




*Building a Telegram bot with
Apache Kafka and ksqlDB*

Robin Moffatt

@rmoff





Addingham
Main St
Main St

Nesfield

Middleton

Denton

Ilkley Golf Club

East Holmes
Fields

Low Park Rd

West Ln

Church St
Ilkley

BEN RHYDDING

Cowpasture Rd

Ben Rhydding Golf Club

Manor Park

ESSCROFT

Ilkley Rd

Ilkley Moor Cow
& Calf Rocks

The Cow and Calf

Moor Rd

Moor Rd

Moor Rd


Moor Ln

Ilkley Moor

Main St

Station Rd

Bur
Whar





NCP Welcome to Hall Ings Car Park

BRADFORD CAPITAL OF CYCLING

*Where's my nearest carpark
with available spaces?*

NCP Welcome to Hall Ings Car Park P

BRADFORD CAPITAL OF CYCLING

BRADFORD CAPITAL OF CYCLING

*How many spaces are available
in this car park?*

NCP Welcome to Hall Ings Car Park P

BRADFORD CAPITAL OF CYCLING

BRADFORD CAPITAL OF CYCLING



💡 *Tell me when a car park with spaces is available*

NCP Welcome to Hall Ings Car Park P

BRADFORD CAPITAL OF CYCLING



How does occupancy vary over time?





Datasets



Products



District Dashboard (visual)



Latest CQC Results (visual)



Other open data sites



Contact

Bradford car parks

City of Bradford Metropolitan District Council



Resources from the City of Bradford Metropolitan District Council (CBMDC) Parking service.

1 - Car park locations

a simple csv containing name and location including latitude / longitude

2 - Car park current status.

API that returns a csv dataset of the current status of 8 Bradford city centre car parks. The dataset returns capacity, empty places, status together with location details.

The dataset is updated every 3 minutes for a live view of spaces in these car parks.

3 - Car park historic status


API that returns a csv dataset building up the historic status of the 8 city centre car parks.

The dataset is updated every 30 minutes.

4 Resources ...

[csv](#) [api](#) [geojson](#)

More Information

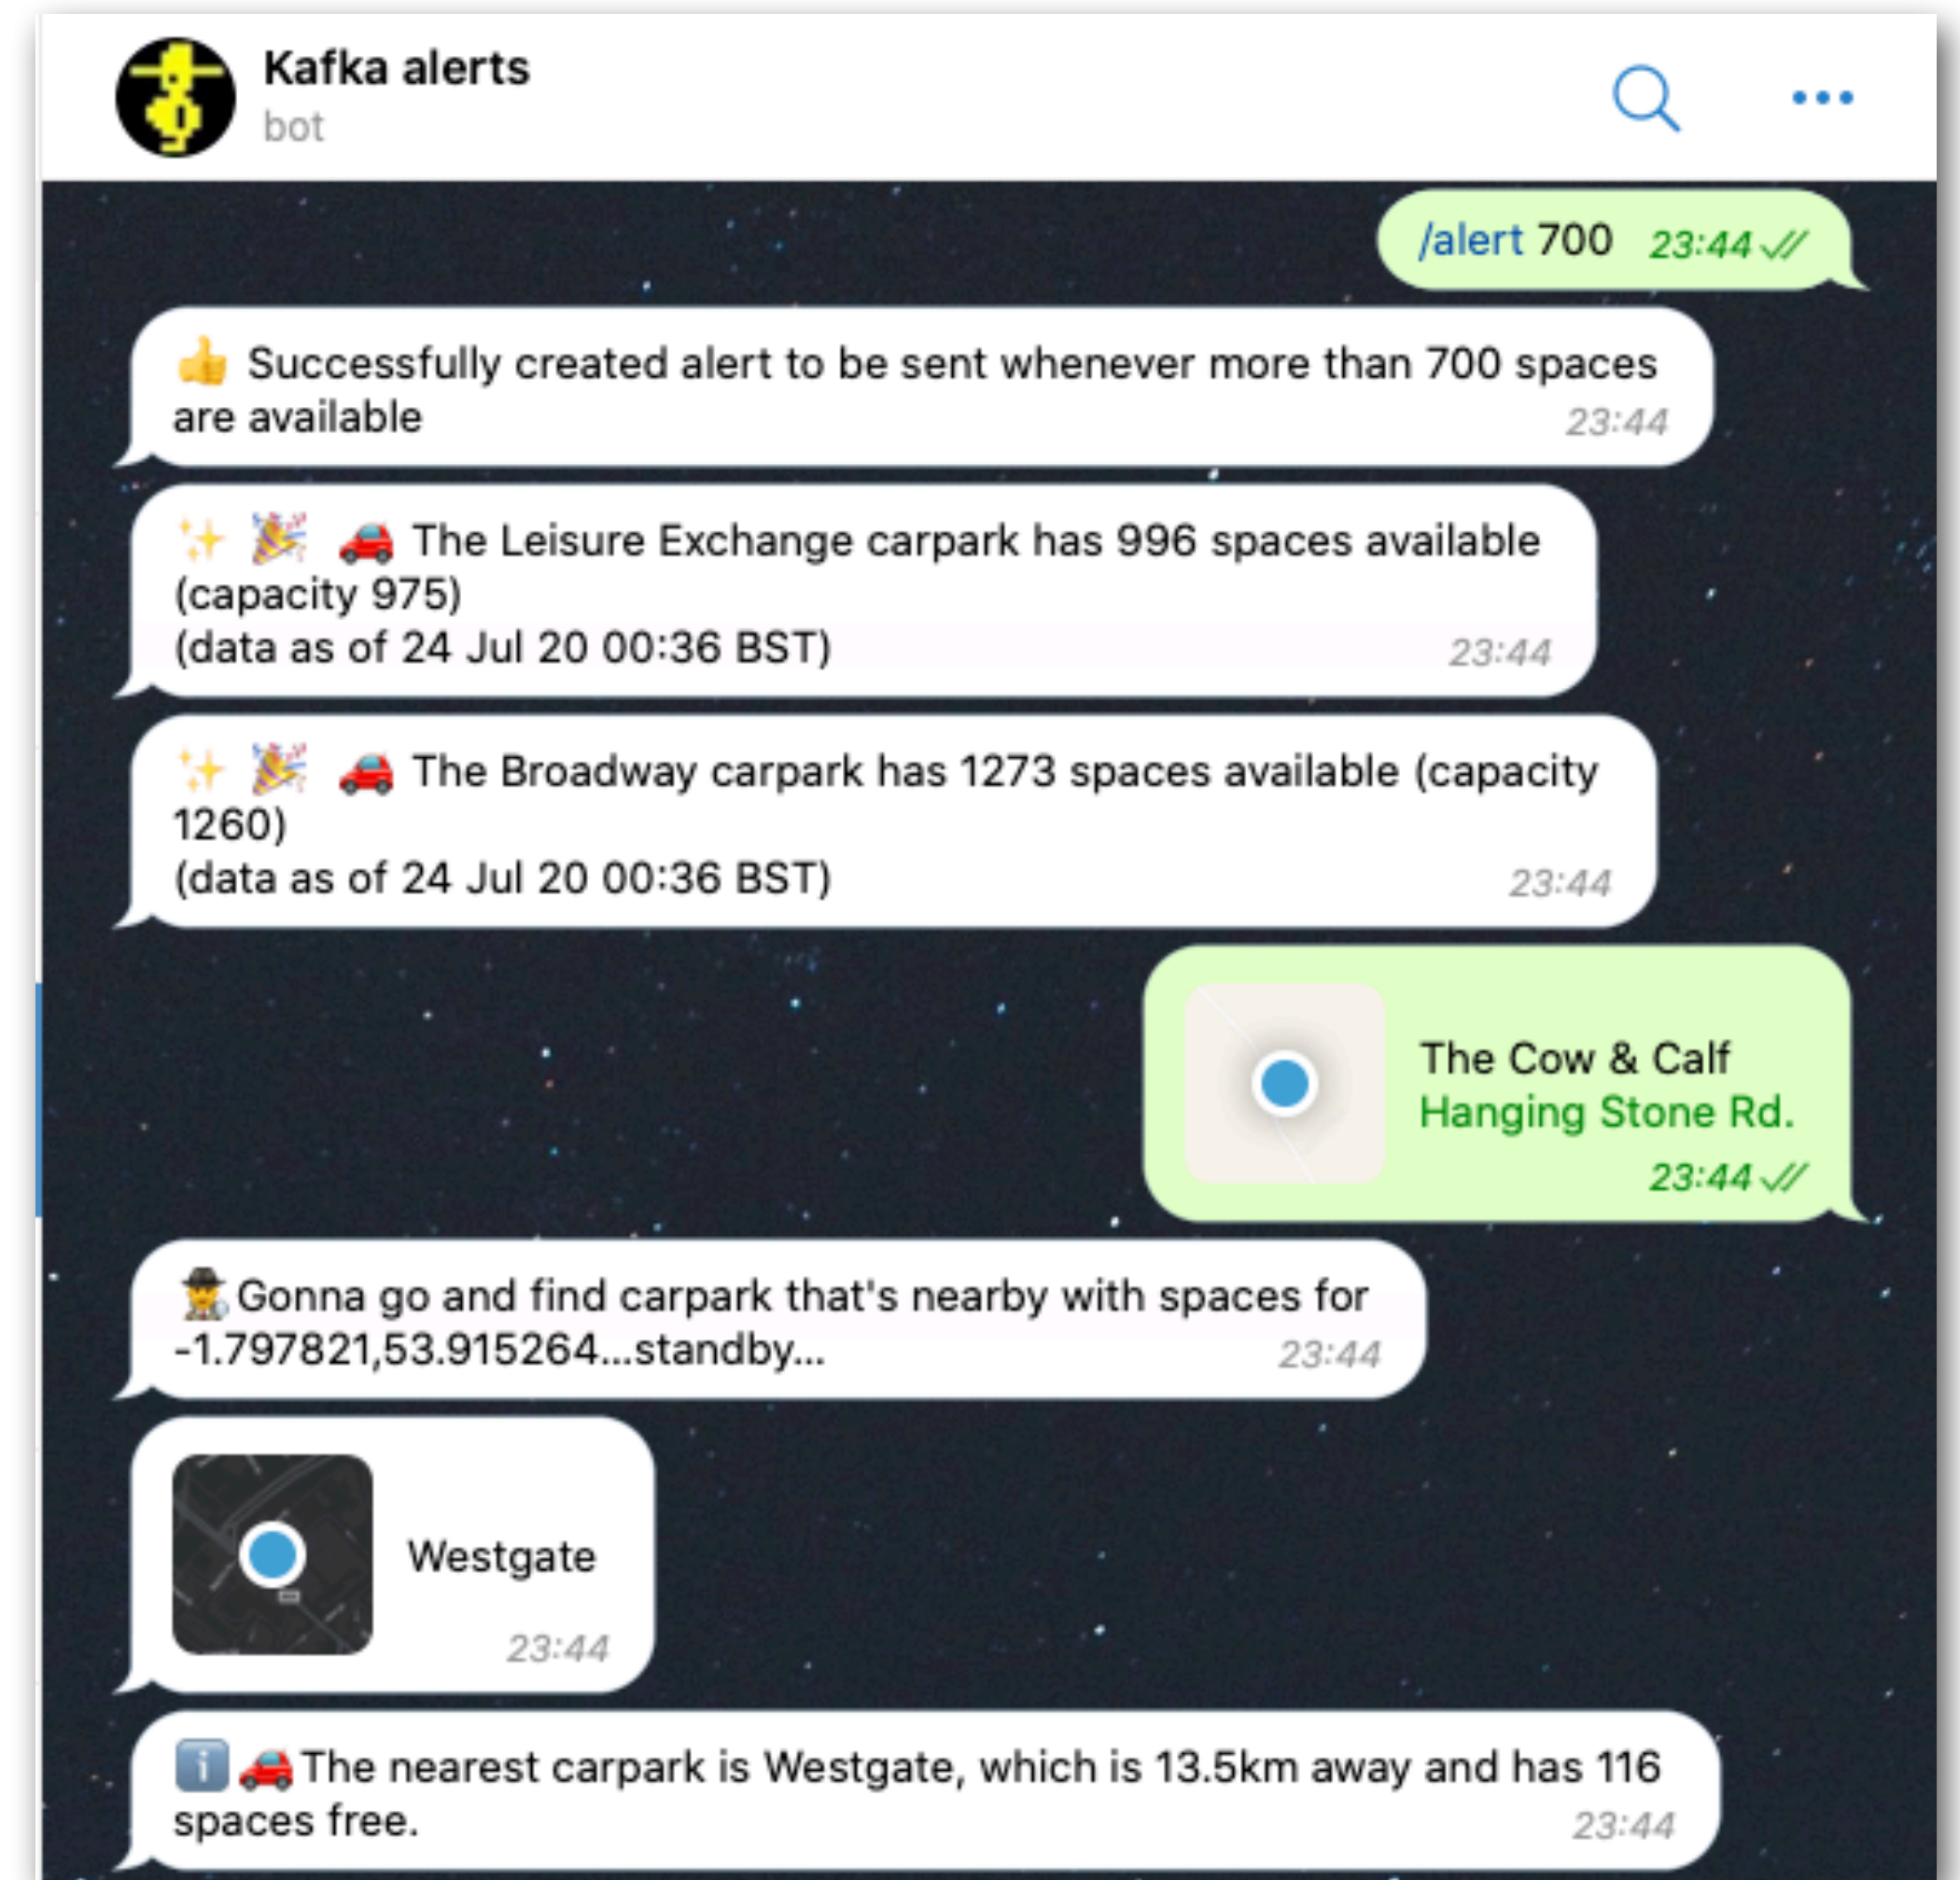
 [Map of car parks across Bradford district](#)

[html](#) [Live map of Bradford car parks](#)

License UK Open Government Licence (OGL v3)

Frequency daily

Telegram






Don't just tell me...

show me!




Demo code: <https://rmoff.dev/carparks>

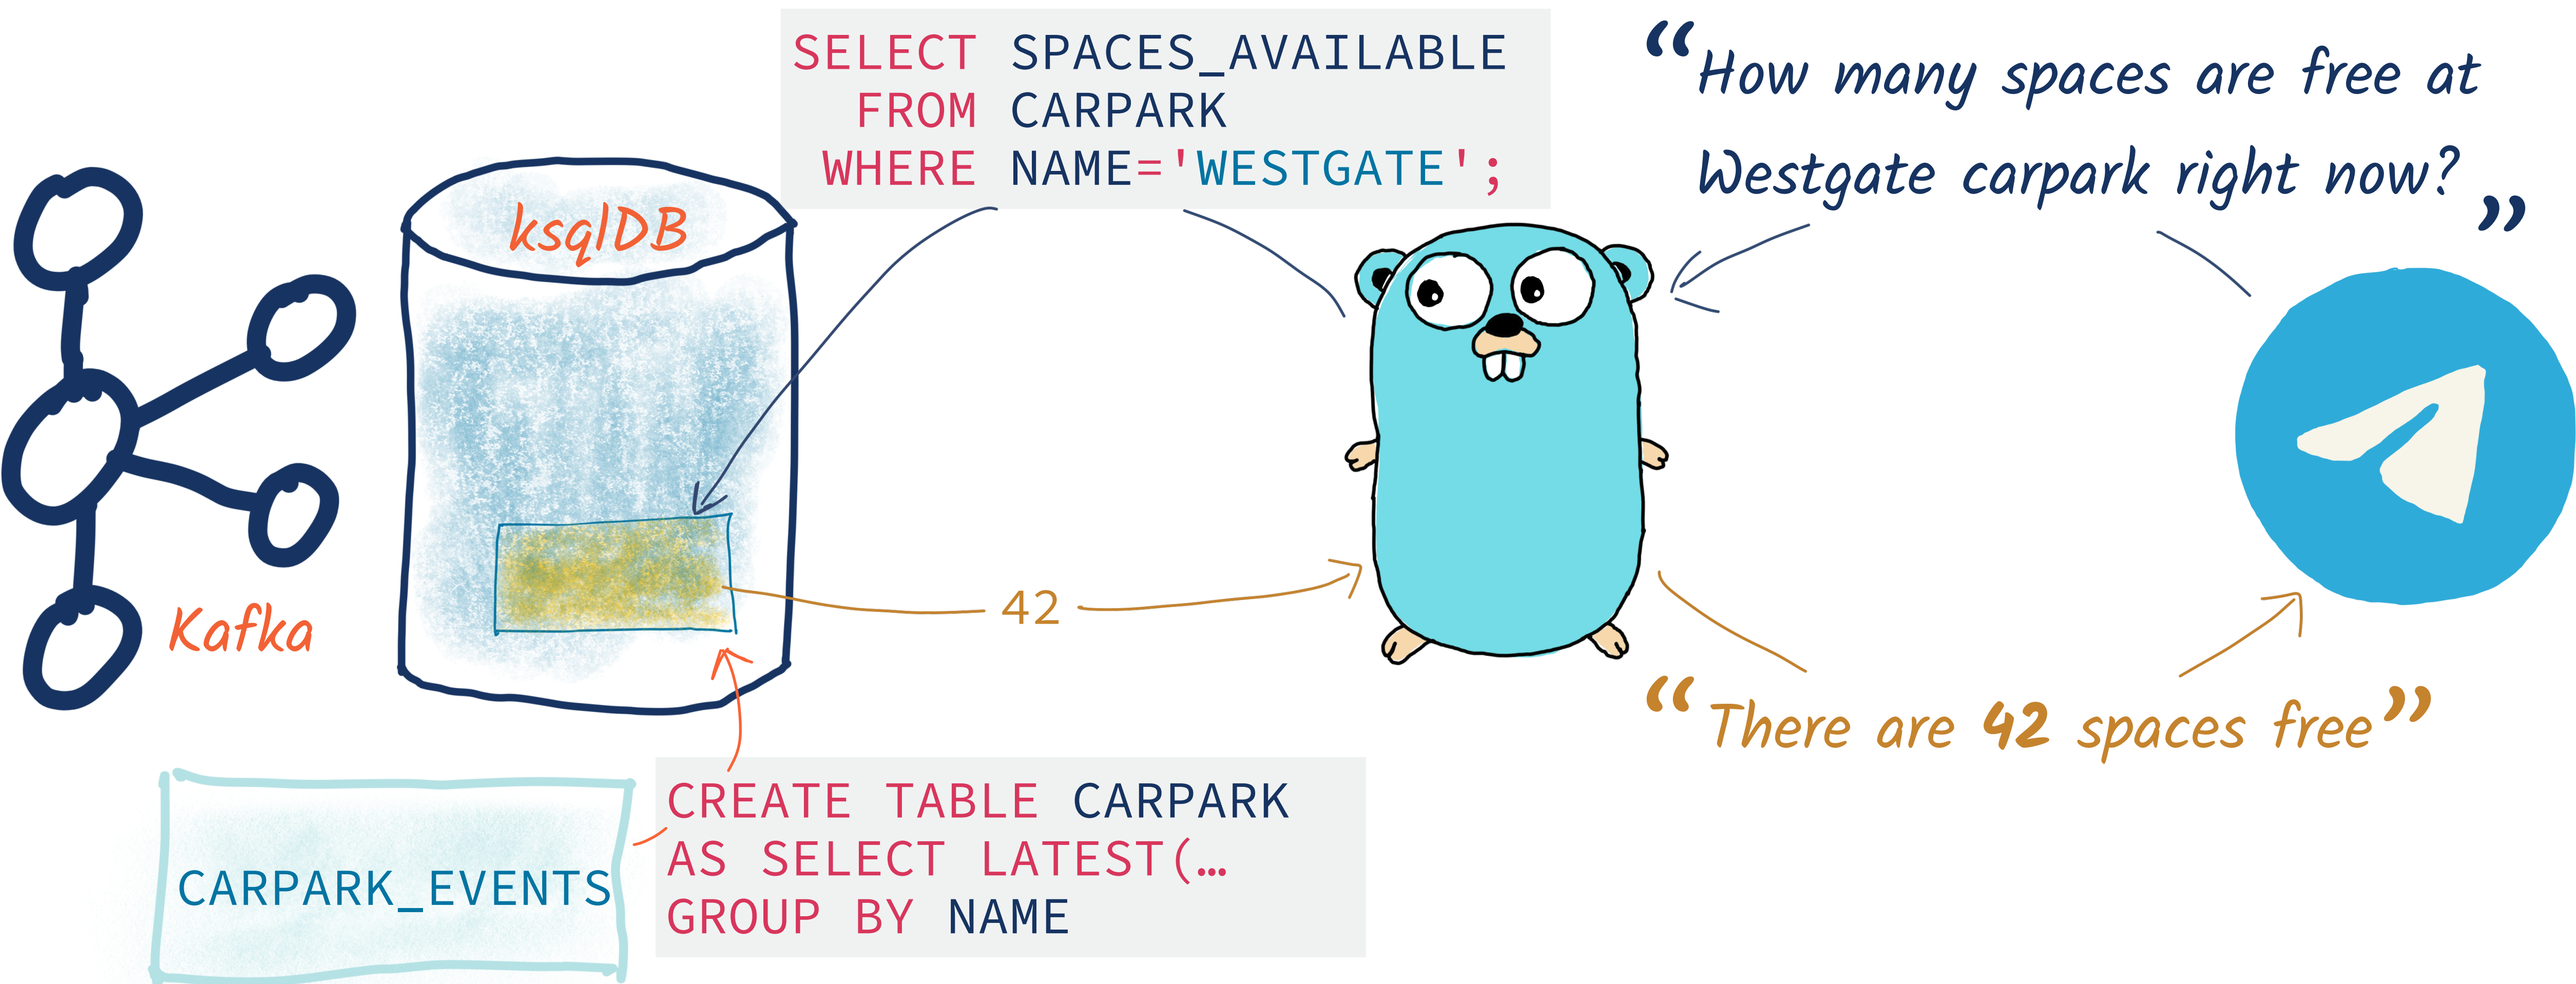


*What are the
key pieces of
the design?*


Event Driven Alerts




KN Lookups (materialised views)





A schema...




A schema...



```
2020-10-14,12:28,Broadway,1132,921  
2020-10-14,12:28,Kirkgate Centre,611,474  
2020-10-14,12:28,Sharpe Street,98,63
```



My kingdom for a schema!

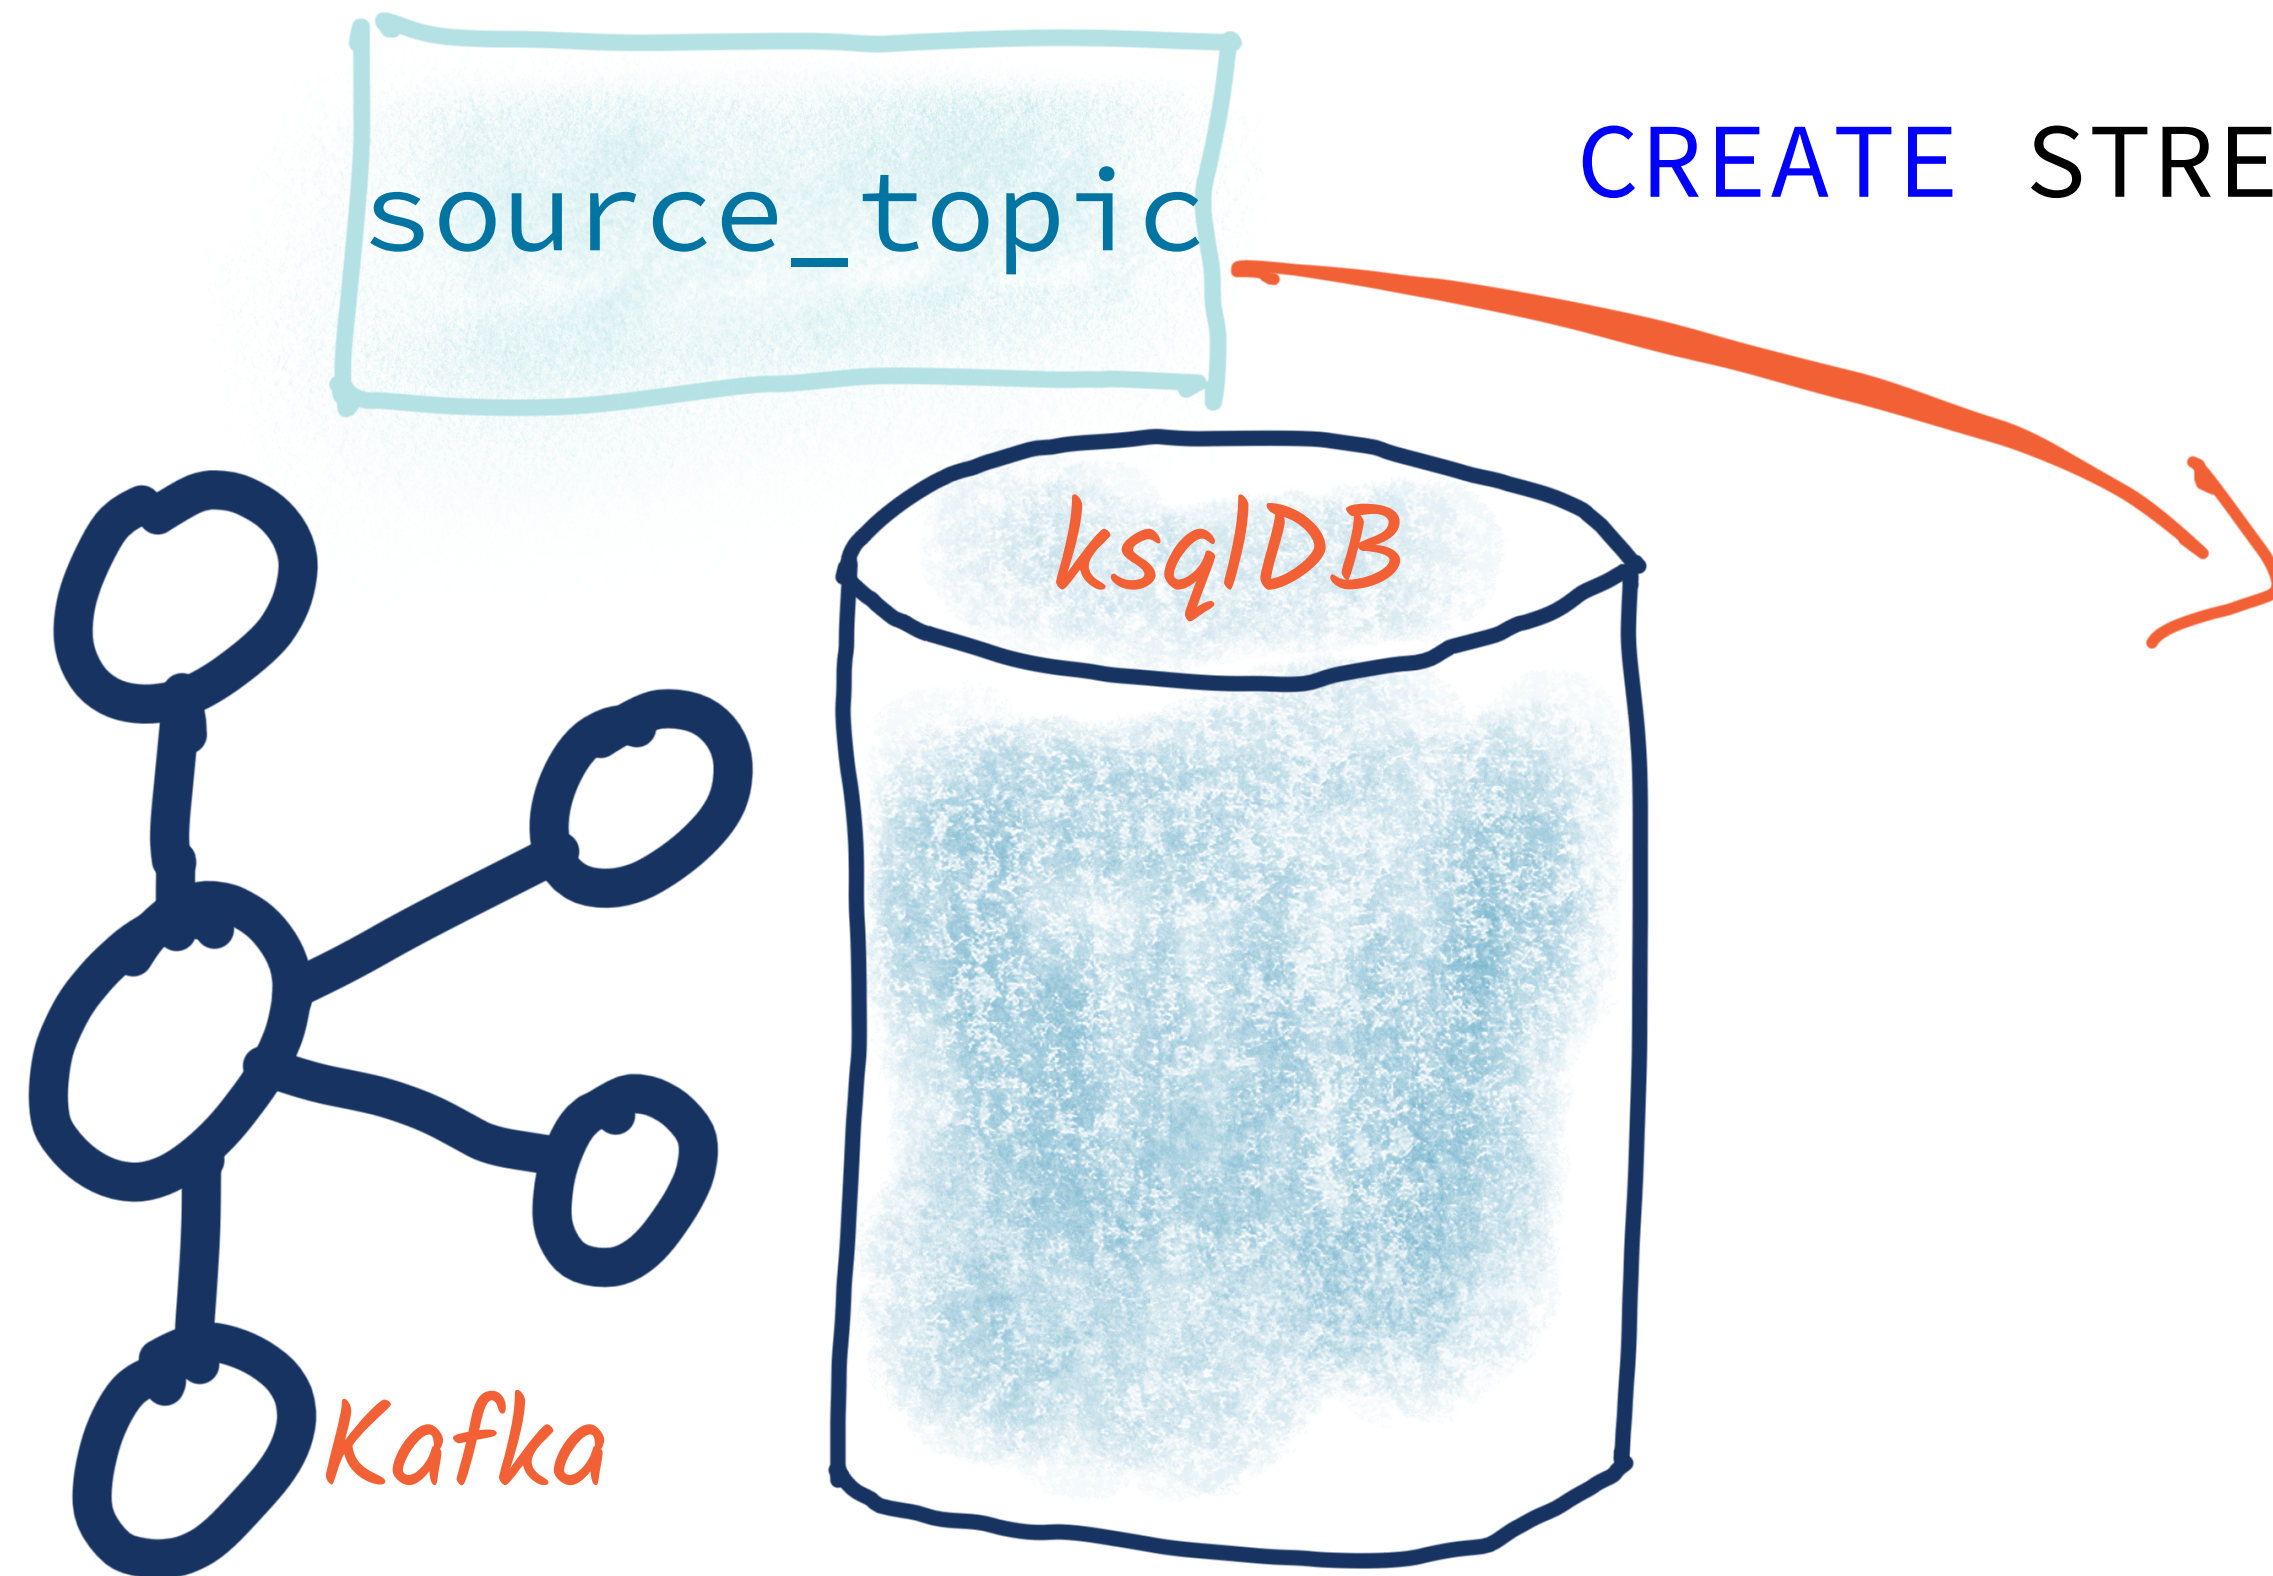


```
2020-10-14,12:28,Broadway,1132,921  
2020-10-14,12:28,Kirkgate Centre,611,474  
2020-10-14,12:28,Sharpe Street,98,63
```



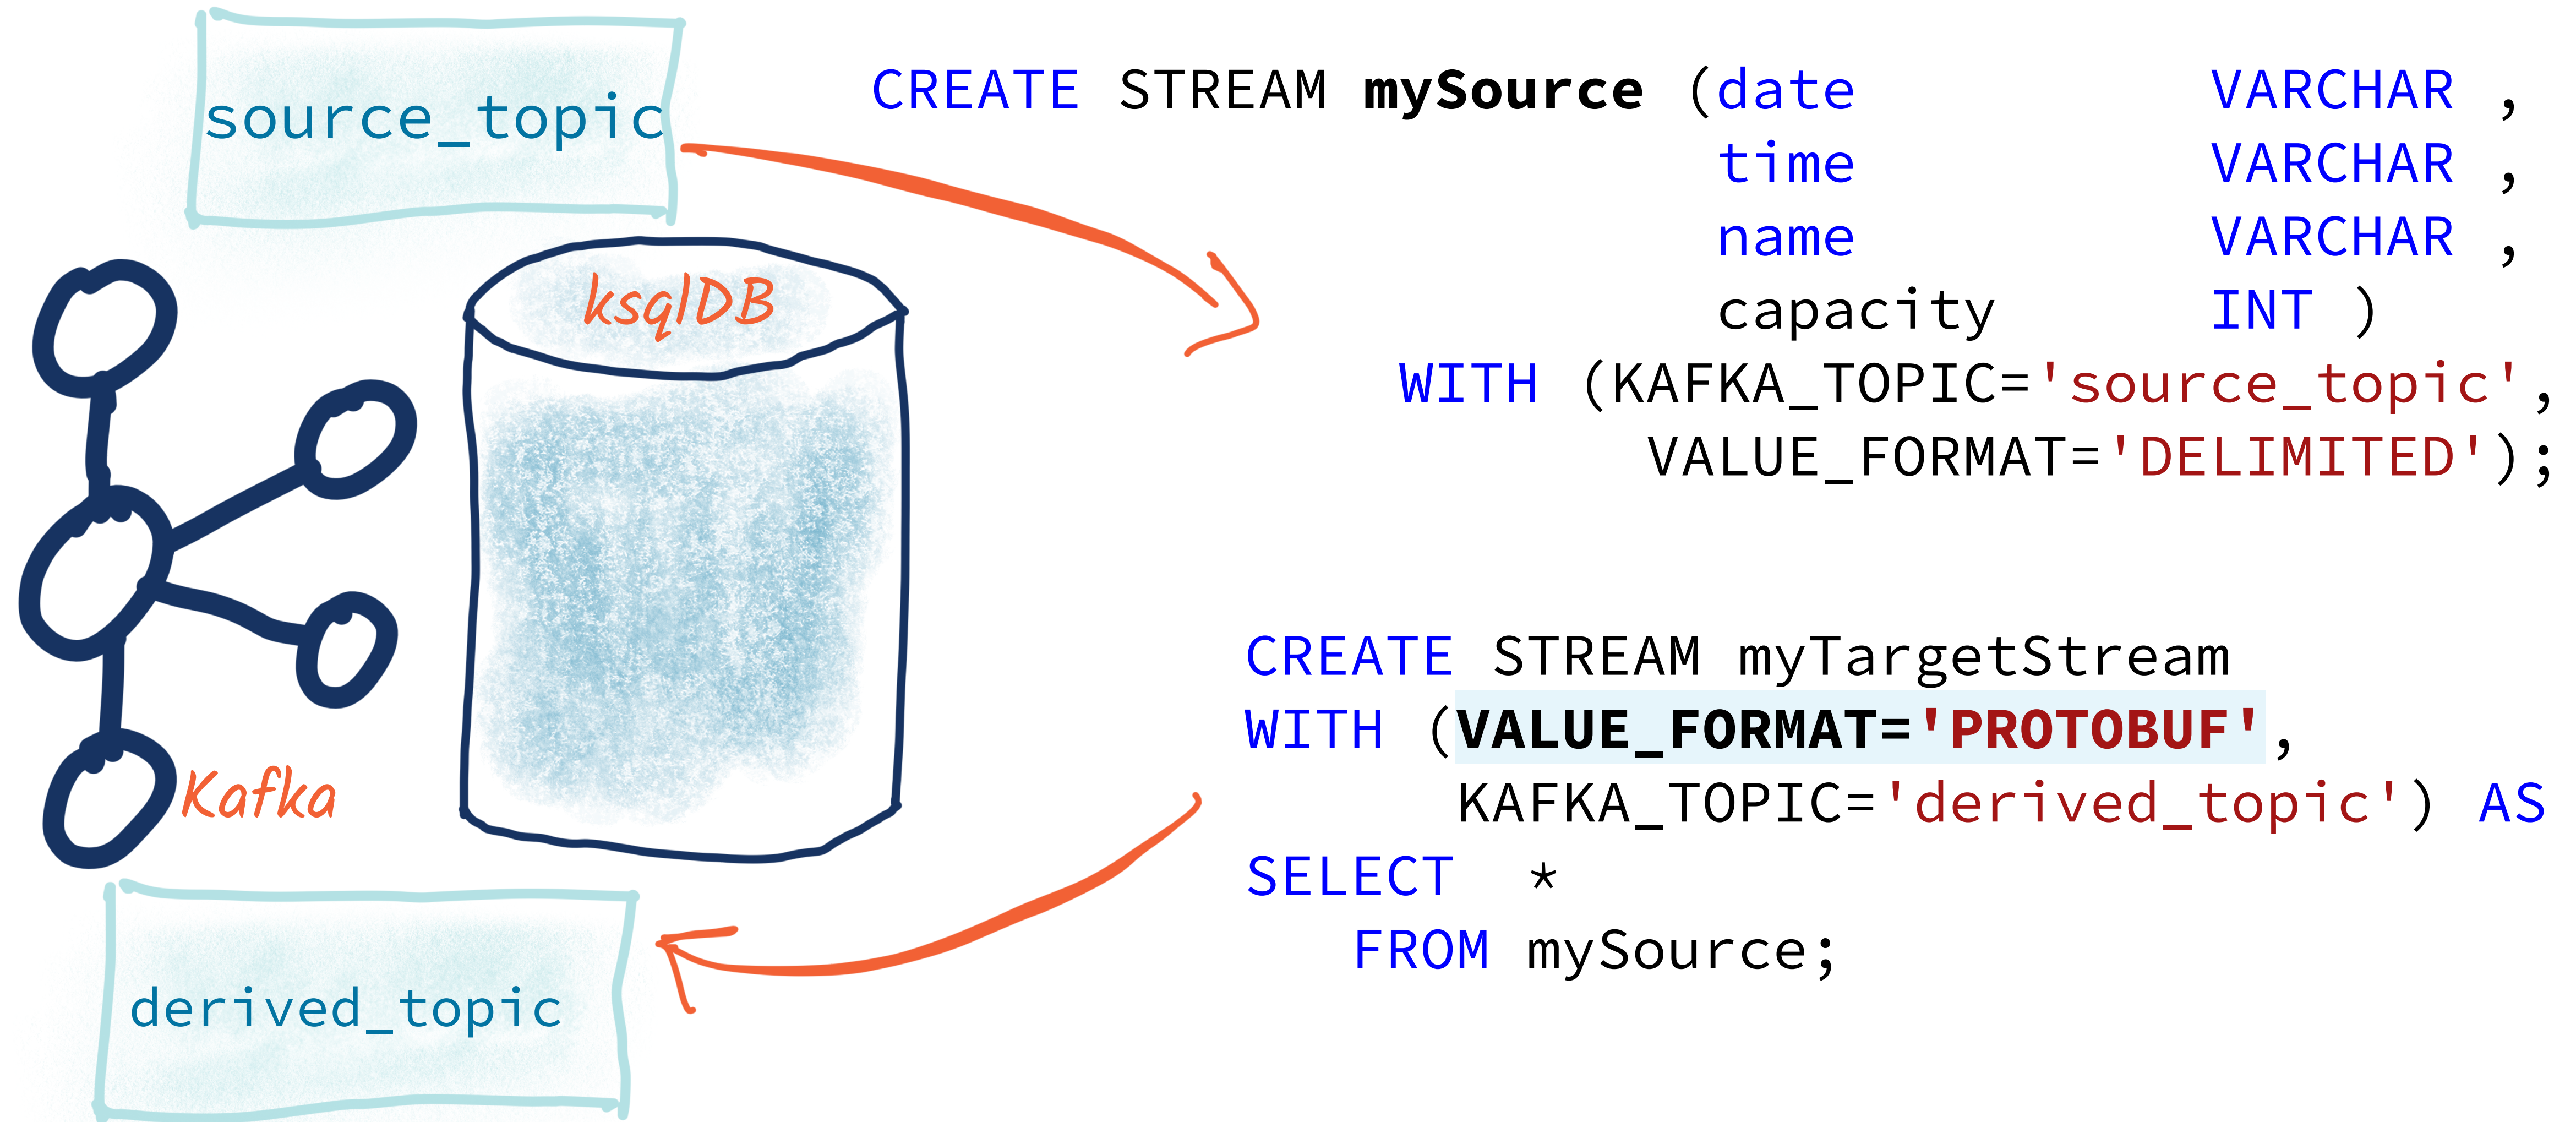
```
{  
  "ts": "2020-10-14T12:28 UTC+1",  
  "name": "Broadway",  
  "capacity": 1132,  
  "empty": 921  
}  
..
```

Applying a schema to streams of data




```
CREATE STREAM mySource (date          VARCHAR ,  
                        time         VARCHAR ,  
                        name          VARCHAR ,  
                        capacity      INT )  
WITH (KAFKA_TOPIC='source_topic',  
      VALUE_FORMAT='DELIMITED');
```


Applying a schema to streams of data




Integration



Streaming Integration with Kafka Connect




Kafka Connect




Kafka Brokers

Streaming Integration with Kafka Connect




Kafka Connect




Streaming Integration with Kafka Connect

Logos of various data sources and connectors: JSON, MQTT, MySQL, Oracle, syslog, http://, CSV, Salesforce, and Microsoft SQL Server.

Logos of various data destinations and connectors: mongoDB, elasticsearch, influxdb, snowflake, Amazon S3, Java JDBC, IBM MQ, Google BigQuery, salesforce, http://, Oracle, neo4j, splunk, hadoop HDFS, and MQTT.

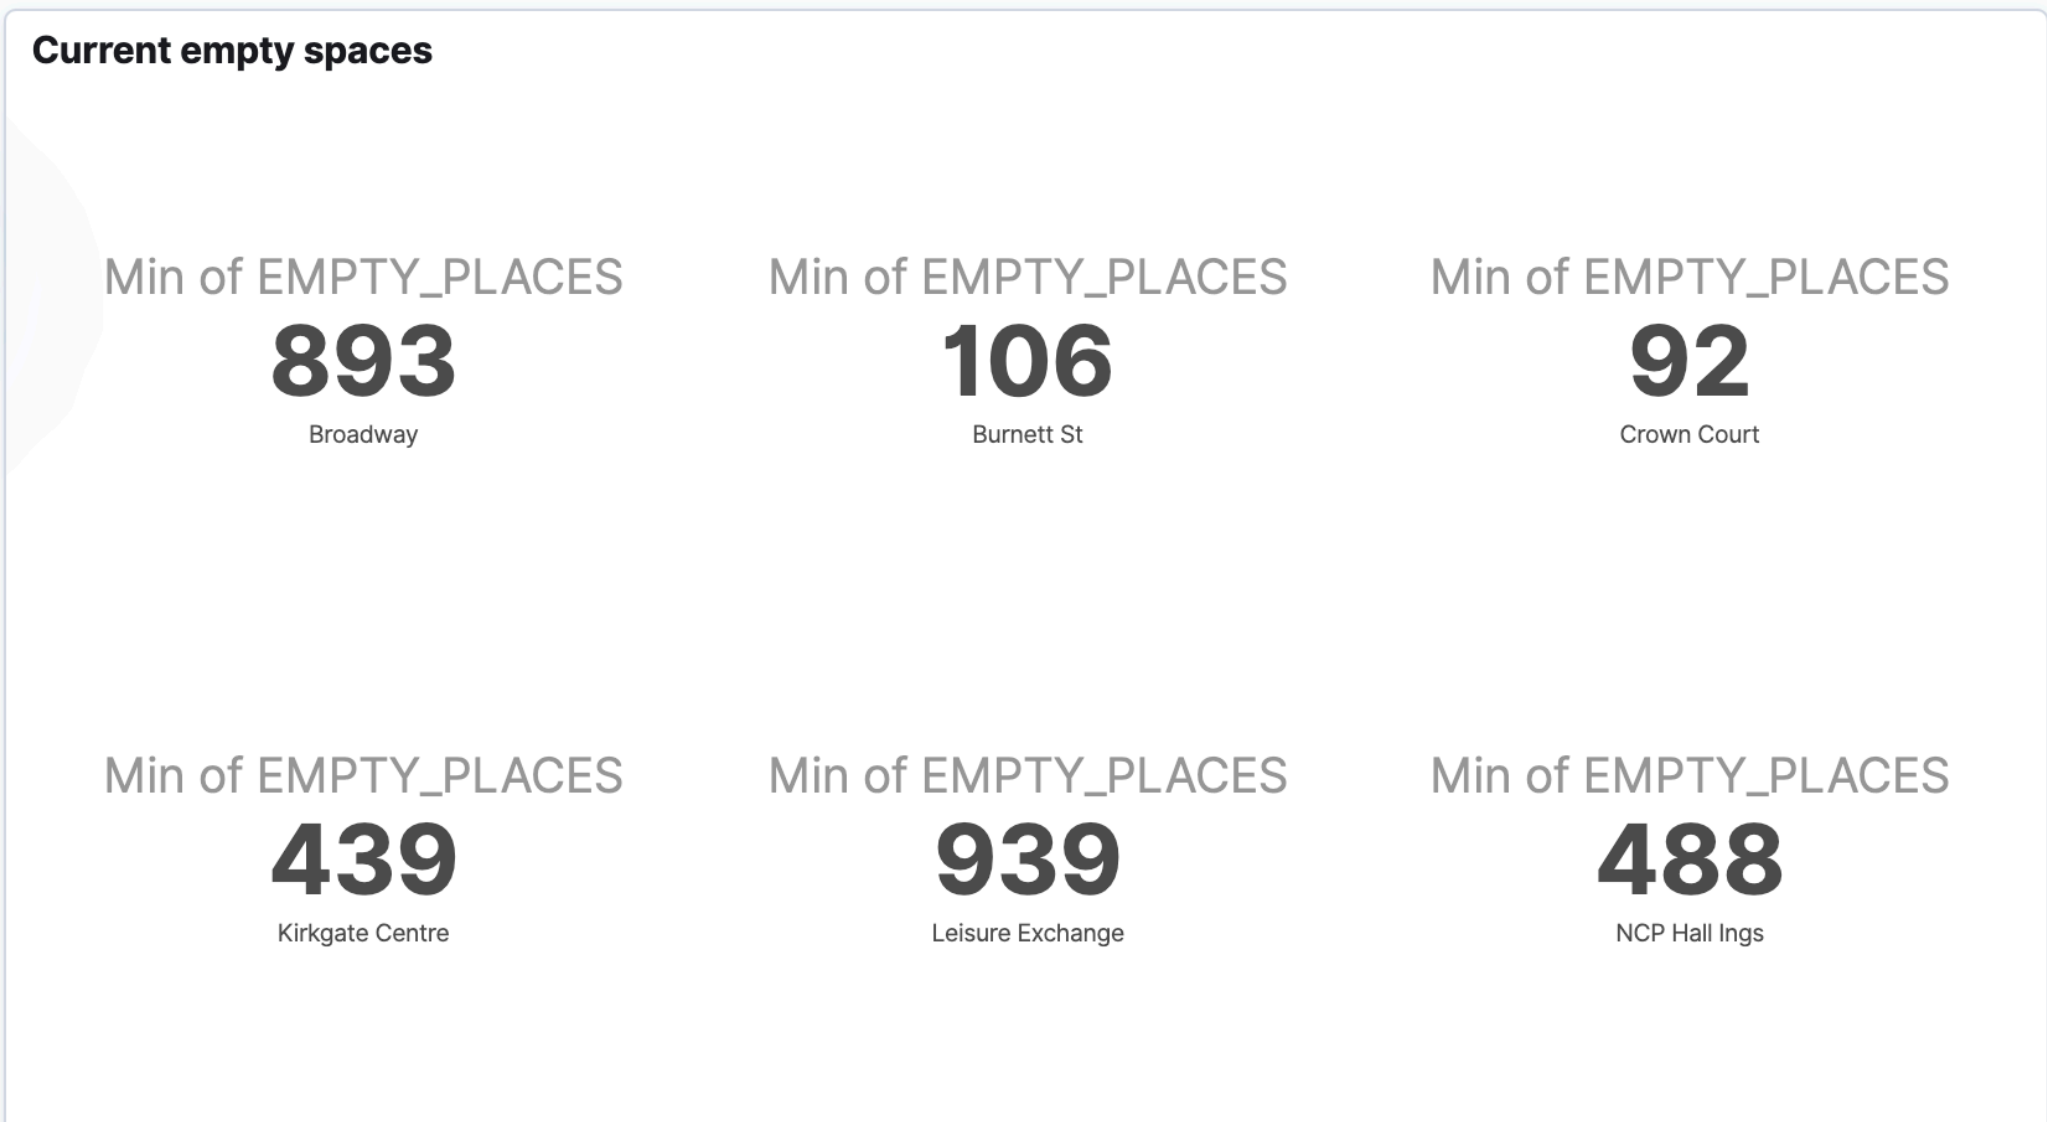
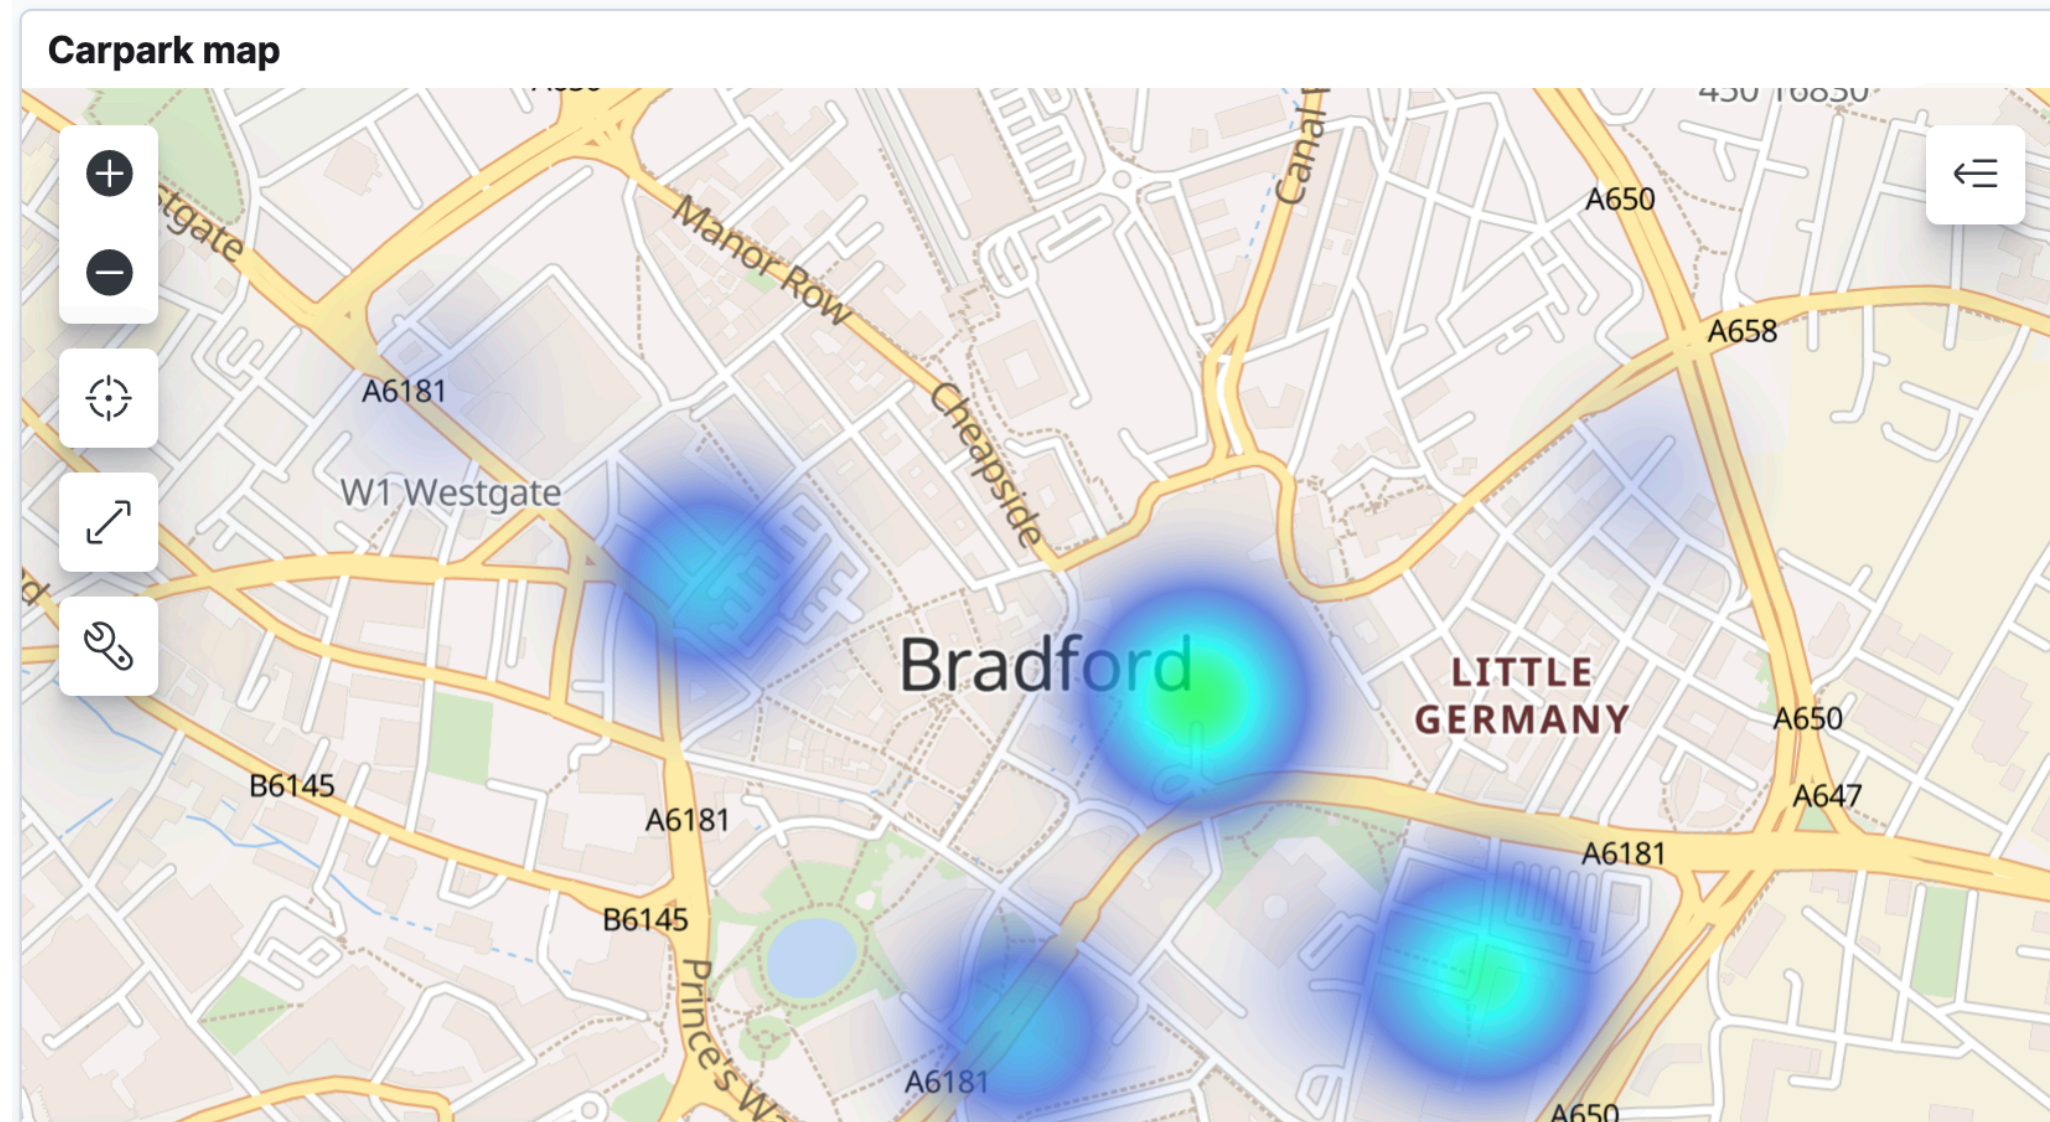



Kafka Connect



Kafka Brokers

Streaming Analytics



Why build
it this way?

Events

Streams

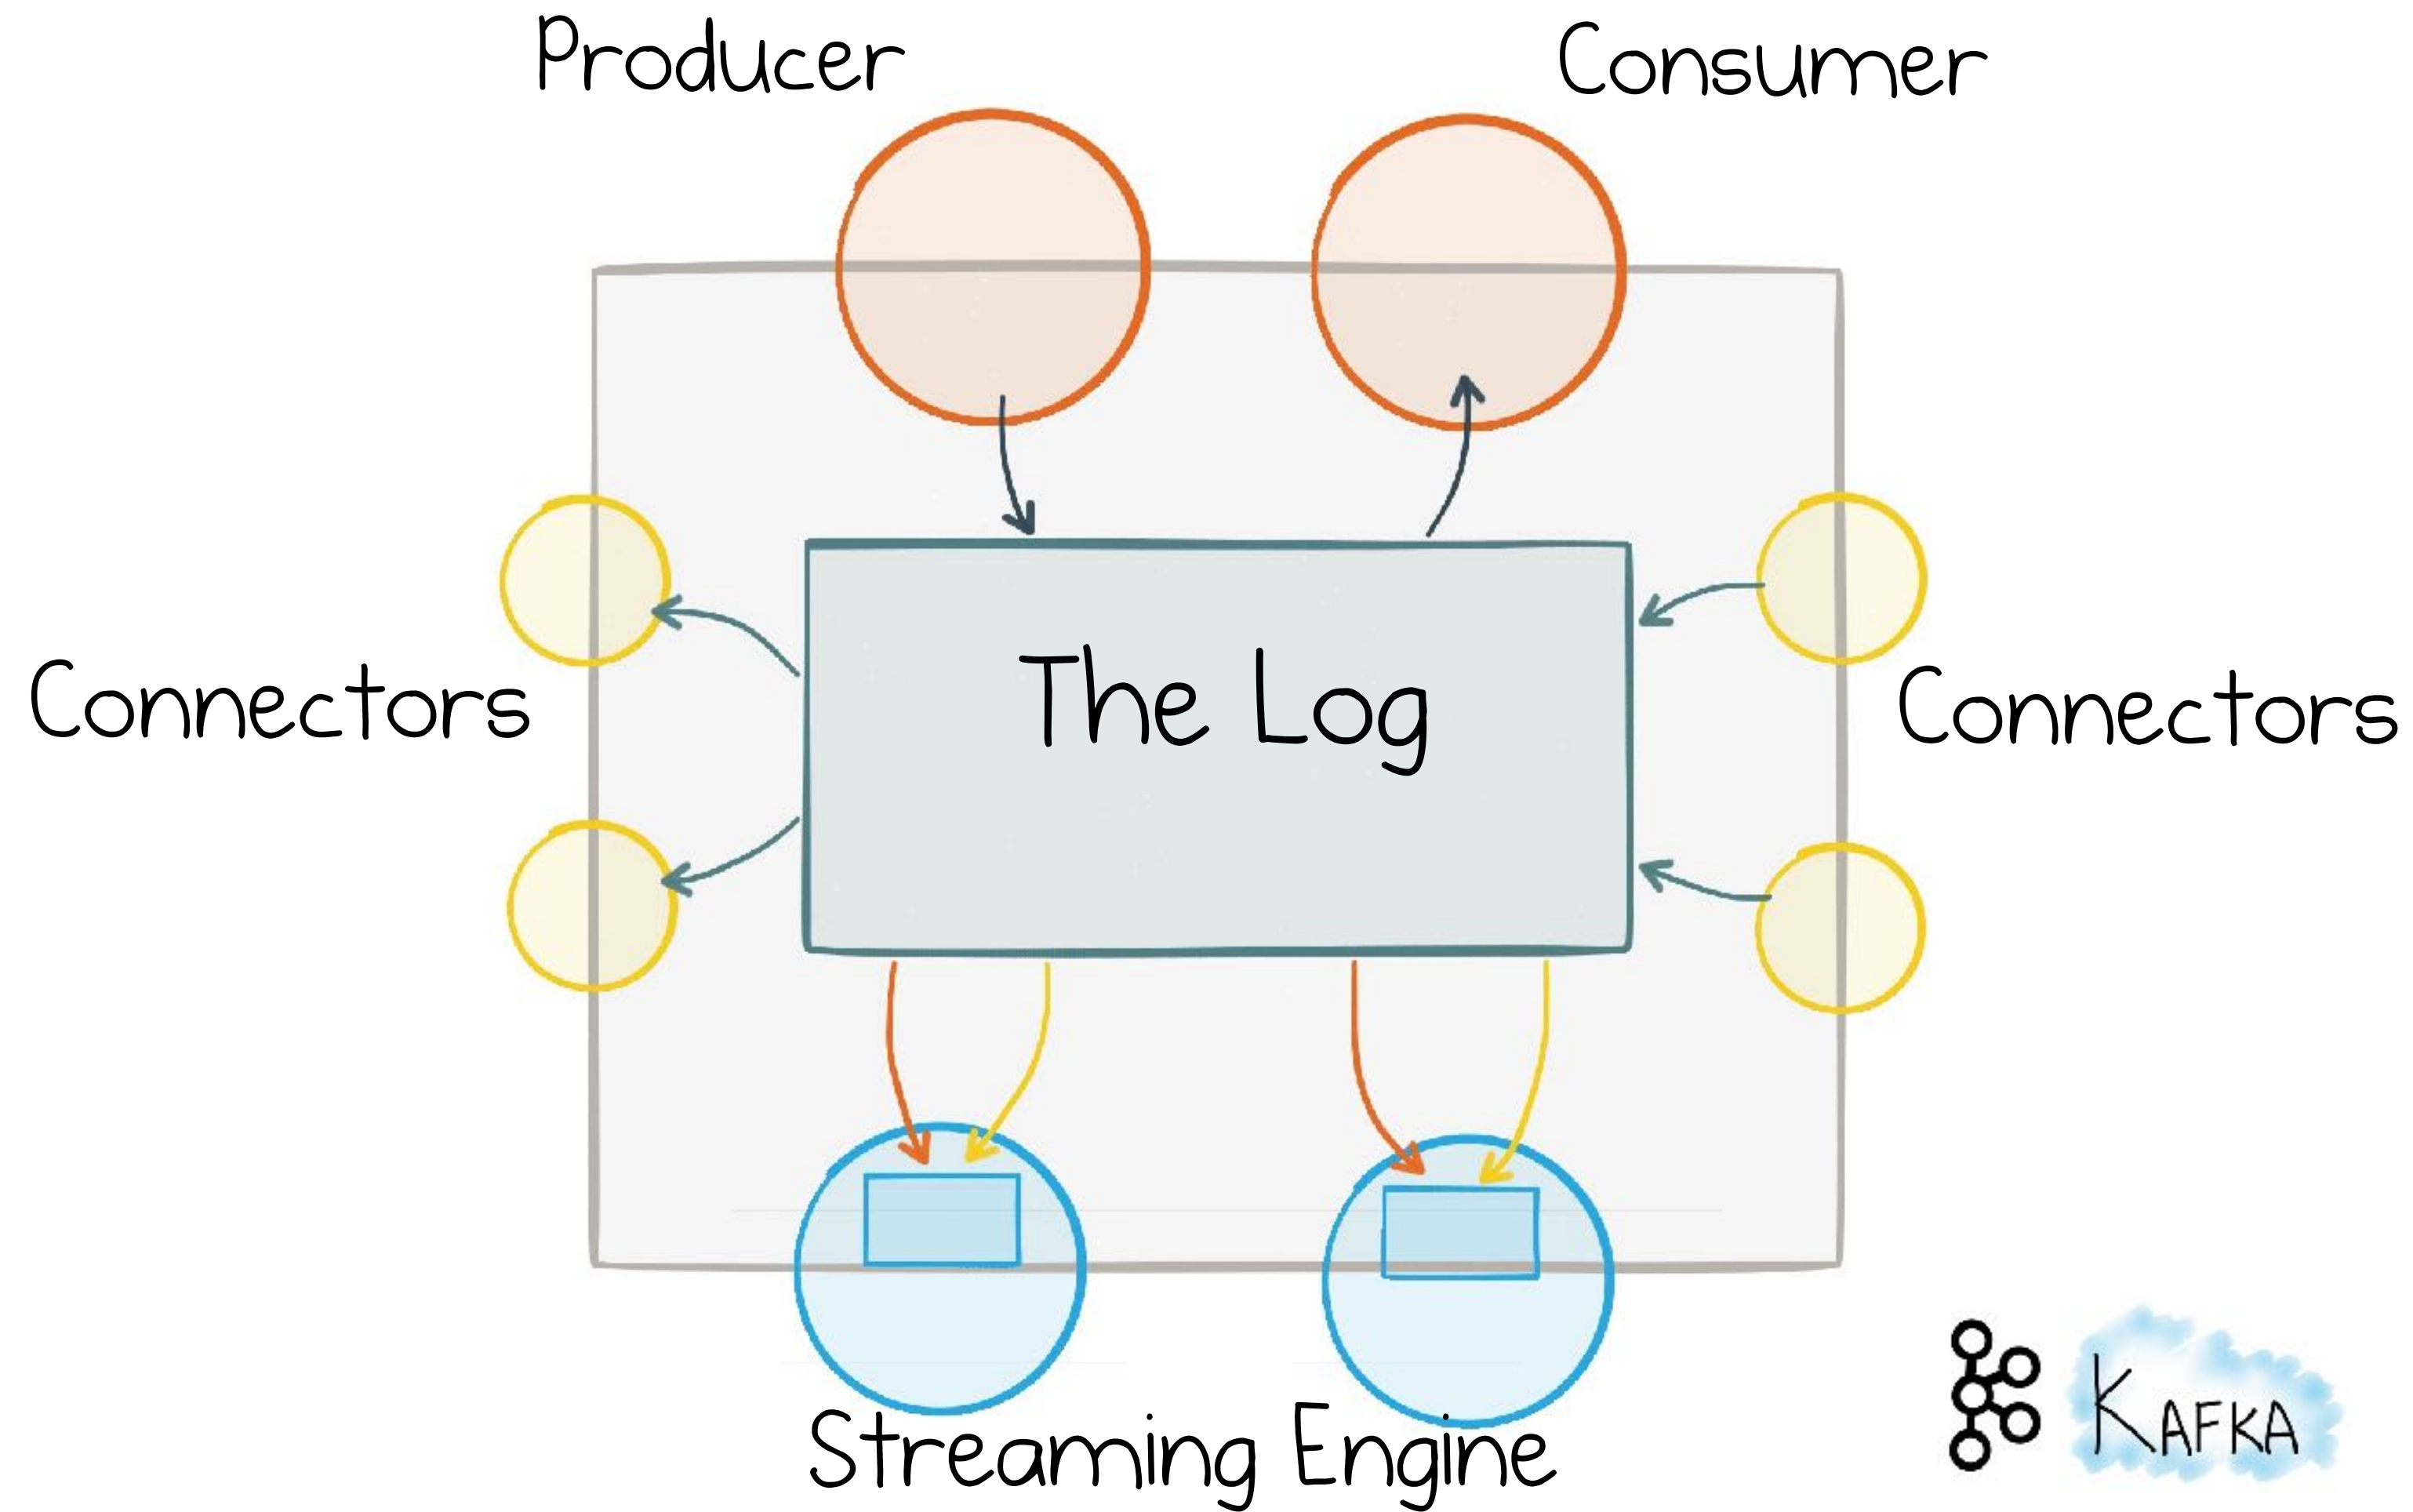
of Events

We want to *react* to
them as they happen

We want to build
state from a stream
of events

We want to provide
the *latest* data in our
analytics

Apache Kafka - an Event Streaming Platform



Why

Kafka?


Distributed, Immutable, Event Log




Events are added at the end of the log

Consumers can seek to any point


Read to offset & scan



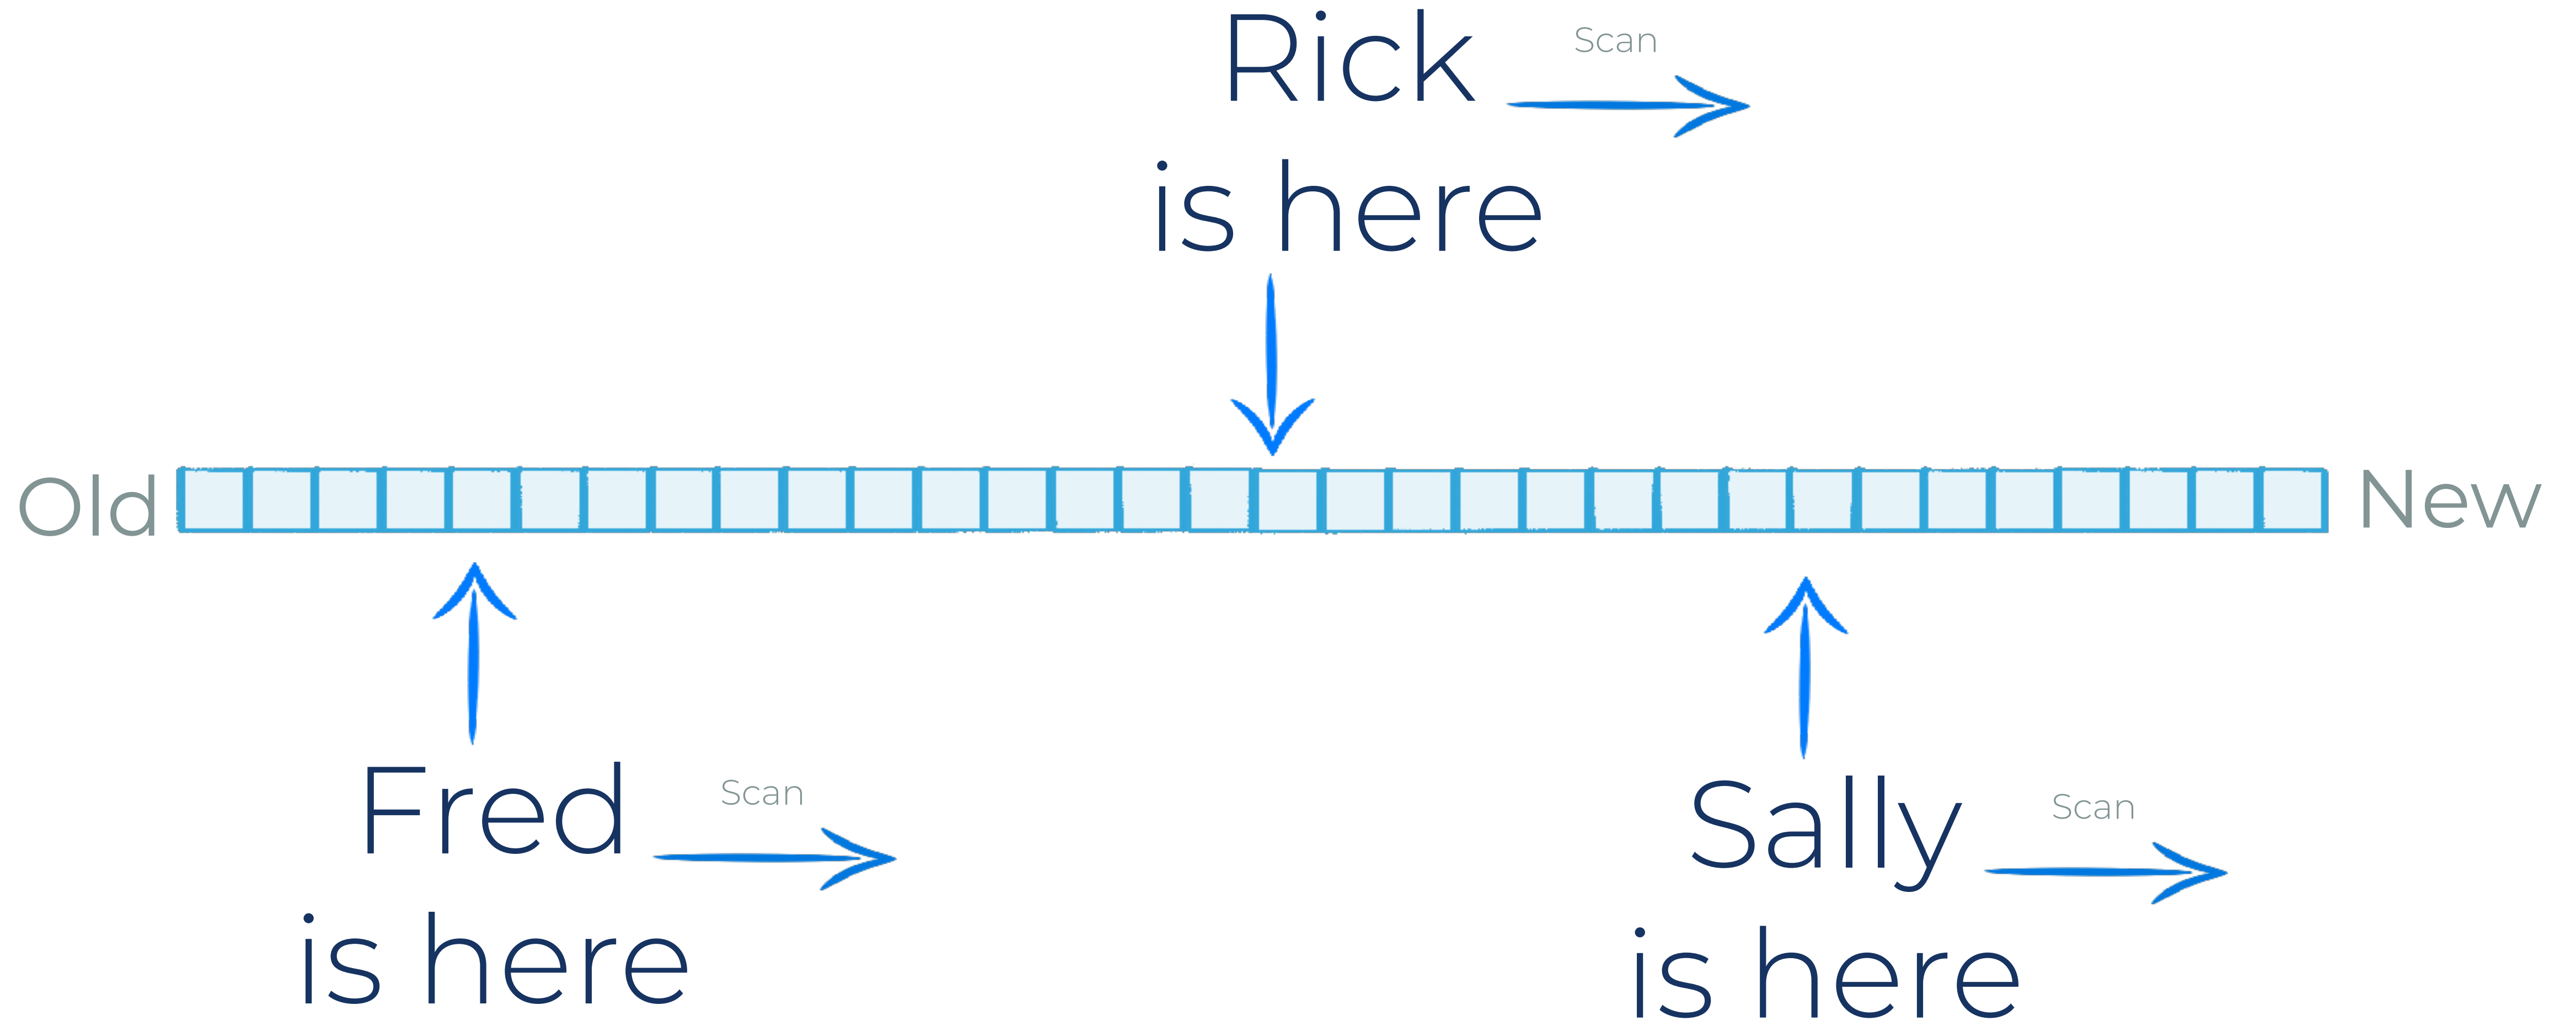
Data is not deleted once read



Consumers are independent of each other




Consumers can be added later







Stream Processing with ksqlDB

Source stream






Stream Processing with ksqlDB

Source stream







Stream Processing with ksqlDB

Source stream




Stream Processing with ksqlDB

Source stream

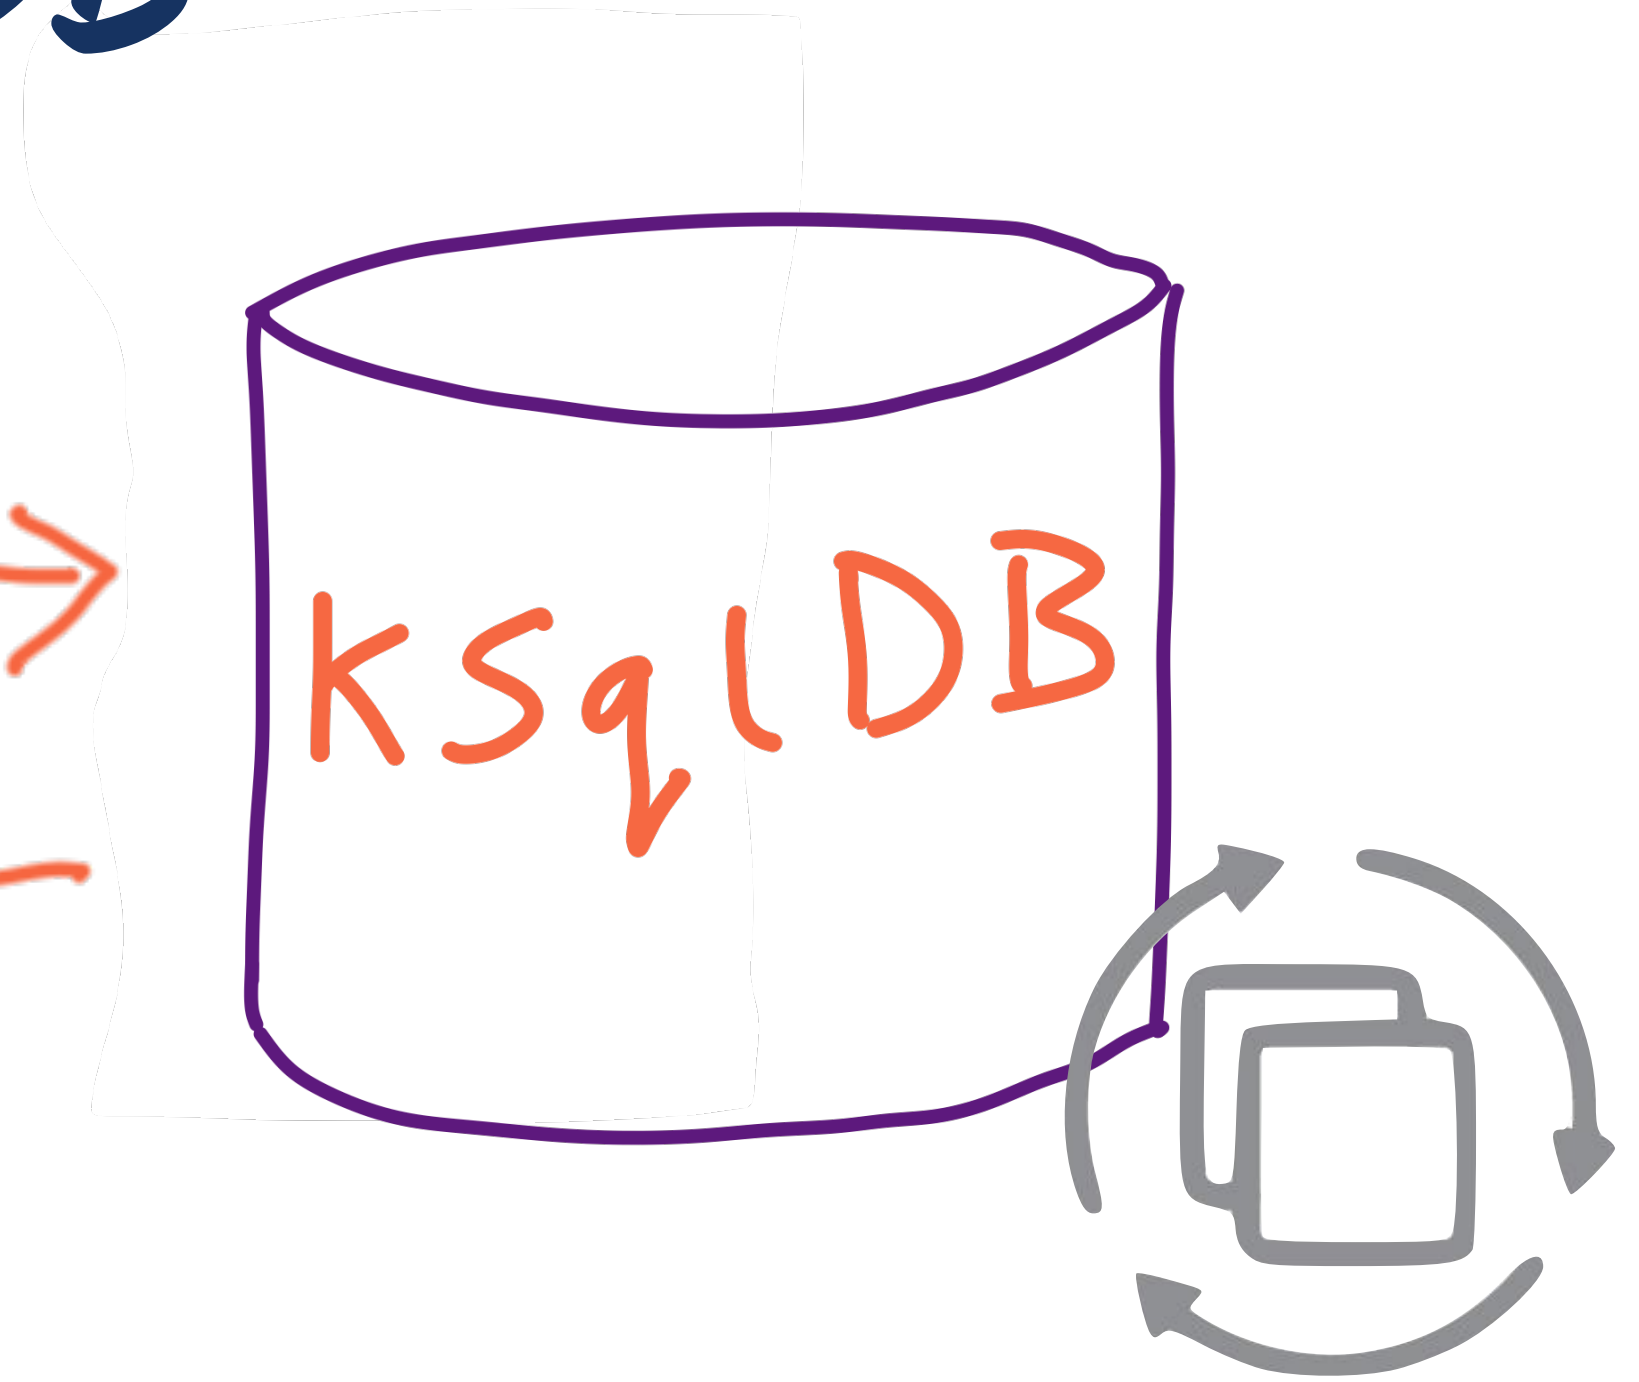




Analytics



Stream Processing with ksqlDB

Source stream






Applications /
Microservices




Stream Processing with ksqlDB

...SUM(TXN_AMT)
GROUP BY AC_ID

Source stream



Applications /
Microservices





AC_ID=42

BALANCE=94.00
AC_ID=42


*Under the
covers of
ksqlDB*



Kafka cluster



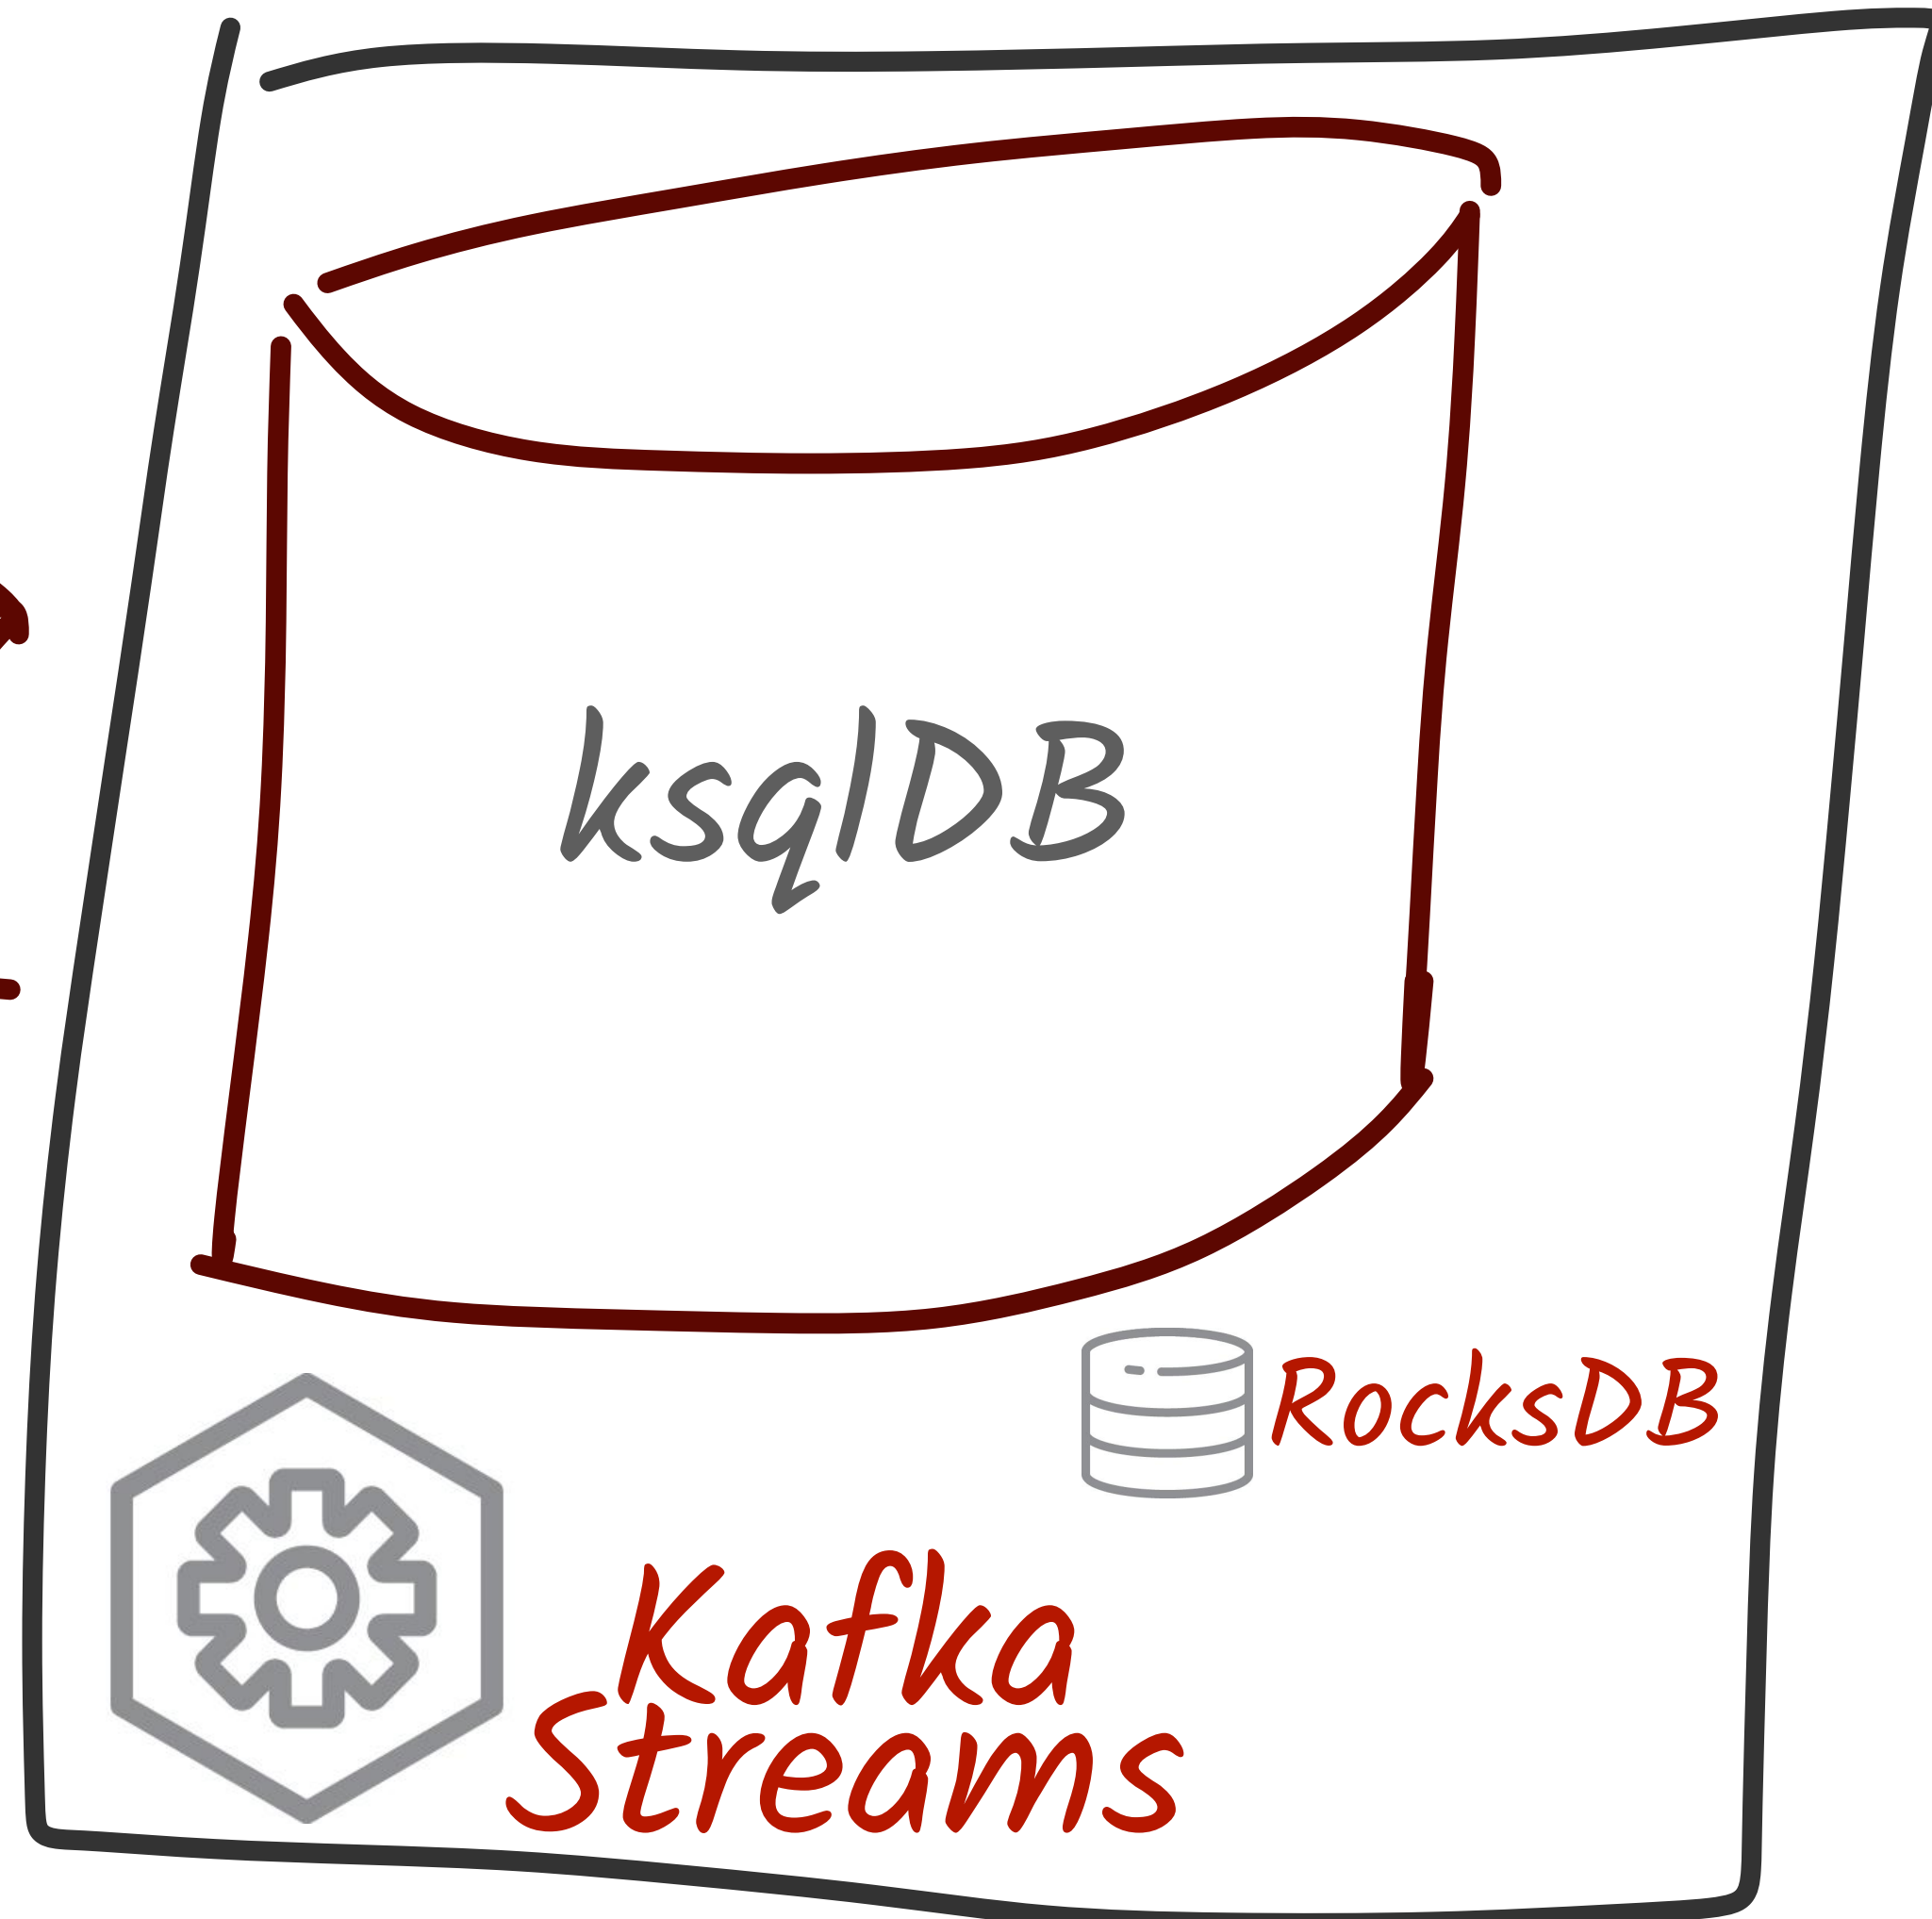
Kafka cluster



consume

produce

JVM



WHERE CAN YOU RUN KSQLDB?




ANYWHERE!



confluent cloud

Fully Managed Kafka ^{& KSQLDB} as a Service

Running ksqlDB - self-managed



ksqlDB Server
(JVM process)

DEB, RPM, ZIP, TAR downloads

<http://confluent.io/download>

Docker images

[confluentinc/ksqldb-server](https://hub.docker.com/r/confluentinc/ksqldb-server)



Physical



docker



kubernetes



openstack®

vmware®



Microsoft
Azure



Google Cloud Platform



amazon
web services

...and many more...

Why

Kafka?

Stream Store

Process Integrate

Stream

Store

Process Integrate

Stream

Store

Process

Integrate

Stream

Store

Process Integrate

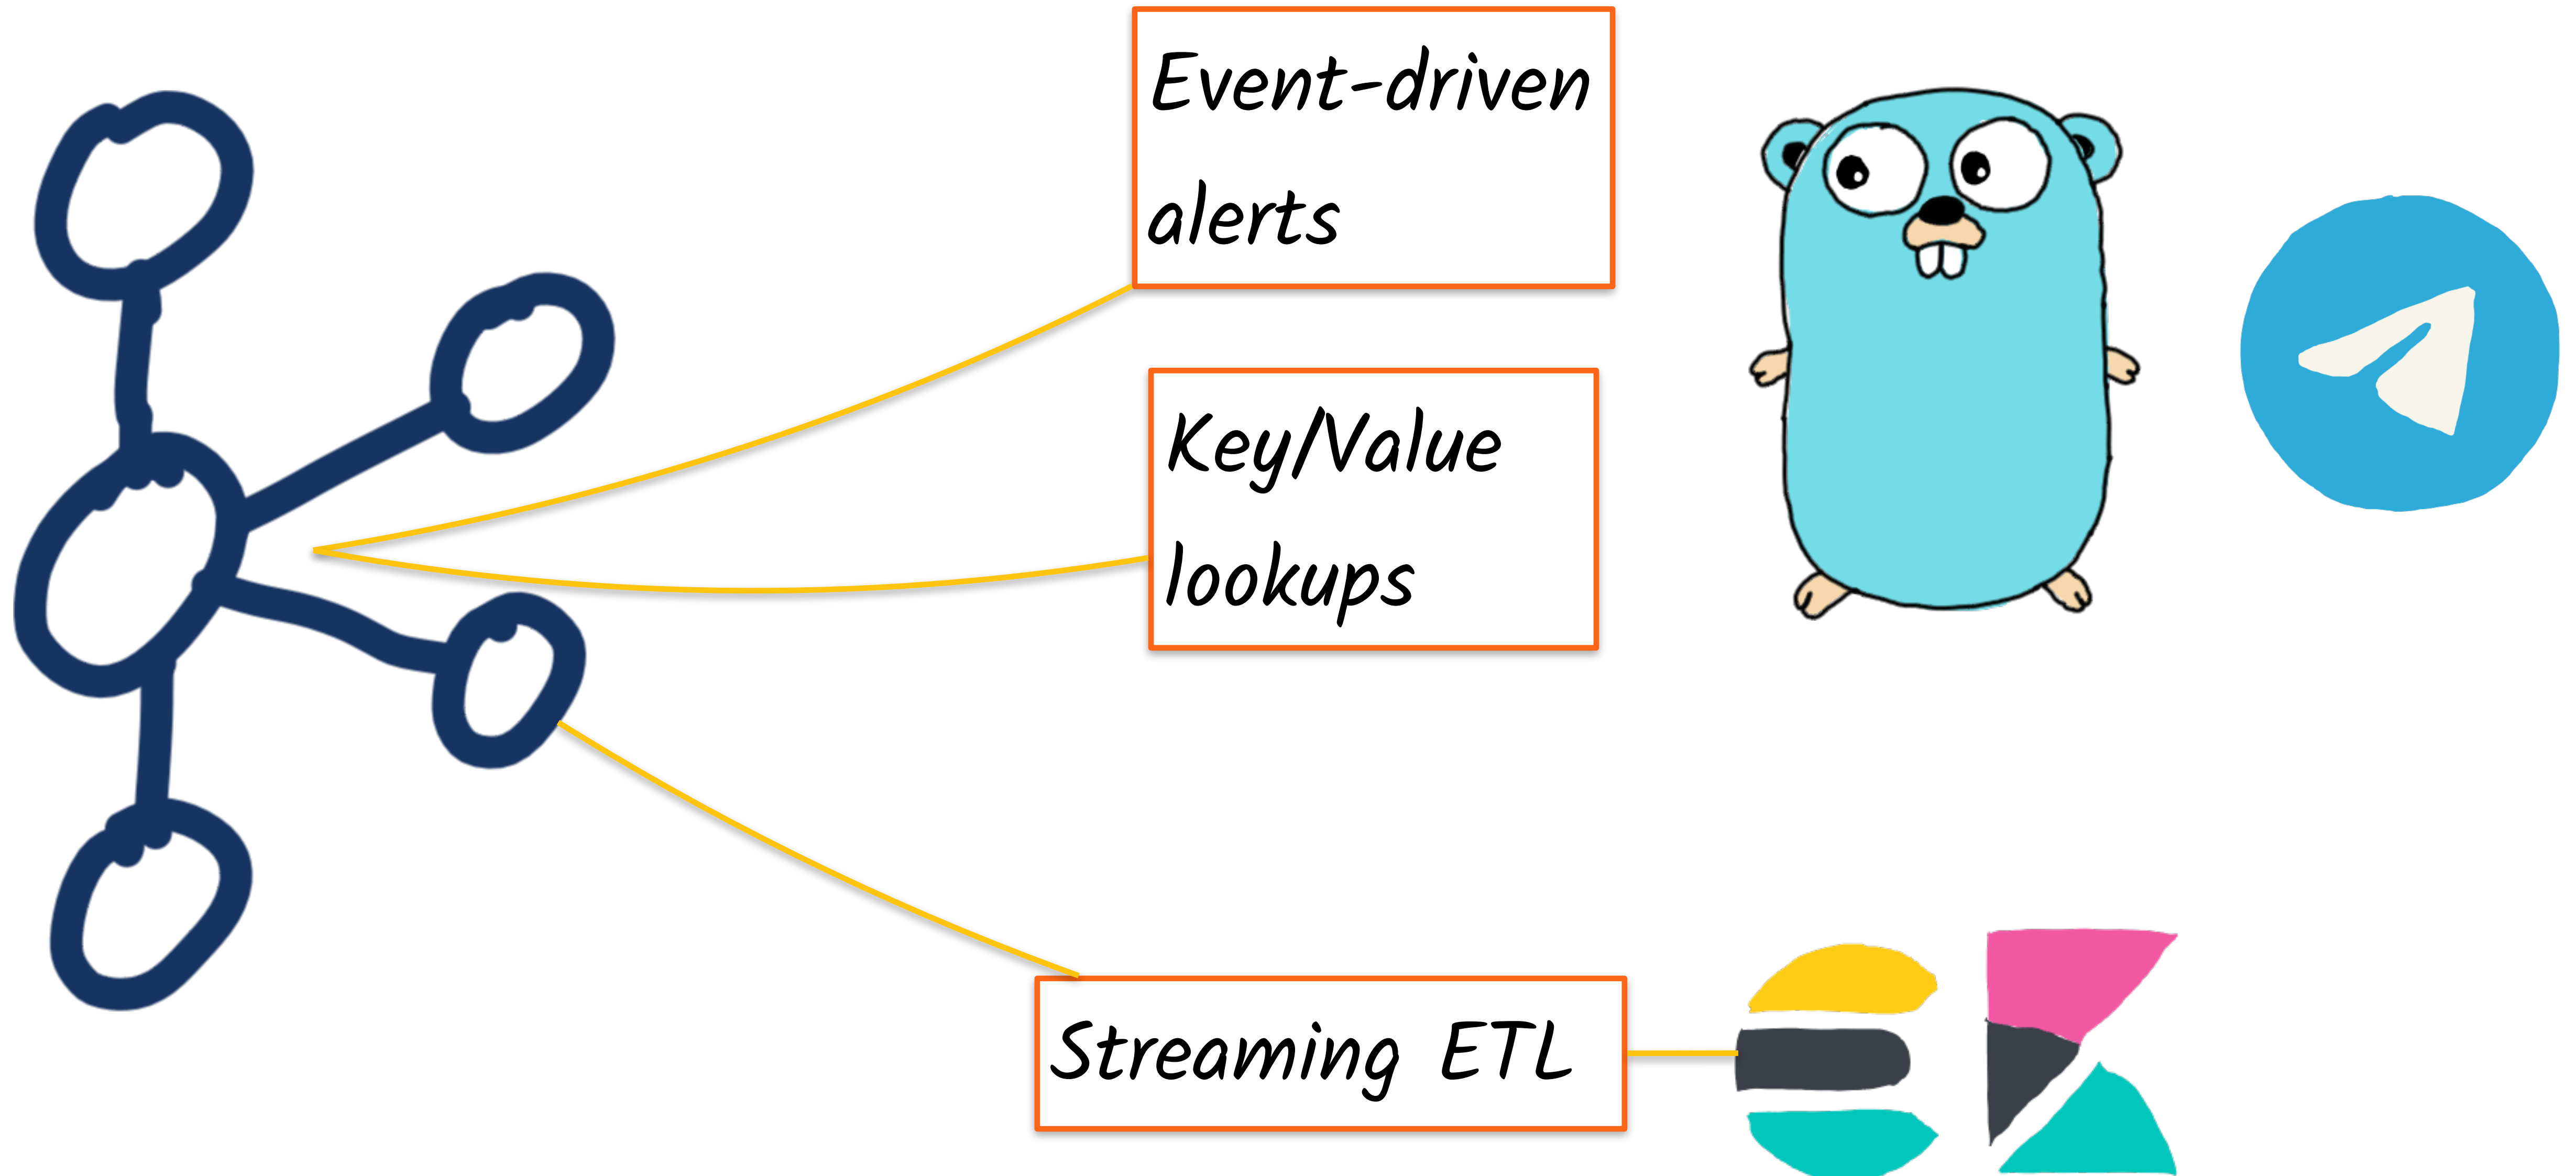
Stream

Store

Process

Integrate

Flexible, event-driven applications



Want to
learn
more?



Try it out for yourself

<https://rmoff.dev/carparks>



Deep-dive articles


- What is Apache Kafka?
- Event Streaming vs. Related Trends
- KRaft: Kafka without ZooKeeper
- Transactions & Guarantees in Kafka
- Processing & Storage Fundamentals

- Kafka Performance
- Cloud-native Kafka
- Streaming Database Systems
- Testing Apache Kafka
- Explore Kafka's Internals

Over 10 hours of free courses

- Apache Kafka 101
- Kafka Connect 101
- Kafka Streams 101
- ksqlDB 101
- Inside ksqlDB

- Spring Framework and Kafka
- Building Data Pipelines with Kafka
- Event Sourcing with Kafka
- Data Mesh 101



CONFLUENT

Developer

Plus: Hands-on Quick Starts and Client Language Guides + Event Streaming Patterns + More

developer.confluent.io

RM0FF200

**\$200 USD off your bill each calendar month
for the first three months when you sign up**

<https://rmoff.dev/cc1oud>

Free money!
(additional \$200 towards
your bill 😊)



confluent cloud

Fully Managed Kafka as a Service



#

EOF

@rmoff

rmoff.dev/talks

youtube.com/rmoff




Further reading / watching

<https://rmoff.dev/kafka-talks>

Real-life examples

Here's a nice example using real data to solve a real problem - is my train late now? What are the routes most likely to be delayed?

 [On Track with Apache Kafka: Building a Streaming Platform solution with Rail Data](#)

Moving from  to  let's take another real data feed and build some realtime location-based notifications 

 [Building a Telegram bot with Go, Apache Kafka, and ksqlDB](#)

Integration and data pipelines

Integration between Kafka and other systems? Kafka Connect has you covered ⚡

 [Kafka Connect in 60 seconds](#)

 [From Zero to Hero with Kafka Connect](#)