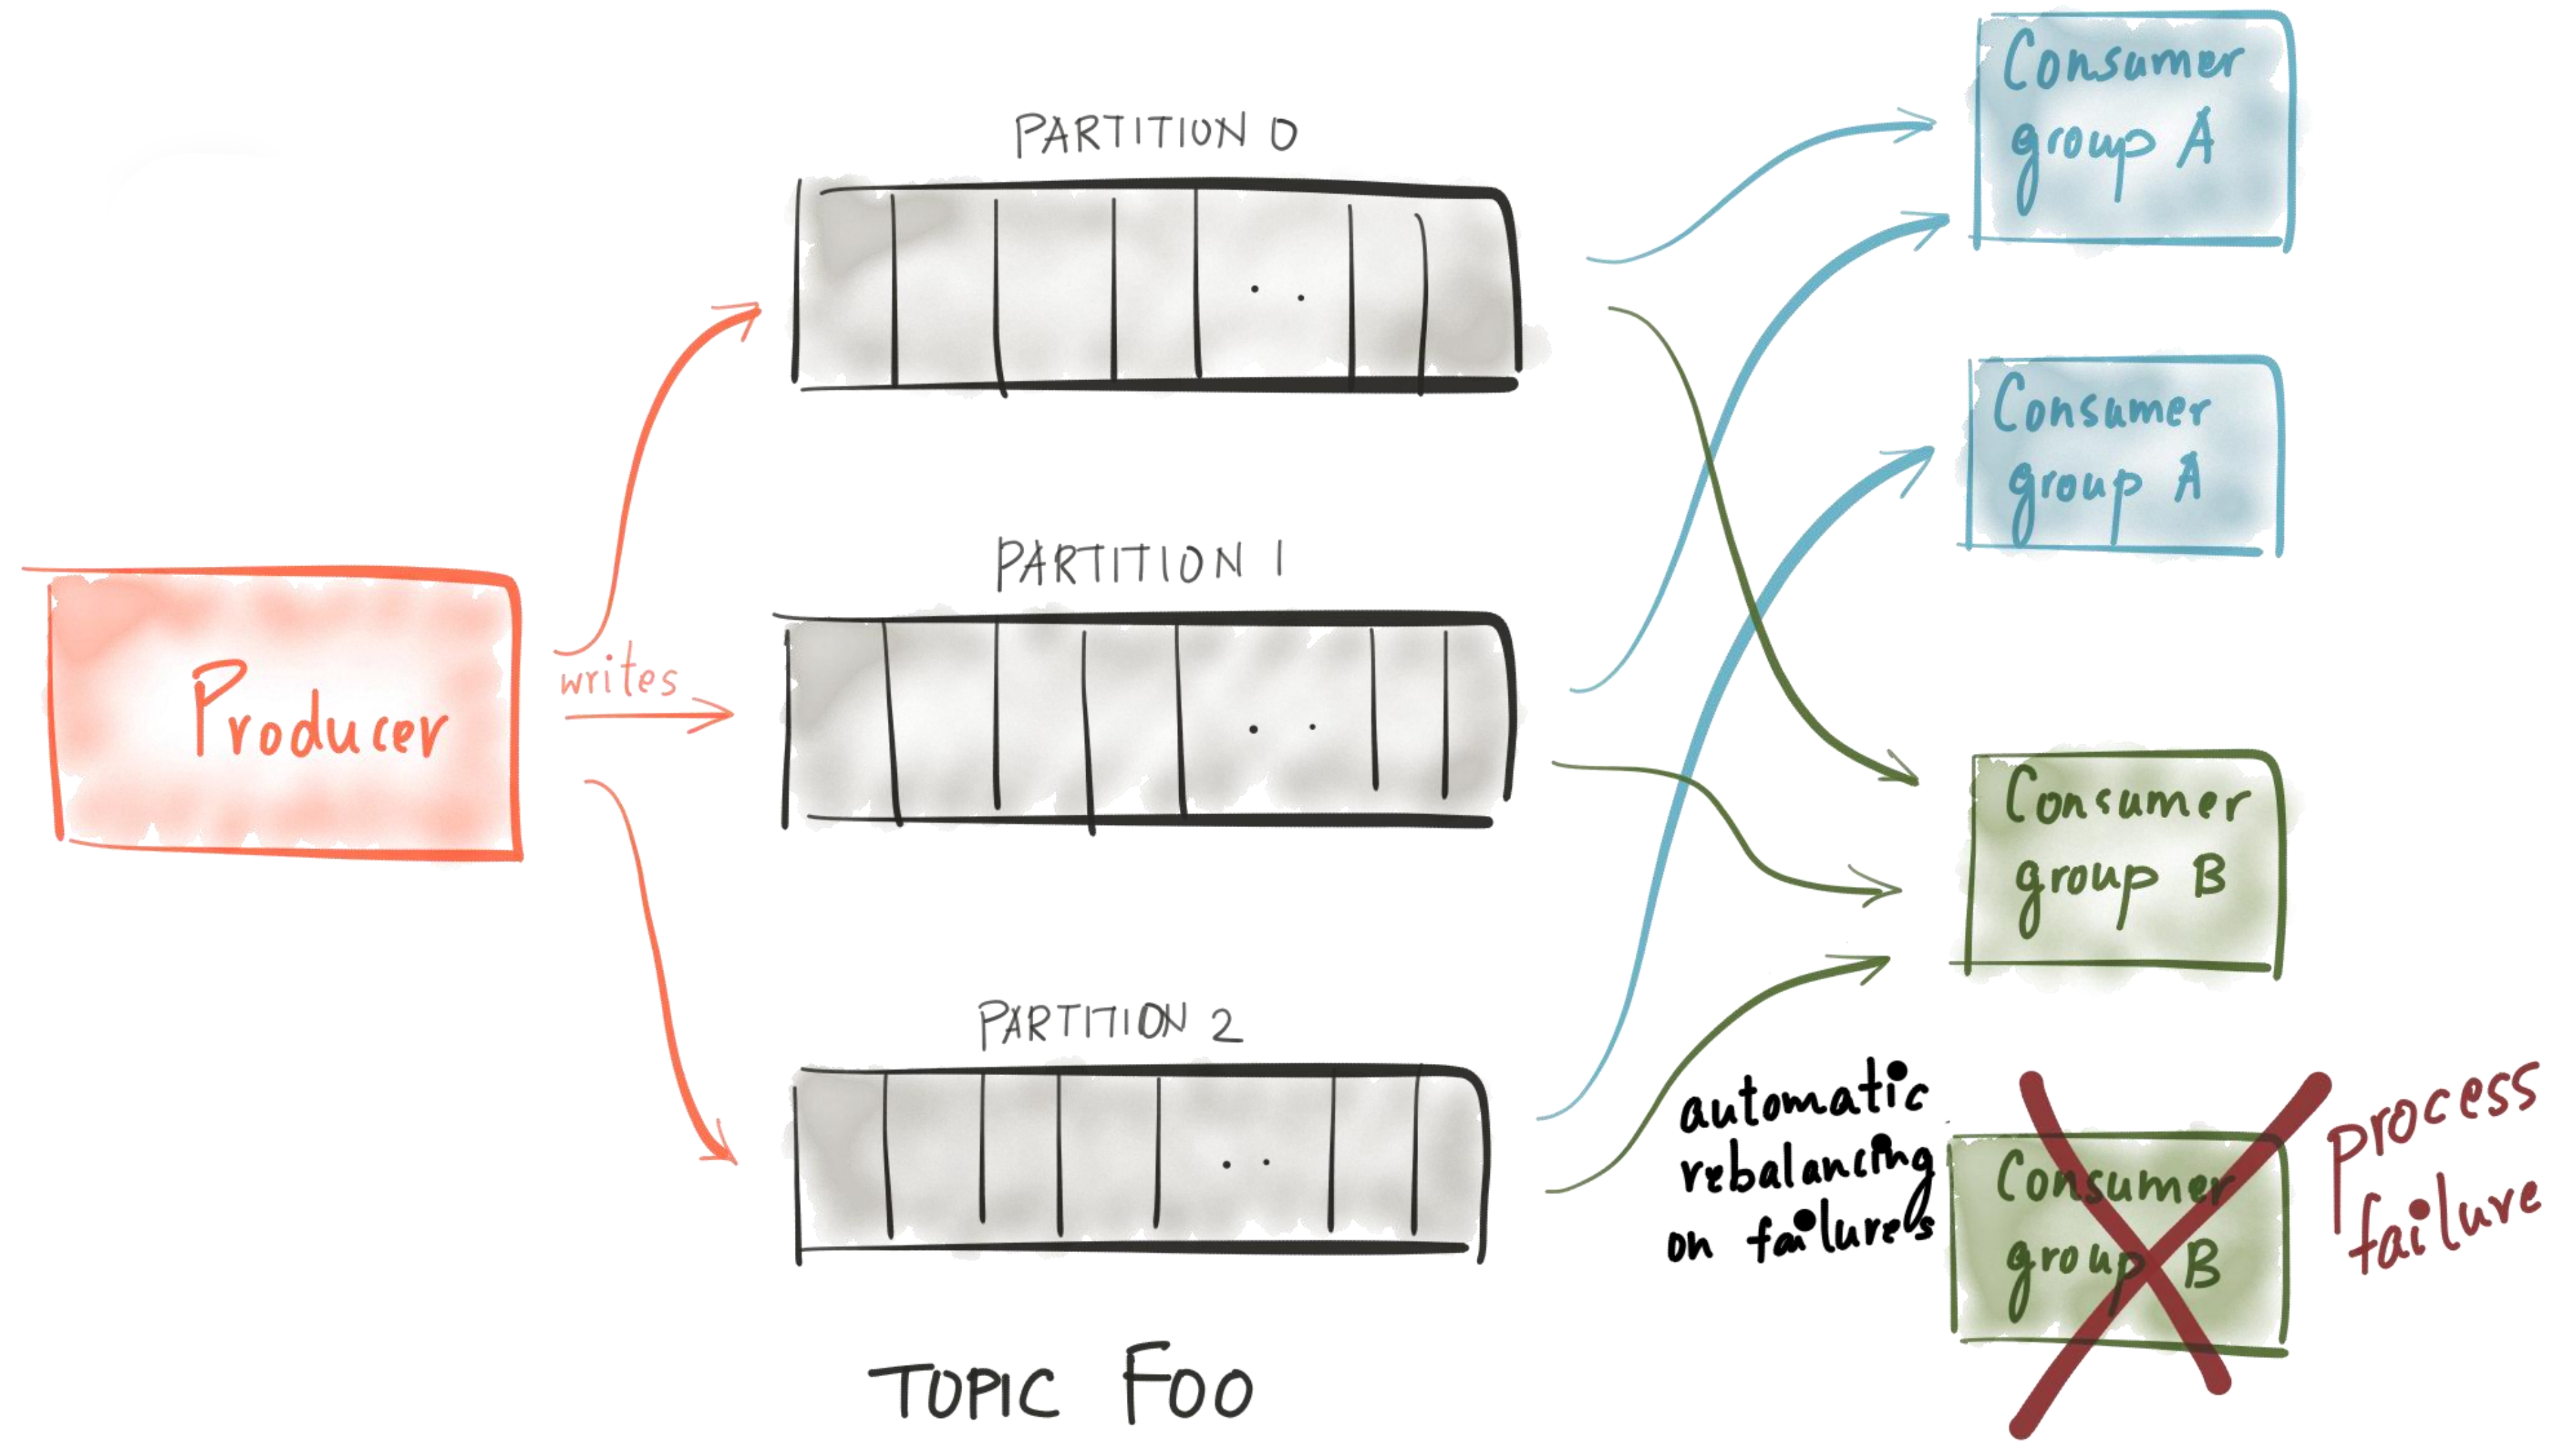
# TOPIC $\approx$ PARTITIONED LOG

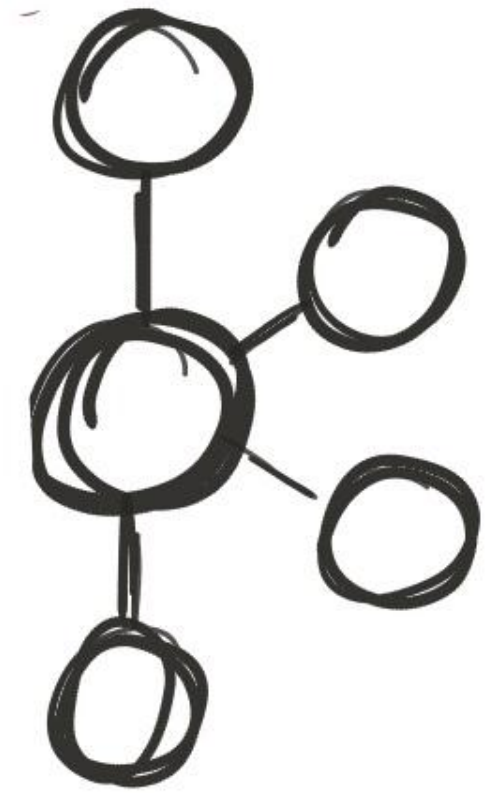
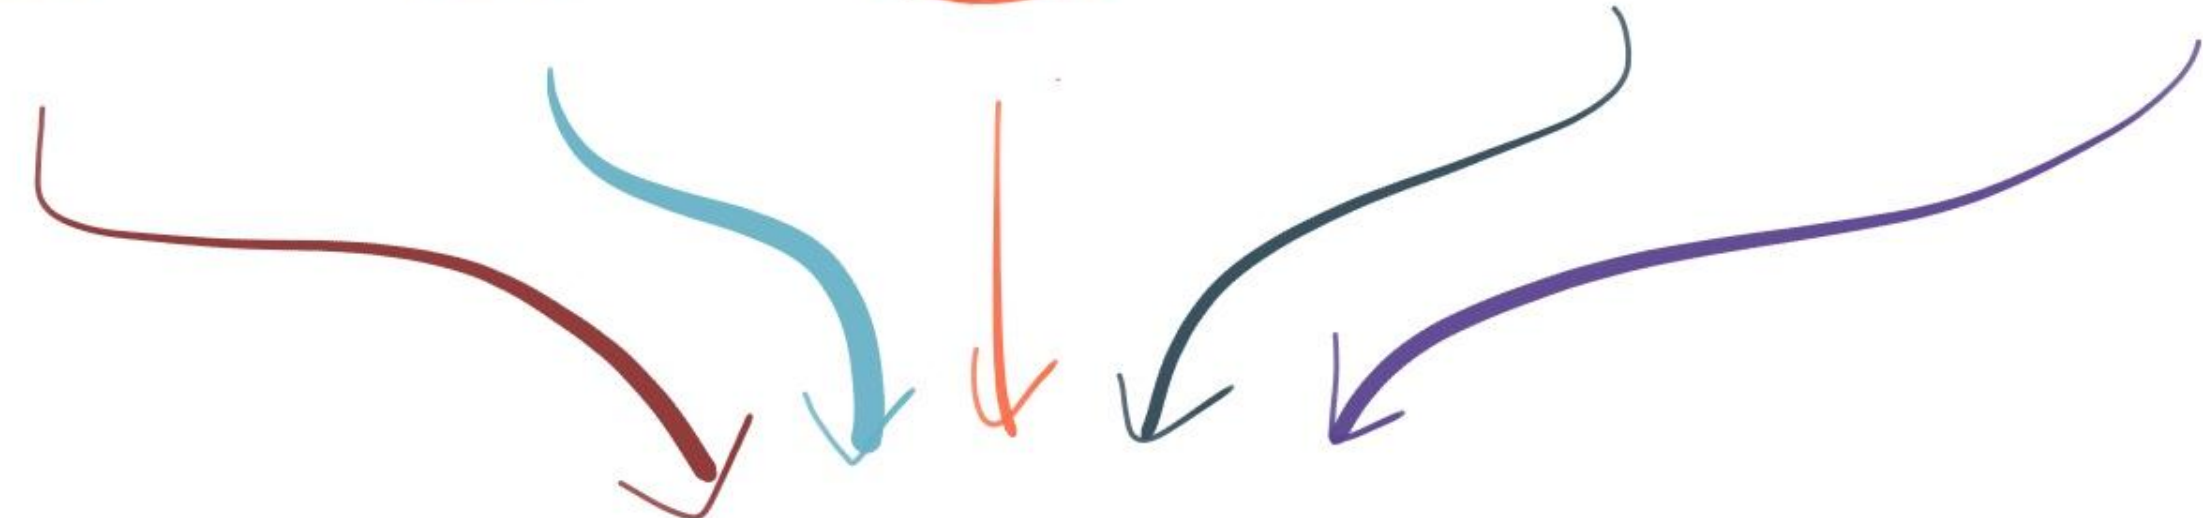




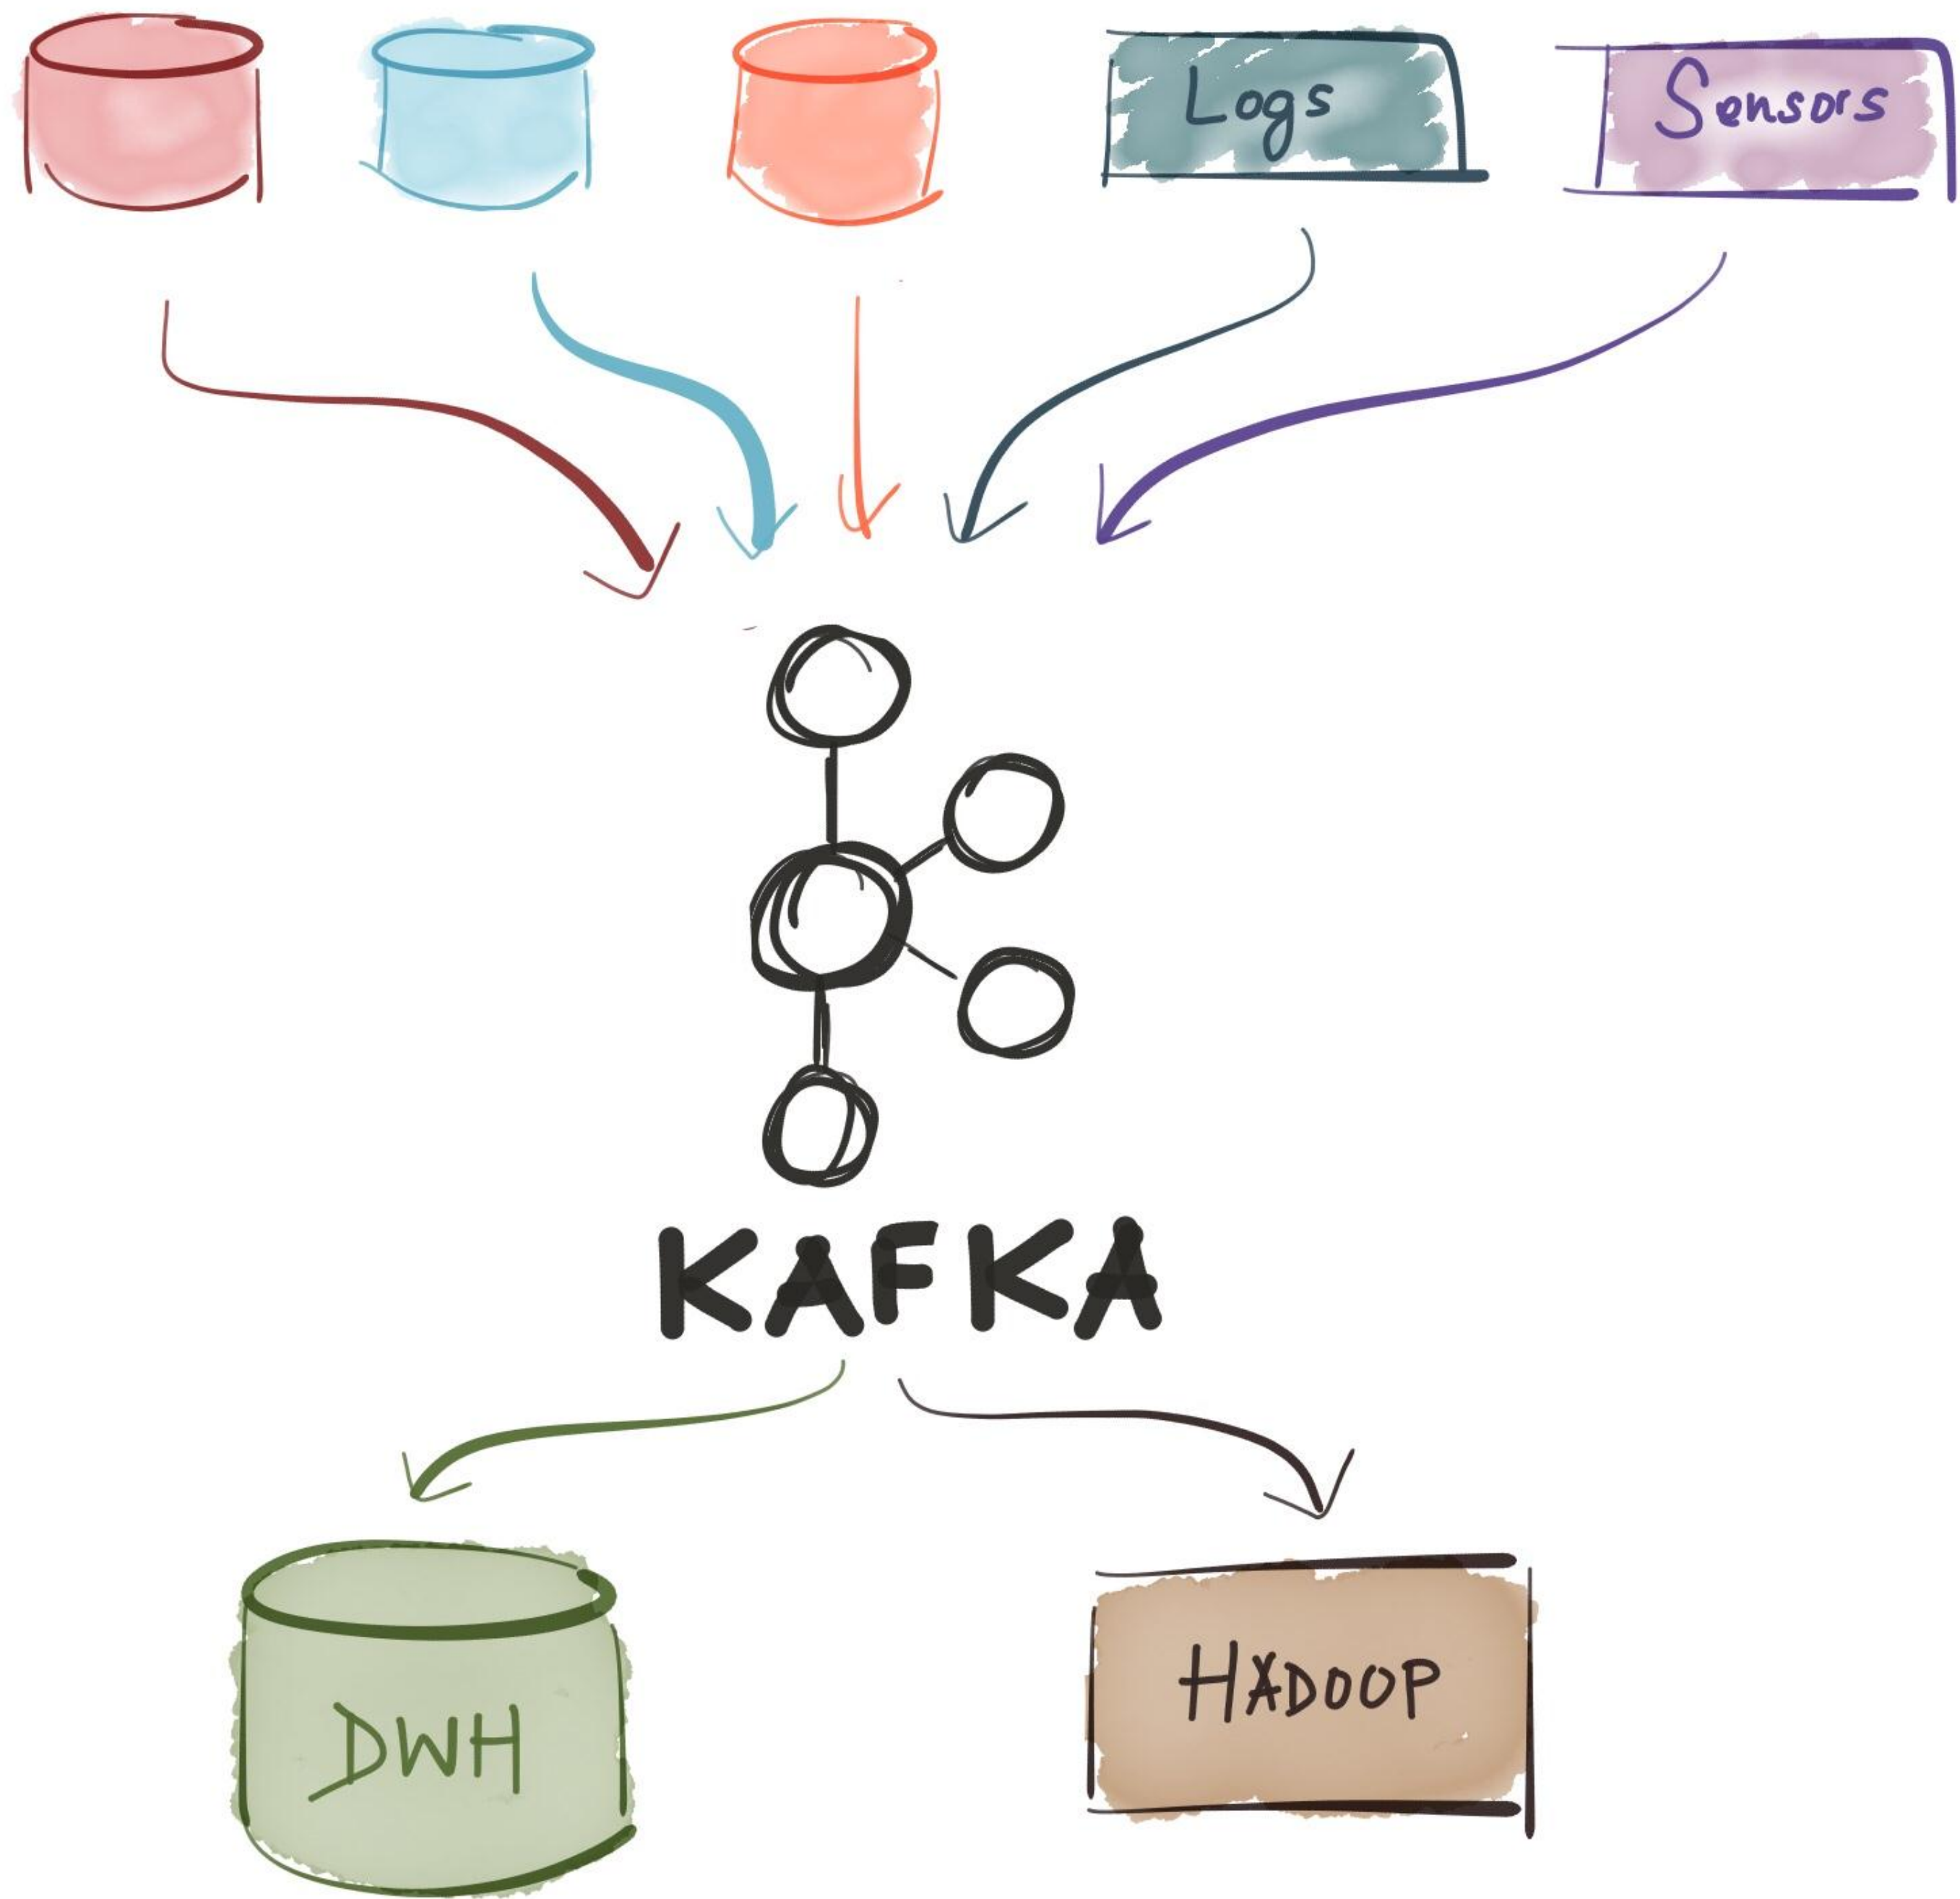


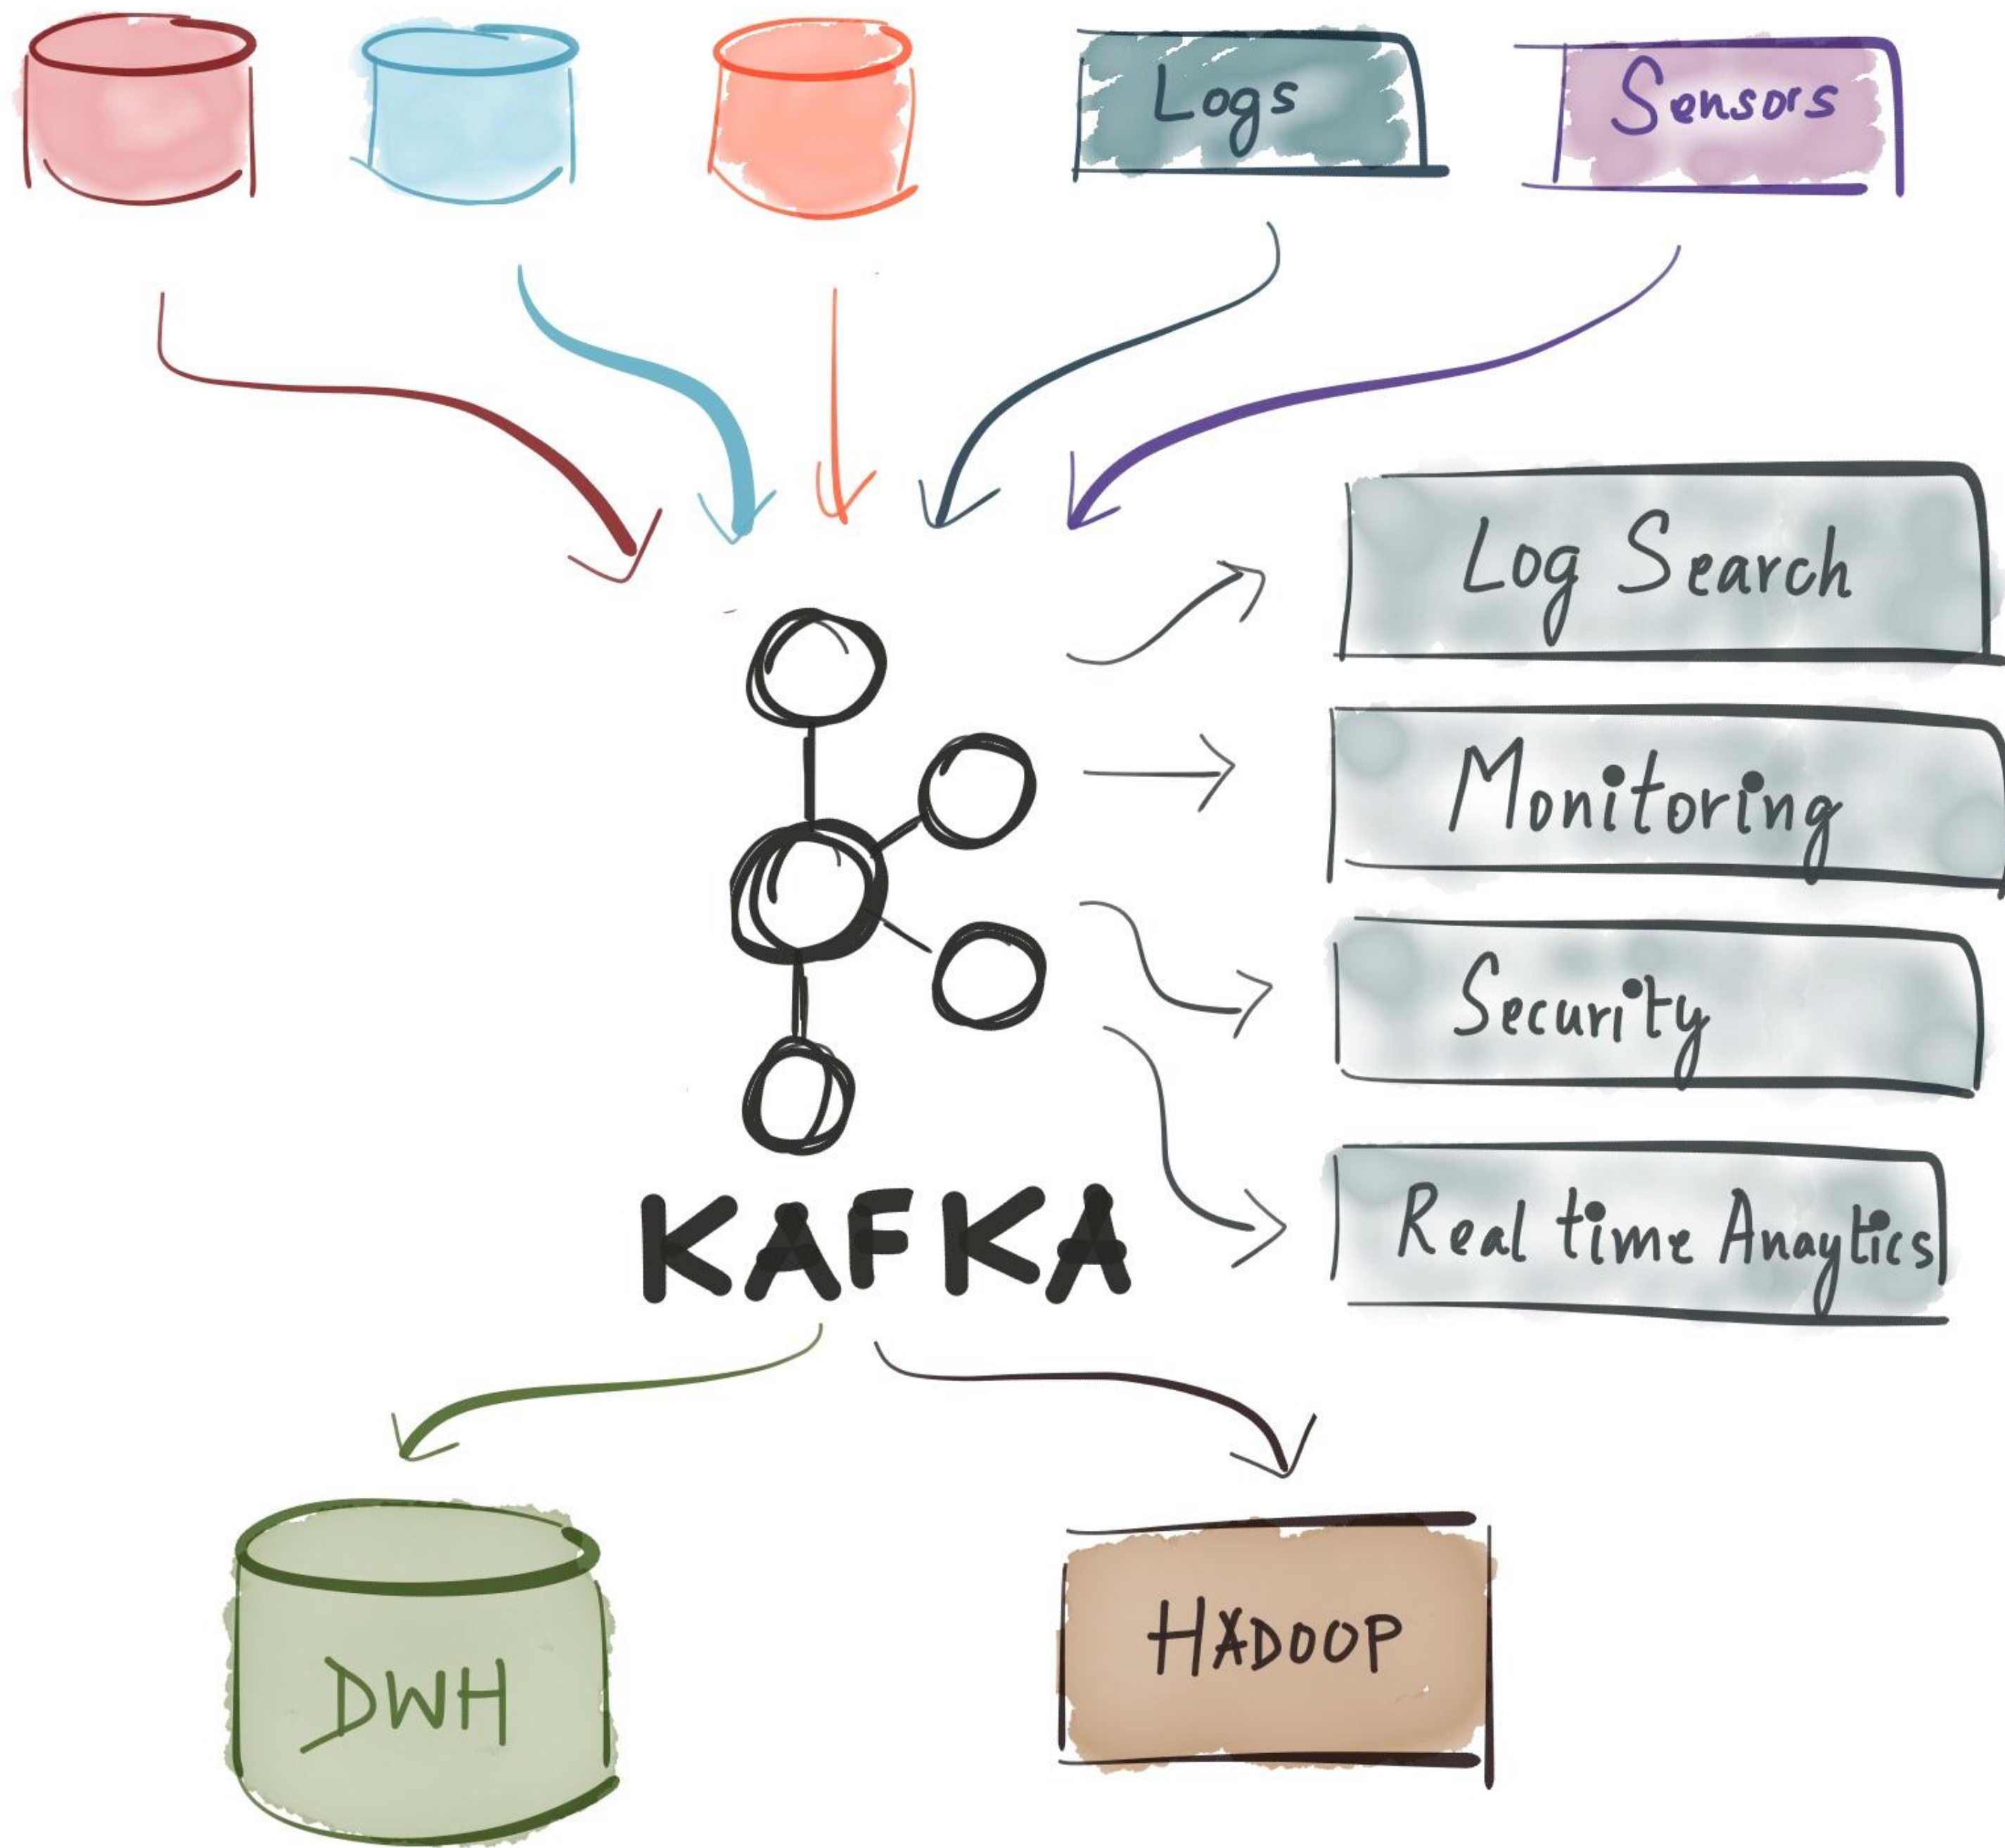
# Topic $\approx$ PARTITIONED LOG

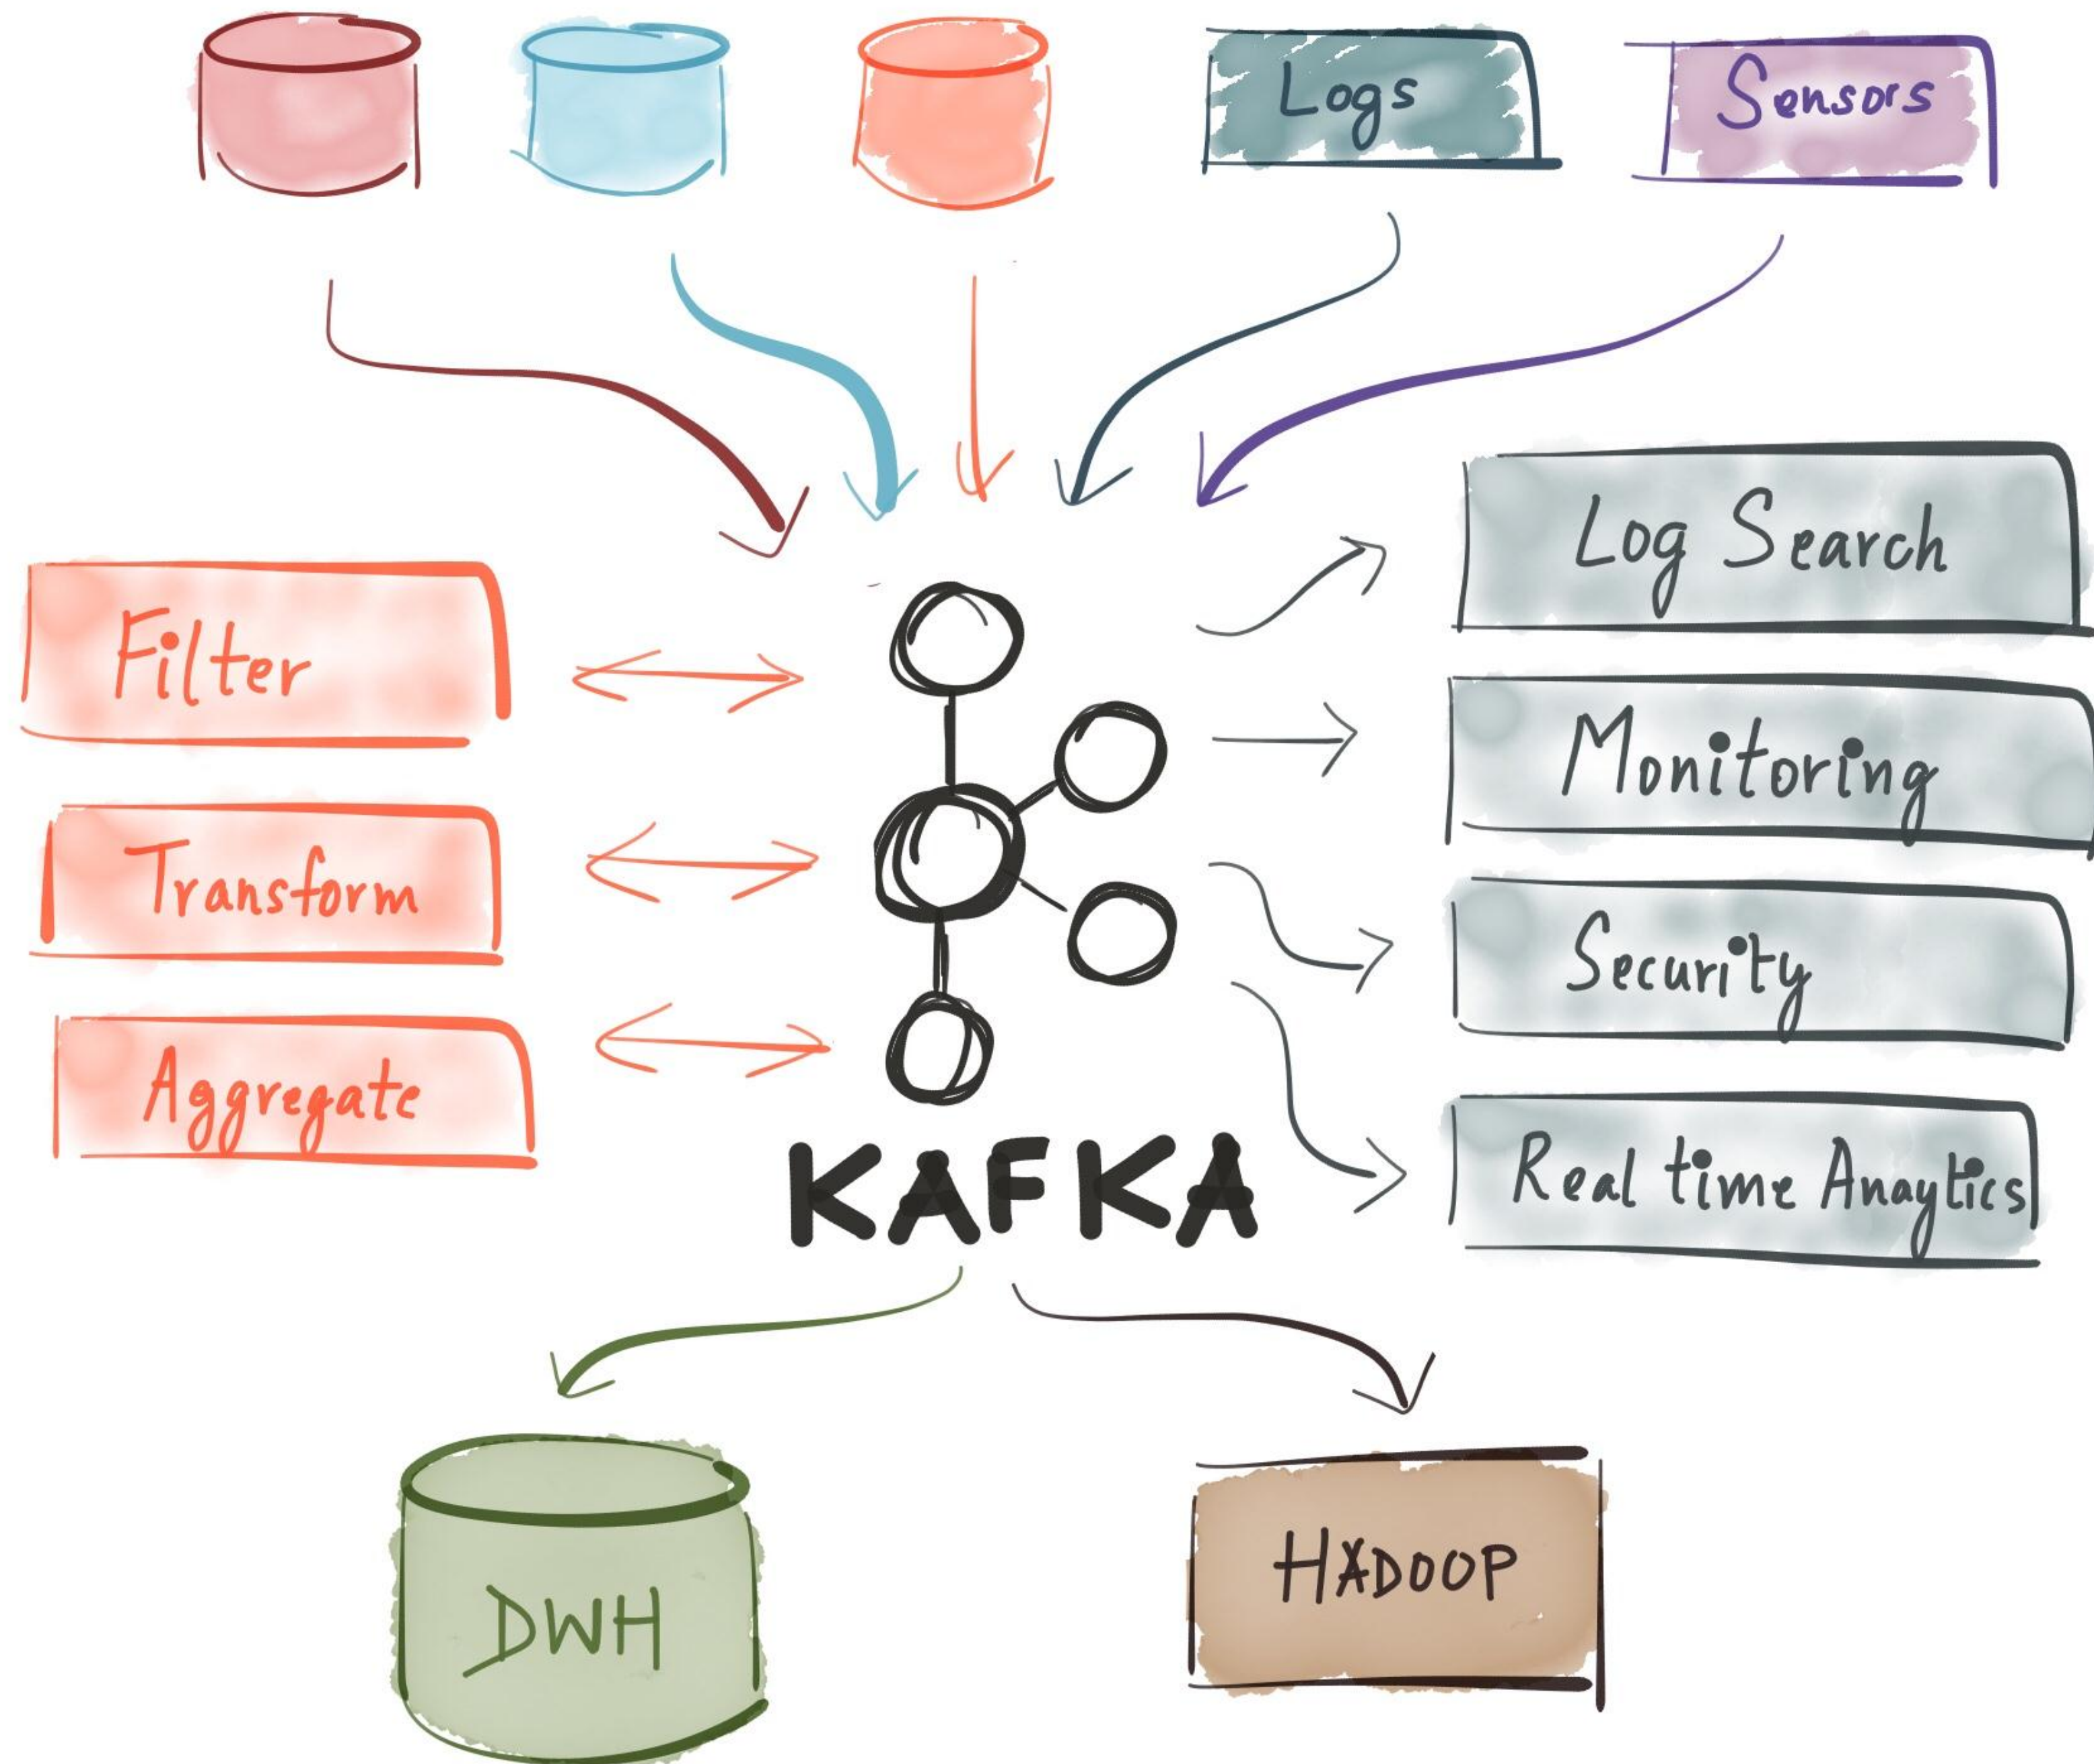




**KAFKA**



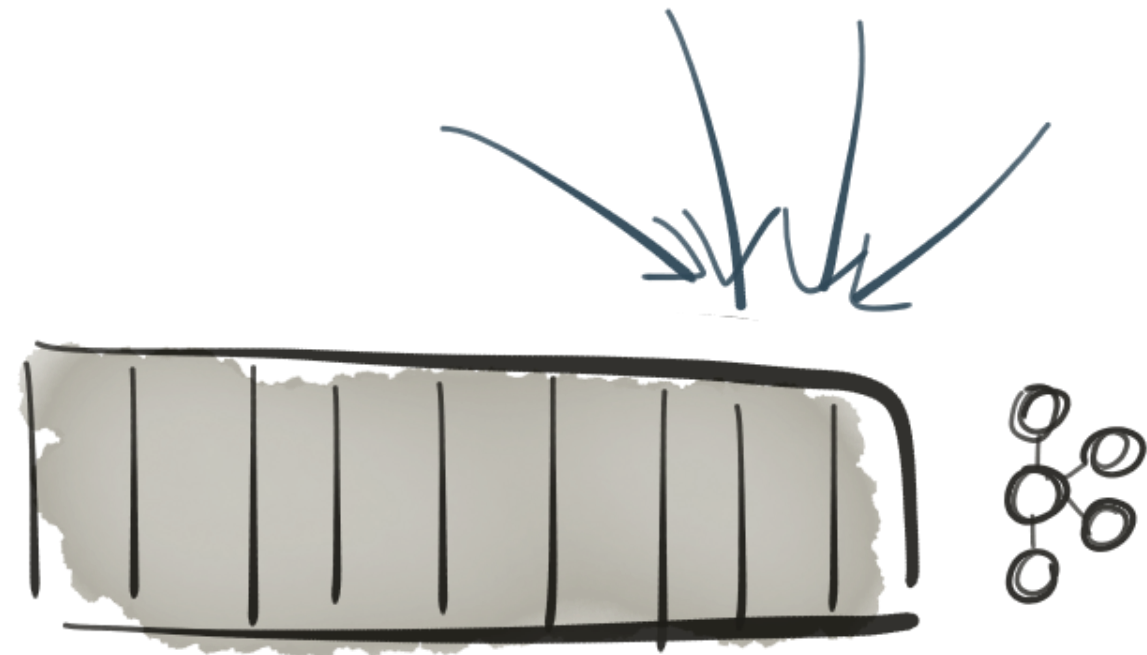






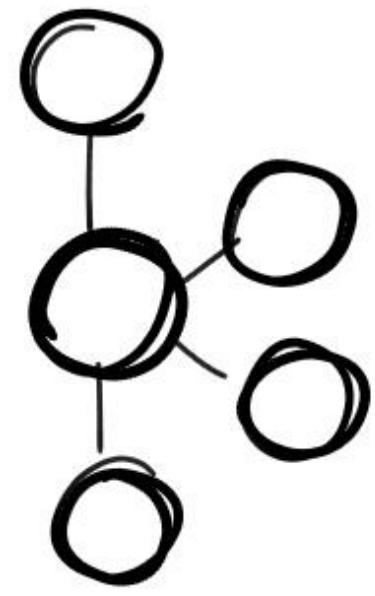
Everything  
is a  
STREAM

800 BILLION



> 2.5 TRILLION

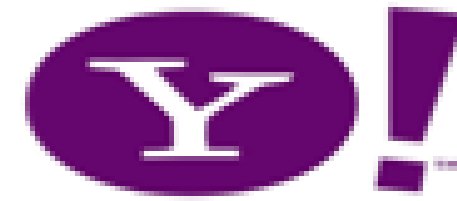
MESSAGES/DAY



# EVERYWHERE ELSE

intuit.

 Cerner

















 Square



UBER





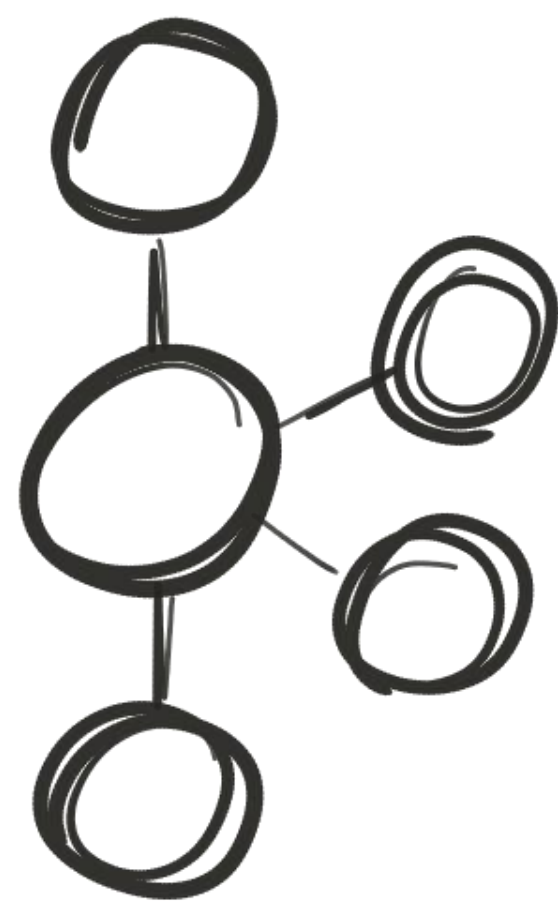
  
WIKIPEDIA  
The Free Encyclopedia









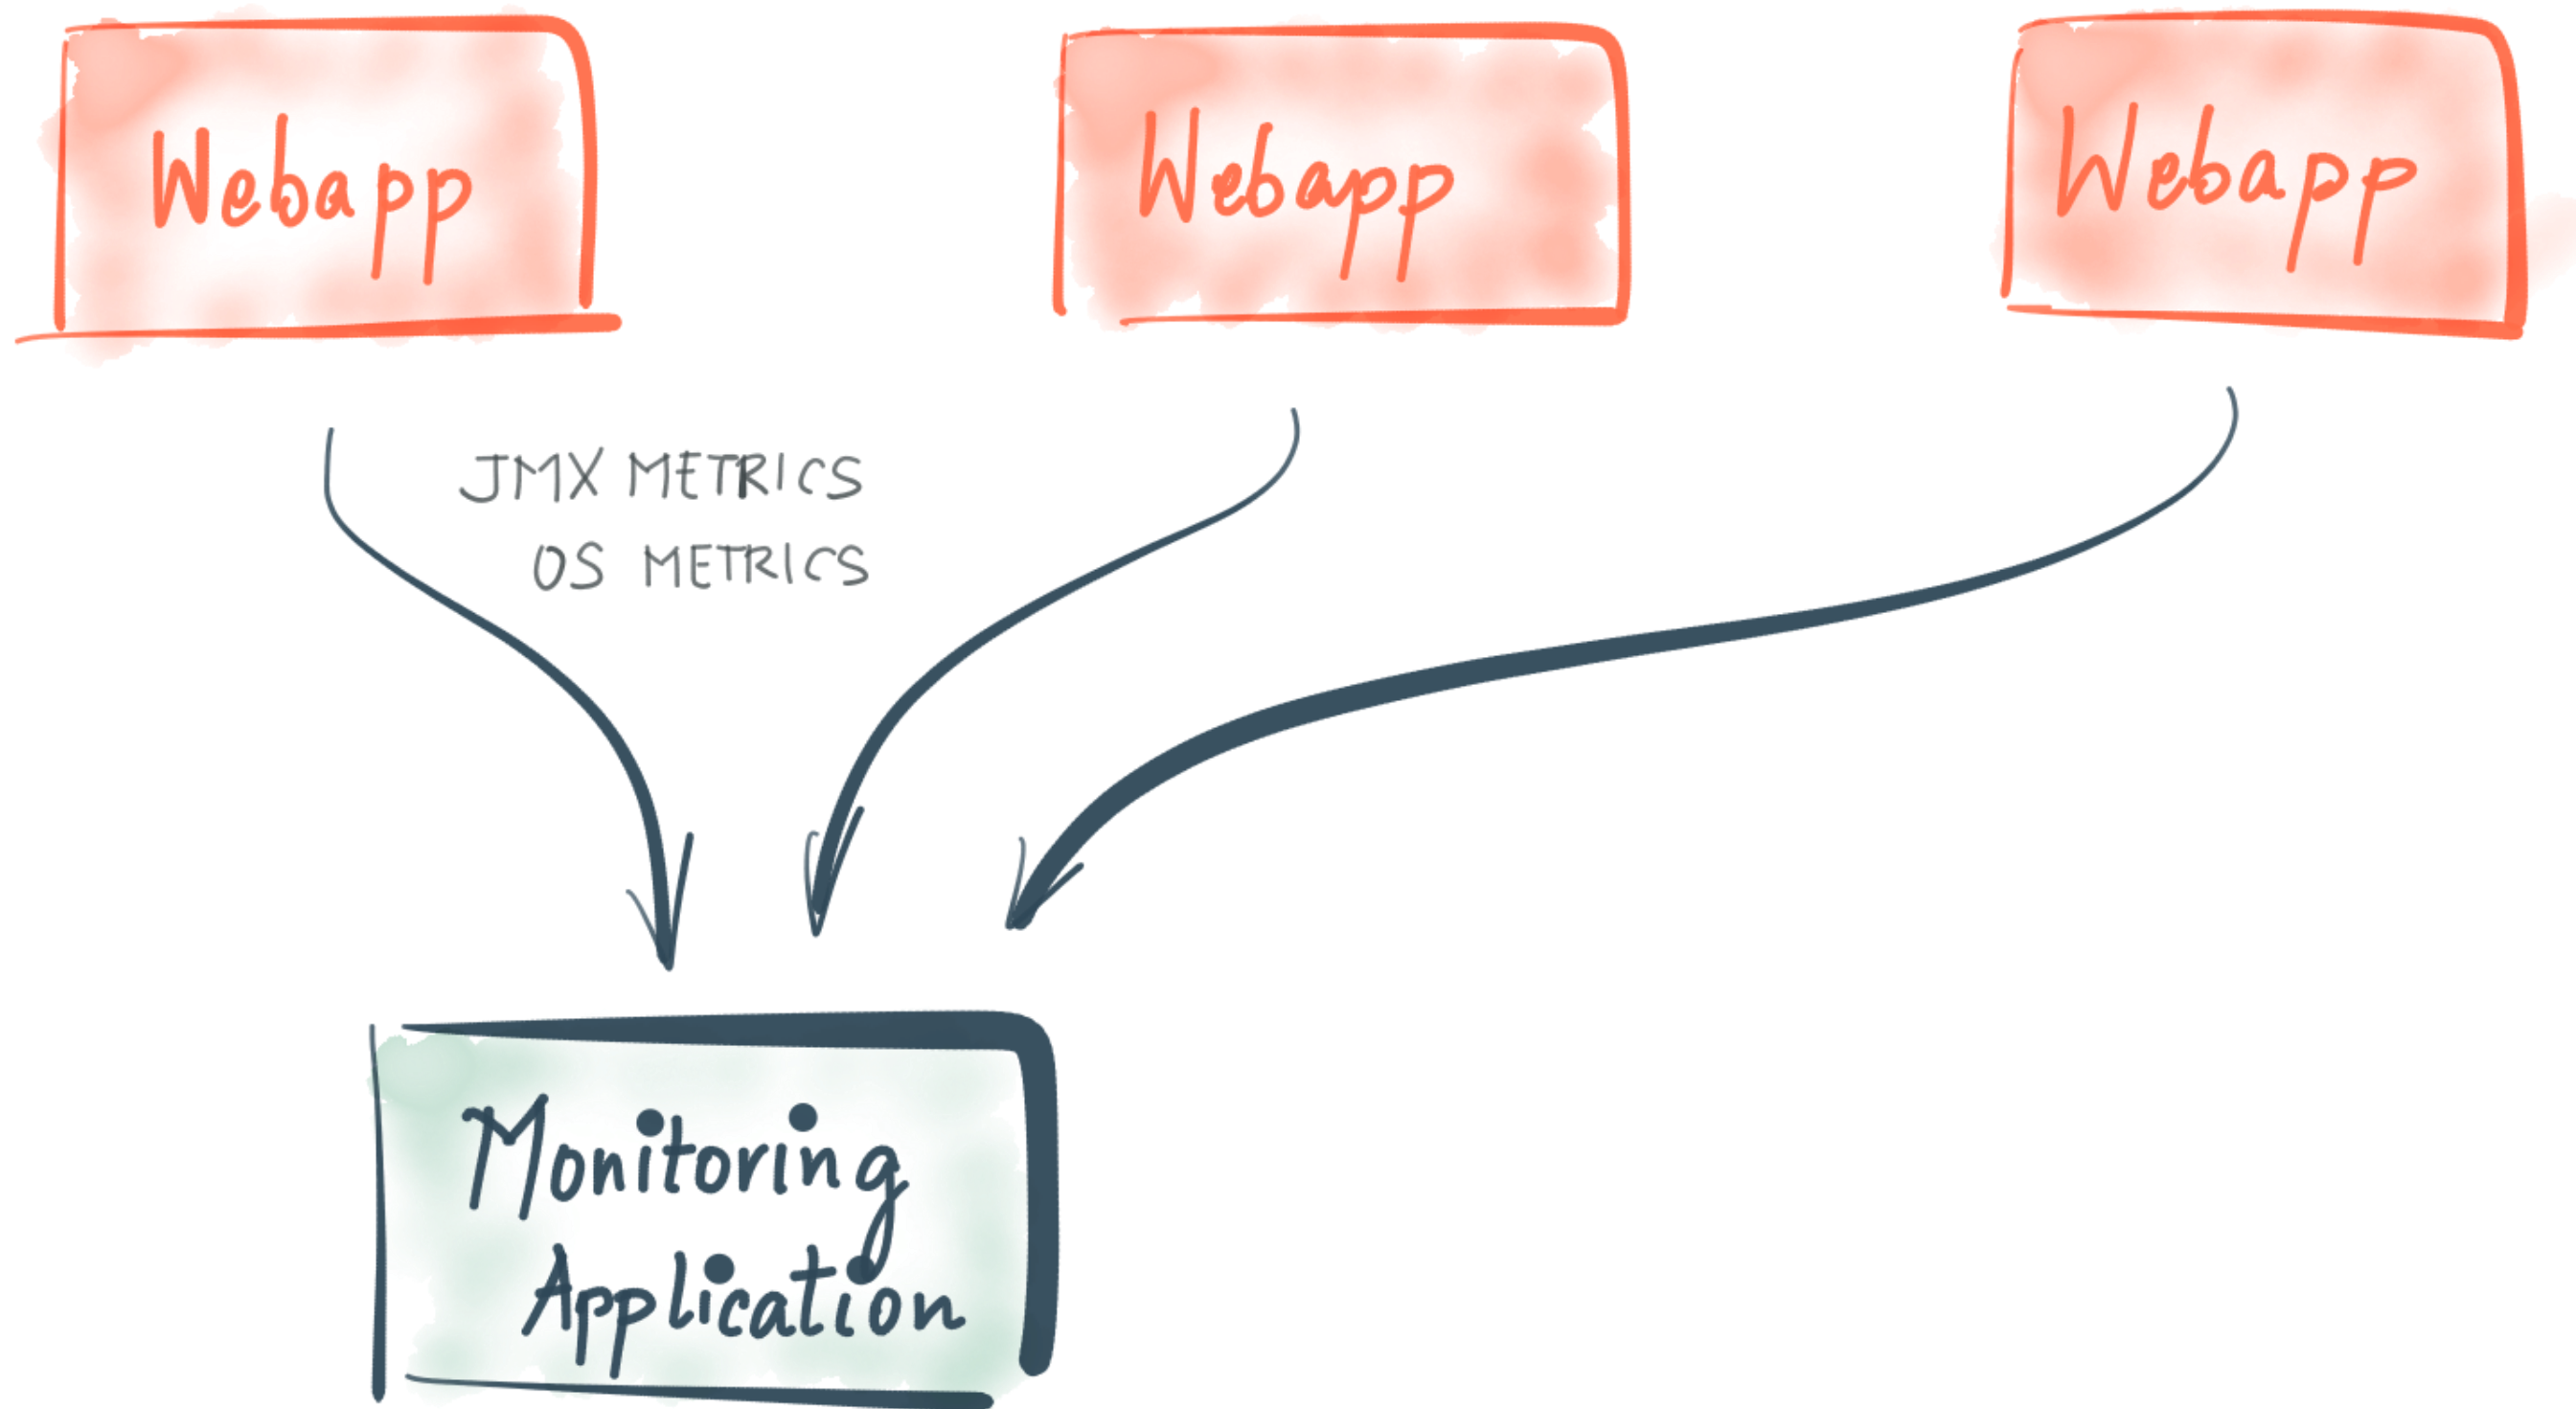


Monitoring



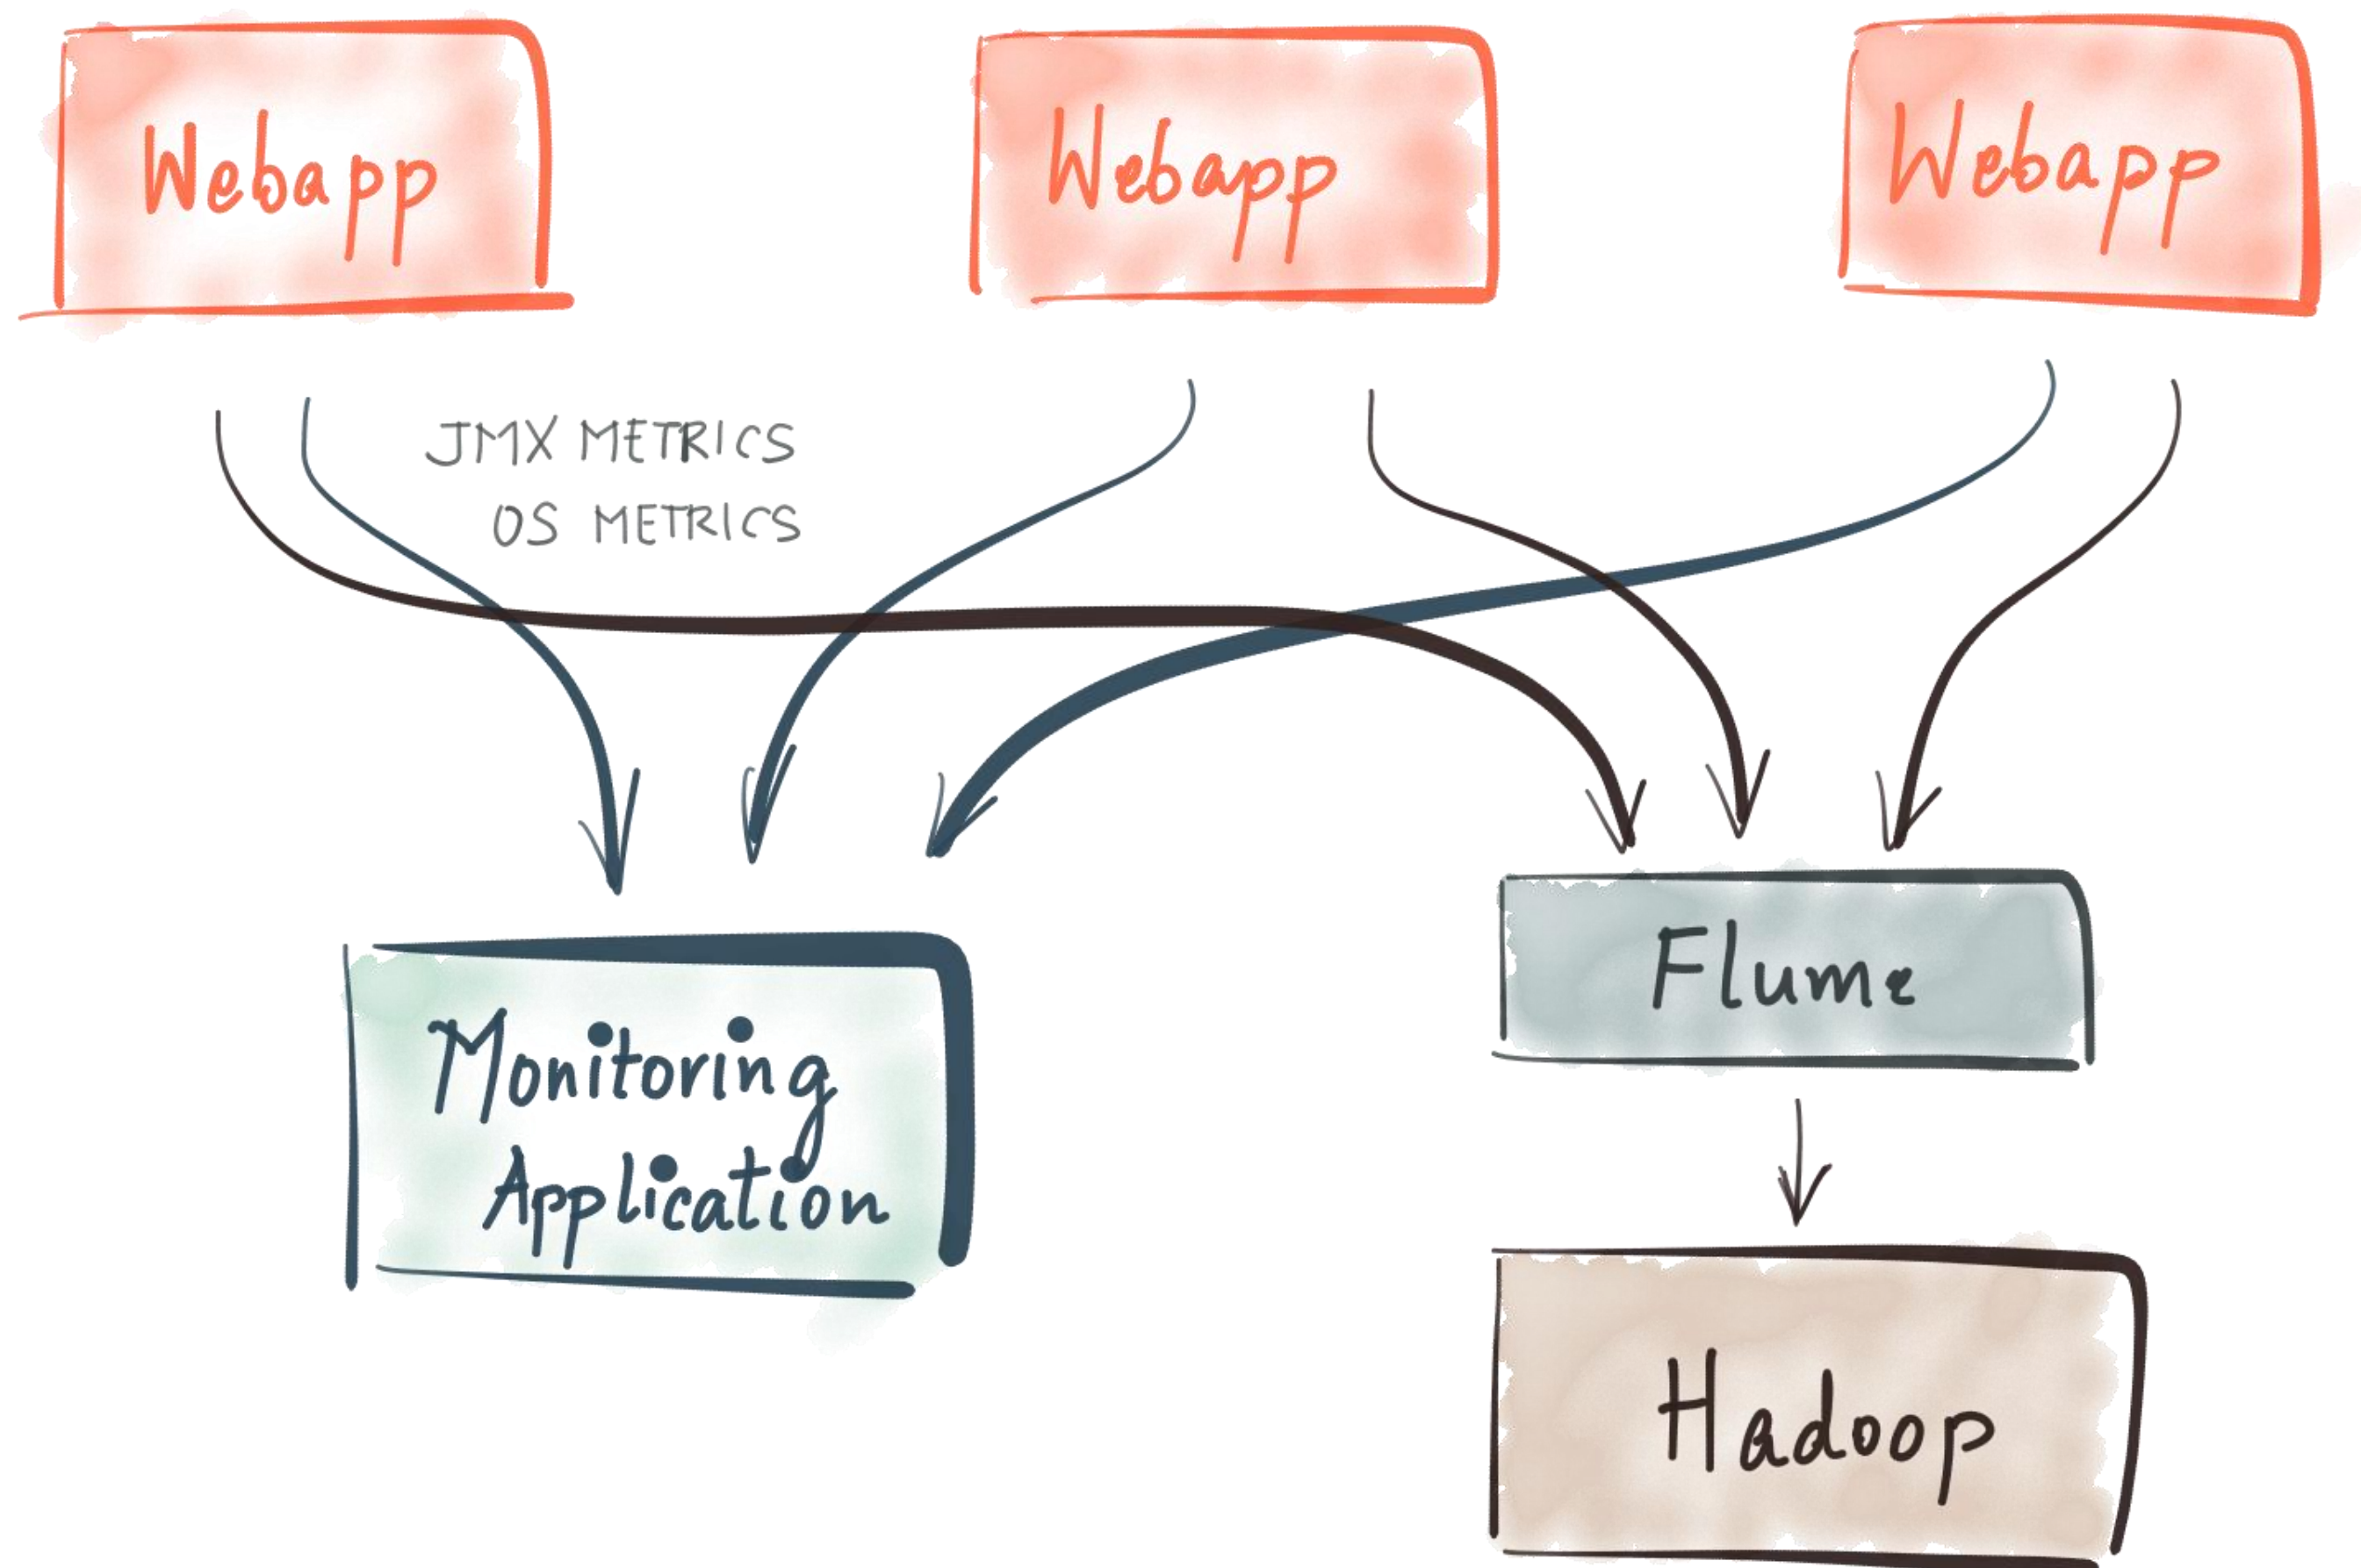


# FOR MONITORING



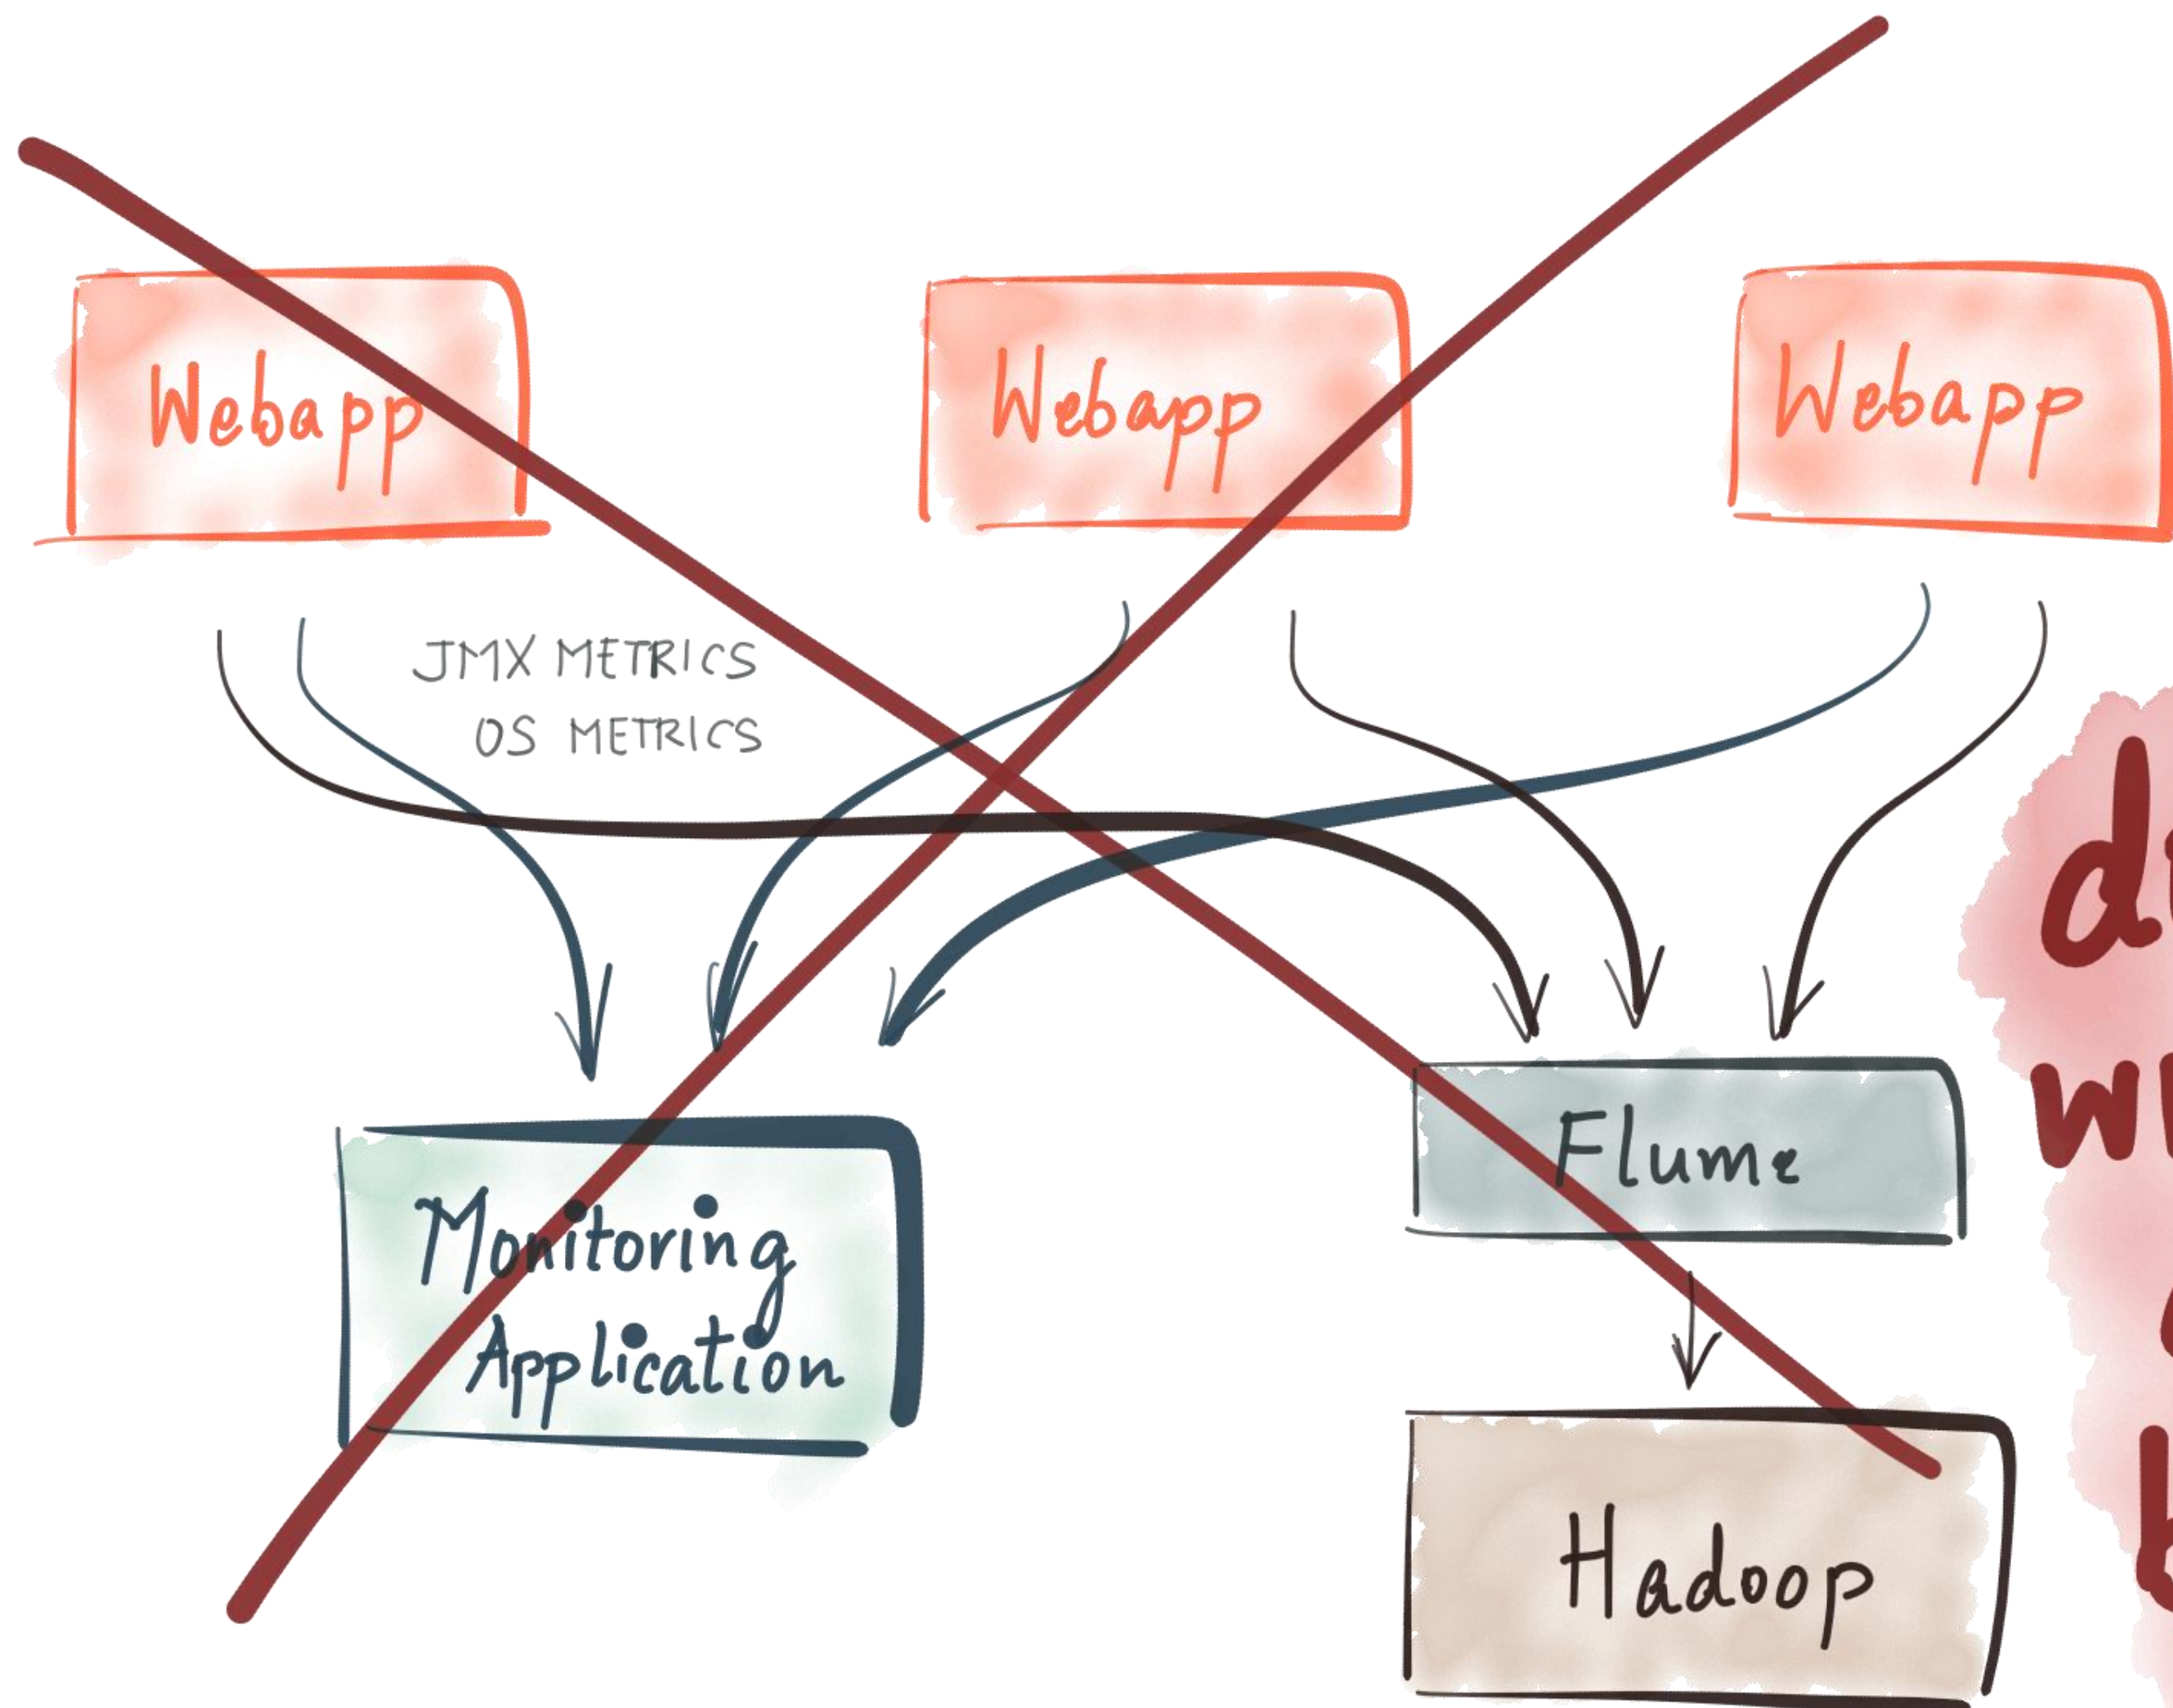


# FOR MONITORING





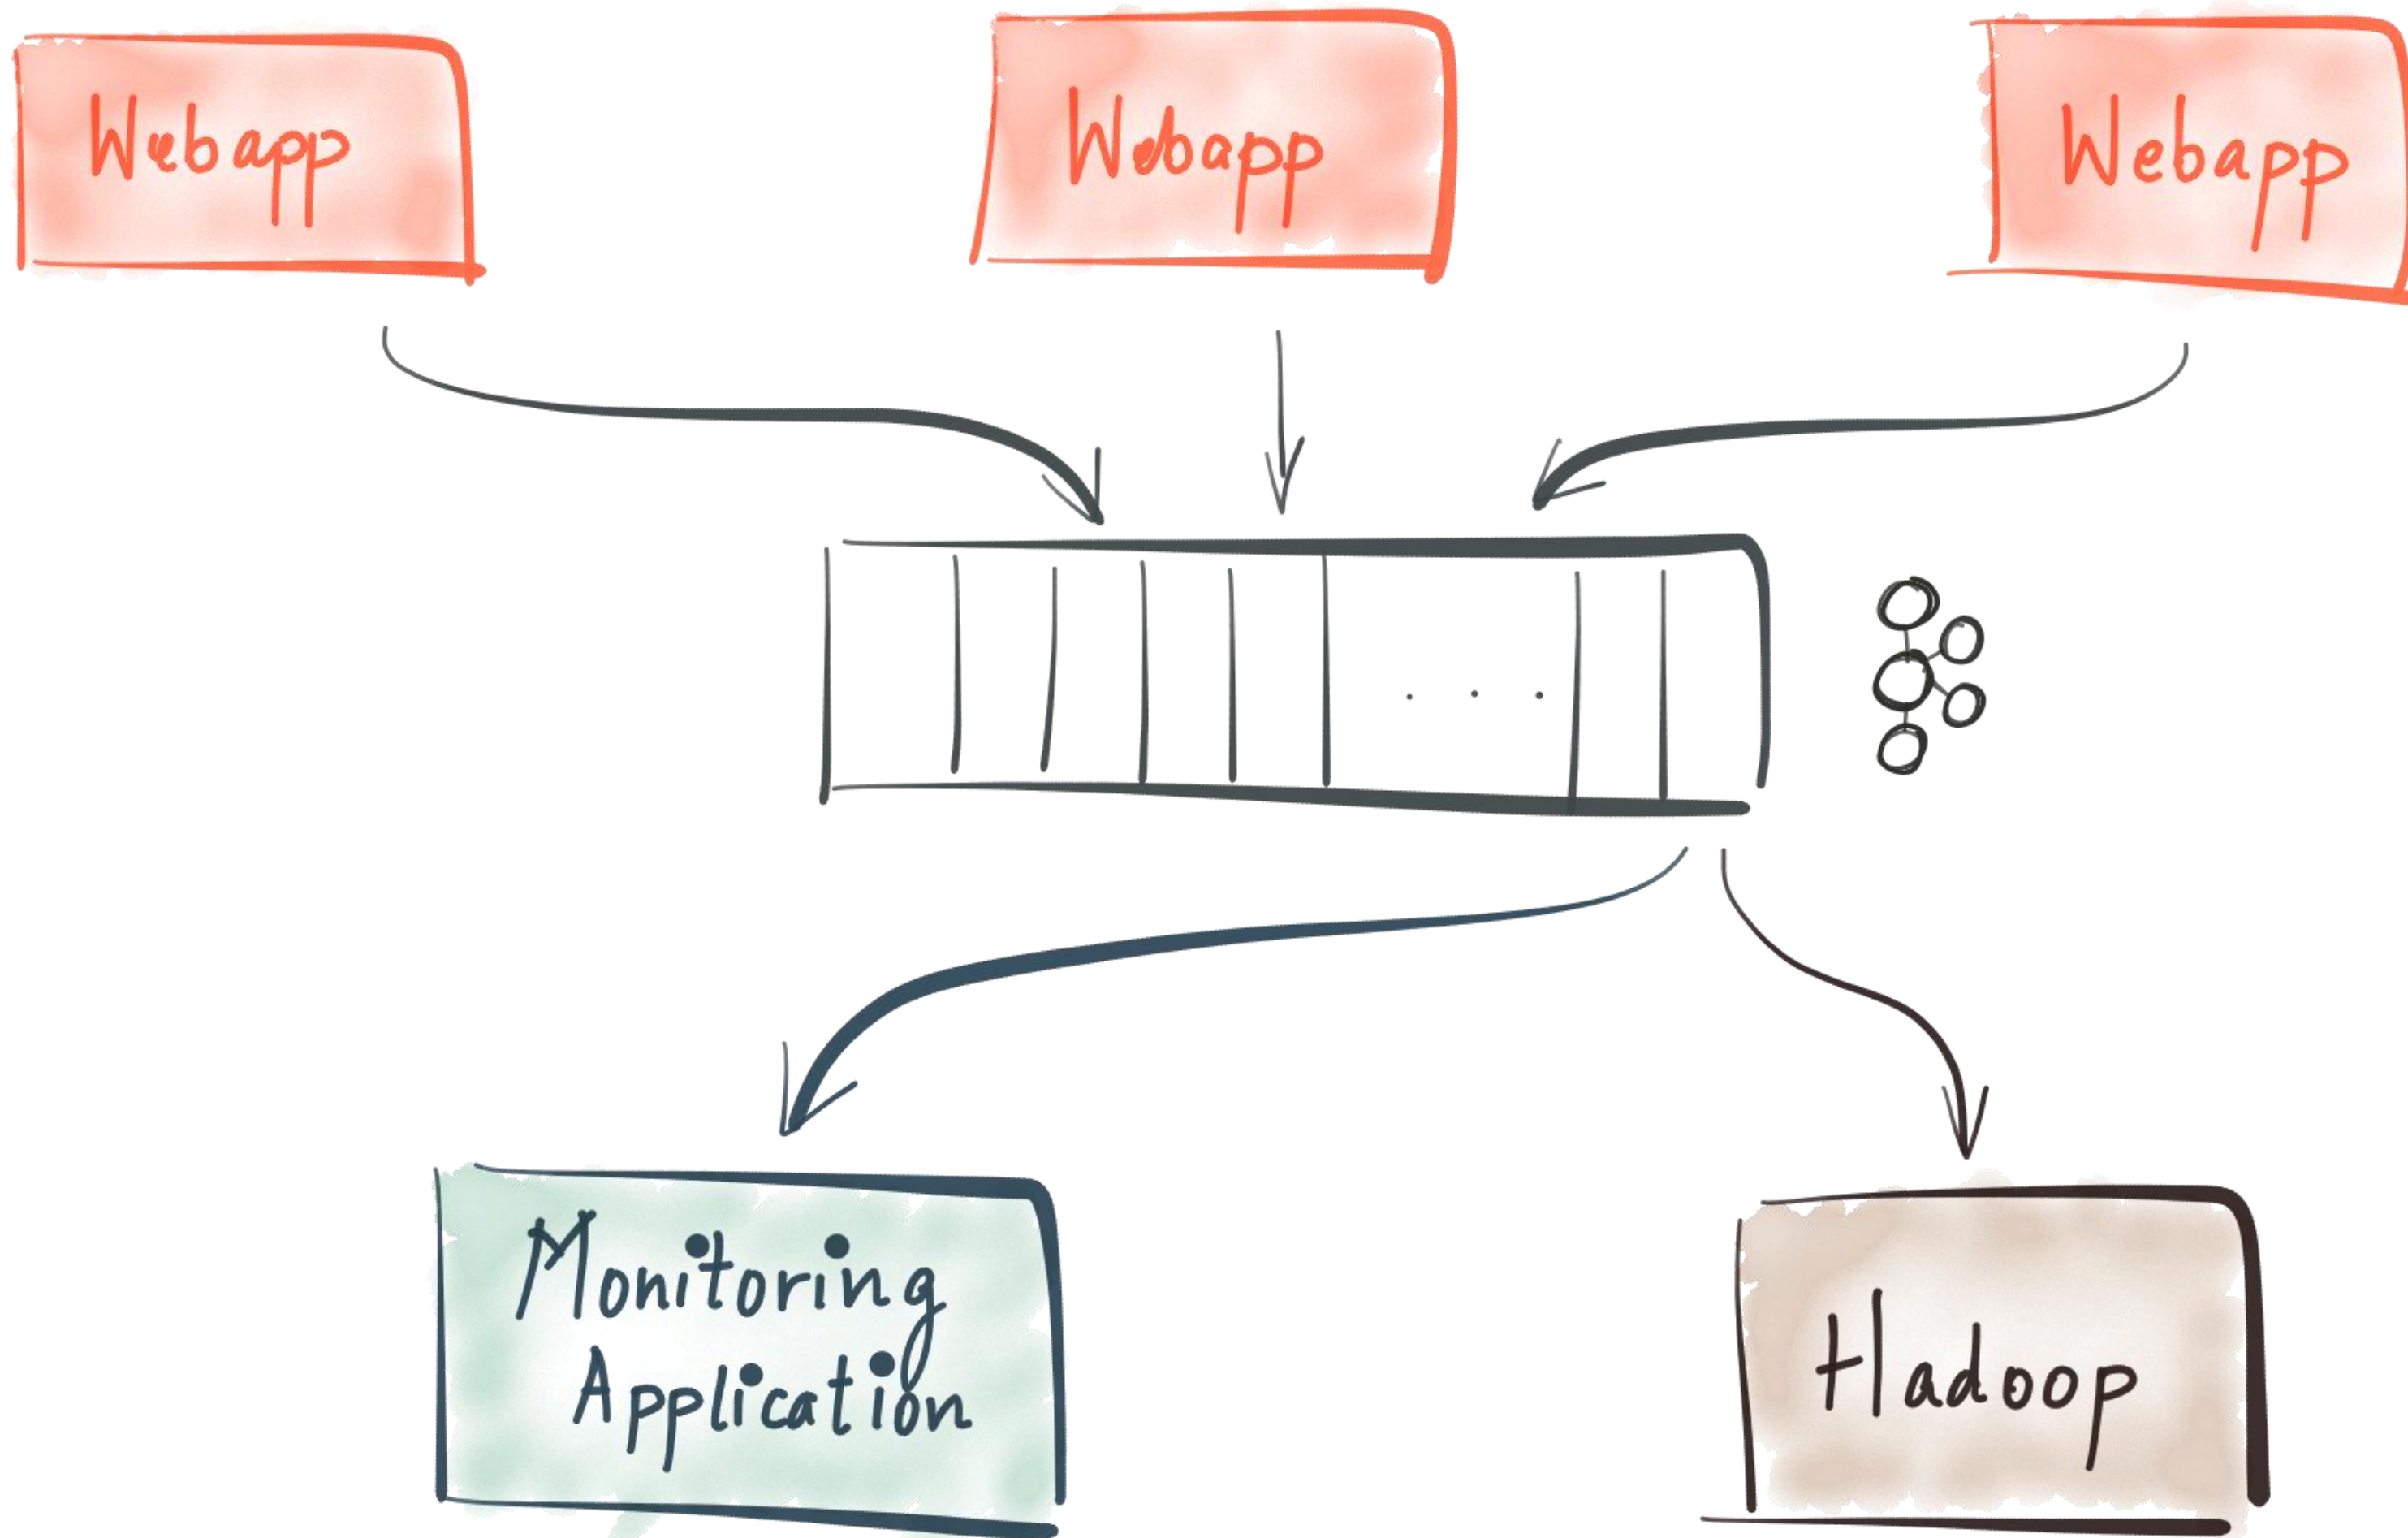
# FOR MONITORING

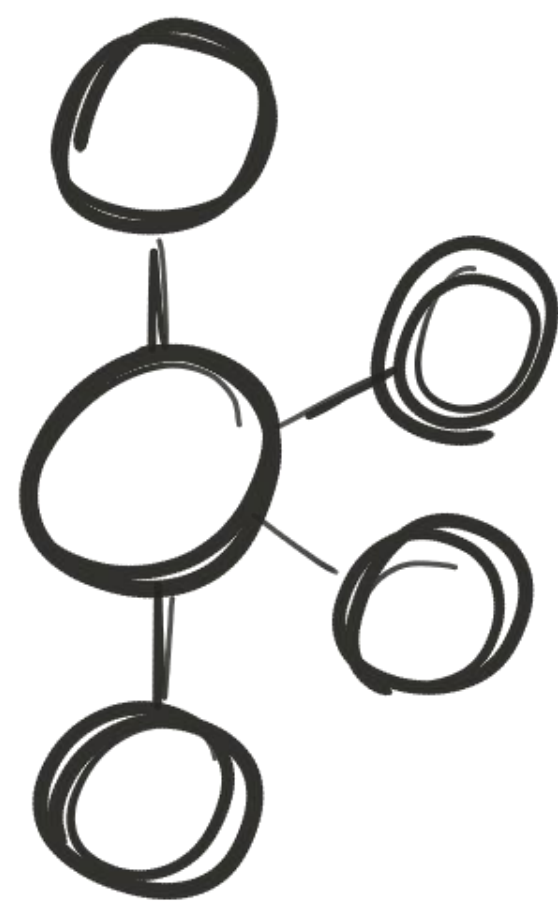


double writes are bad!



# FOR MONITORING

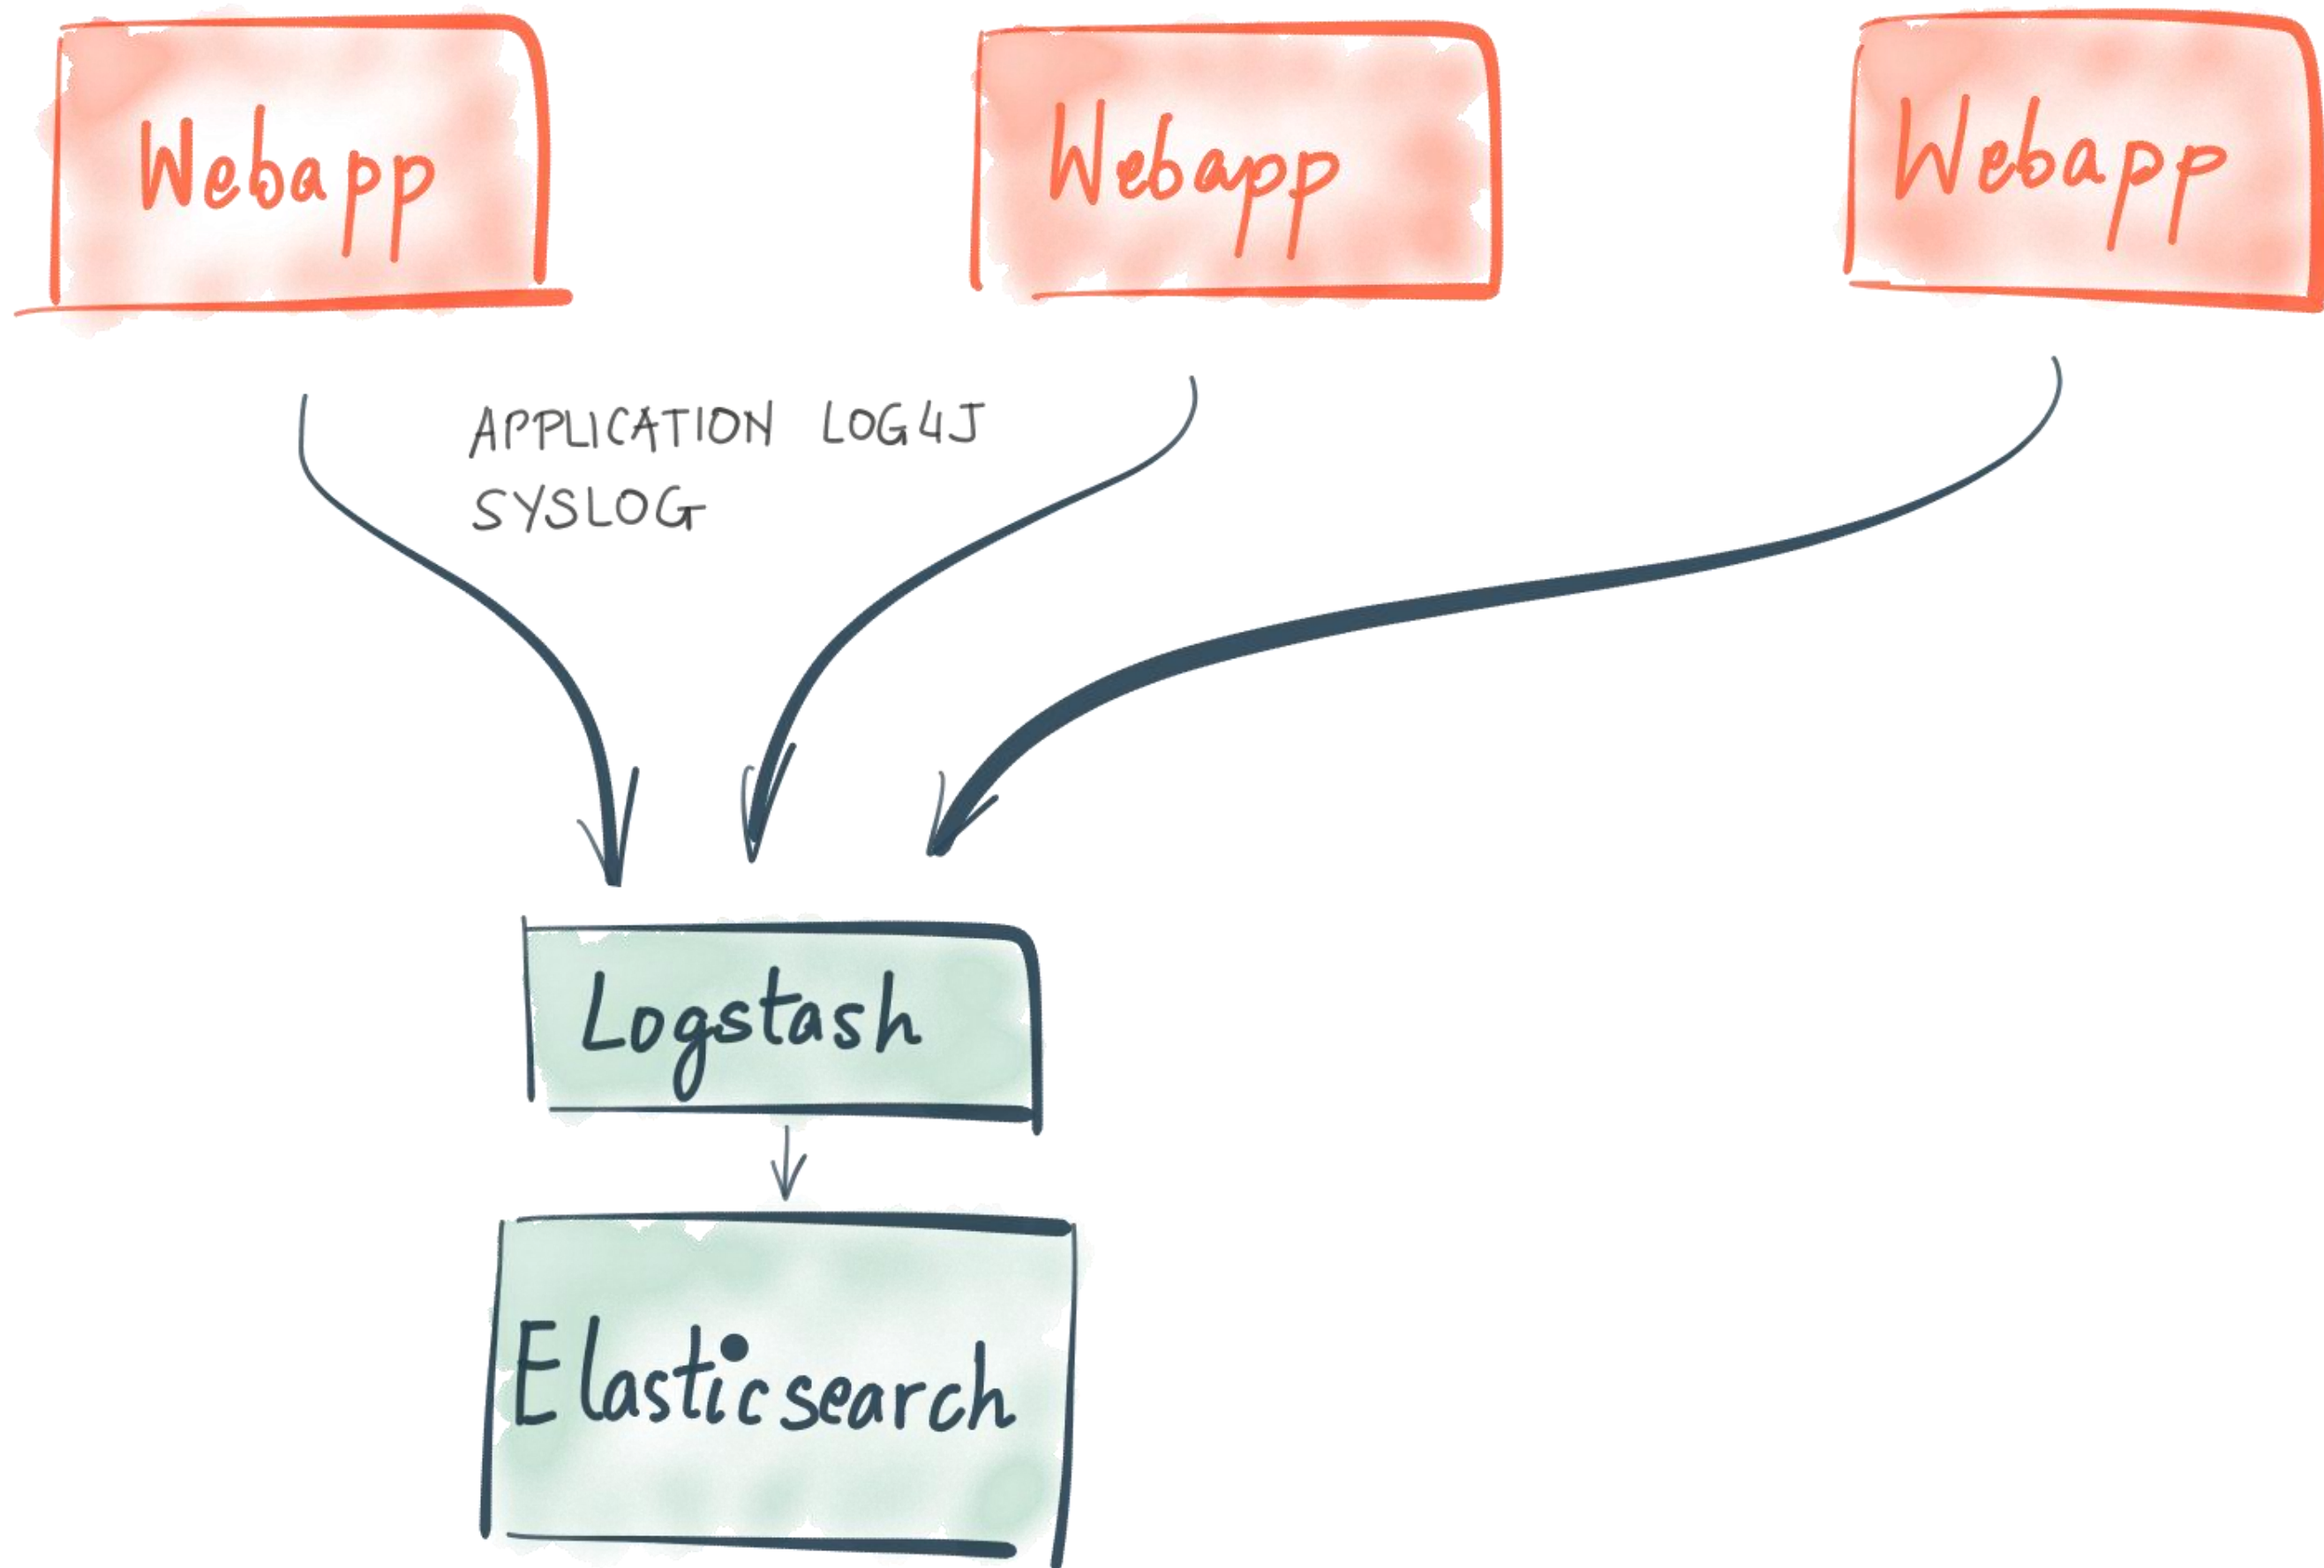




Log Collection

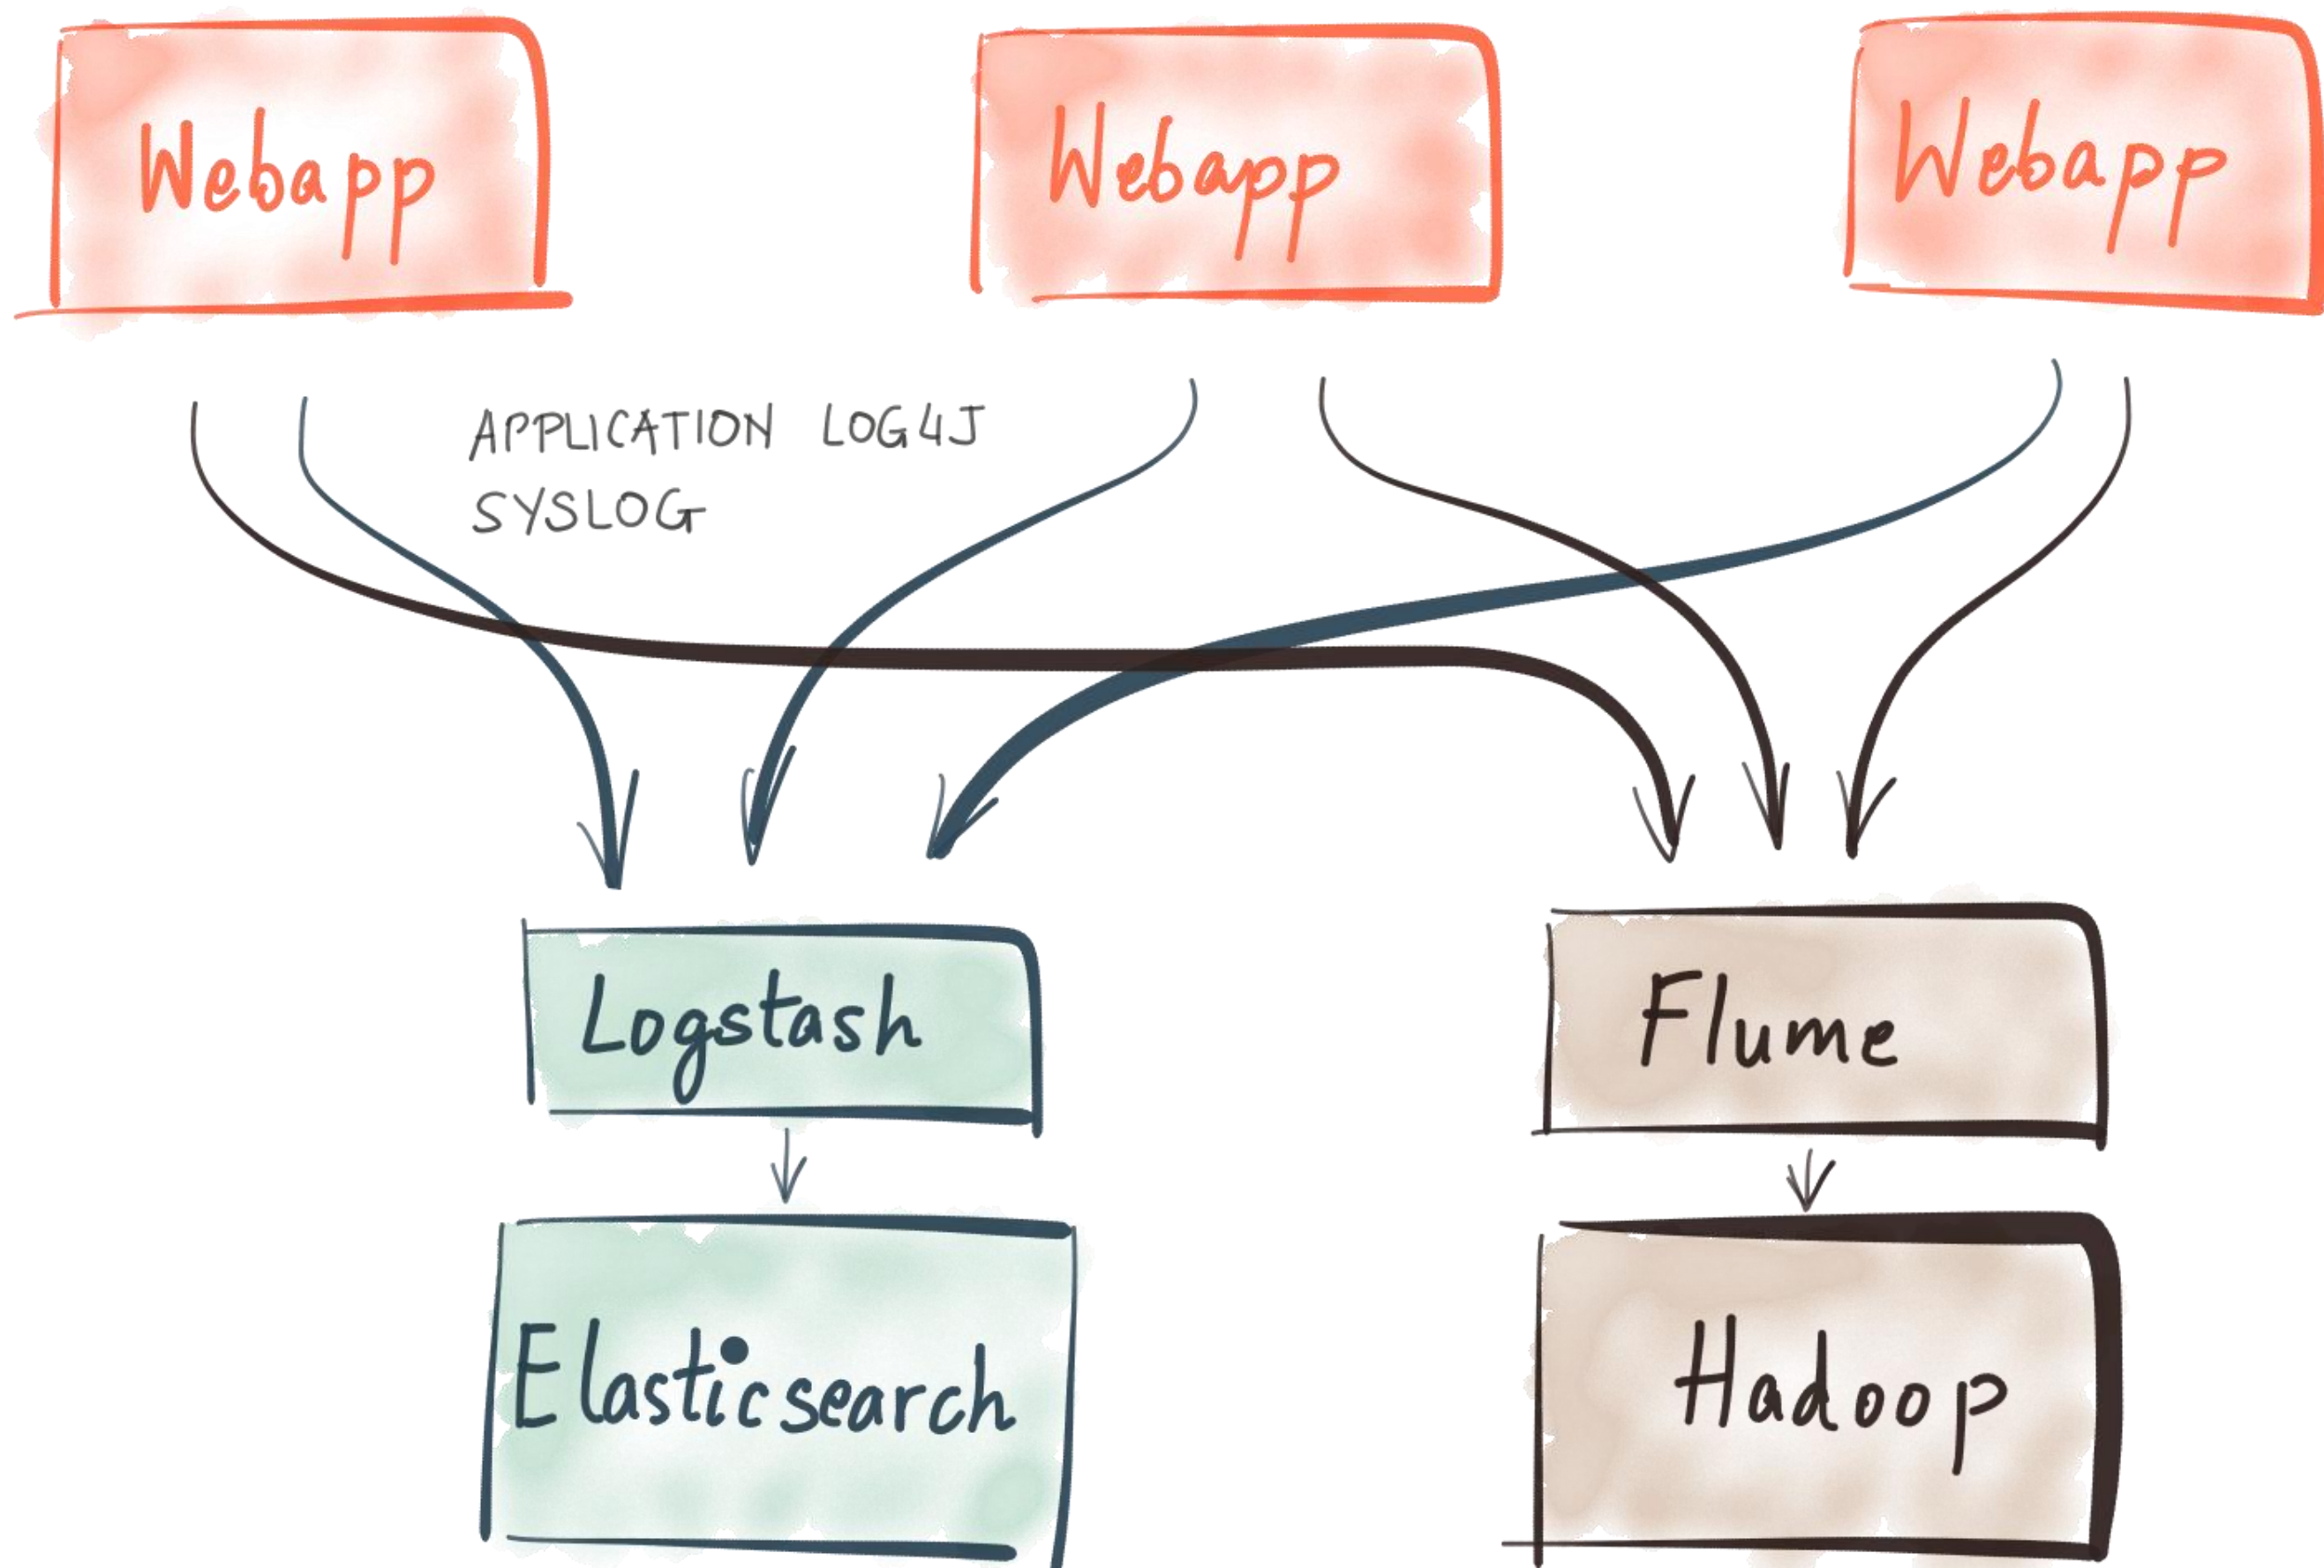


# FOR LOGGING



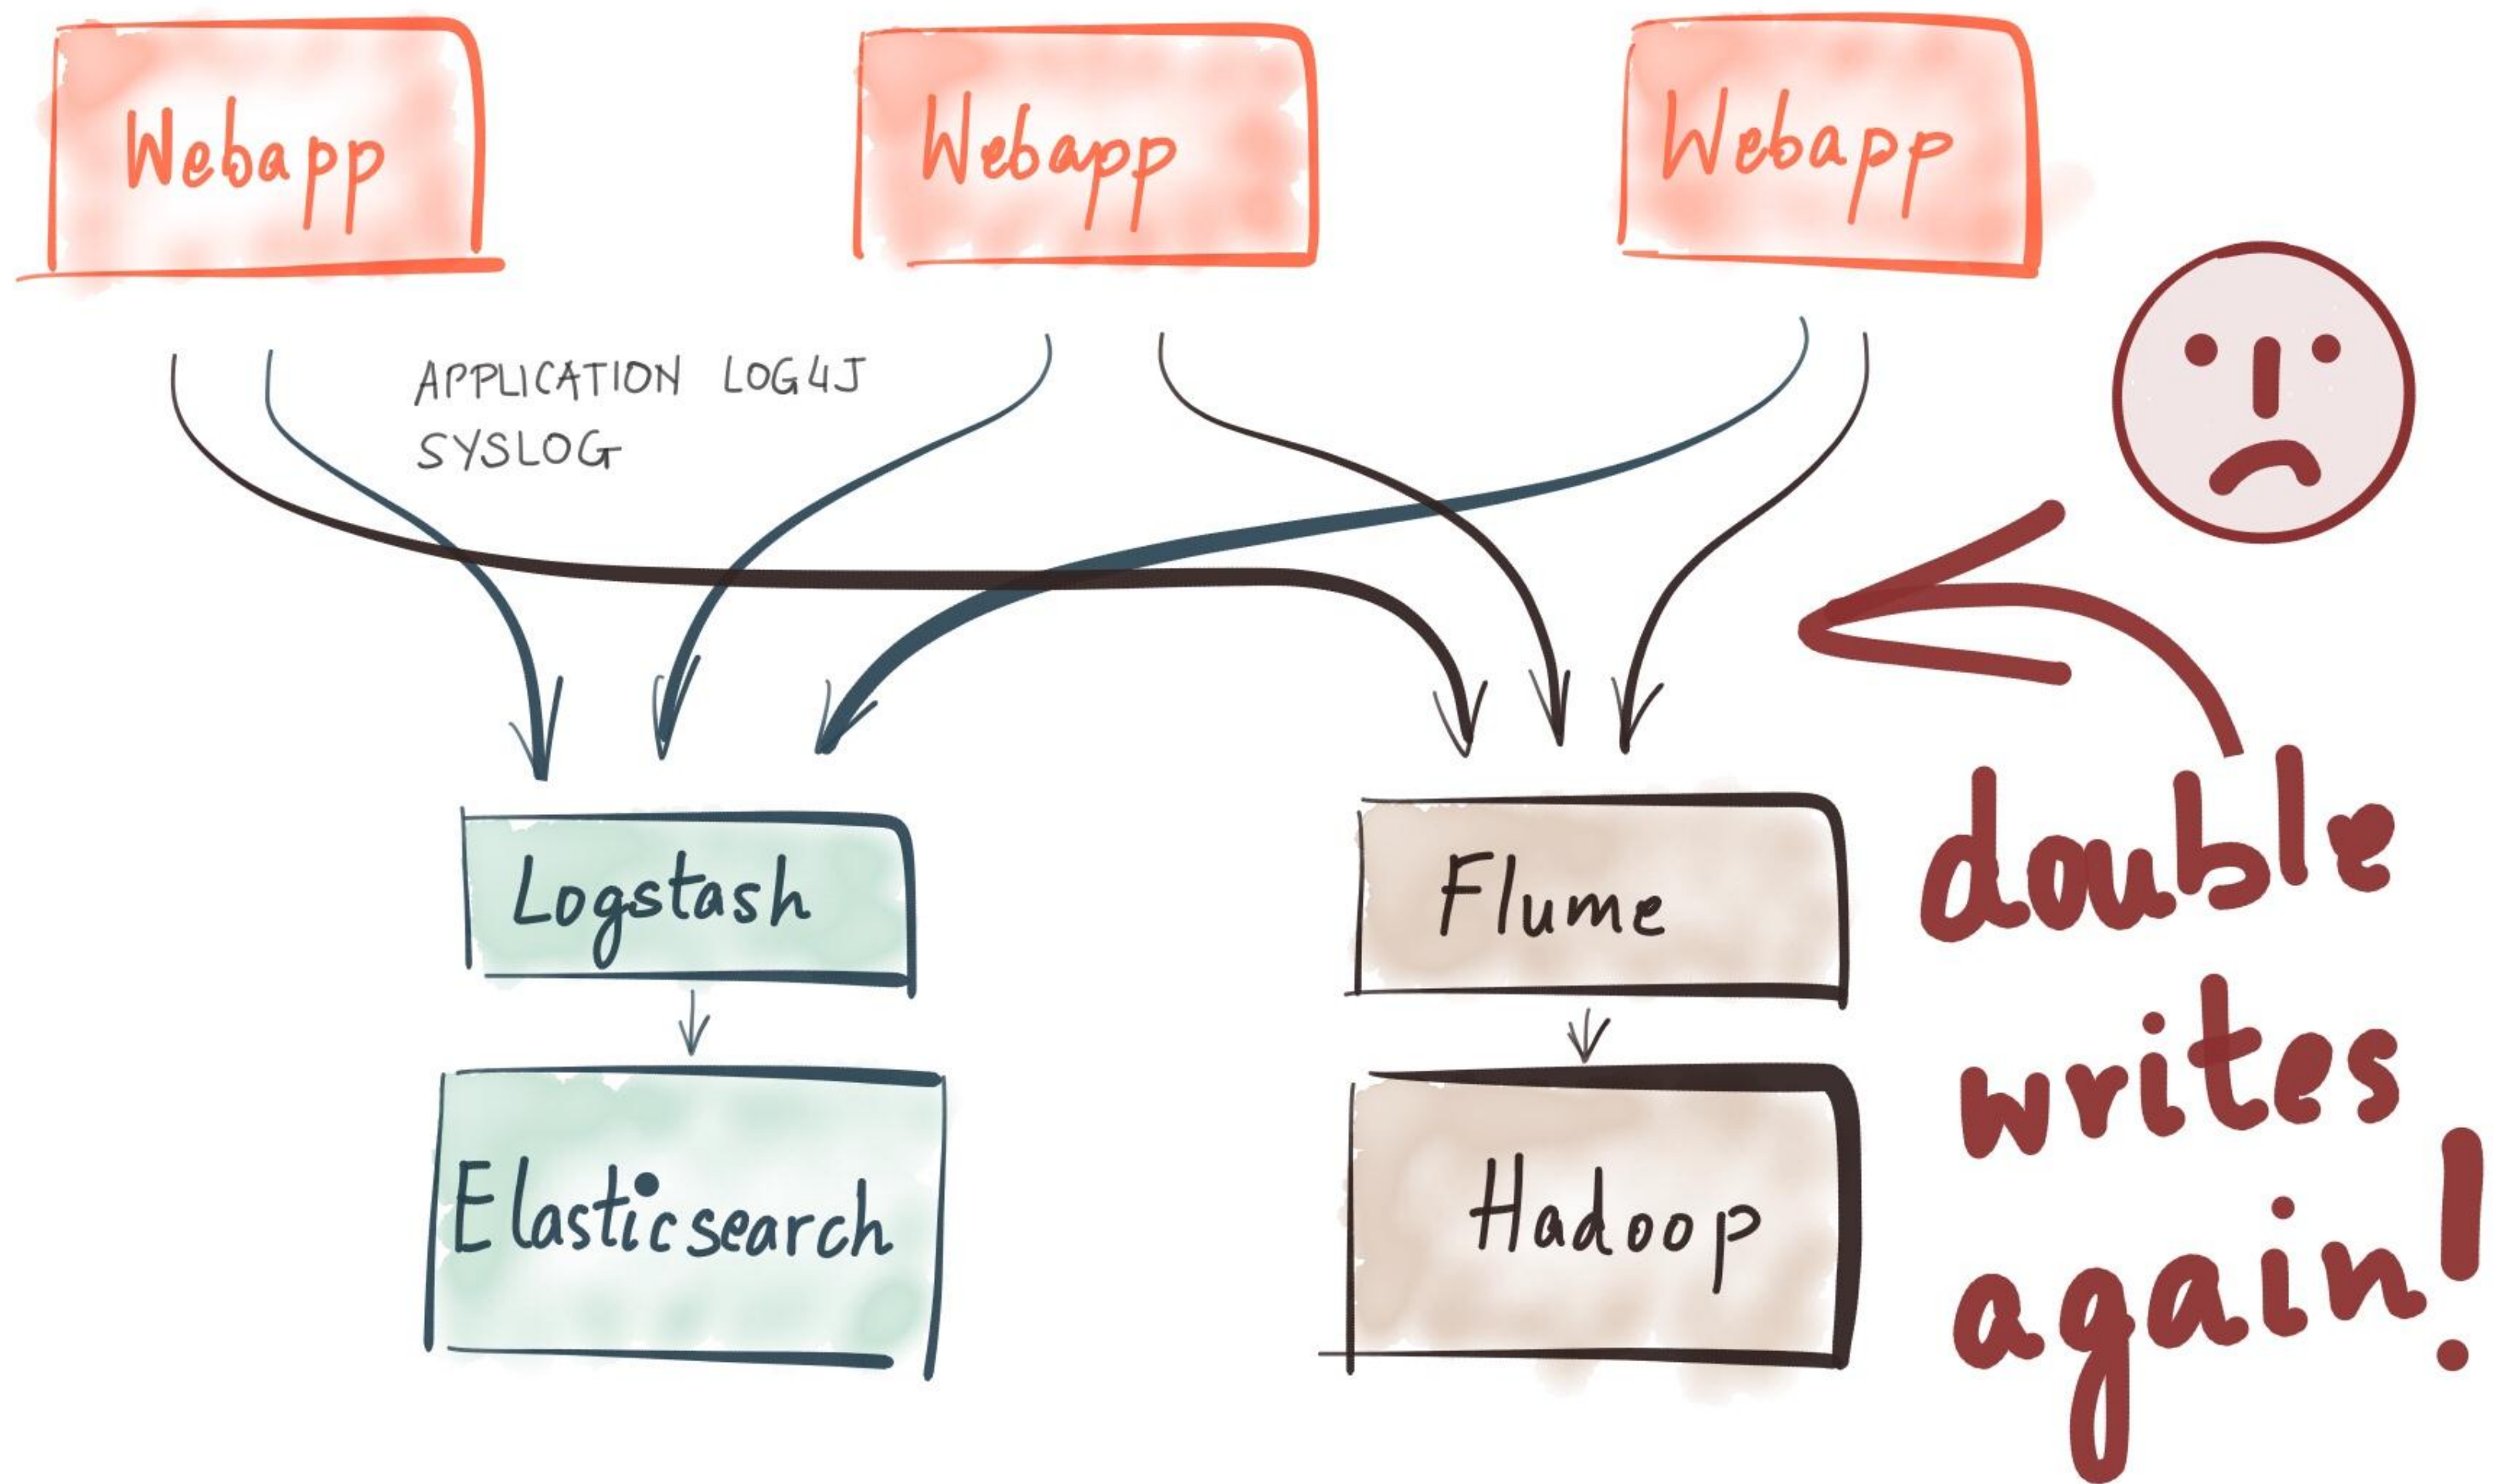


# FOR LOGGING





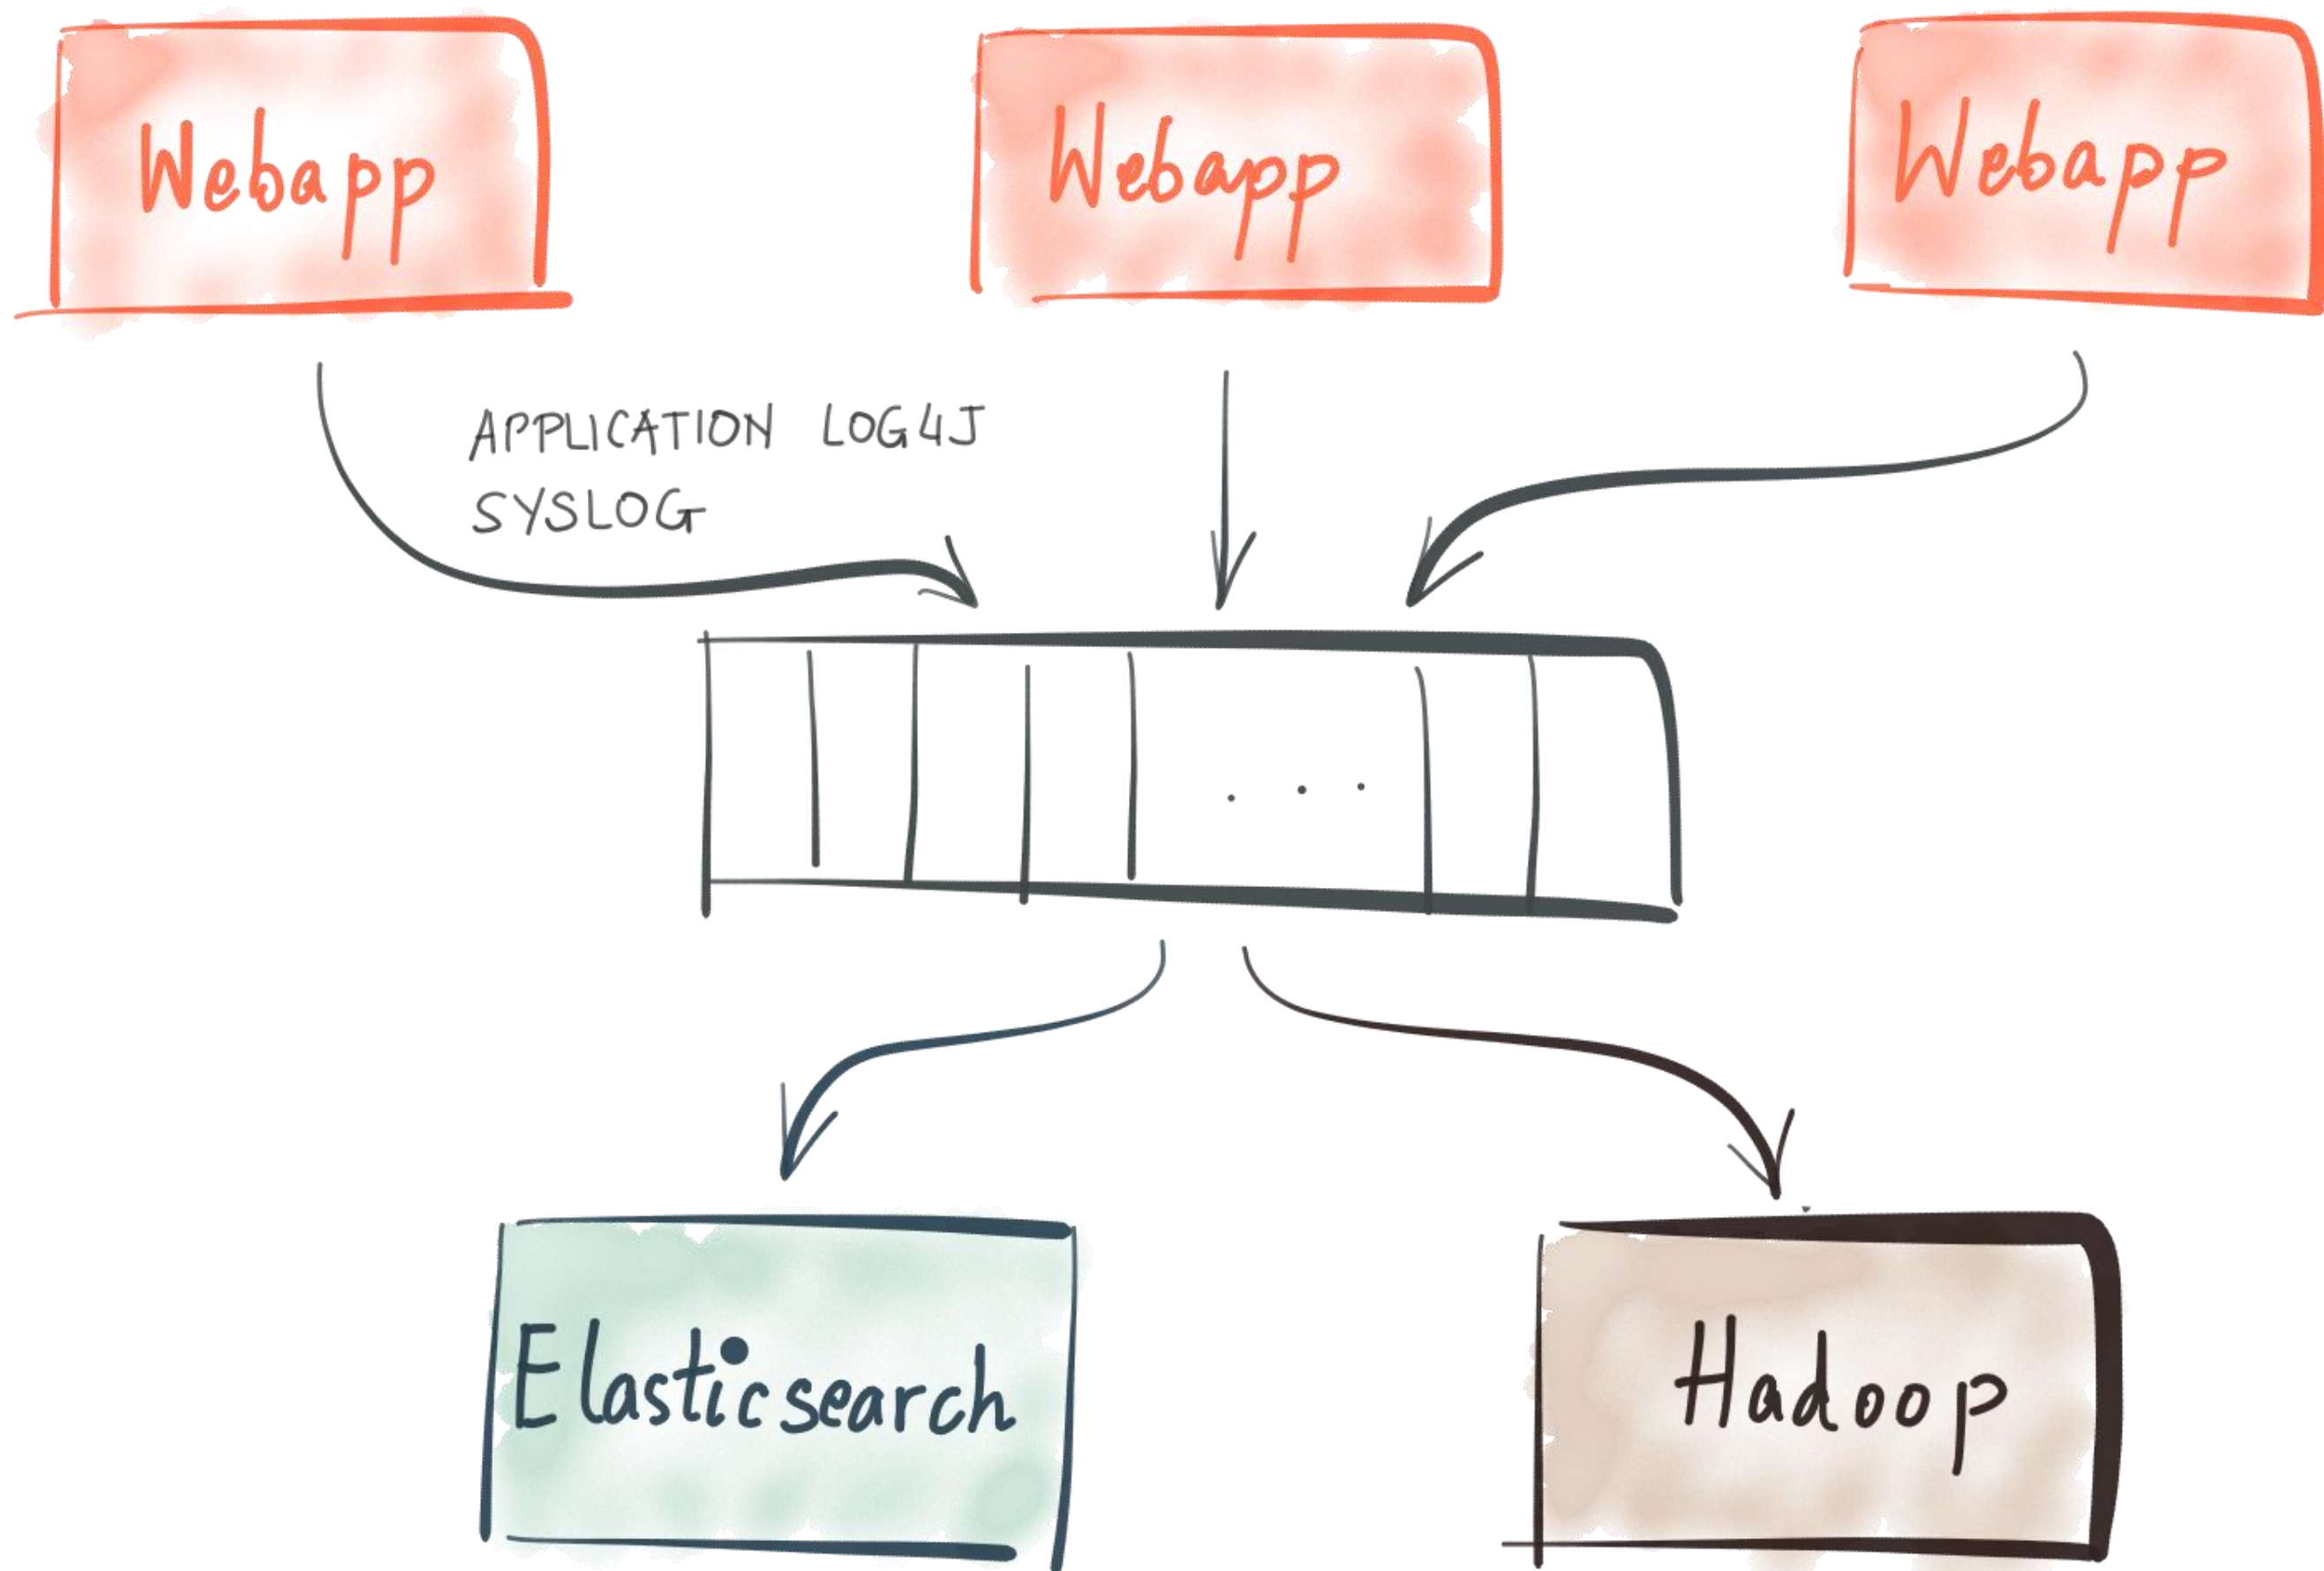
# FOR LOGGING

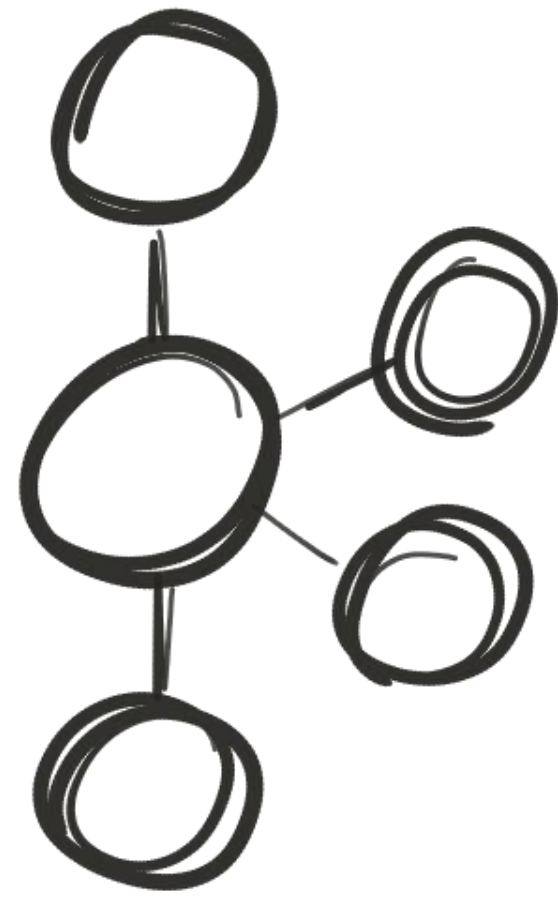




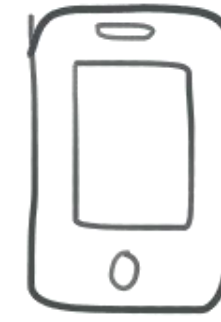


# FOR LOGGING

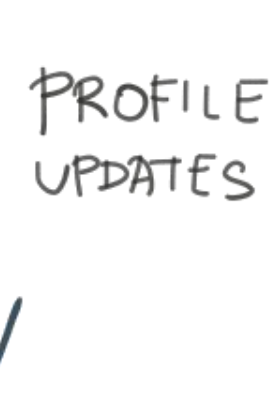




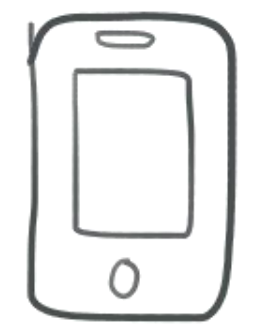
Data Denormalization



Profile Webapp



HBASE /  
CASSANDRA /  
RDBMS



Profile Webapp

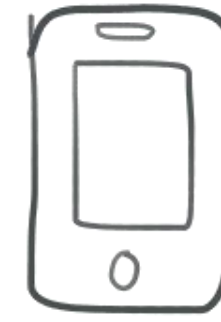


MEMCACHED

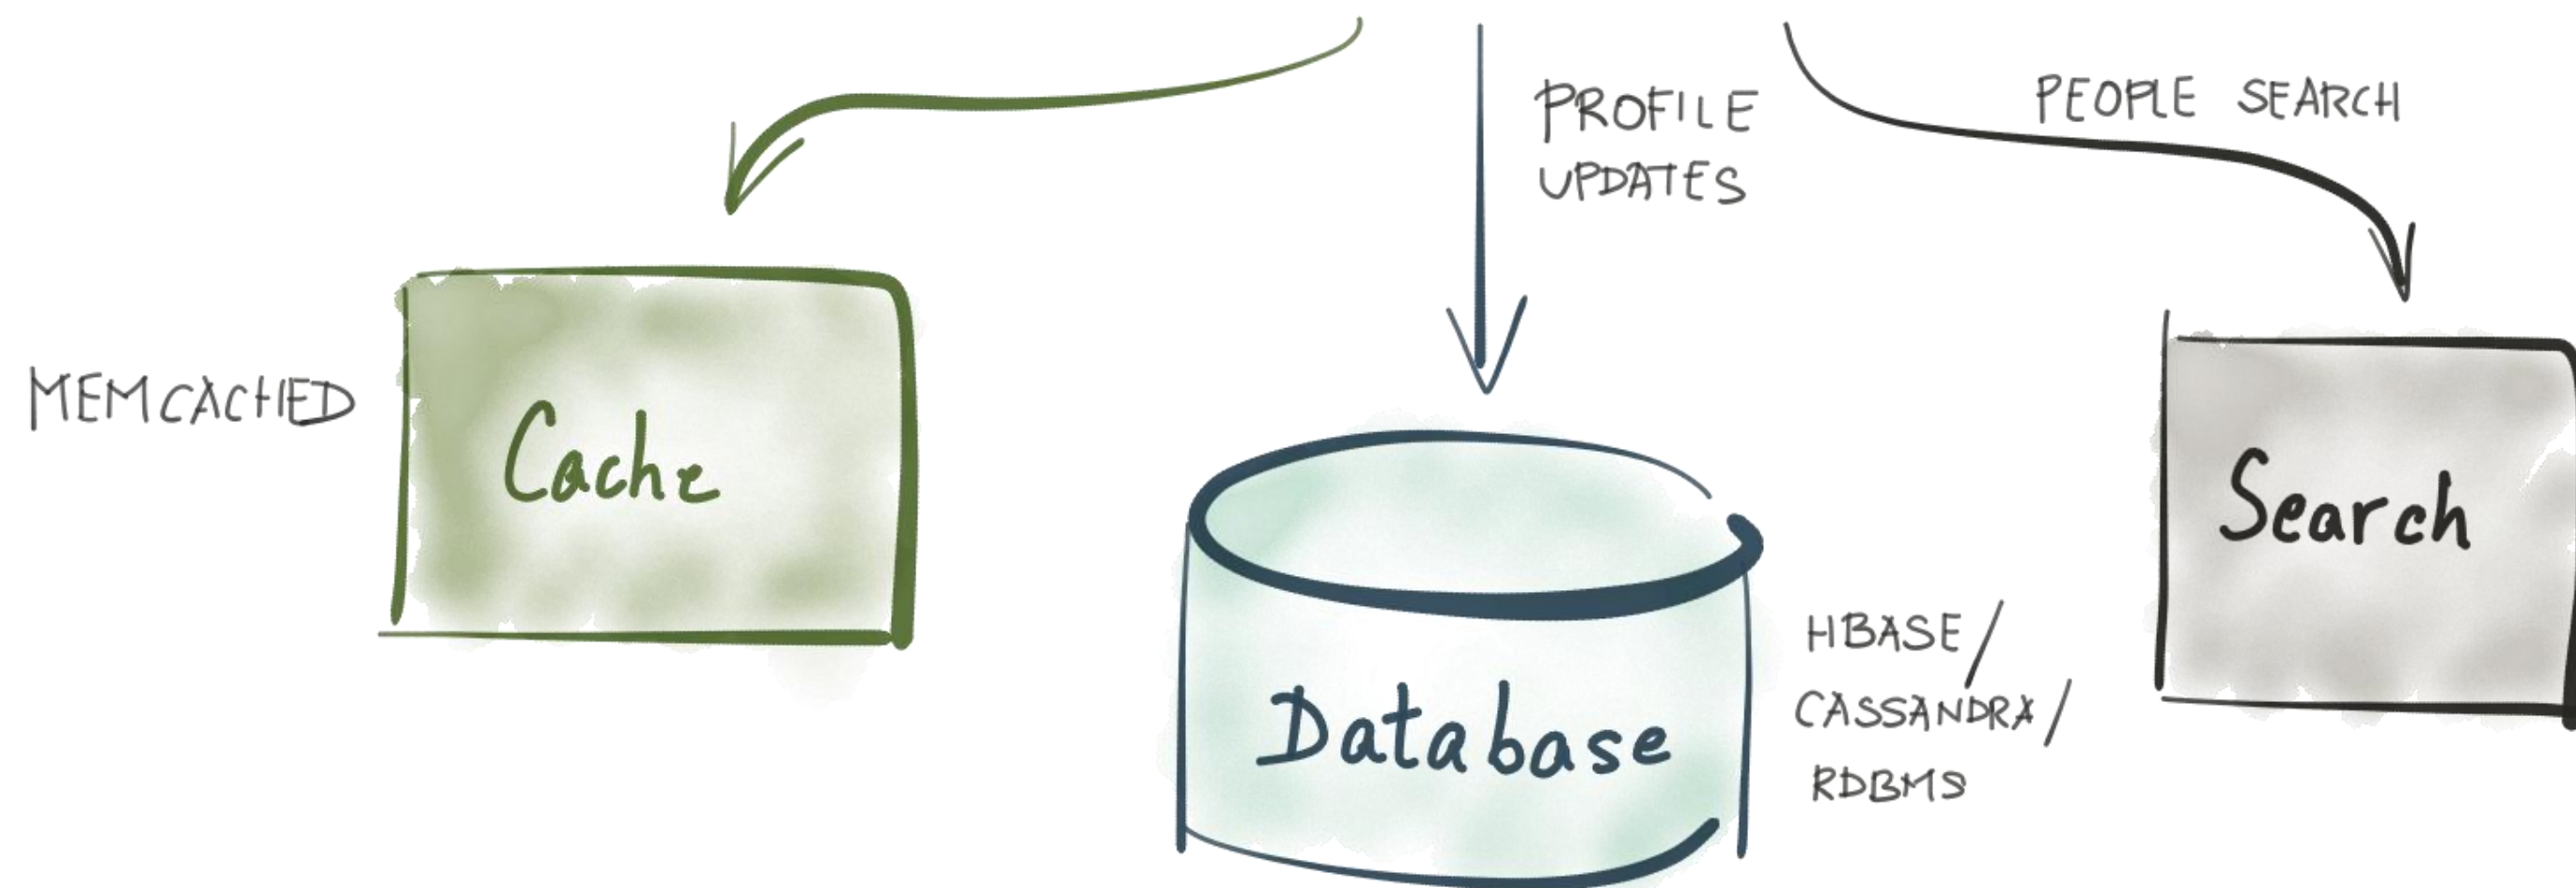
Cache

Database

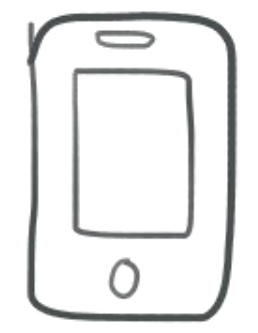
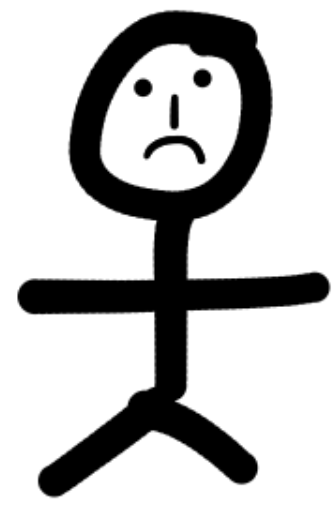
HBASE /  
CASSANDRA /  
RDBMS



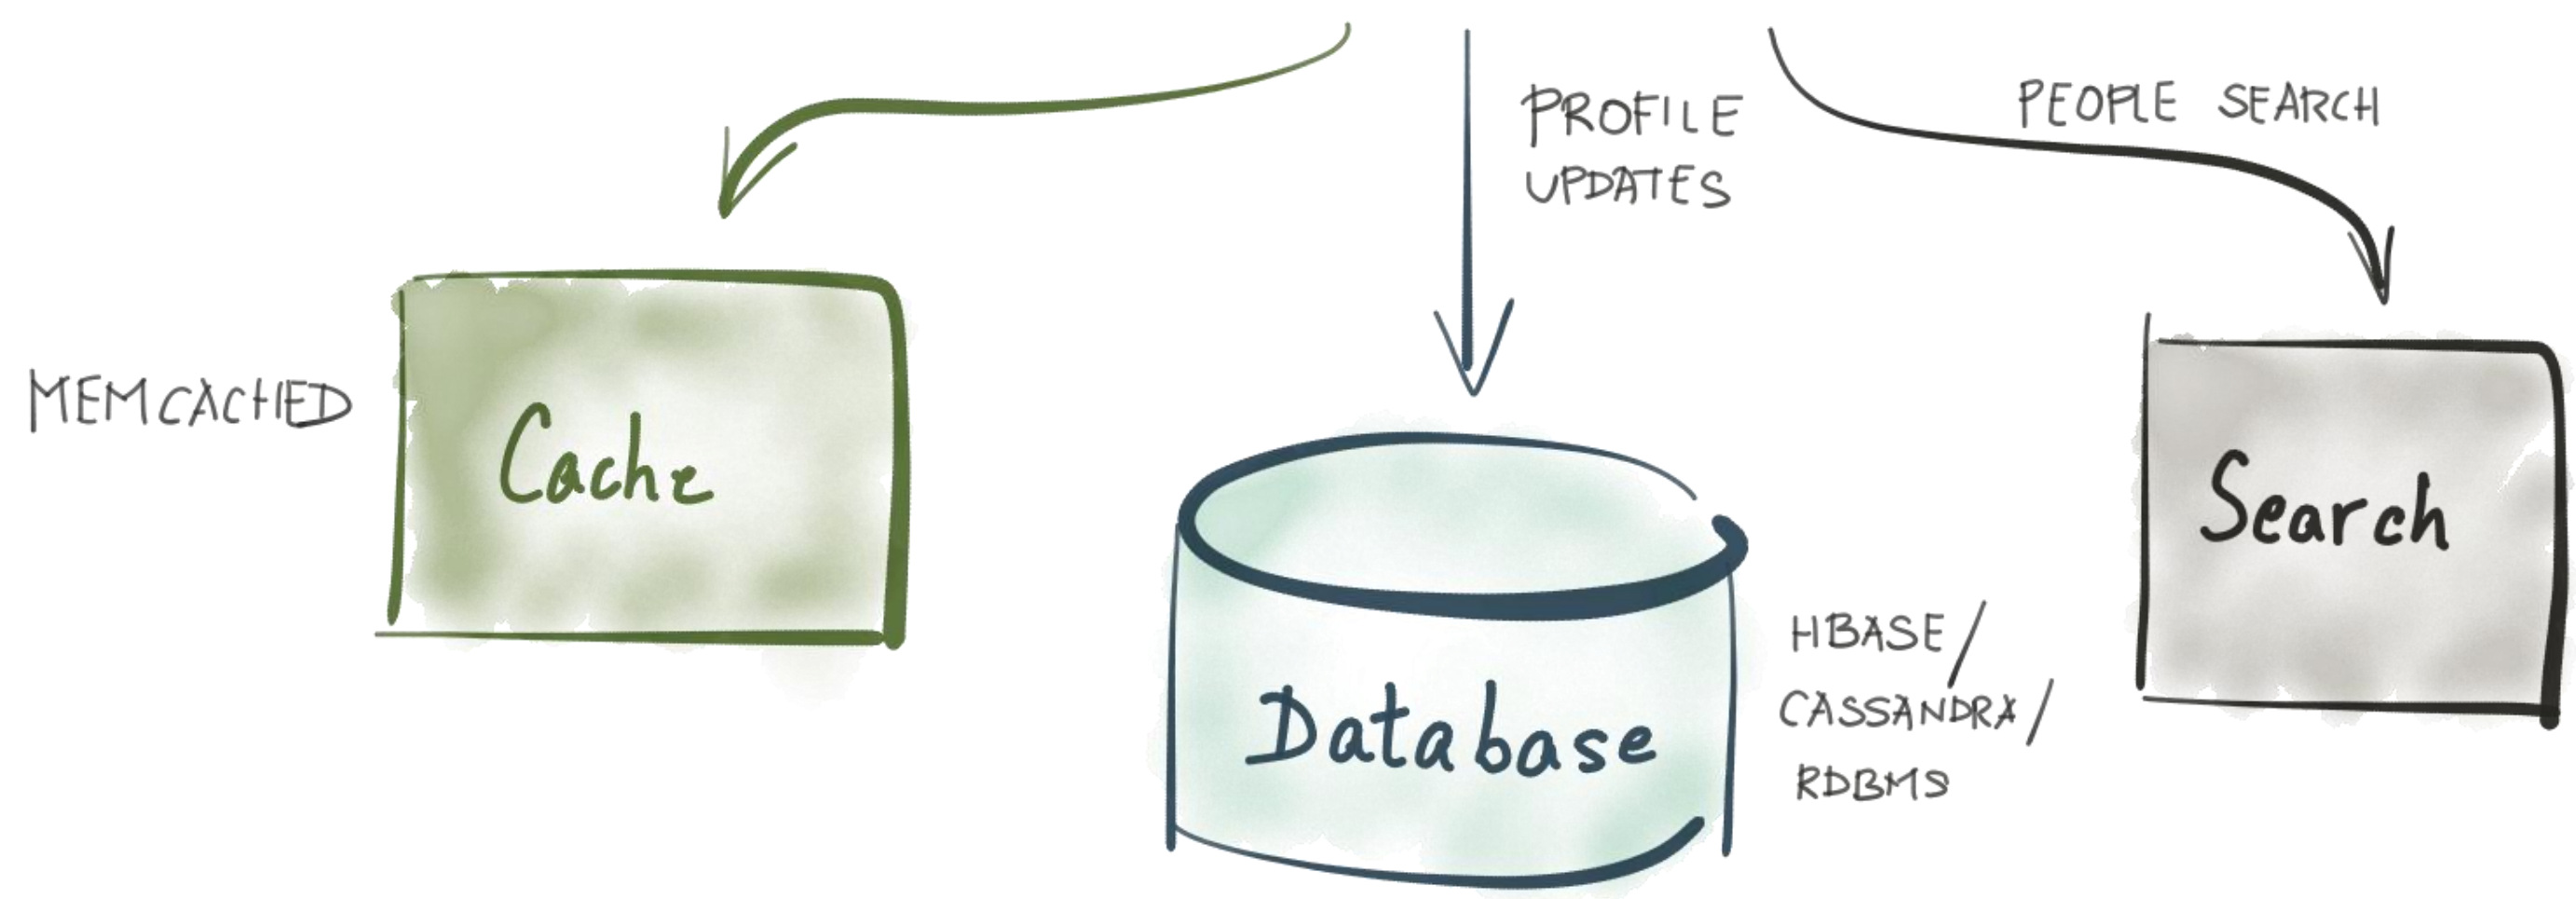
Profile Webapp

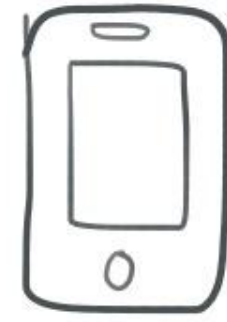
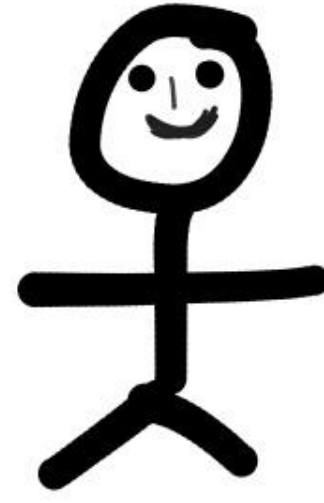


App is slow!



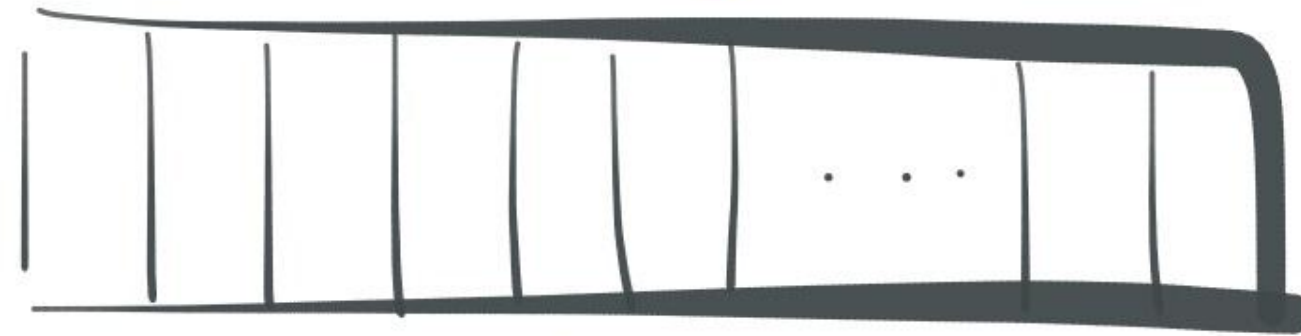
Profile Webapp





Profile Webapp

WRITE ONCE



MEMCACHED

Cache

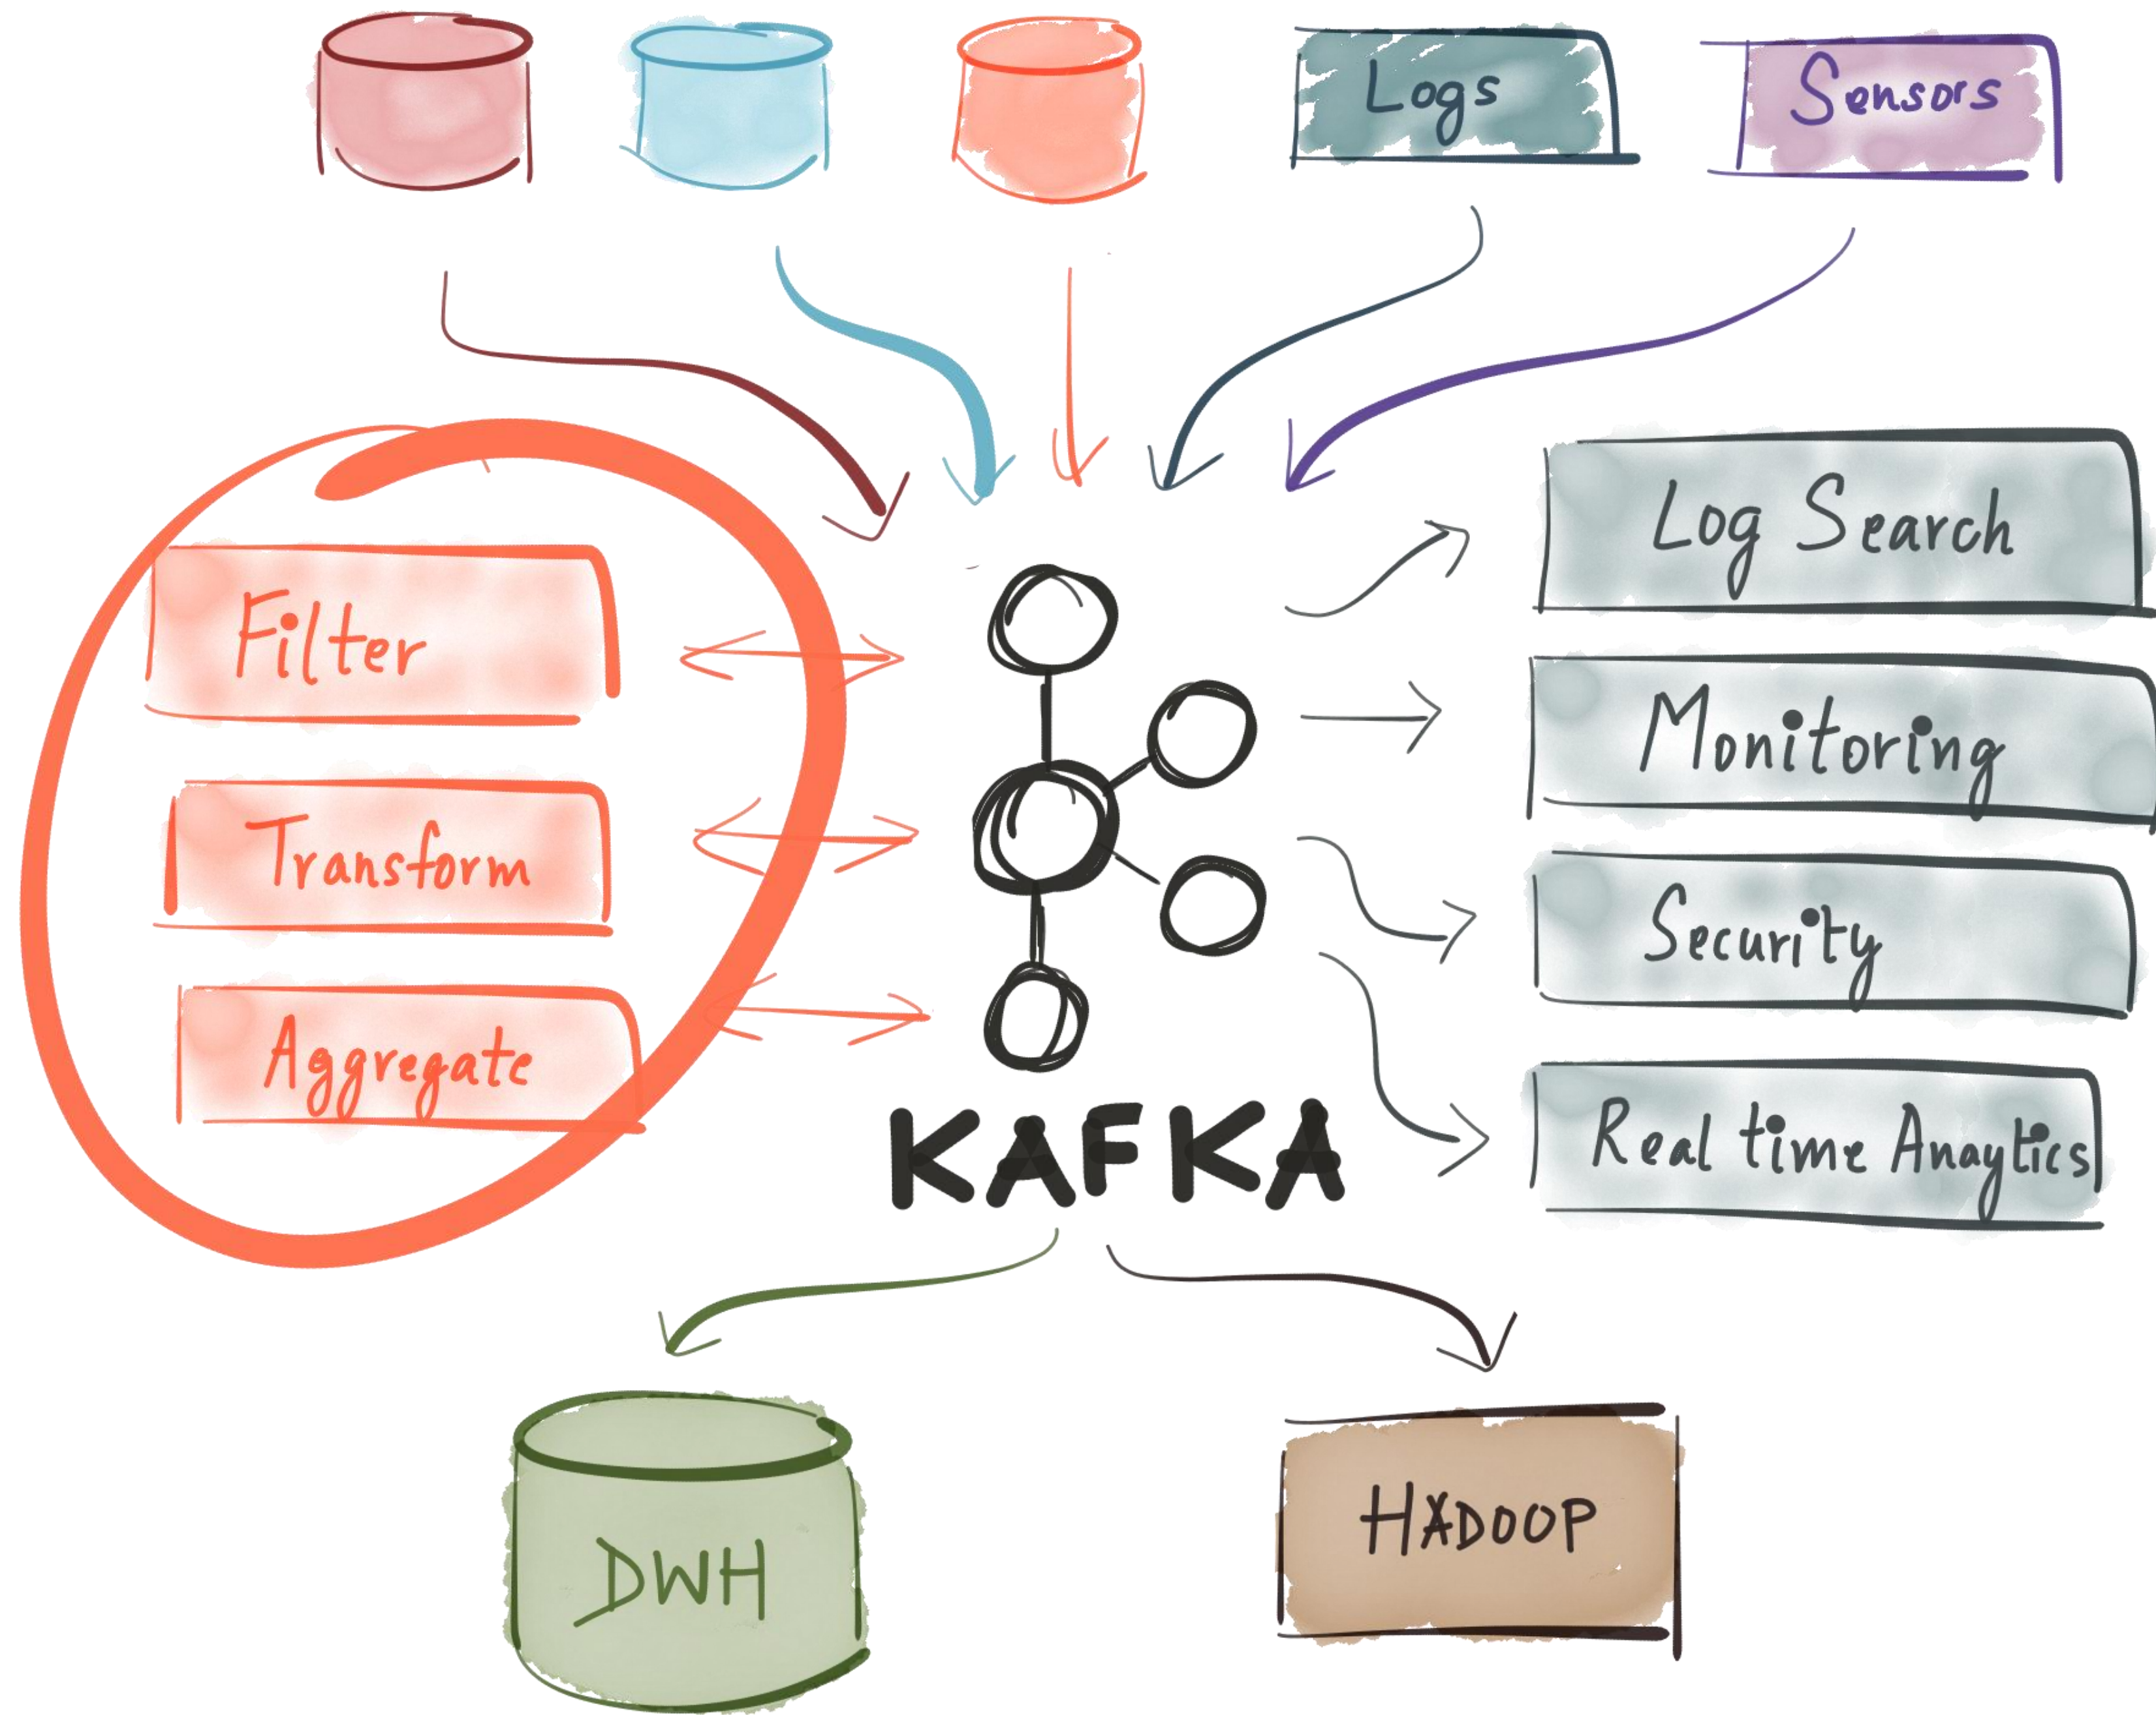
READ MANY TIMES

Database

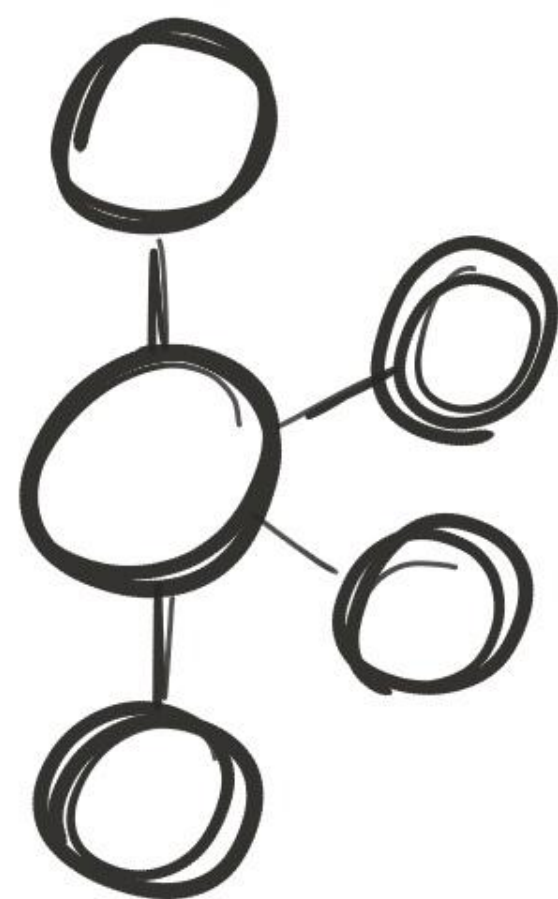
HBASE /  
CASSANDRA /  
RDBMS

Search

ELASTICSEARCH /  
SOLR



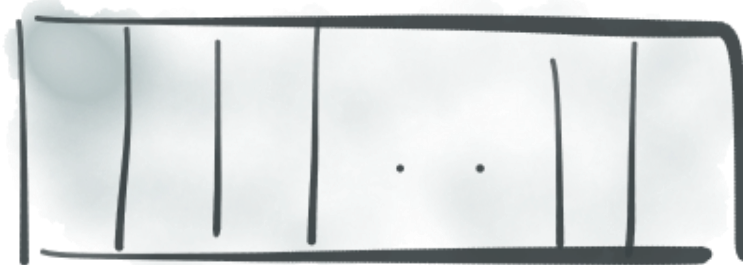




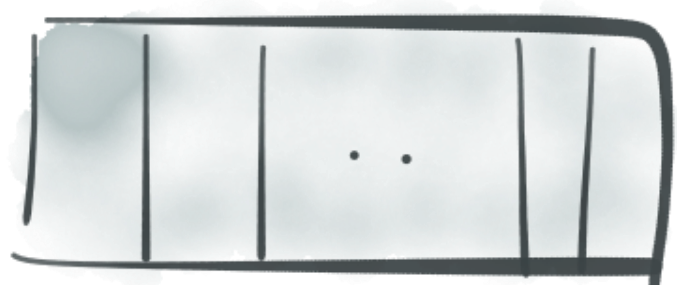
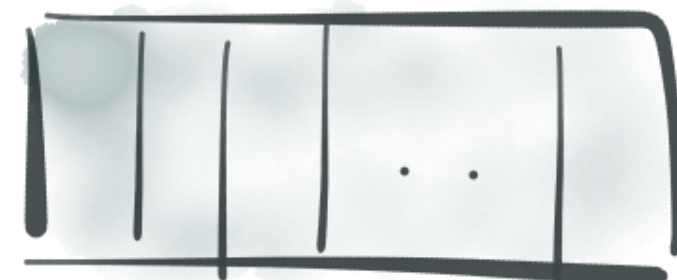
# Stream Processing

# STREAM PROCESSING

Credit Card Txns



Valid Txns

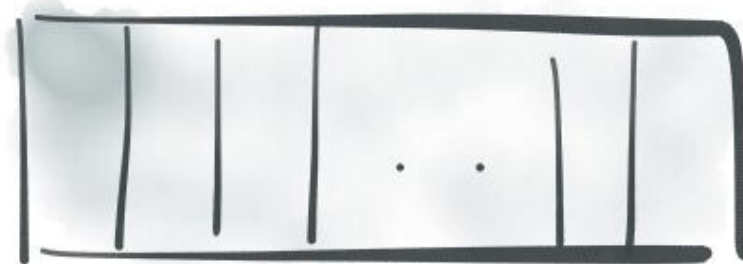


Txns On Hold

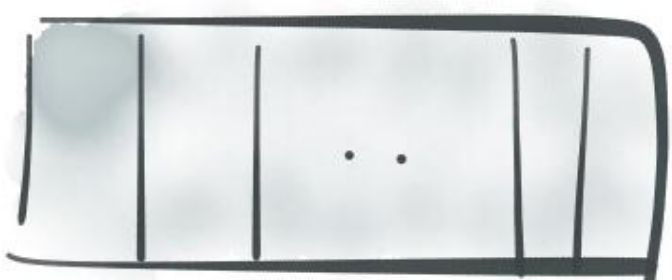
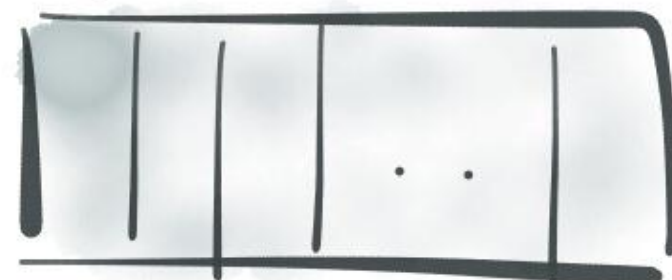


# STREAM PROCESSING

Credit Card Txns



Valid Txns



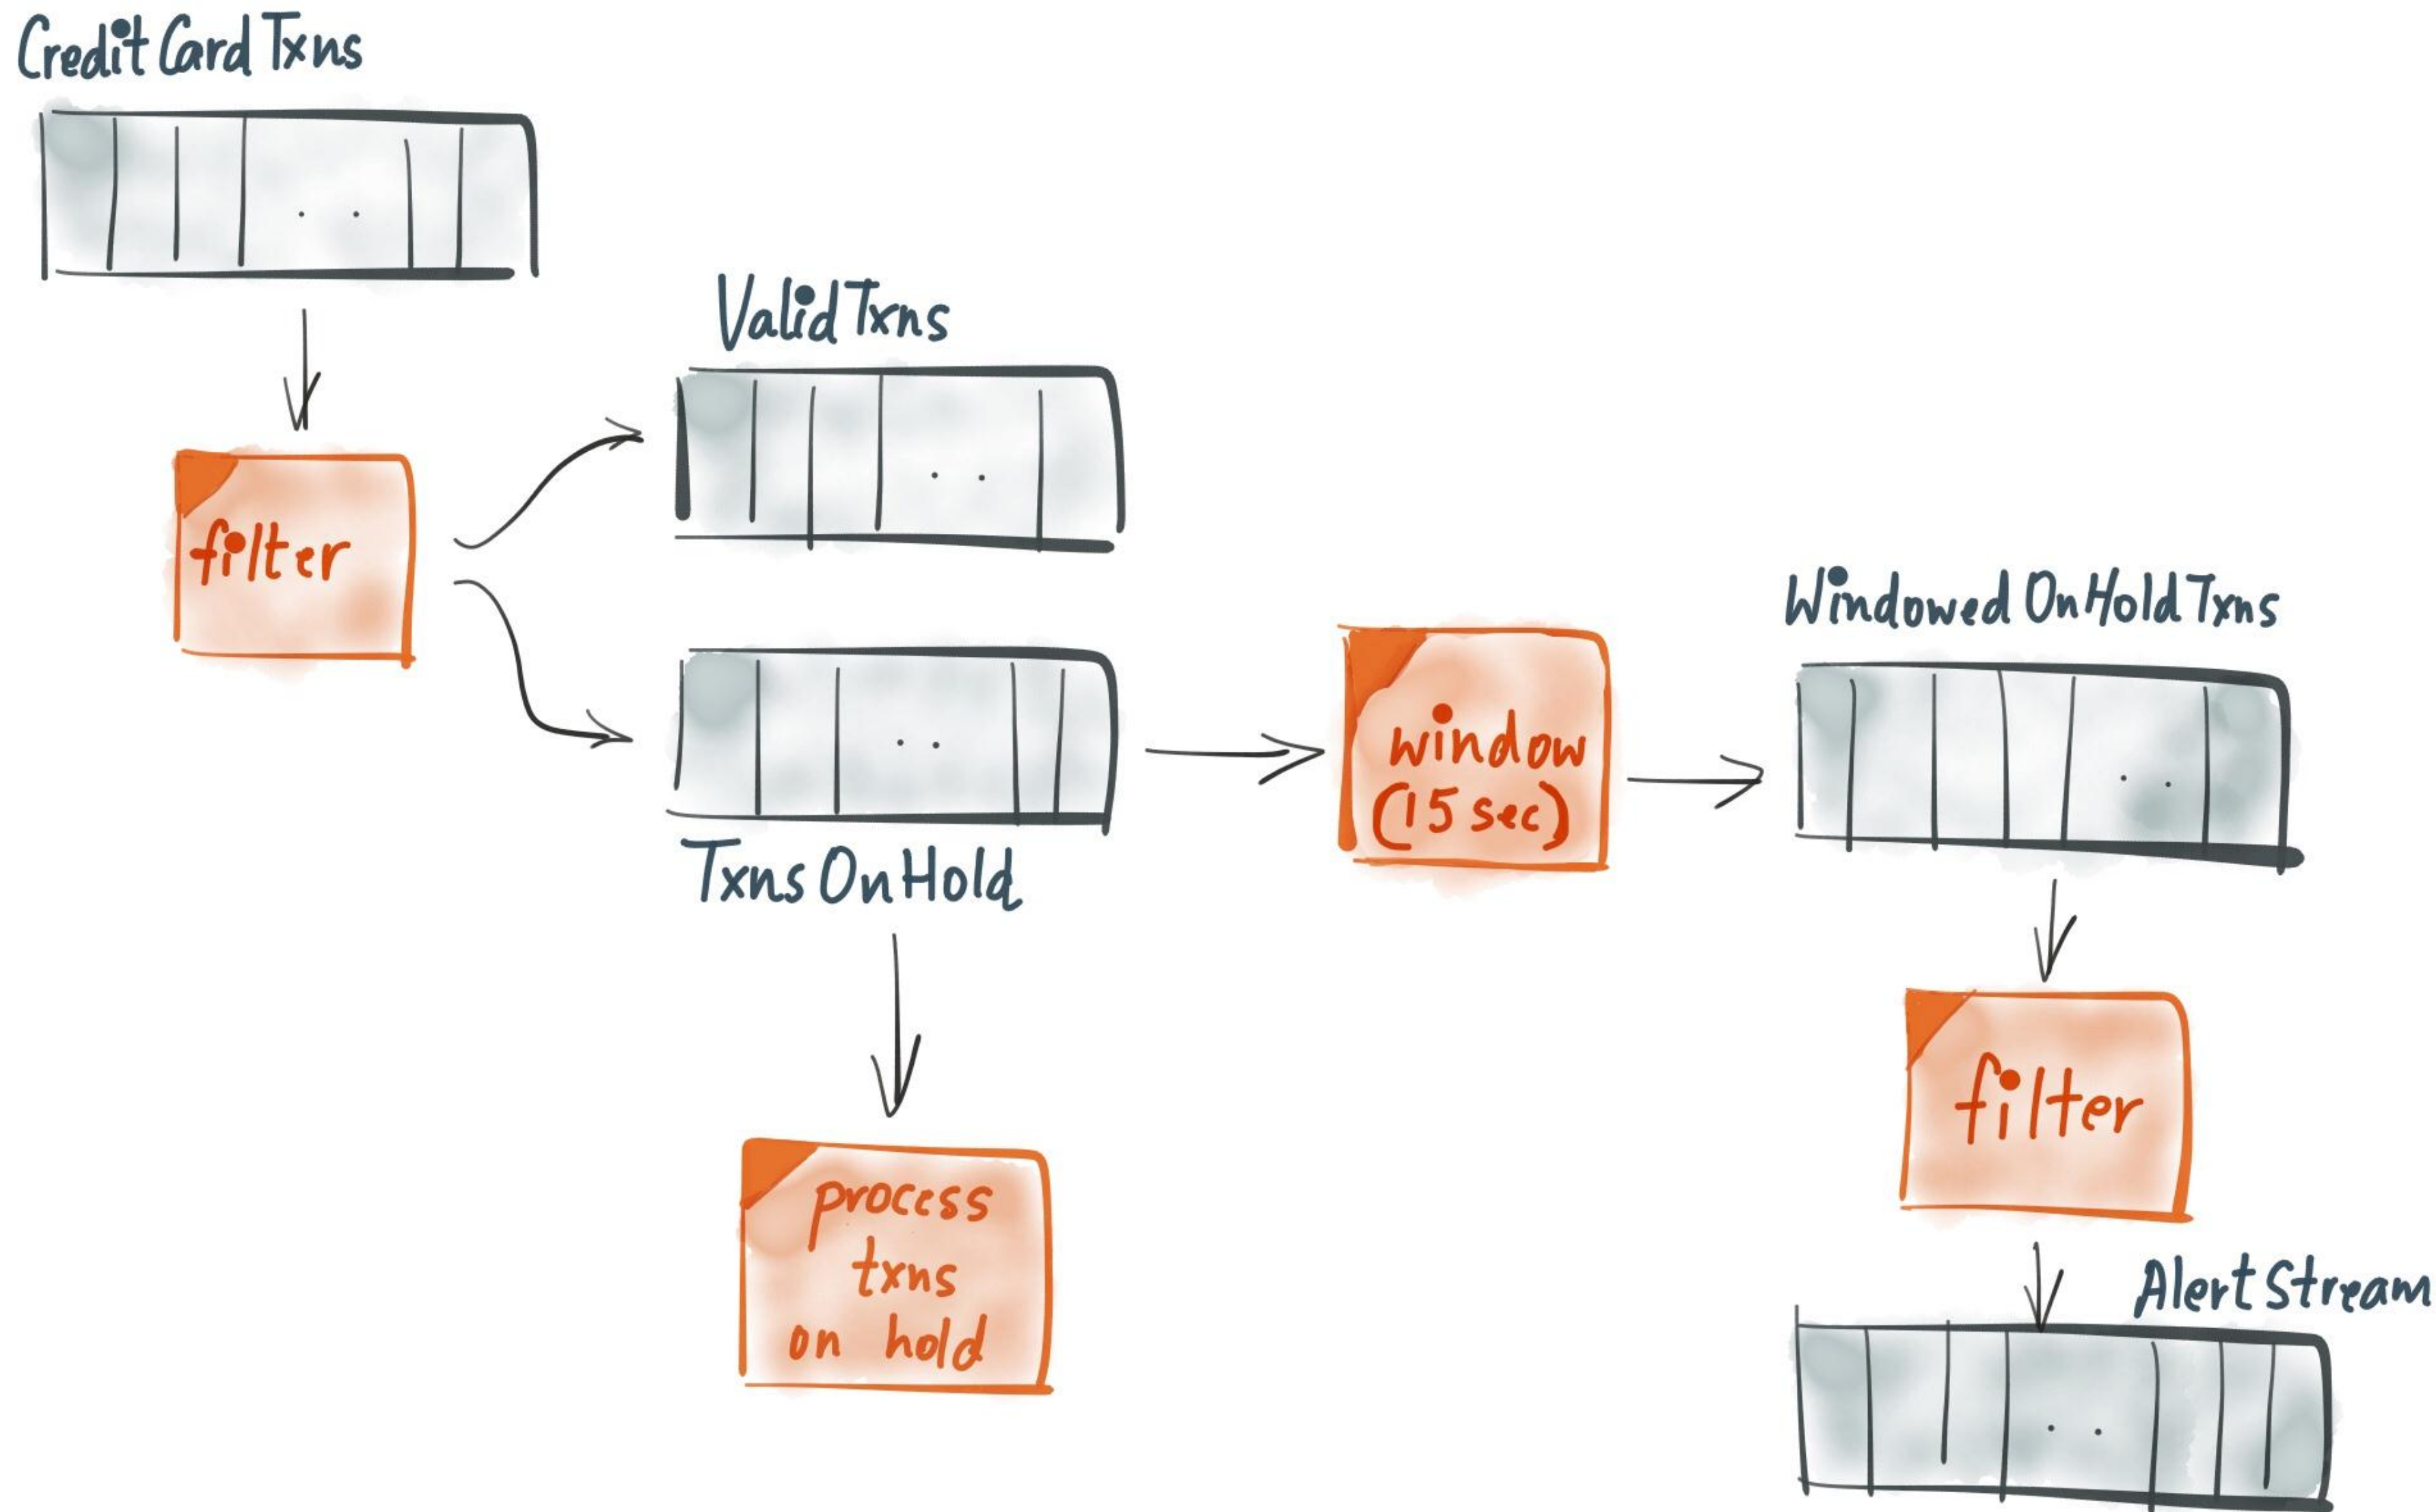
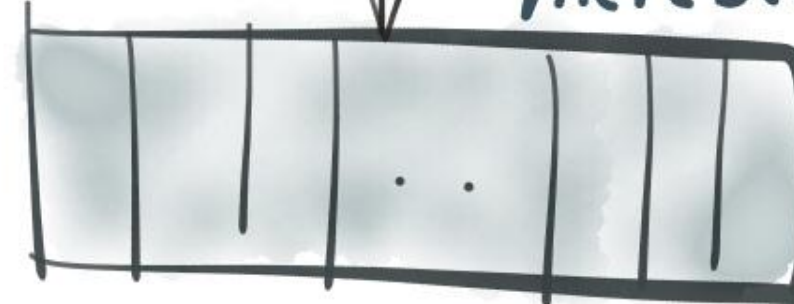
Txns On Hold



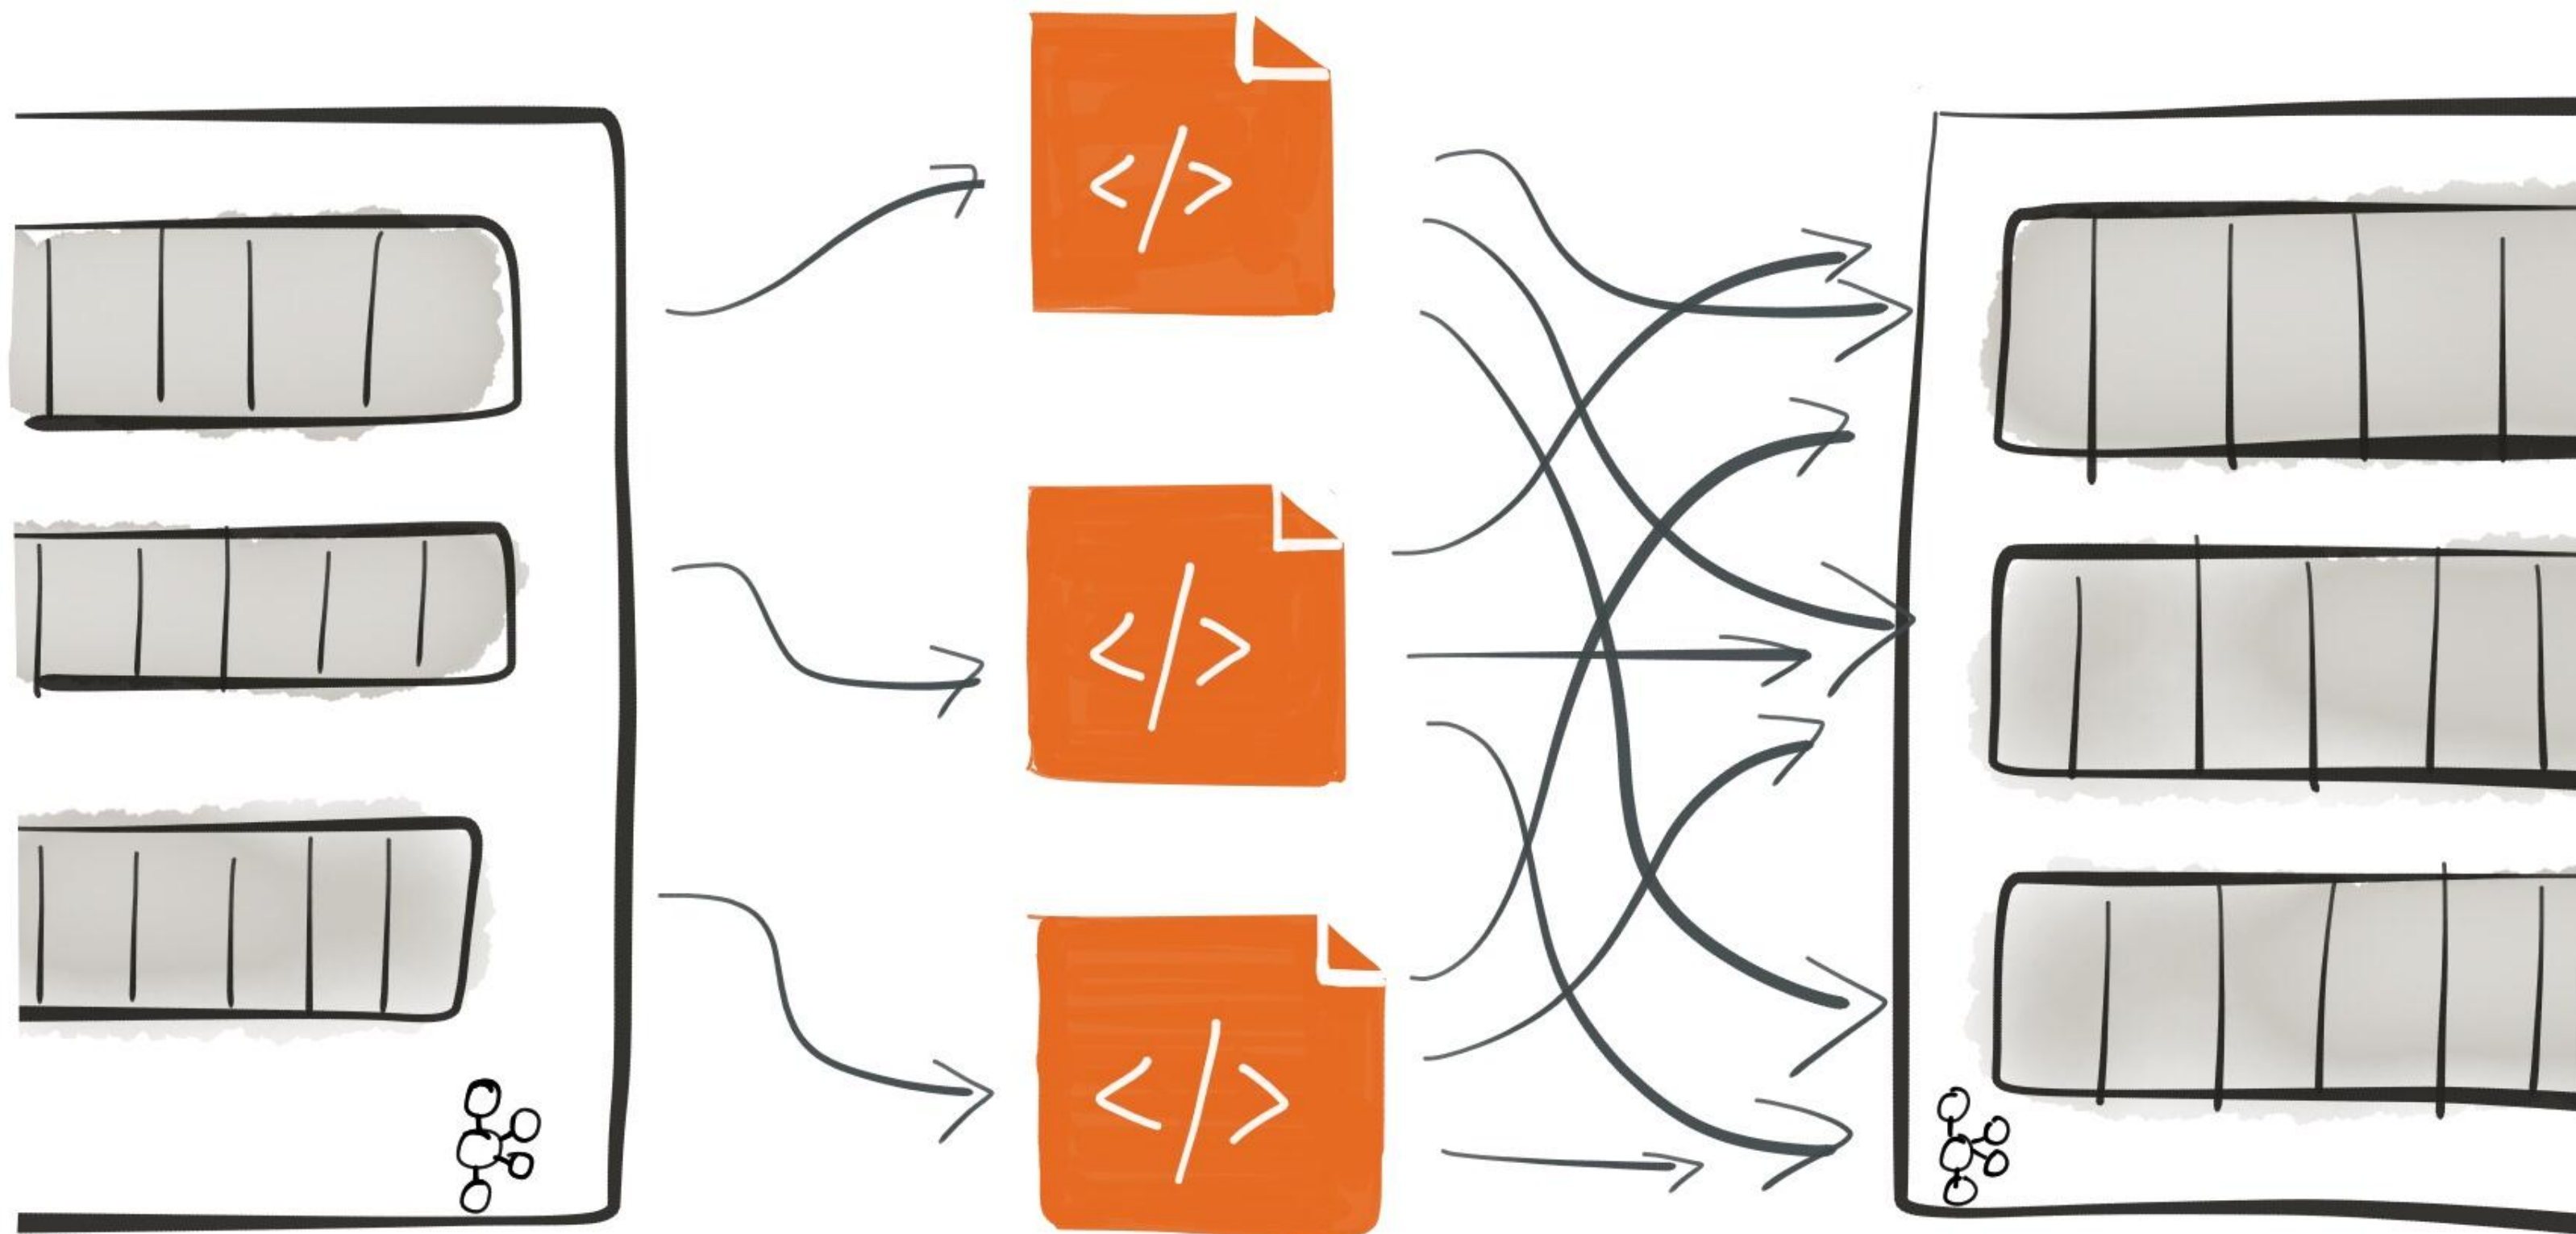
Windowed On Hold Txns



Alert Stream

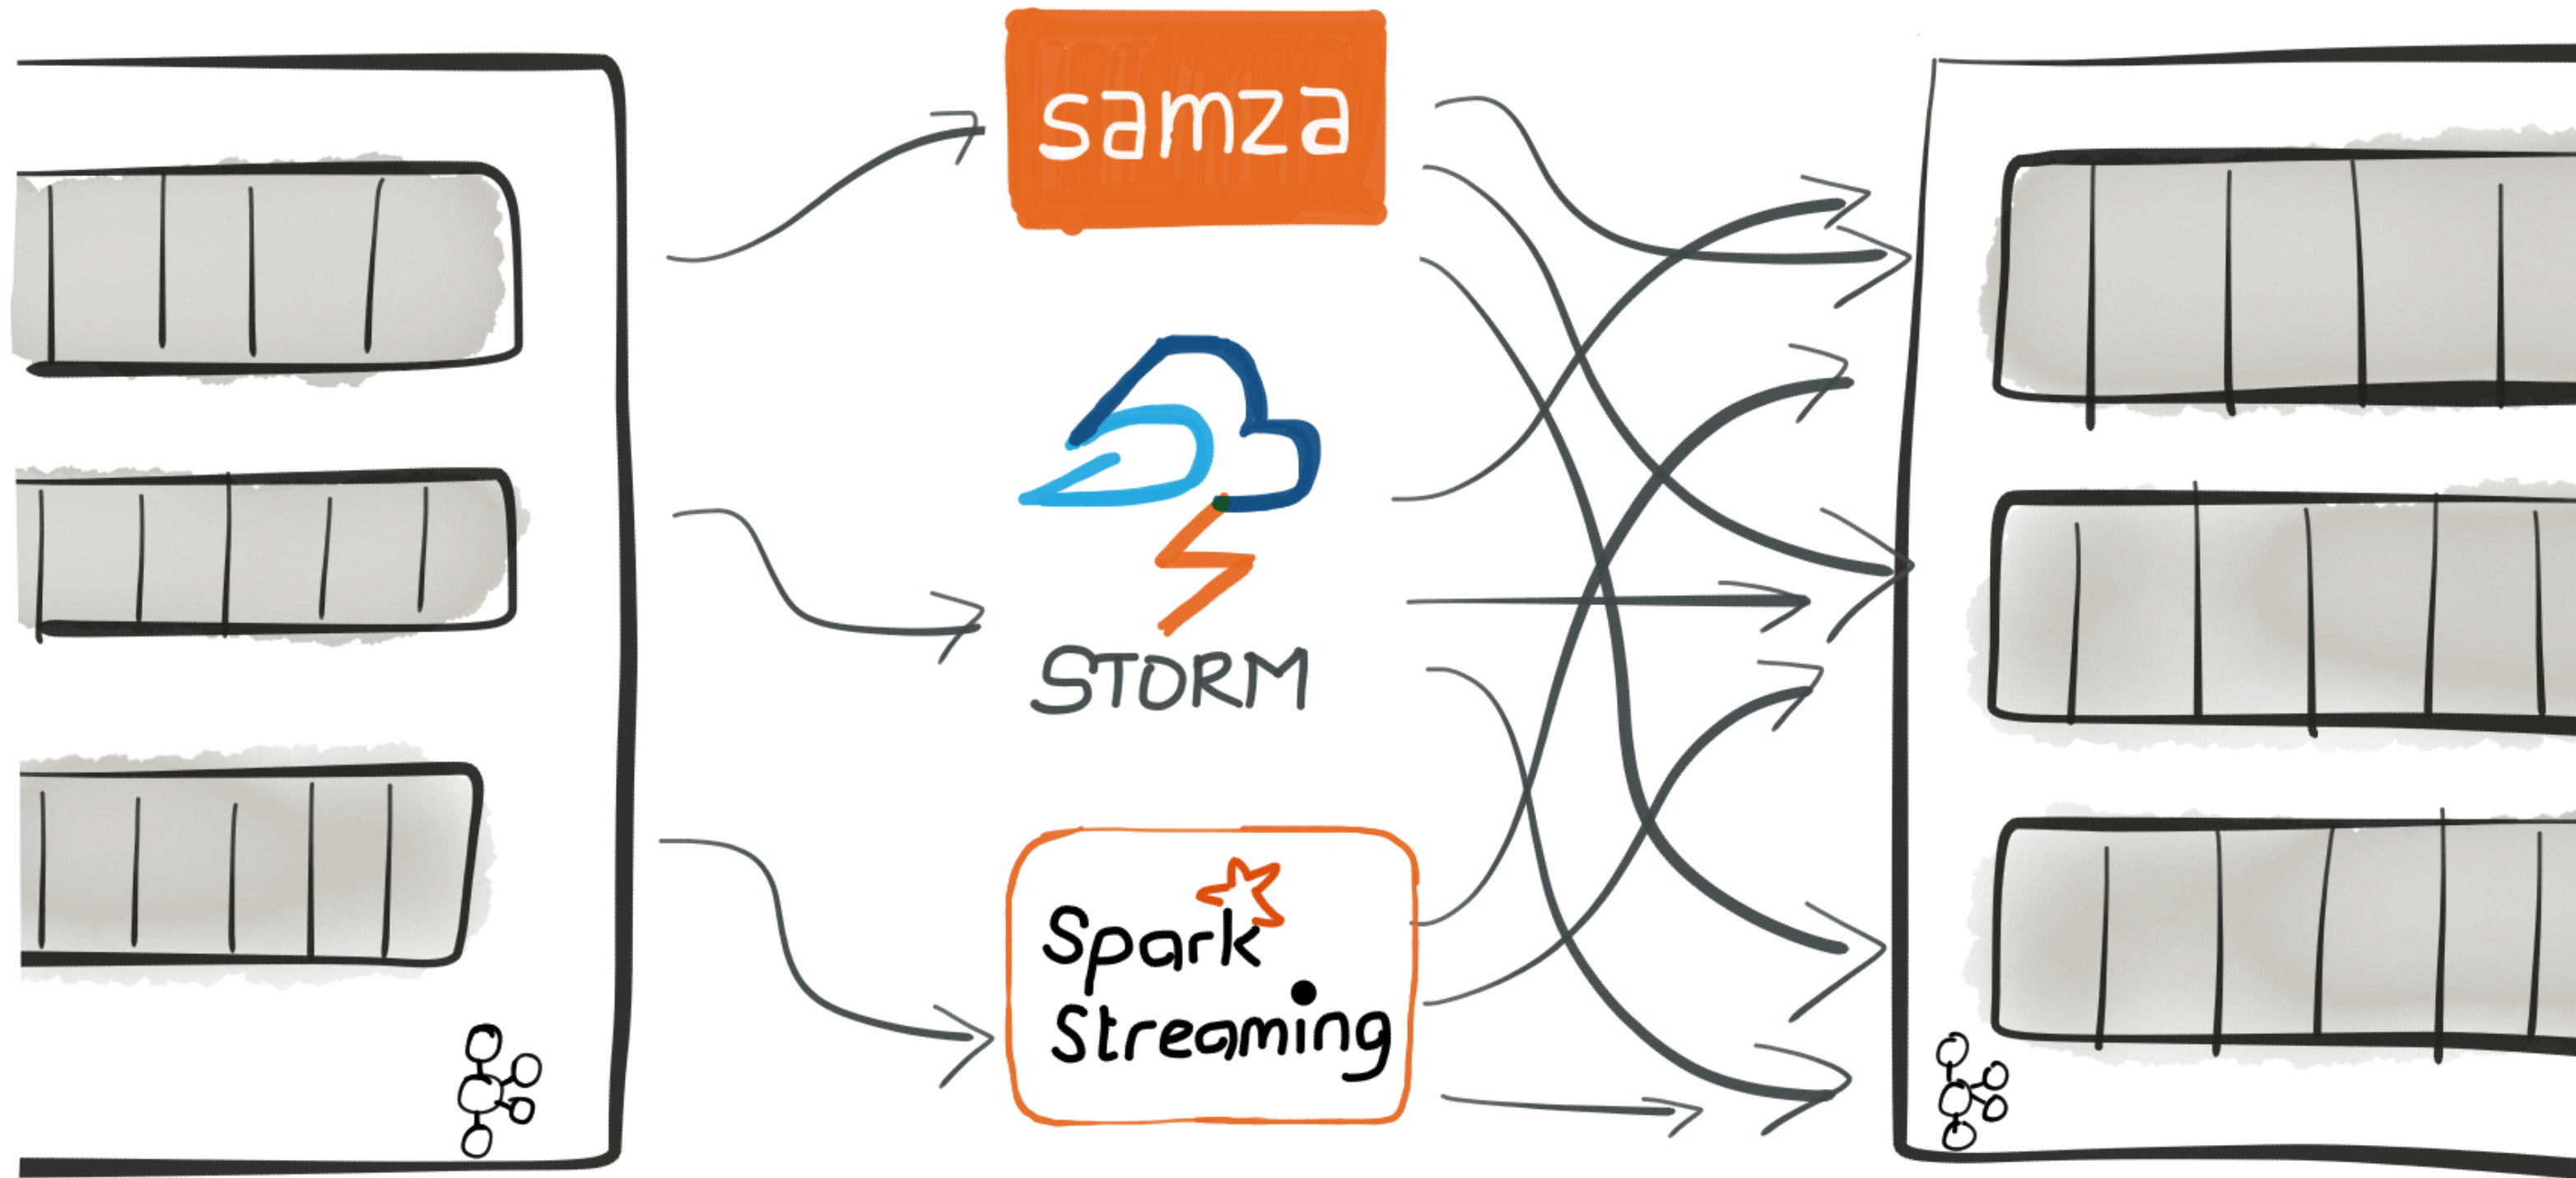


# STREAM PROCESSING



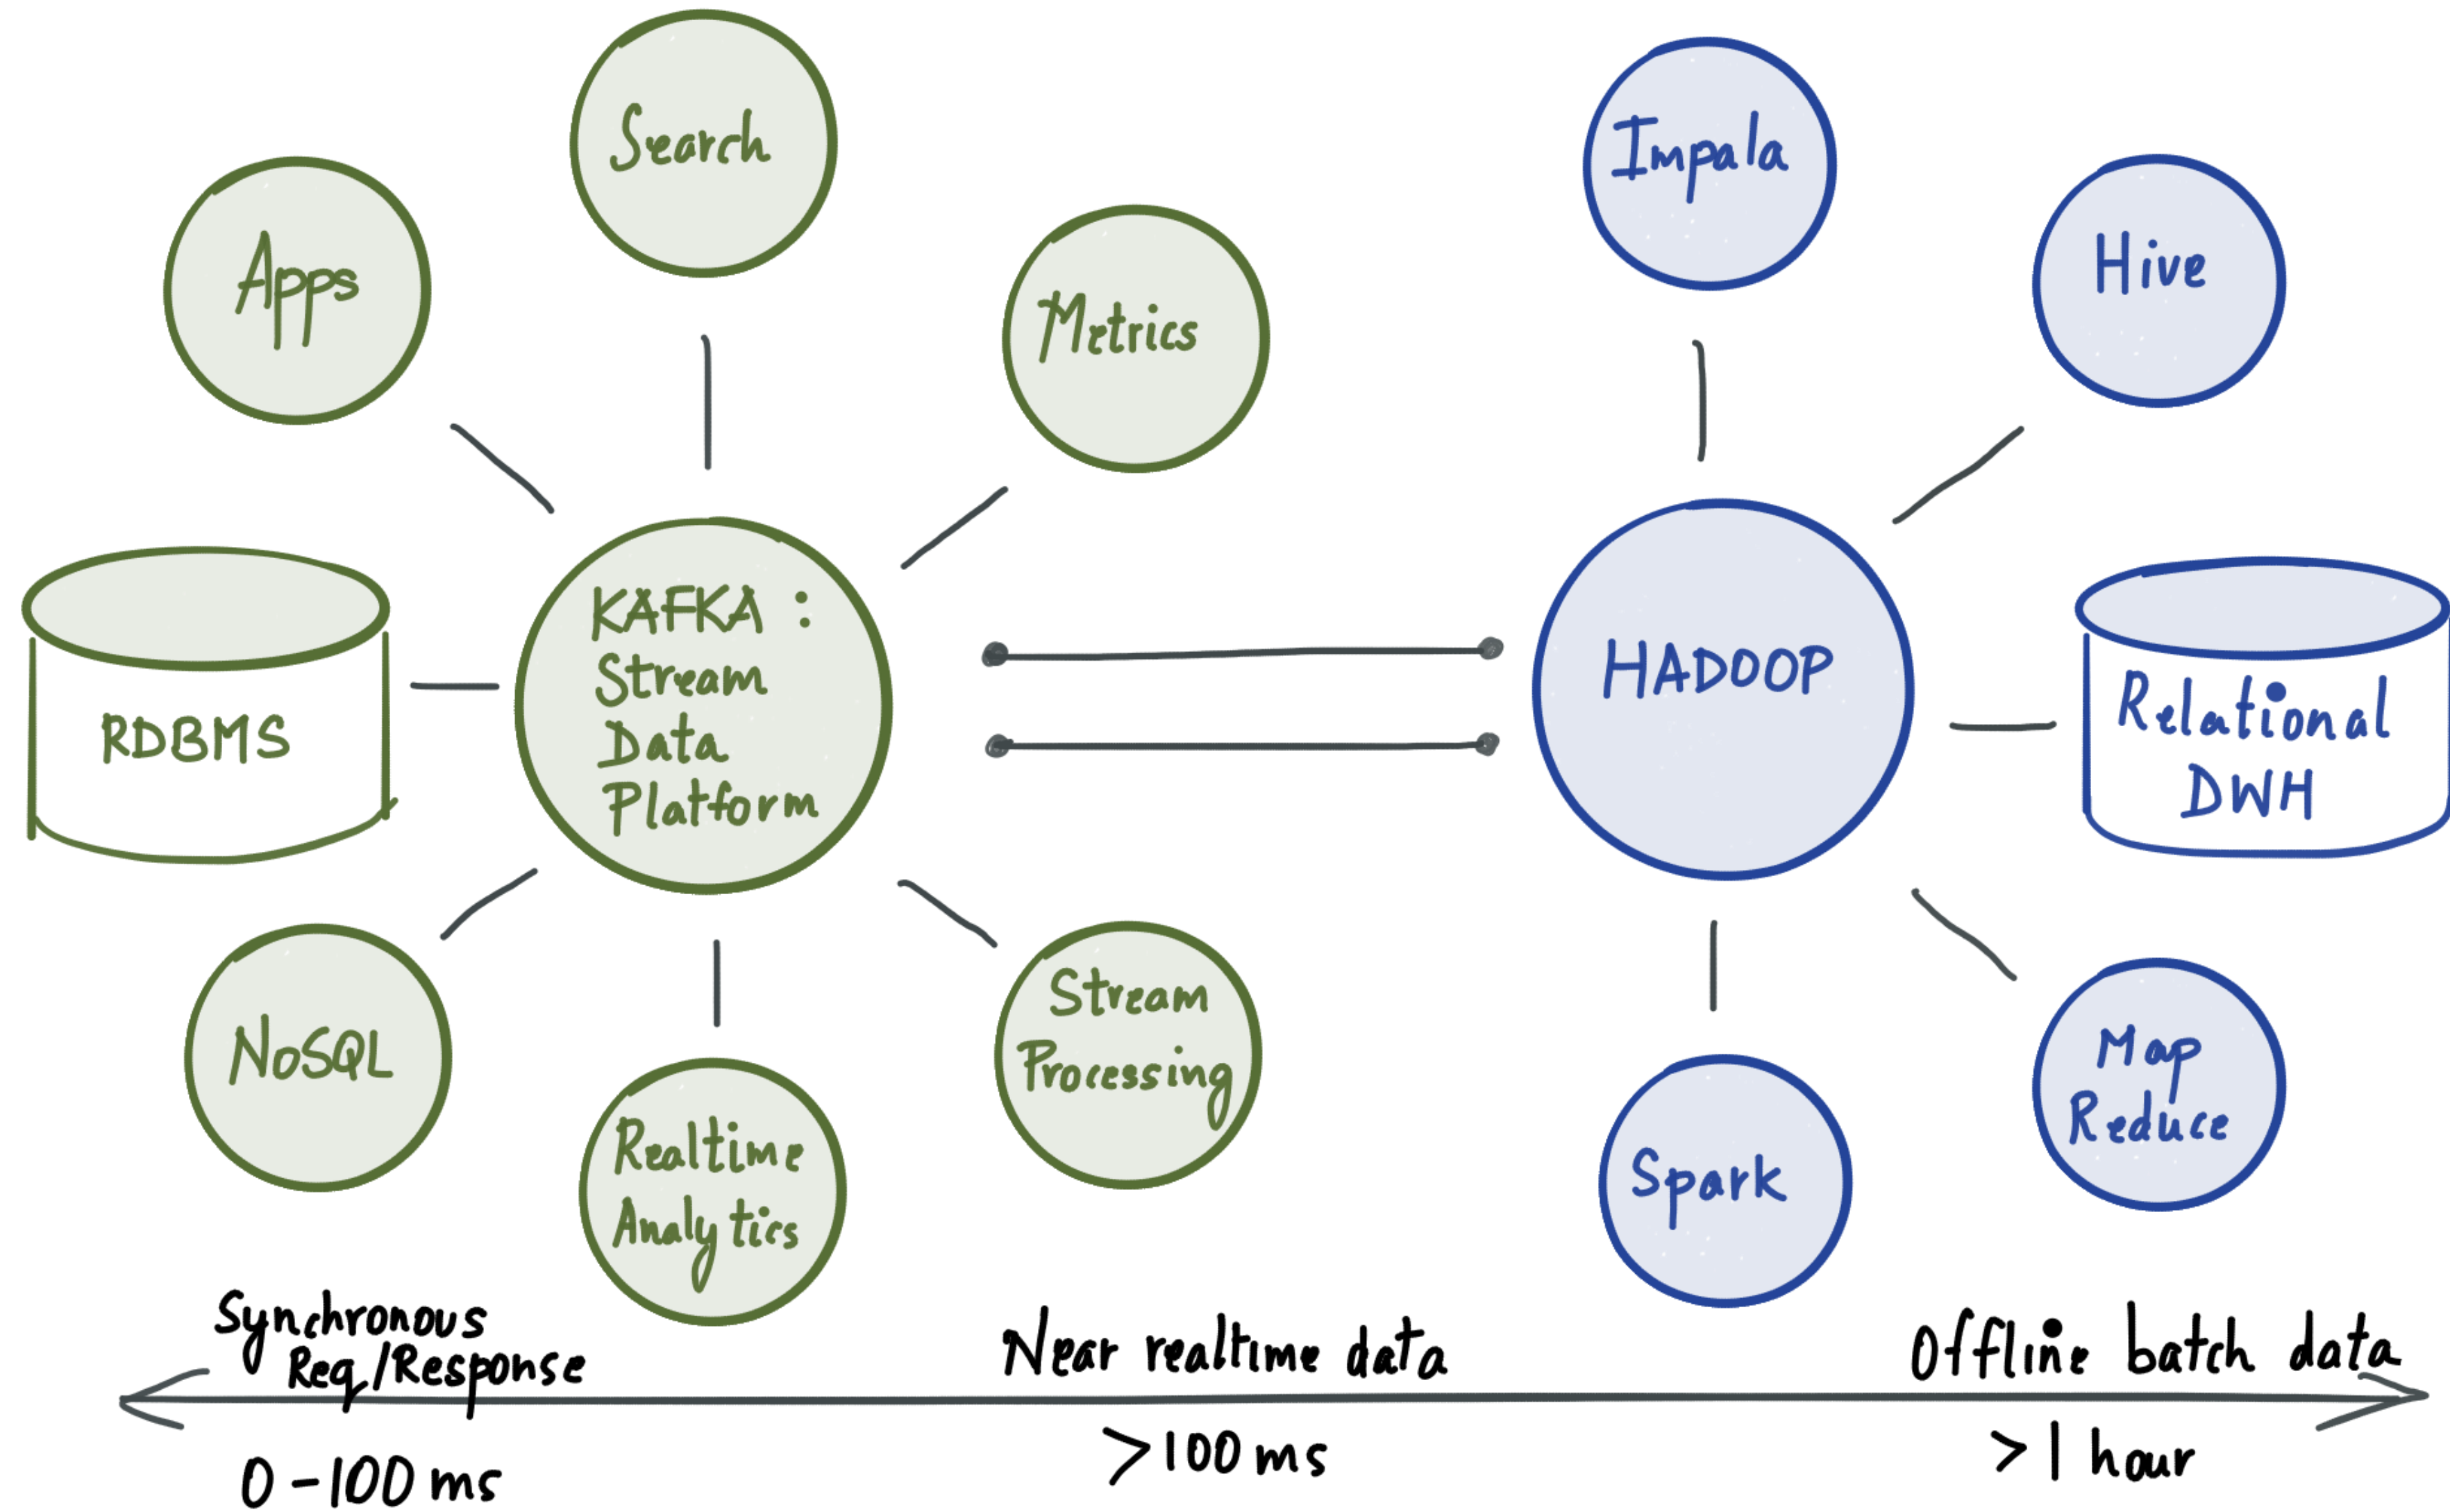
CONSUMER + CUSTOM + PRODUCER  
CODE

# STREAM PROCESSING



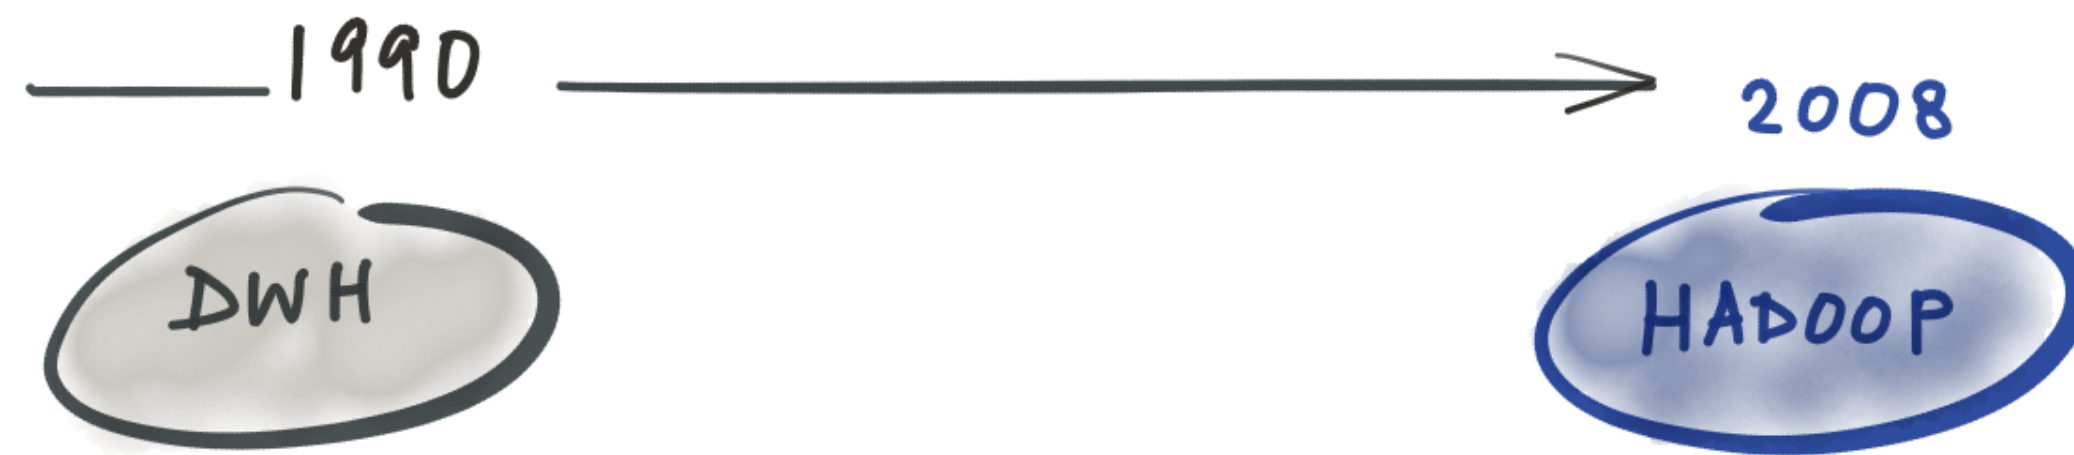
CONSUMER + CUSTOM + PRODUCER  
CODE

# STREAM DATA PLATFORM



1990

DWH



- Centralized  $\Rightarrow$  Distributed Computing
- Closed  $\Rightarrow$  Open Source





- Centralized  $\Rightarrow$  Distributed Computing
- Batch  $\Rightarrow$  Realtime
- Closed  $\Rightarrow$  Open Source
- Daily  $\Rightarrow$  Continuous



- Mission: Make Kafka-based stream data platform a practical reality everywhere
- Product
  - Stream processing integration for Kafka
  - Connectors for streaming data flow for common systems
  - Monitor end-to-end data flow
  - Schemas and metadata management
  - First release - Confluent Platform 1.0

# Thank you

[@nehanarkhede](#)

<http://confluent.io/careers>