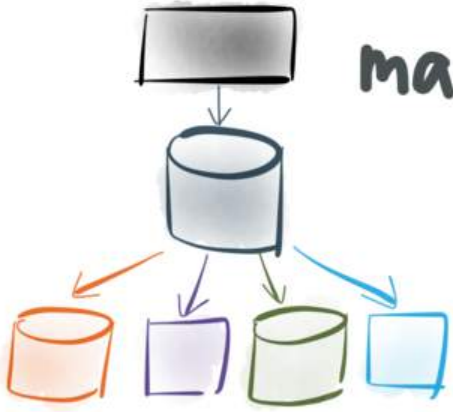



CHANGE DATA CAPTURE

THE
MAGIC
WAND
WE FORGOT



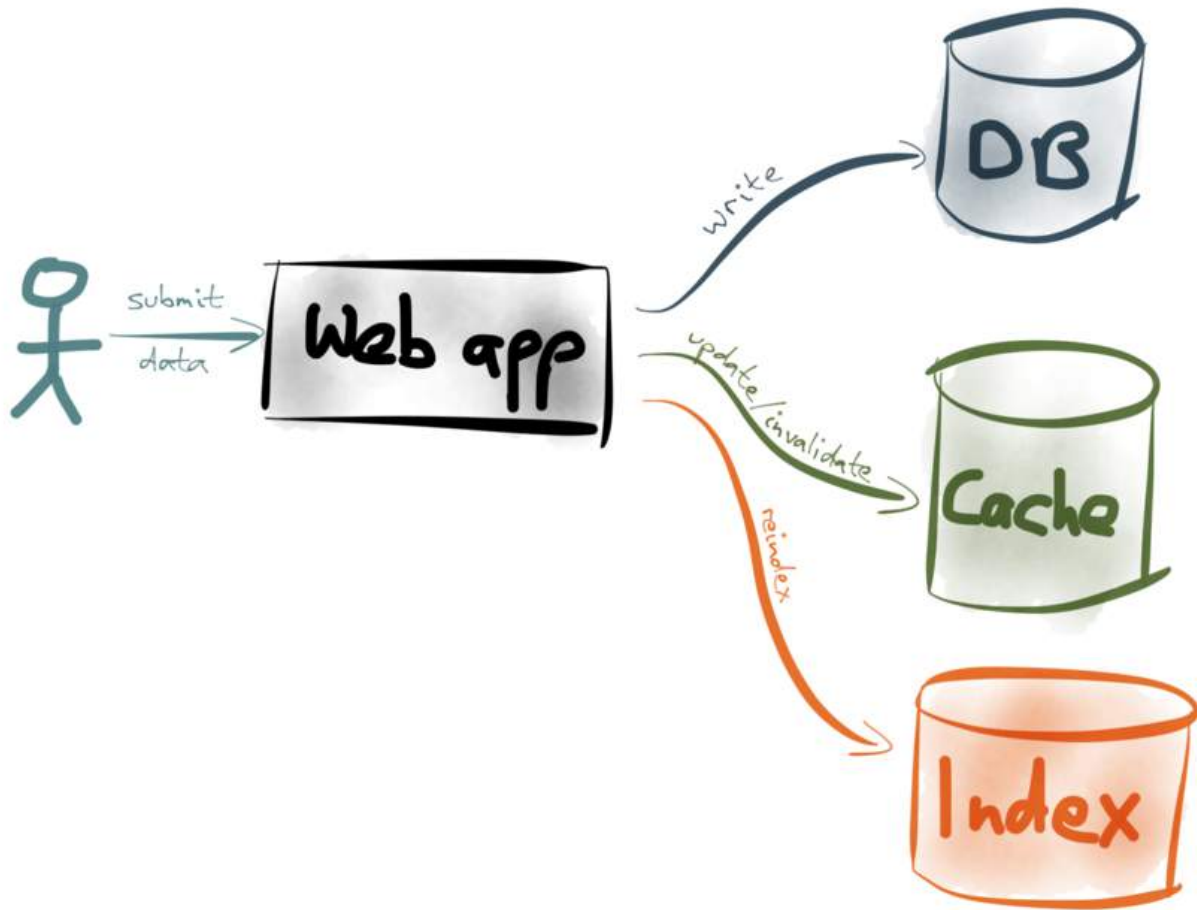
martinkl.com/bbuzz

MARTIN
KLEPPMANN
@martinkl



dataintensive.net

@martinkl



SAME
DATA



in

DIFFERENT
FORM

SAME
DATA

in


DIFFERENT
FORM

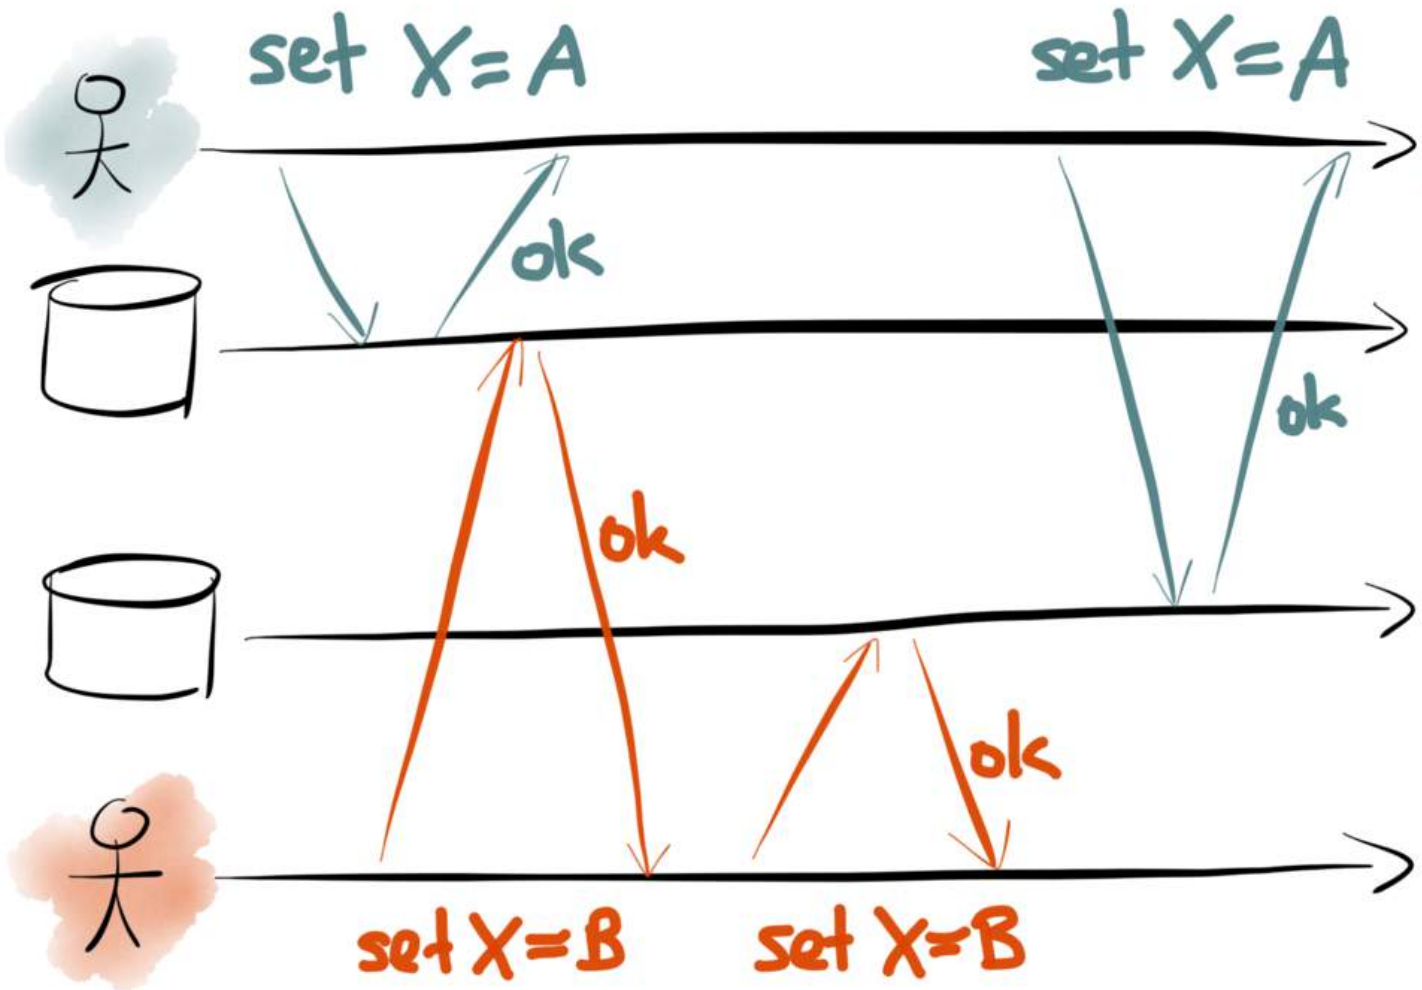
Denormalisation

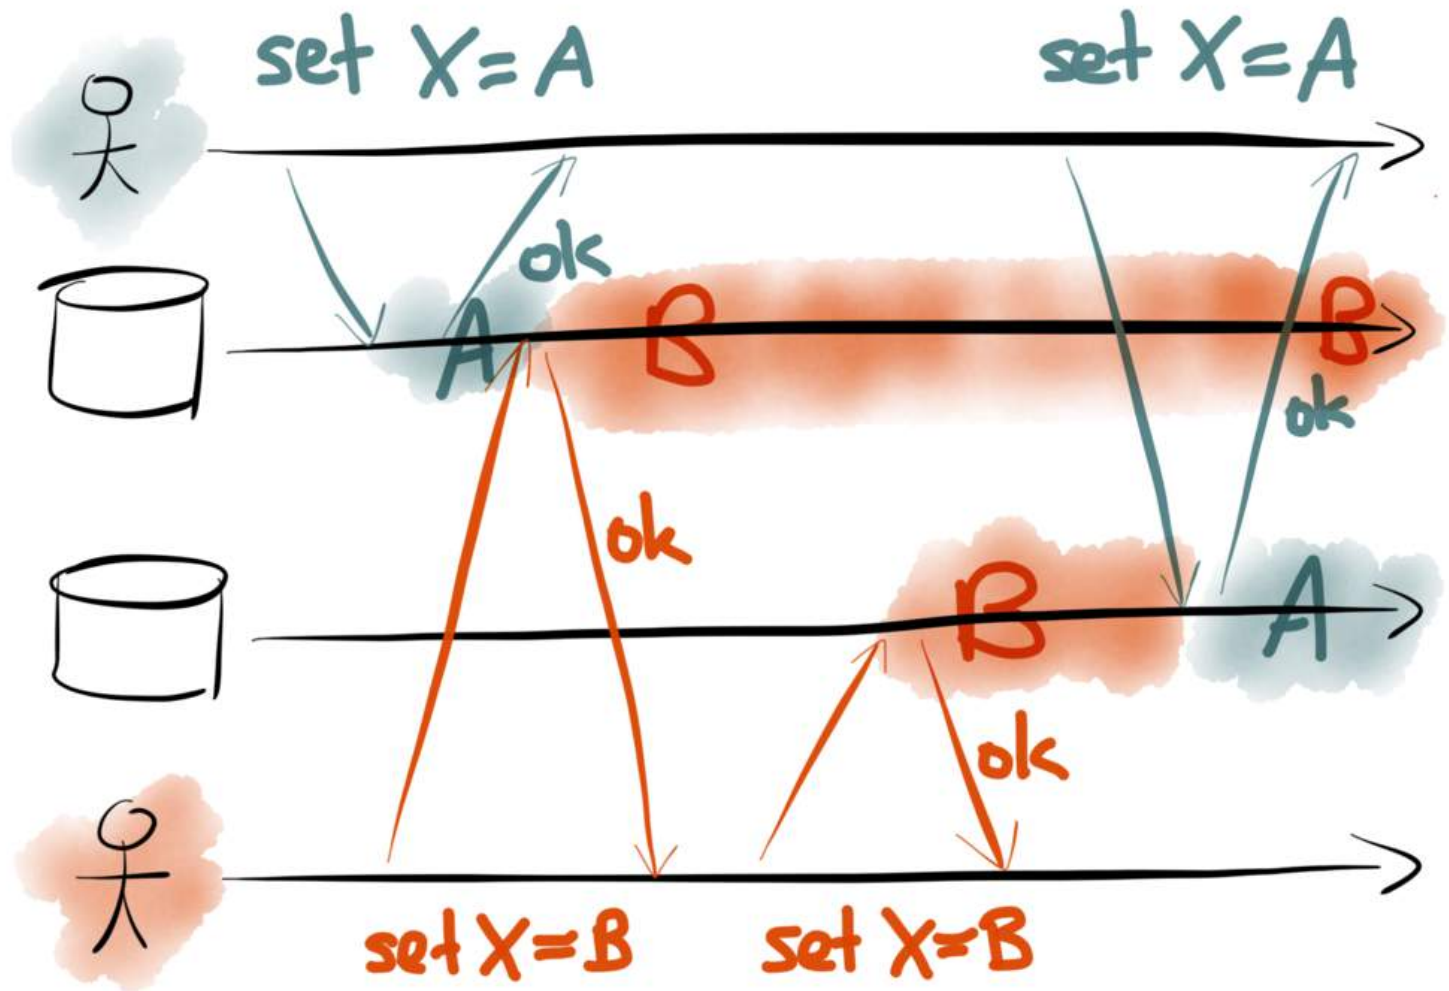
Caching

Indexes

Aggregations








OUR CHALLENGE:

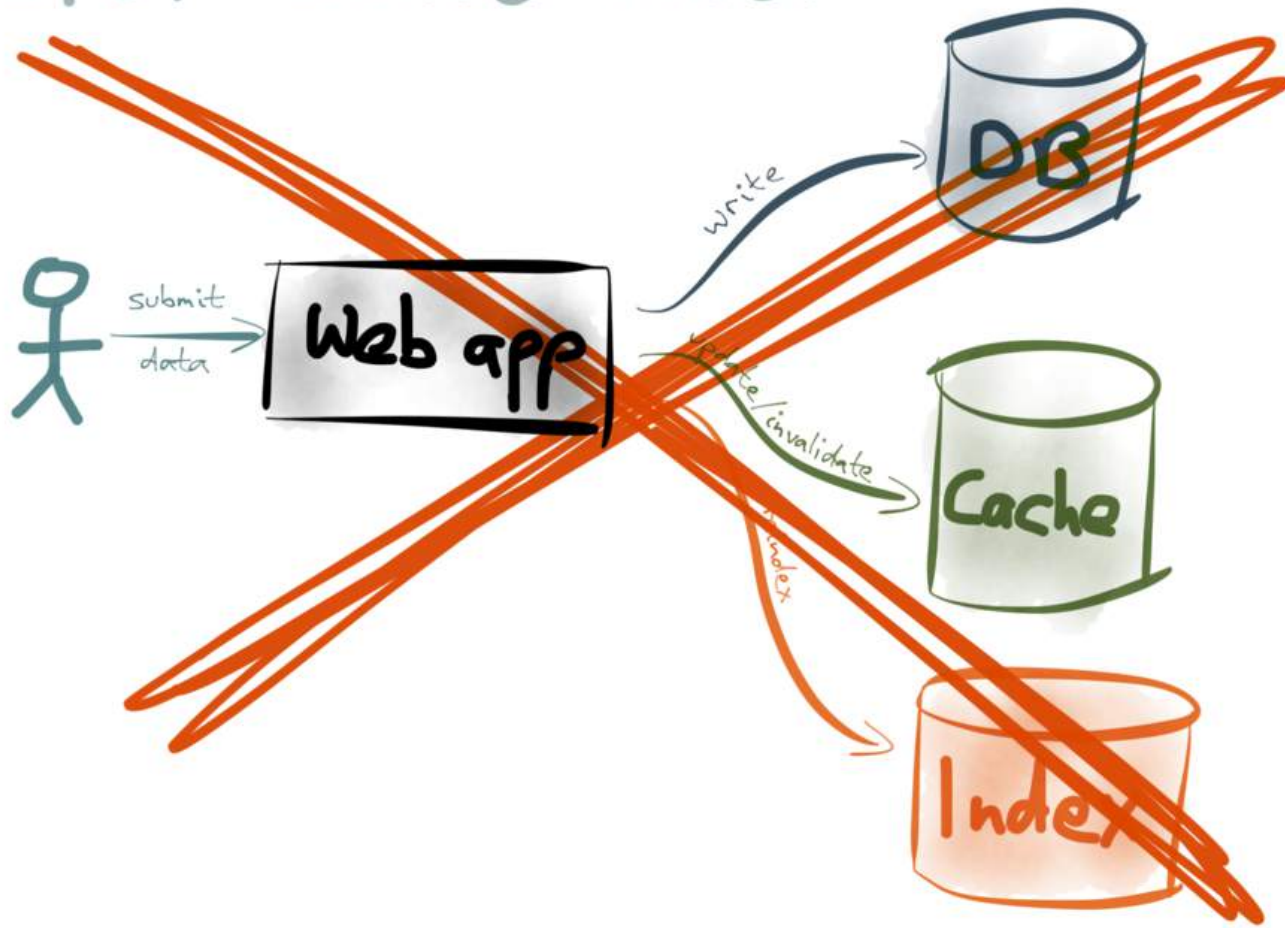
Data Integration

"Making sure the data
ends up in all the right places"

STOP DOING THIS.



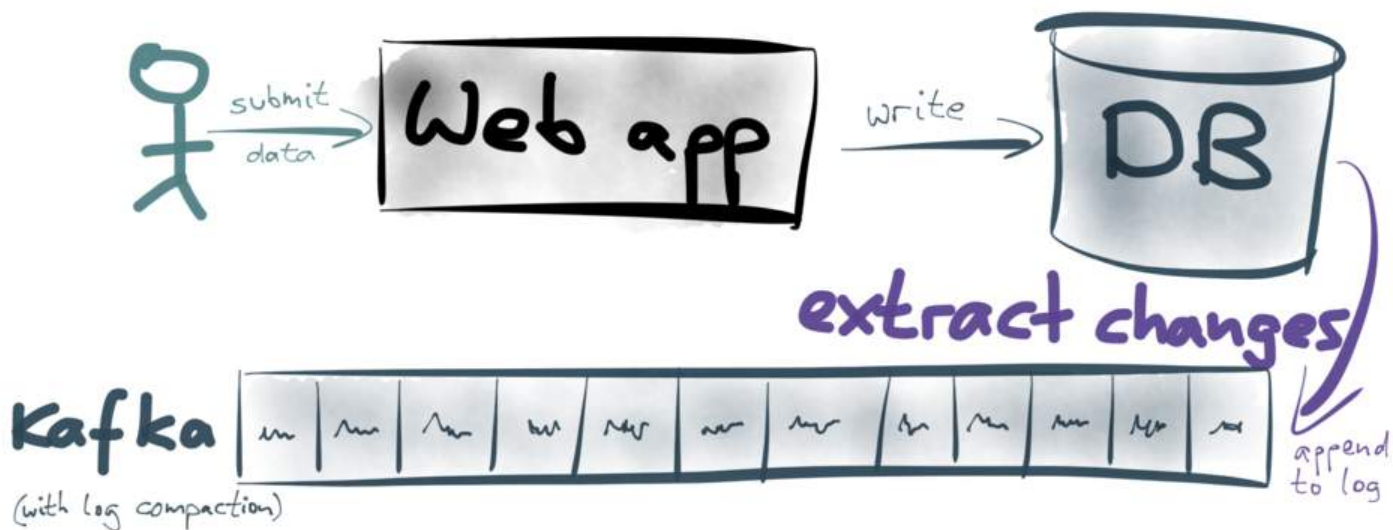
STOP DOING THIS.



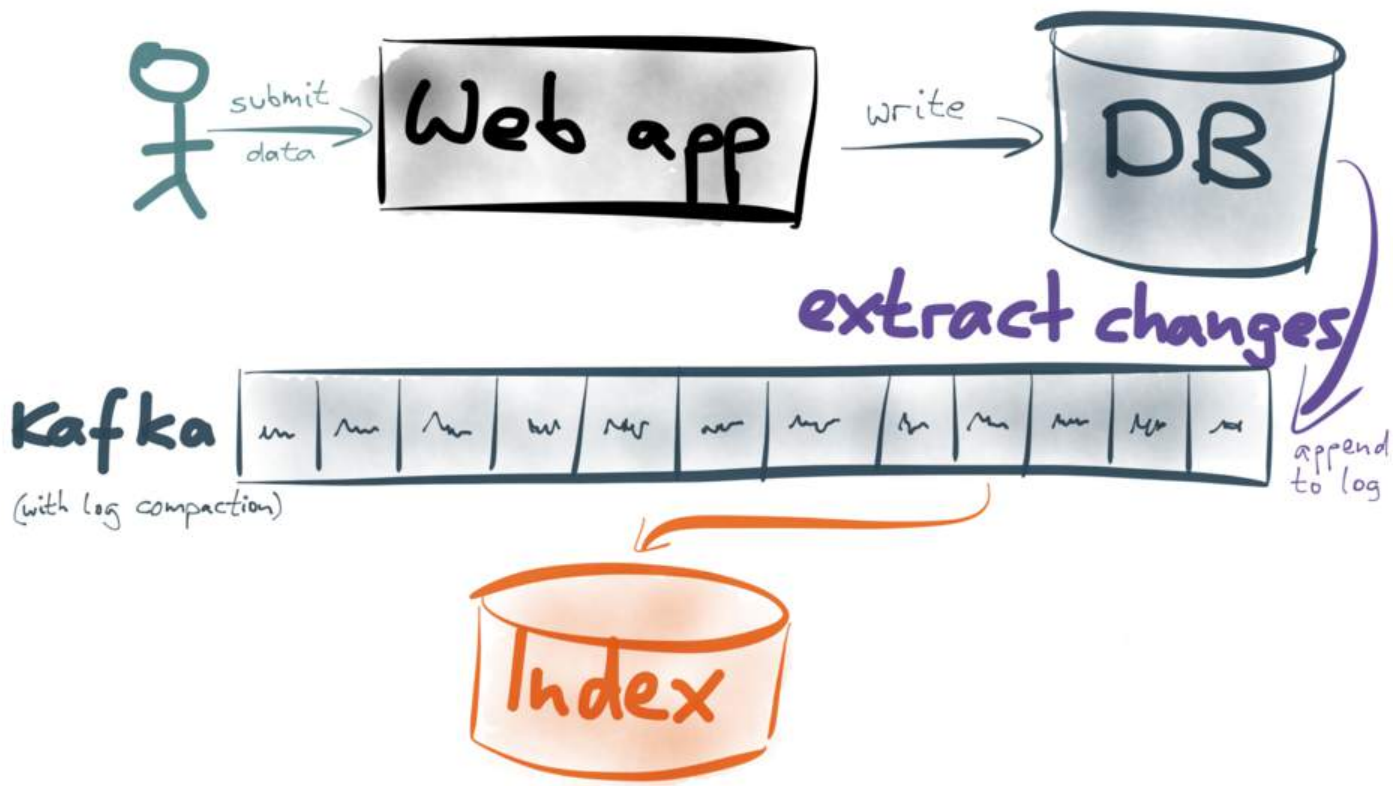
USING CHANGE CAPTURE



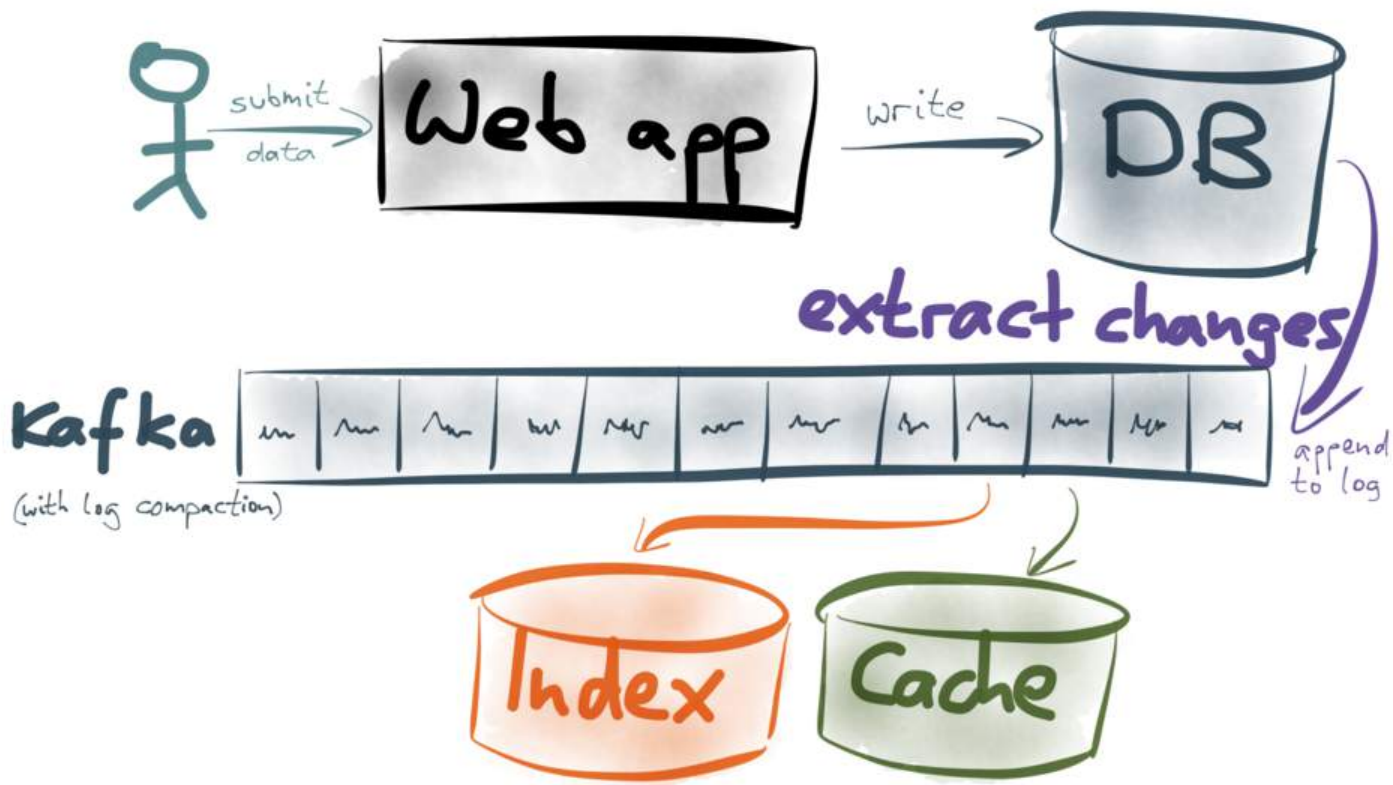
USING CHANGE CAPTURE



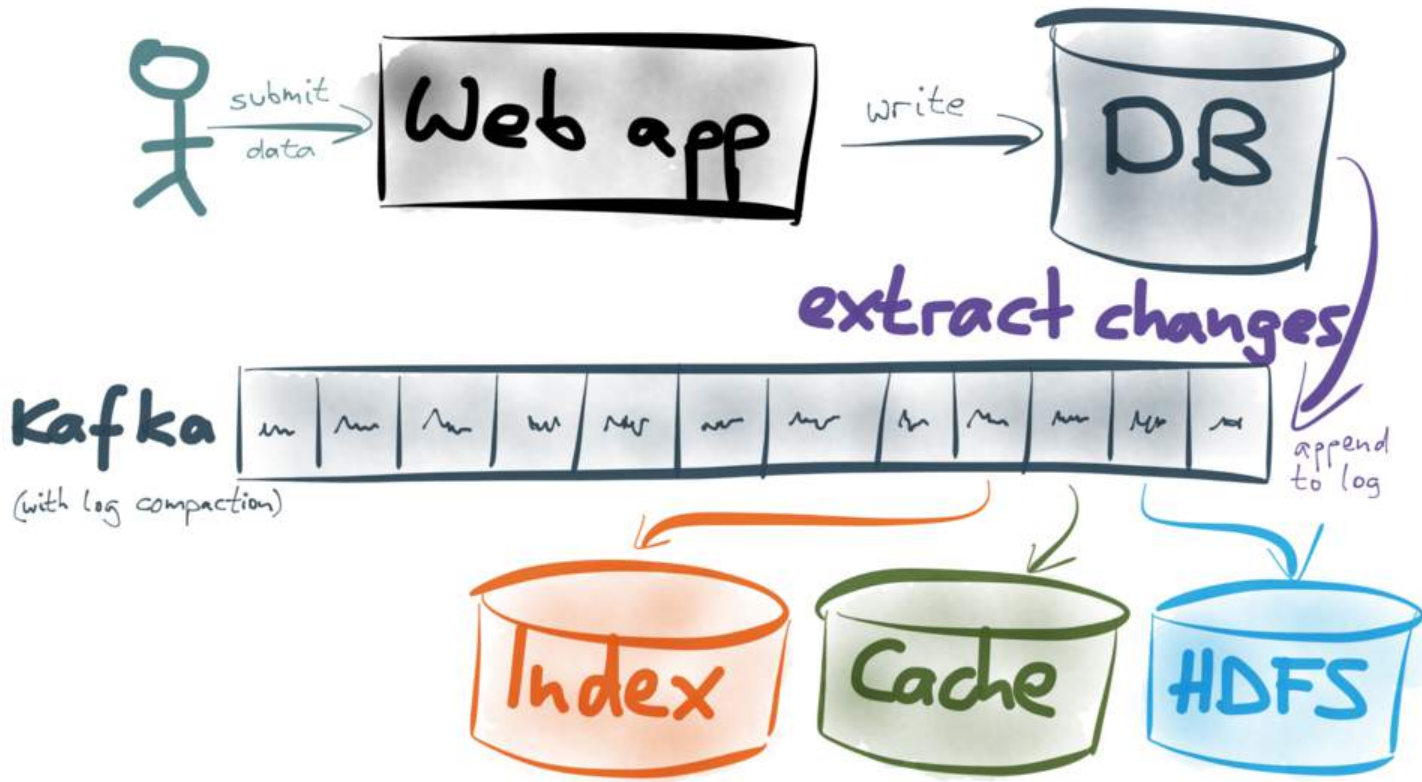
USING CHANGE CAPTURE



USING CHANGE CAPTURE

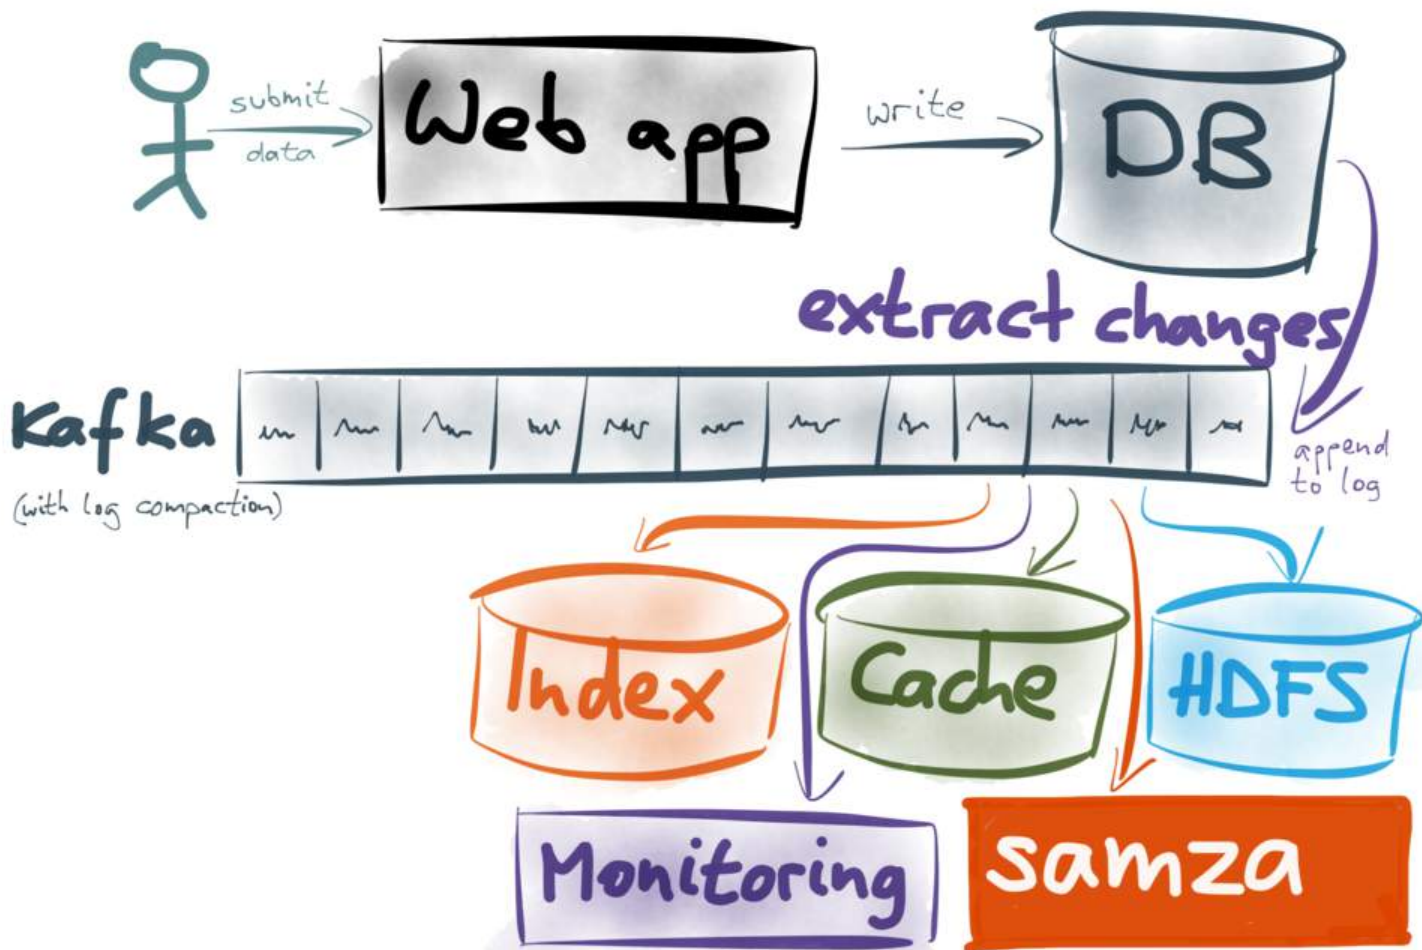


USING CHANGE CAPTURE



consumers independently apply writes

USING CHANGE CAPTURE



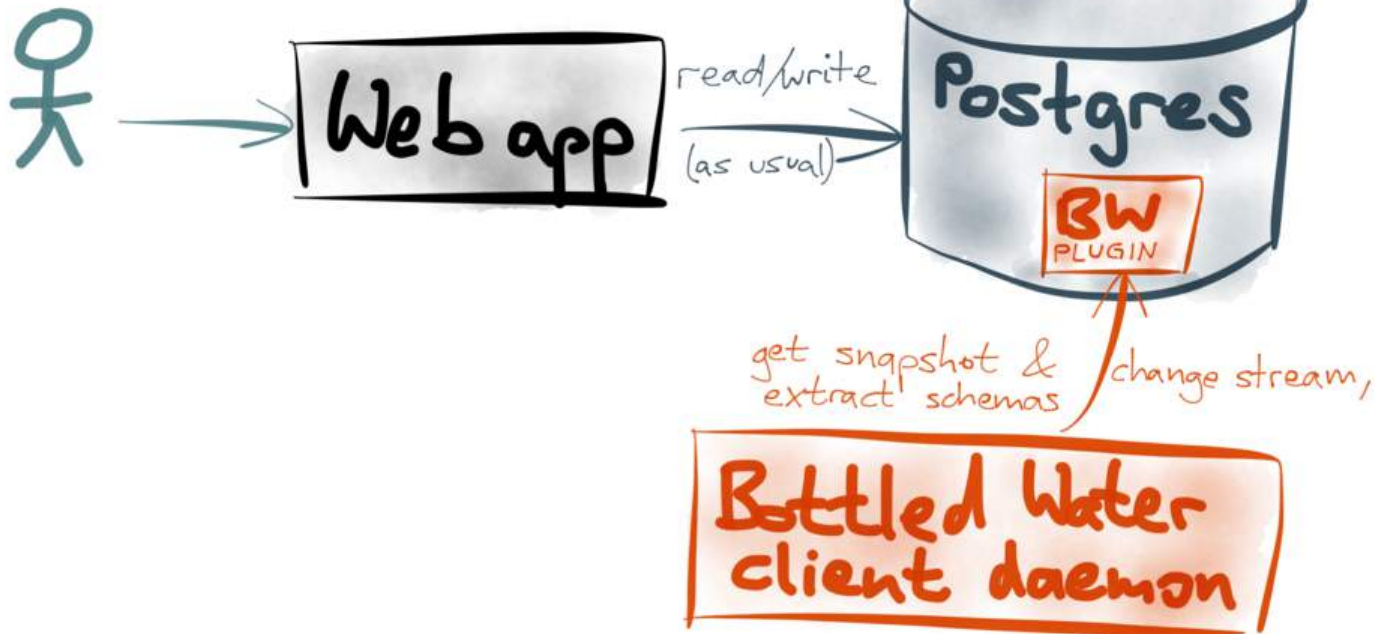


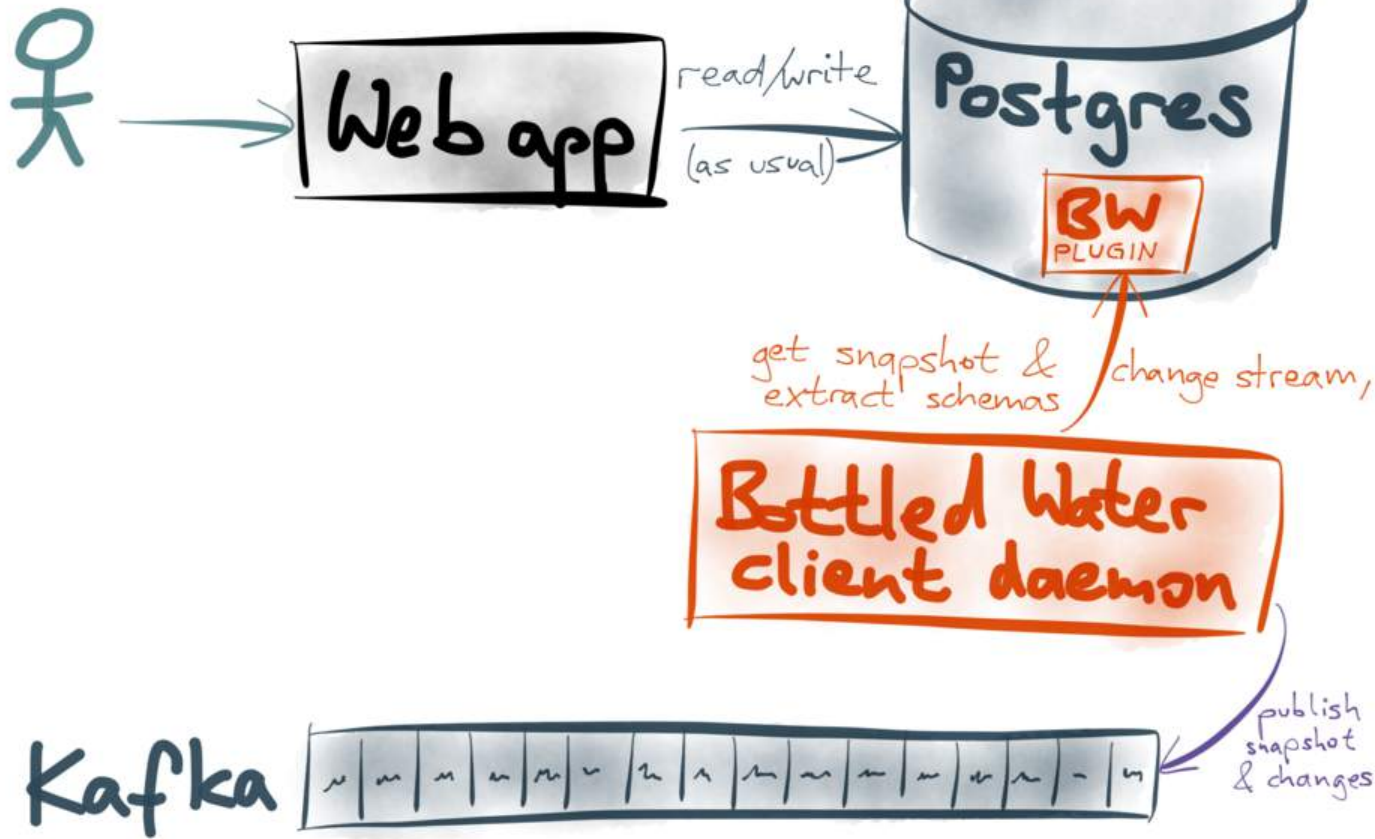
**DATA STREAMS - FRESHLY
BOTTLED AT SOURCE**

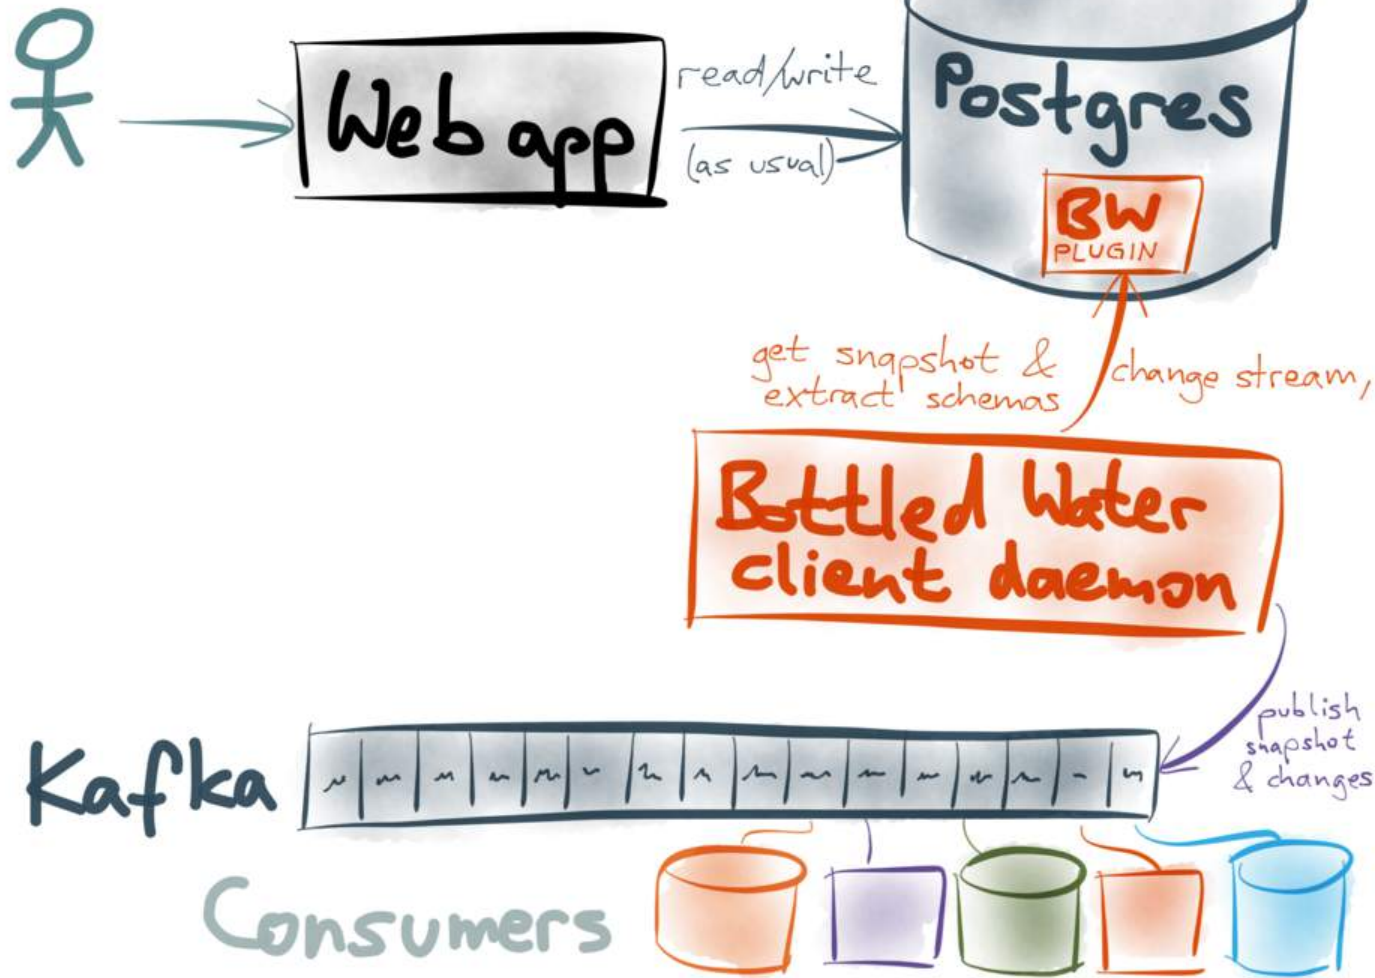


A Stream Data Platform powered by Apache Kafka











get snapshot & extract schemas

change stream,



register
← schemas



Kafka




publish snapshot & changes


Consumers

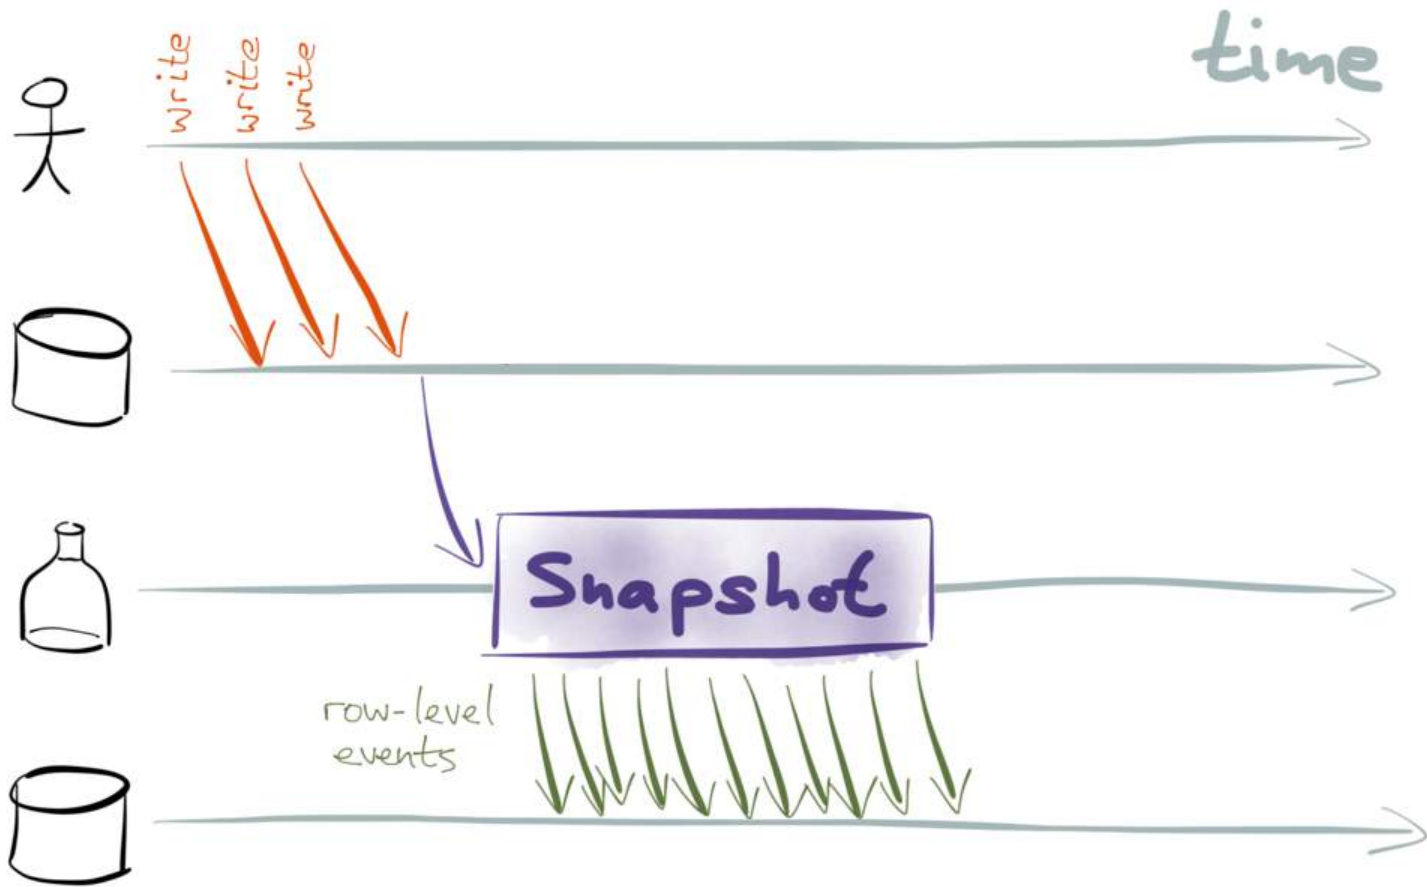



Demo




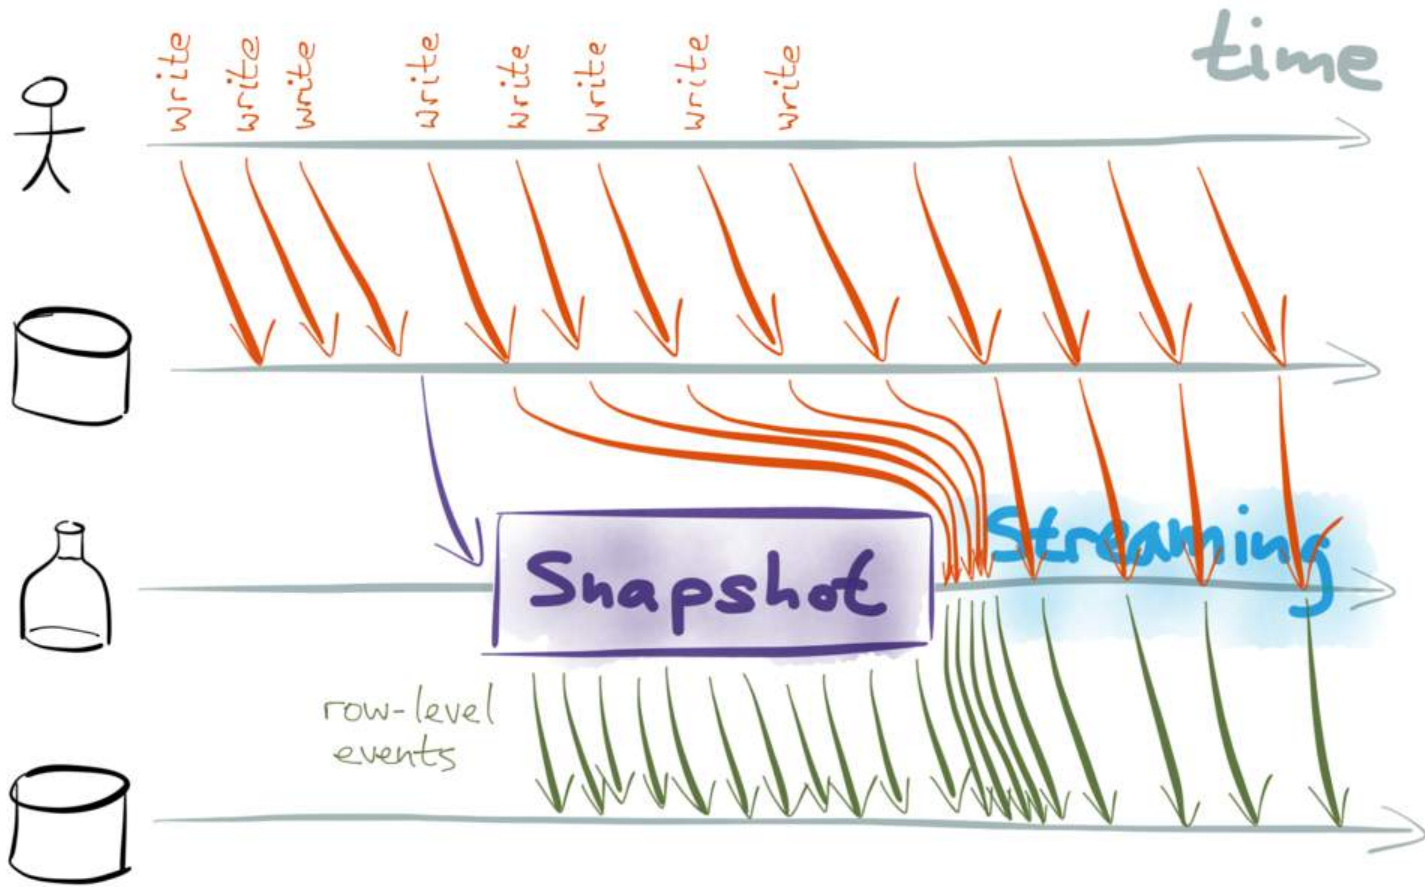













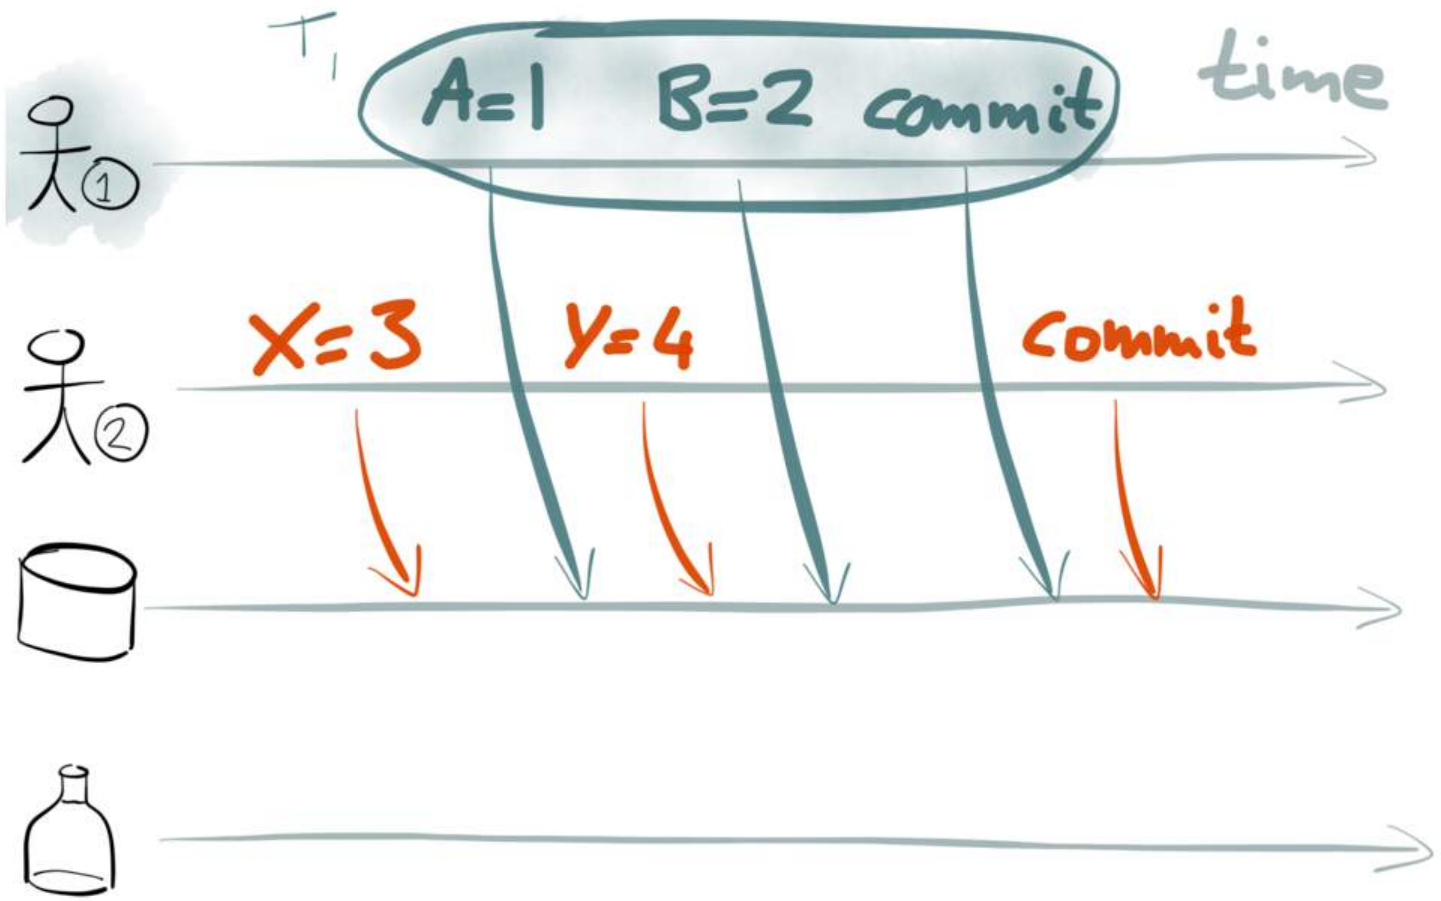
Transactions
& concurrency


time

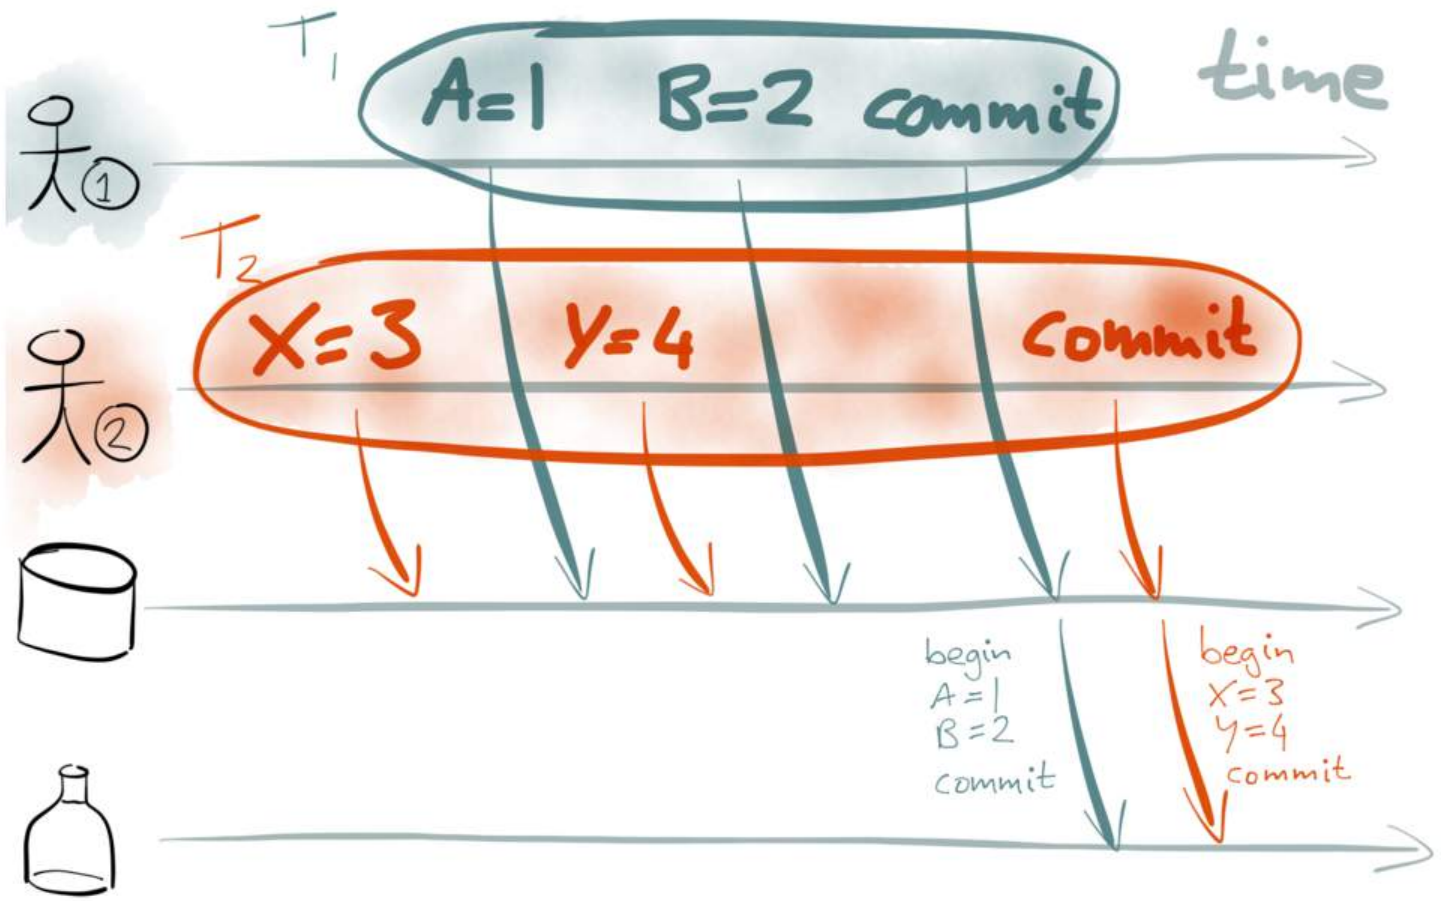















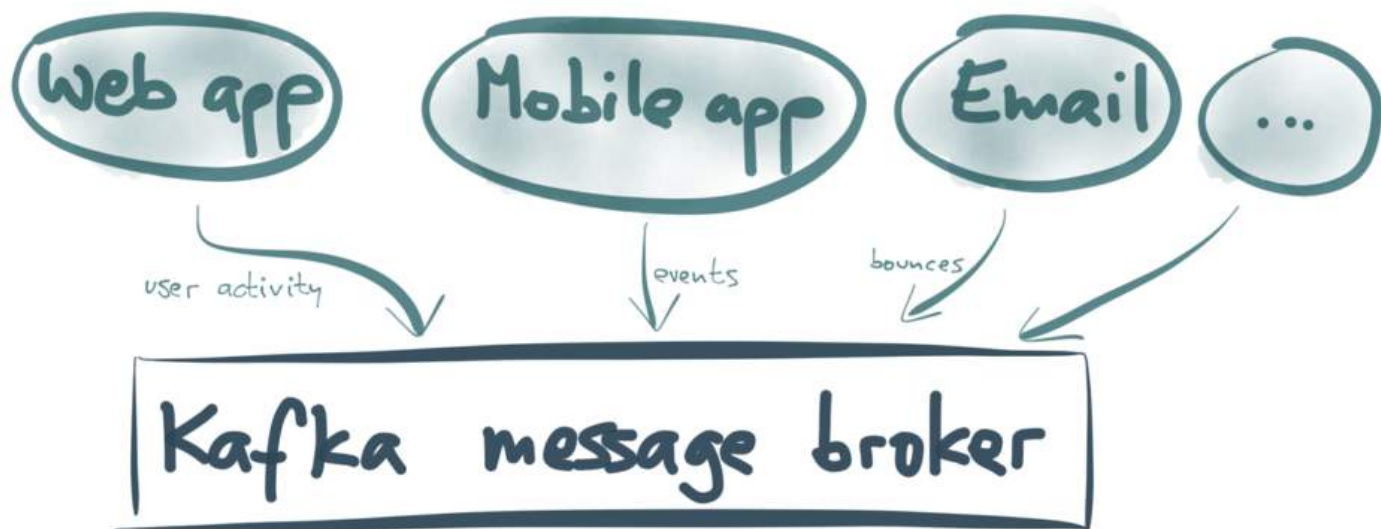
Kafka



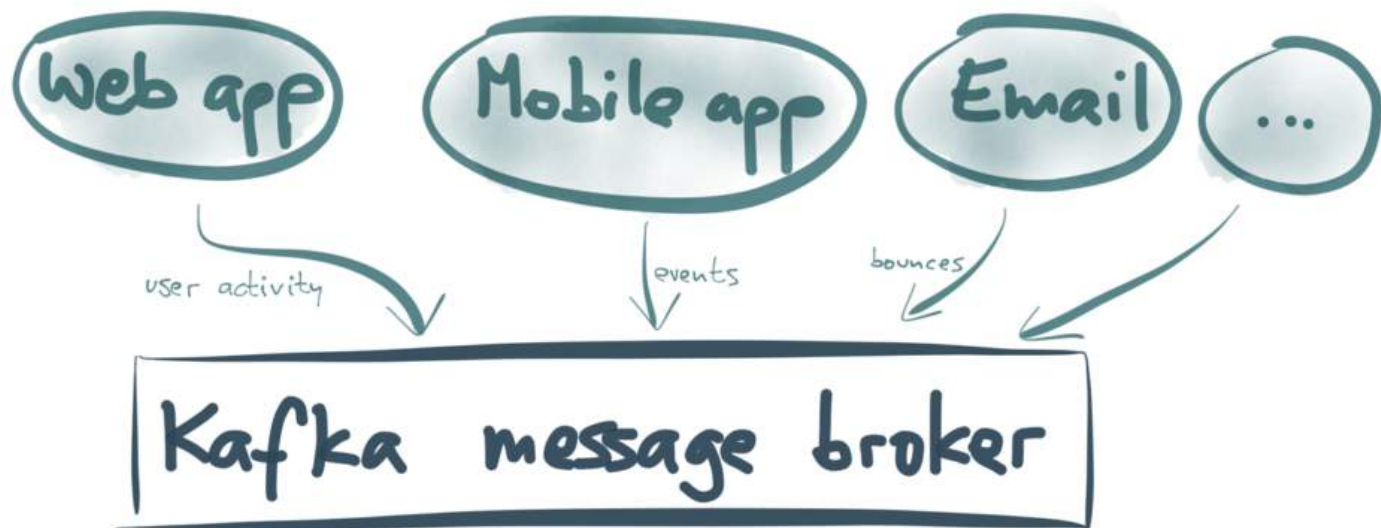
APACHE
SOFTWARE
FOUNDATION

Kafka message broker

PRODUCERS



PRODUCERS




CONSUMERS



stream

← oldest events

most recent events →




stream

new events added here

← oldest events

most recent events →




stream

new events added here

← oldest events

most recent events →



real-time
consumer
position (close to
head of
stream)

PostgreSQL

Bottled Water

PostgreSQL

Bottled Water

Table

Kafka topic

PostgreSQL

Bottled Water

Table

DDL

Kafka topic

Avro schema

PostgreSQL

Bottled Water

Table

Kafka topic

DDL

Avro schema

Row insert/update

Message (new value)

PostgreSQL

Bottled Water

Table

Kafka topic

DDL

Avro schema

Row insert/update

Message (new value)

Row delete

Message (null value)

A	42
---	----


B	21
---	----

A	44
---	----

C	12
---	----

A	45
---	----

Kafka changelog compaction




stream

new events added here

← oldest events

most recent events →



re-processing
historical events
starts here


real-time
consumer
position
(close to
head of
stream)

stream

new events added here

← oldest events

most recent events →



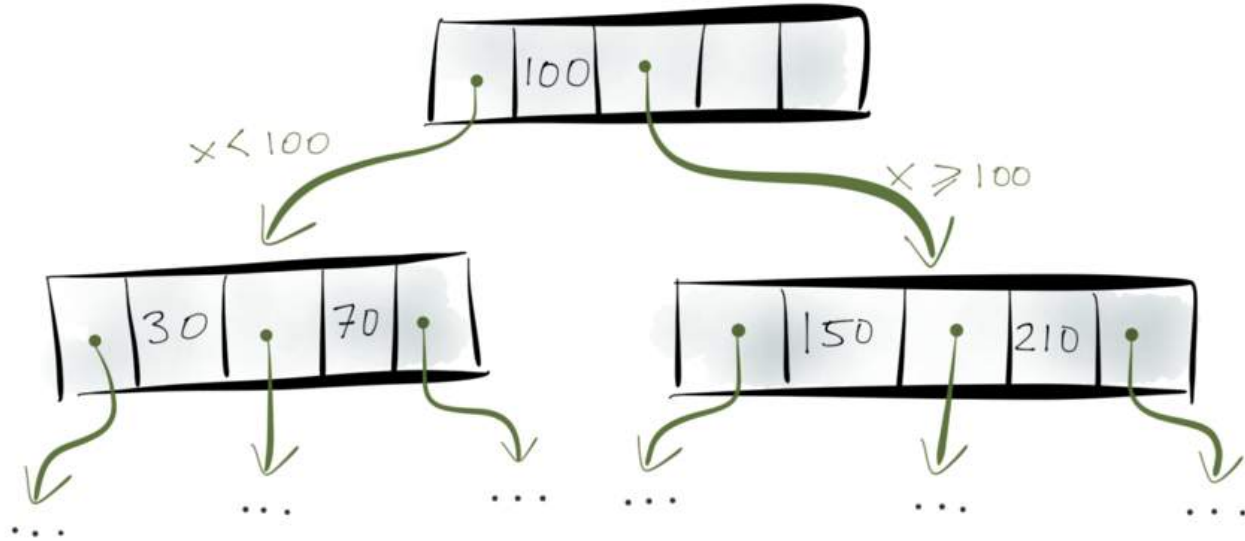
complete history

(using compaction to collect garbage)

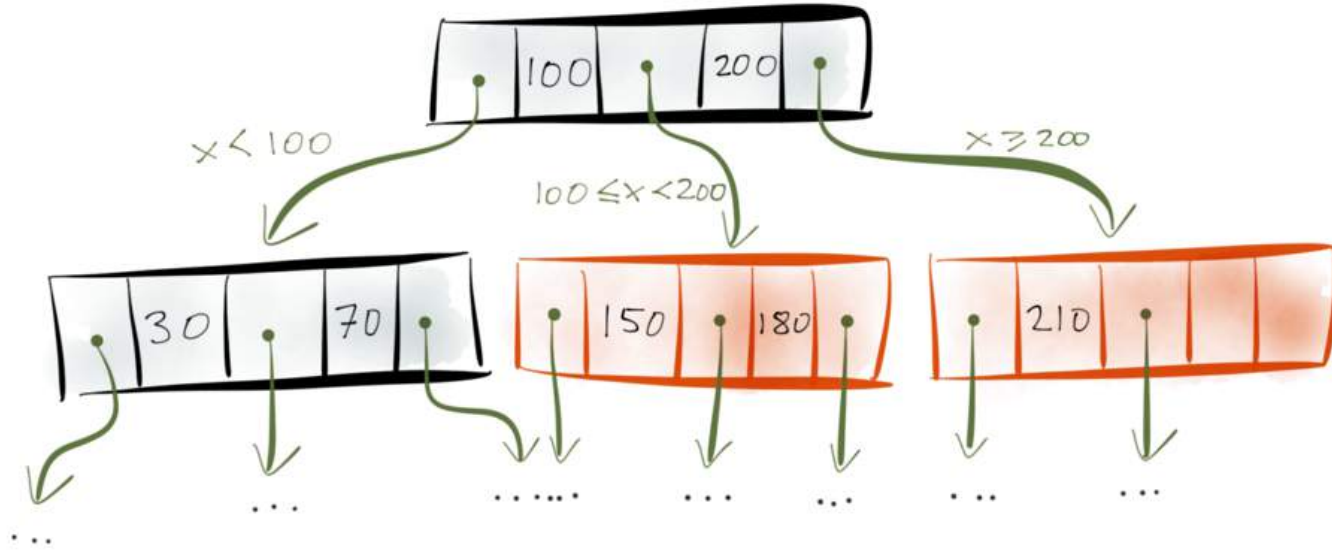


Inside
the bottle
factory

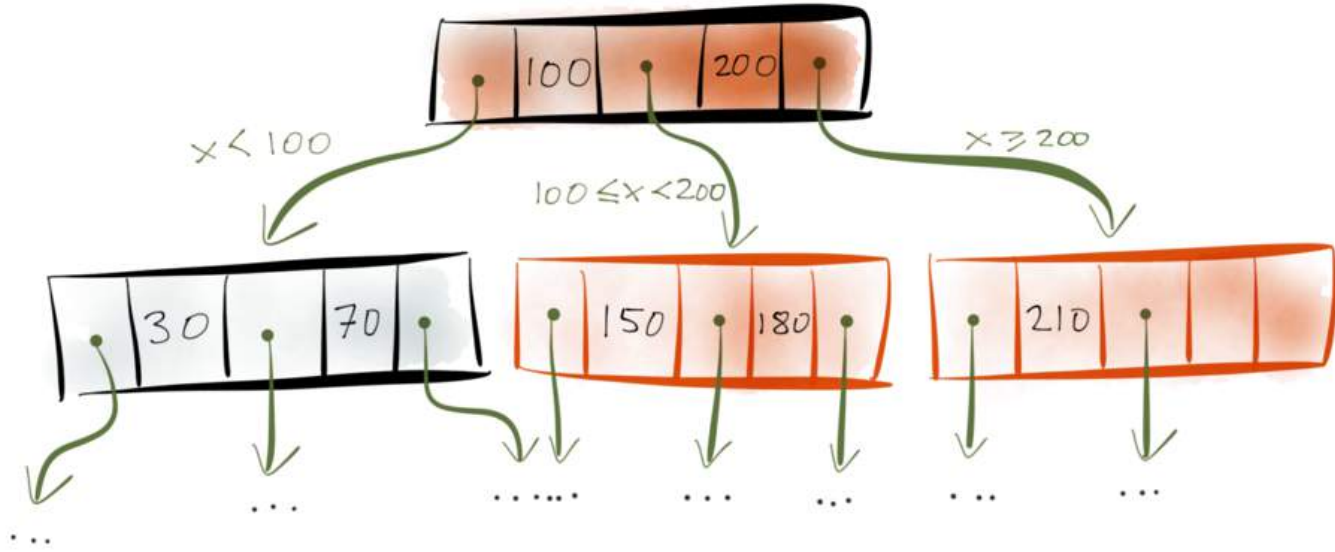
B-TREE



B-TREE NODE SPLIT




B-TREE NODE SPLIT



Atomically overwrite 3 pages?



B-TREE NODE SPLIT



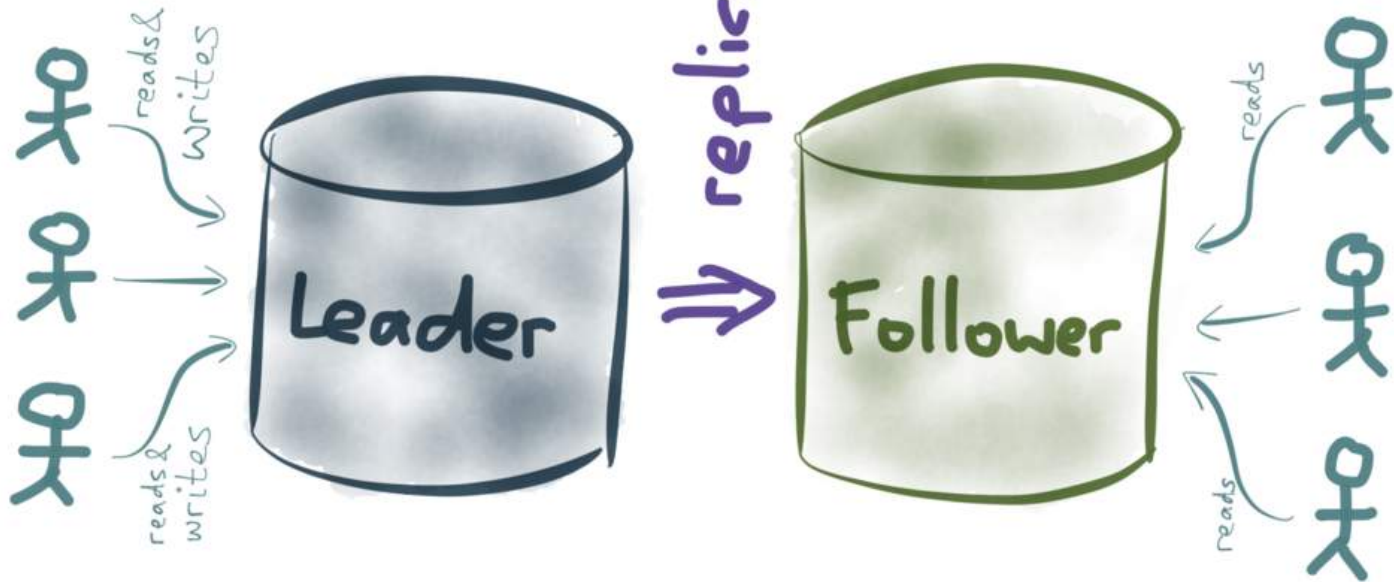
WRITE-AHEAD LOG (WAL)



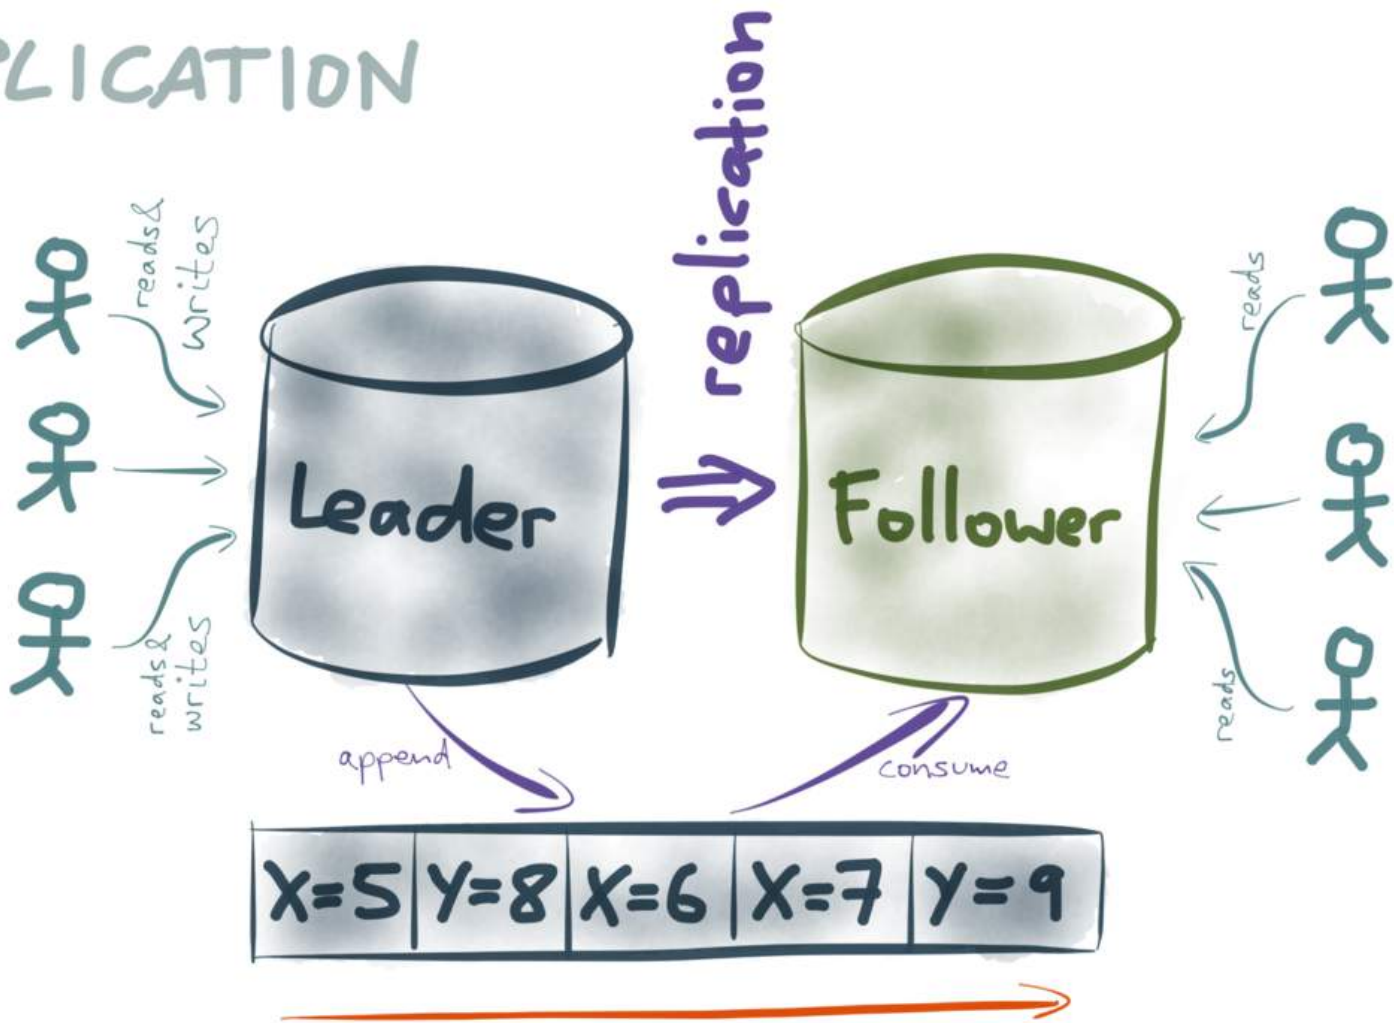
REPLICATION



REPLICATION



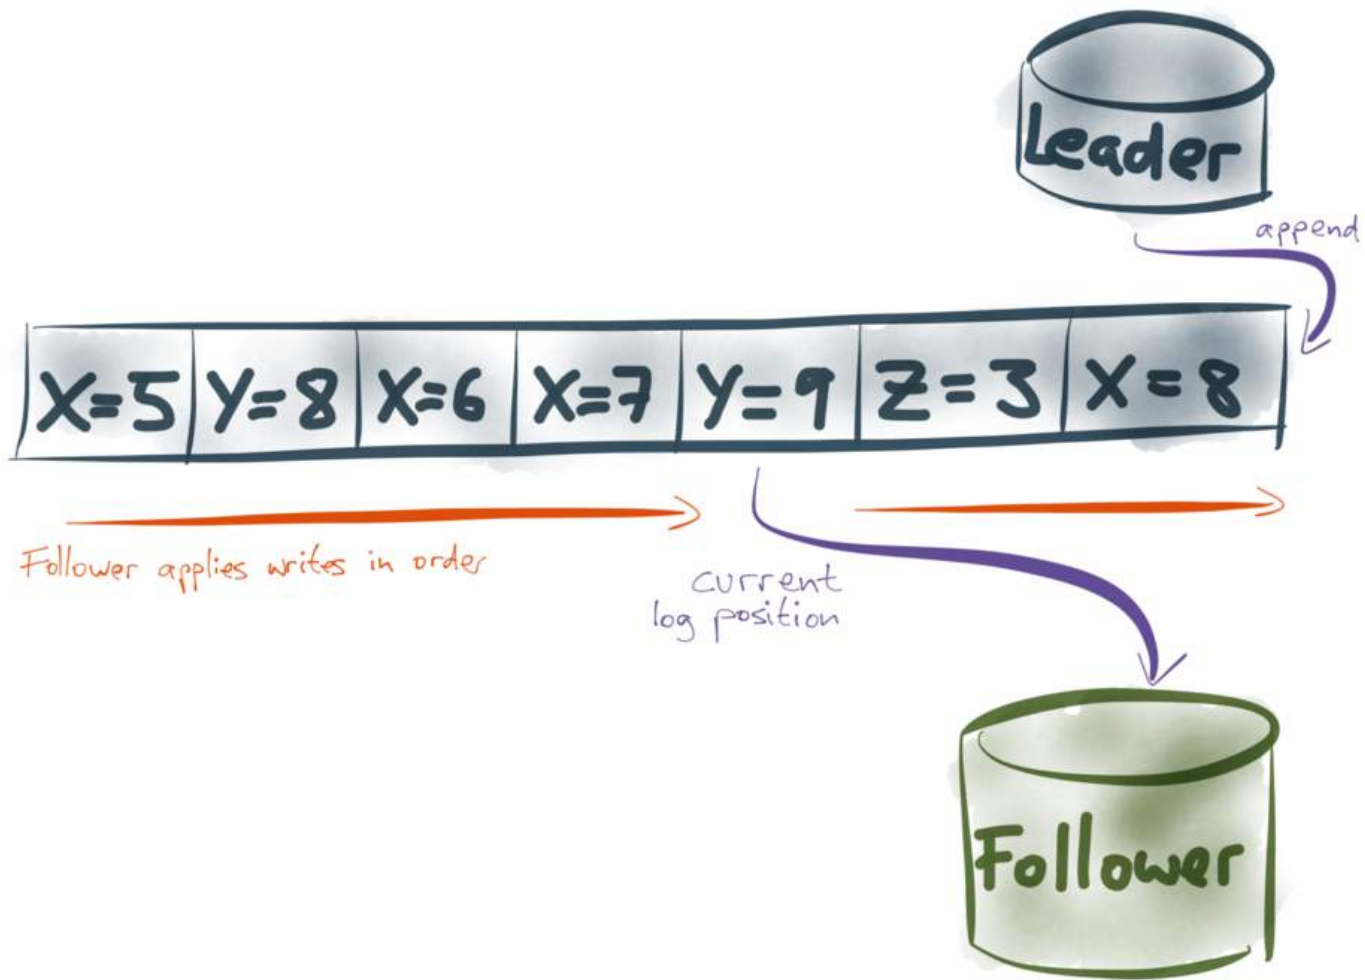
REPLICATION





append







get snapshot & extract schemas

change stream,



register
← schemas



Kafka




publish
snapshot
& changes

Consumers



Further reading

1. Martin Kleppmann: “Bottled Water: Real-time integration of PostgreSQL and Kafka.” 23 April 2015. <http://blog.confluent.io/2015/04/23/bottled-water-real-time-integration-of-postgresql-and-kafka/>
2. Shirshanka Das, Chavdar Botev, Kapil Surlaker, et al.: “All Aboard the Databus!” at *ACM Symposium on Cloud Computing (SoCC)*, October 2012. <http://www.socc2012.org/slides-das.pdf>
3. Yogeshwer Sharma, Philippe Ajoux, Petchean Ang, et al.: “Wormhole: Reliable Pub-Sub to Support Geo-replicated Internet Services,” at *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, May 2015. <https://www.usenix.org/system/files/conference/nsdi15/nsdi15-paper-sharma.pdf>
4. Jay Kreps: “I ♥ Logs.” O’Reilly Media, September 2014. <http://shop.oreilly.com/product/0636920034339.do>
5. Martin Kleppmann: “Designing data-intensive applications.” O’Reilly Media, to appear in 2015. <http://dataintensive.net>
6. Martin Kleppmann: “Turning the database inside-out with Apache Samza.” 4 March 2015. <http://blog.confluent.io/2015/03/04/turning-the-database-inside-out-with-apache-samza/>
7. Pat Helland: “Immutability Changes Everything,” at *7th Biennial Conference on Innovative Data Systems Research (CIDR)*, January 2015. http://www.cidrdb.org/cidr2015/Papers/CIDR15_Paper16.pdf



dataintensive.net

Bottled Water
github.com
/confluentinc
/bottledwater-pg