

***I DON'T
ALWAYS TEST
MY STREAMS,
BUT WHEN I
DO I DO IT IN
PRODUCTION***

**[j] conf
.dev**

online 2020



@gamussa

| @jconf.dev

| @confluentinc



@gamussa | @jconf.dev | @confluentinc

SPECIAL THANKS



@inponomarev

@gamussa

|

@jconf.dev

|



@bsideup

@confluentinc

<https://gamov.dev/jdevconf2020>



Agenda



01

Kafka Apps

What are those?

02

Testing basics

Run Kafka and CP locally

03

Can we better?

Test with TestContainers

04

Testing Streams

TopologyTestDriver vs TestContainers

05

Testing ksqlDB

Ksql-test-runner vs TestContainers

06

Conclusion

Always Be Testing

PREFACE



Senior Developer:

- Where did you get this code?

Junior Developer:

- From Stackoverflow

Senior Developer:

- from question or answer?



An example from my life:

Kafka Tutorials website

Literate Programming...

Do you speak it...

https://en.wikipedia.org/wiki/Literate_programming



@gamussa

|

@jconf.dev

|

@confluentinc



Say Hello World to Event Streaming.

Apache Kafka is a powerful, scalable, fault-tolerant distributed streaming platform. This site features full code examples using Kafka, Kafka Streams, and ksqlDB to demonstrate real use cases.



Kafka 101

Console Producer and Consumer Basics

Console Consumer with primitive keys and values

Using the Console Consumer to read from a specific offset and partition



Apply a function to data.

Transform a stream of events

Flatten deeply nested events

Filter a stream of events

Rekey a stream with a value

Rekey a stream with a function

More



Aggregate data.

Count a stream of events

Sum a stream of events

Find the min/max in a stream of events

Calculate a running average

Cogrouping multiple streams of aggregates



Join data.

Join a stream and a table together

Join a stream and a stream together

Join a table and a table together

Join two tables with a foreign key



Collect data over time.

Create tumbling windows

Create hopping windows

Create session windows

Emit a final result from a time window



Connect data sources & sinks.

Add a key to a stream from a JDBC source



@gamussa



@jconf.dev



@confluentinc

↑ Merge pull request #410 from confluentinc/master c2fbbd8 by bbejeck

Workflow Artifacts

Rerun

Edit Workflow

18:17 · Kafka Tutorials pipeline · Stopped · Jun 3 4:35 PM

00:41 · Kafka Tutorials live site deployment · Passed · Jun 3 5:29 PM

18:17 · Kafka Tutorials pipeline · Stopped

Run the tests		
01:44	· KSQL transforming tests	✓
01:47	· KStreams transforming tests	✓
04:08	· Kafka transforming tests	✓
01:50	· KSQL filtering tests	✓
01:41	· KStreams filtering tests	✓
01:48	· KSQL splitting tests	✓
17:29	· KStreams splitting tests	⊖
02:08	· KSQL merging tests	✓
01:55	· KStreams merging tests	✓
02:15	· KSQL join stream to stream tests	✓
02:09	· KSQL join stream to table tests	✓
02:19	· KStreams join stream to table tests	✓
02:23	· KStreams table to table foreign key join tests	✓
01:42	· KSQL join table to table tests	✓
02:09	· KStreams tumbling windows tests	✓
01:47	· KSQL tumbling windows tests	✓
01:51	· KSQL session windows tests	✓
01:46	· KSQL aggregation count tests	✓
03:38	· KSQL aggregation MIN/MAX tests	✓
02:10	· KSQL aggregation sum tests	✓
01:40	· KStreams aggregation sum tests	✓
01:40	· KSQL serialization tests	✓
01:59	· KStreams Serialization tests	✓
01:54	· KSQL rekey stream tests	✓
01:57	· KSQL rekey stream with function tests	✓

Build the website

00:35 · Compile with Jekyll ✓

Deploy to live site

Promote →

Jun 3 5:29 PM ✓

Deploy to staging site

Promote →



↑ Merge pull request #410 from confluentinc/master c2fbbd8 by bbejck

Workflow Artifacts

Rerun

Edit Workflow



18:17 · Kafka Tutorials pipeline · Stopped · Jun 3 4:35 PM

00:41 · Kafka Tutorials live site deployment · Passed · Jun 3 5:29 PM

18:17 · Kafka Tutorials pipeline · Stopped

Run the tests		
01:44	· KSQL transforming tests	✓
01:47	· KStreams transforming tests	✓
04:08	· Kafka transforming tests	✓
01:50	· KSQL filtering tests	✓
01:41	· KStreams filtering tests	✓
01:48	· KSQL splitting tests	✓
17:29	· KStreams splitting tests	⊖
02:08	· KSQL merging tests	✓
01:55	· KStreams merging tests	✓
02:15	· KSQL join stream to stream tests	✓
02:09	· KSQL join stream to table tests	✓
02:19	· KStreams join stream to table tests	✓
02:23	· KStreams table to table foreign key join tests	✓
01:42	· KSQL join table to table tests	✓
02:09	· KStreams tumbling windows tests	✓
01:47	· KSQL tumbling windows tests	✓
01:51	· KSQL session windows tests	✓
01:46	· KSQL aggregation count tests	✓
03:38	· KSQL aggregation MIN/MAX tests	✓
02:10	· KSQL aggregation sum tests	✓
01:40	· KStreams aggregation sum tests	✓
01:40	· KSQL serialization tests	✓
01:59	· KStreams Serialization tests	✓
01:54	· KSQL rekey stream tests	✓
01:57	· KSQL rekey stream with function tests	✓

Deploy to live site

Promote →

Jun 3 5:29 PM ✓

Deploy to staging site

Promote →

Build the website

00:35 · Compile with Jekyll ✓

@gamussa

@jconf.dev

@confluentinc

↑ Merge pull request #410 from confluentinc/master c2fbbd8 by bbejeck

Workflow Artifacts

Rerun

Edit Workflow

18:17 · Kafka Tutorials pipeline · Stopped : Jun 3 4:35 PM

00:41 · Kafka Tutorials live site deployment · Passed : Jun 3 5:29 PM

18:17 · Kafka Tutorials pipeline · Stopped

Run the tests	
01:44 · KSQL transforming tests	✓
01:47 · KStreams transforming tests	✓
04:08 · Kafka transforming tests	✓
01:50 · KSQL filtering tests	✓
01:41 · KStreams filtering tests	✓
01:48 · KSQL splitting tests	✓
17:29 · KStreams splitting tests	⊖
02:08 · KSQL merging tests	✓
01:55 · KStreams merging tests	✓
02:15 · KSQL join stream to stream tests	✓
02:09 · KSQL join stream to table tests	✓
02:19 · KStreams join stream to table tests	✓
02:23 · KStreams table to table foreign key join tests	✓
01:42 · KSQL join table to table tests	✓
02:09 · KStreams tumbling windows tests	✓
01:47 · KSQL tumbling windows tests	✓
01:51 · KSQL session windows tests	✓
01:46 · KSQL aggregation count tests	✓
03:38 · KSQL aggregation MIN/MAX tests	✓
02:10 · KSQL aggregation sum tests	✓
01:40 · KStreams aggregation sum tests	✓
01:40 · KSQL serialization tests	✓
01:59 · KStreams Serialization tests	✓
01:54 · KSQL rekey stream tests	✓
01:57 · KSQL rekey stream with function tests	✓

Deploy to live site

Promote →

Jun 3 5:29 PM ✓

Deploy to staging site

Promote →

Build the website

00:35 · Compile with Jekyll ✓



↑ Merge pull request #410 from confluentinc/master c2fbbd8 by bbejeck

Workflow Artifacts

Rerun

Edit Workflow

18:17 · Kafka Tutorials pipeline · Stopped : Jun 3 4:35 PM

00:41 · Kafka Tutorials live site deployment · Passed : Jun 3 5:29 PM

18:17 · Kafka Tutorials pipeline · Stopped

Run the tests

- 01:44 · KSQL transforming tests ✓
- 01:47 · KStreams transforming tests ✓
- 04:08 · Kafka transforming tests ✓
- 01:50 · KSQL filtering tests ✓
- 01:41 · KStreams filtering tests ✓
- 01:48 · KSQL splitting tests ✓
- 17:29 · KStreams splitting tests ⚪
- 02:08 · KSQL merging tests ✓
- 01:55 · KStreams merging tests ✓
- 02:15 · KSQL join stream to stream tests ✓
- 02:09 · KSQL join stream to table tests ✓
- 02:19 · KStreams join stream to table tests ✓
- 02:23 · KStreams table to table foreign key join tests ✓
- 01:42 · KSQL join table to table tests ✓
- 02:09 · KStreams tumbling windows tests ✓
- 01:47 · KSQL tumbling windows tests ✓
- 01:51 · KSQL session windows tests ✓
- 01:46 · KSQL aggregation count tests ✓
- 03:38 · KSQL aggregation MIN/MAX tests ✓
- 02:10 · KSQL aggregation sum tests ✓
- 01:40 · KStreams aggregation sum tests ✓
- 01:40 · KSQL serialization tests ✓
- 01:59 · KStreams Serialization tests ✓
- 01:54 · KSQL rekey stream tests ✓
- 01:57 · KSQL rekey stream with function tests ✓

Build the website
00:35 · Compile with Jekyll ✓

Deploy to live site

Promote →

Jun 3 5:29 PM ✓

Deploy to staging site

Promote →



KAFKA APPS AKA STREAM PROCESSING

STREAM PROCESSING



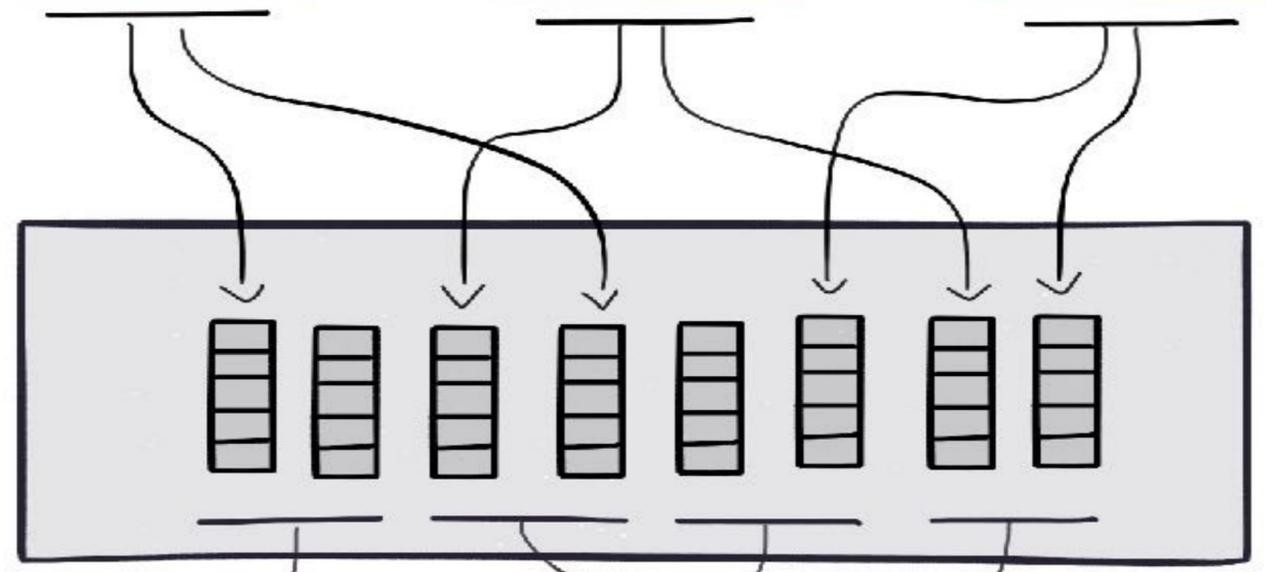
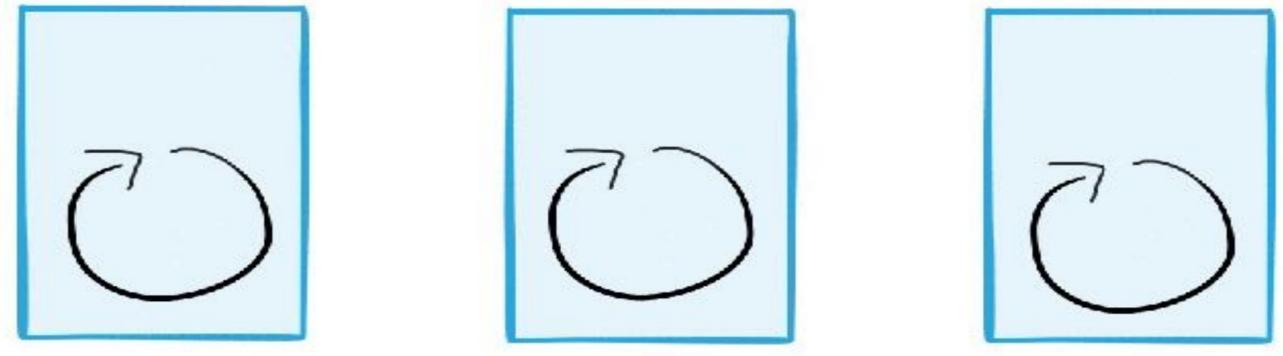
is the **toolset** for
dealing with
events

as they **move!**



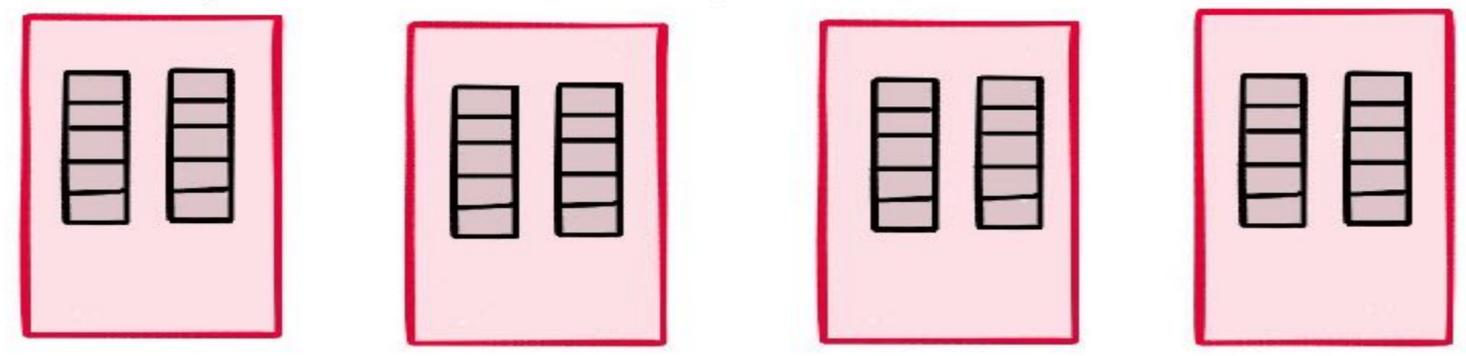


Producers



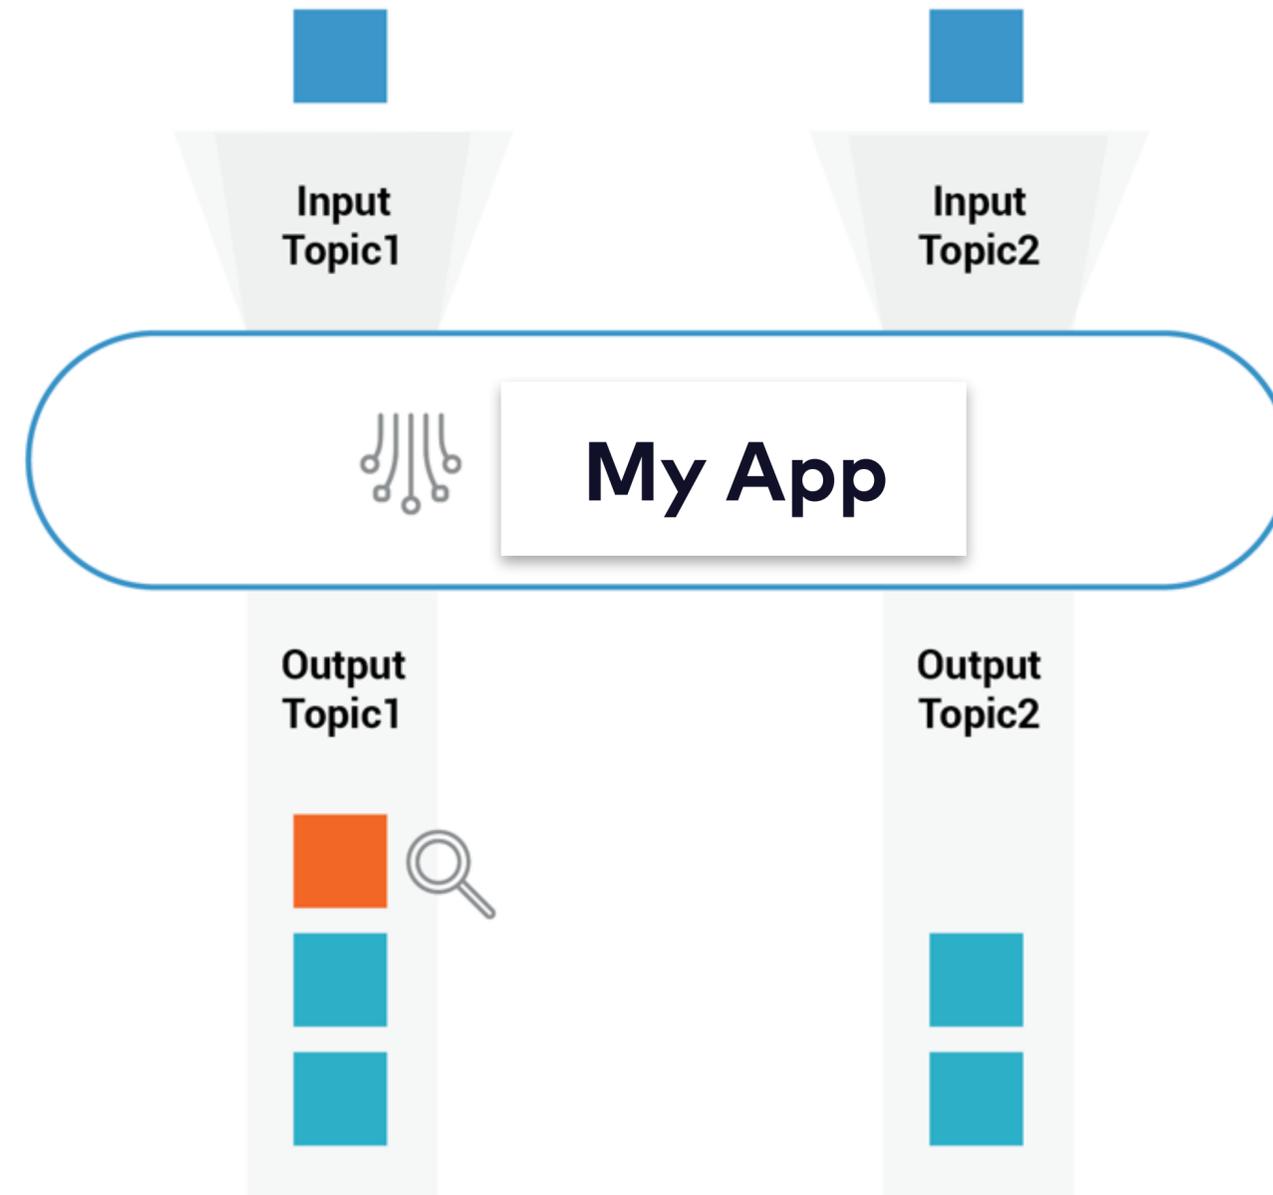
Kafka

Consumers

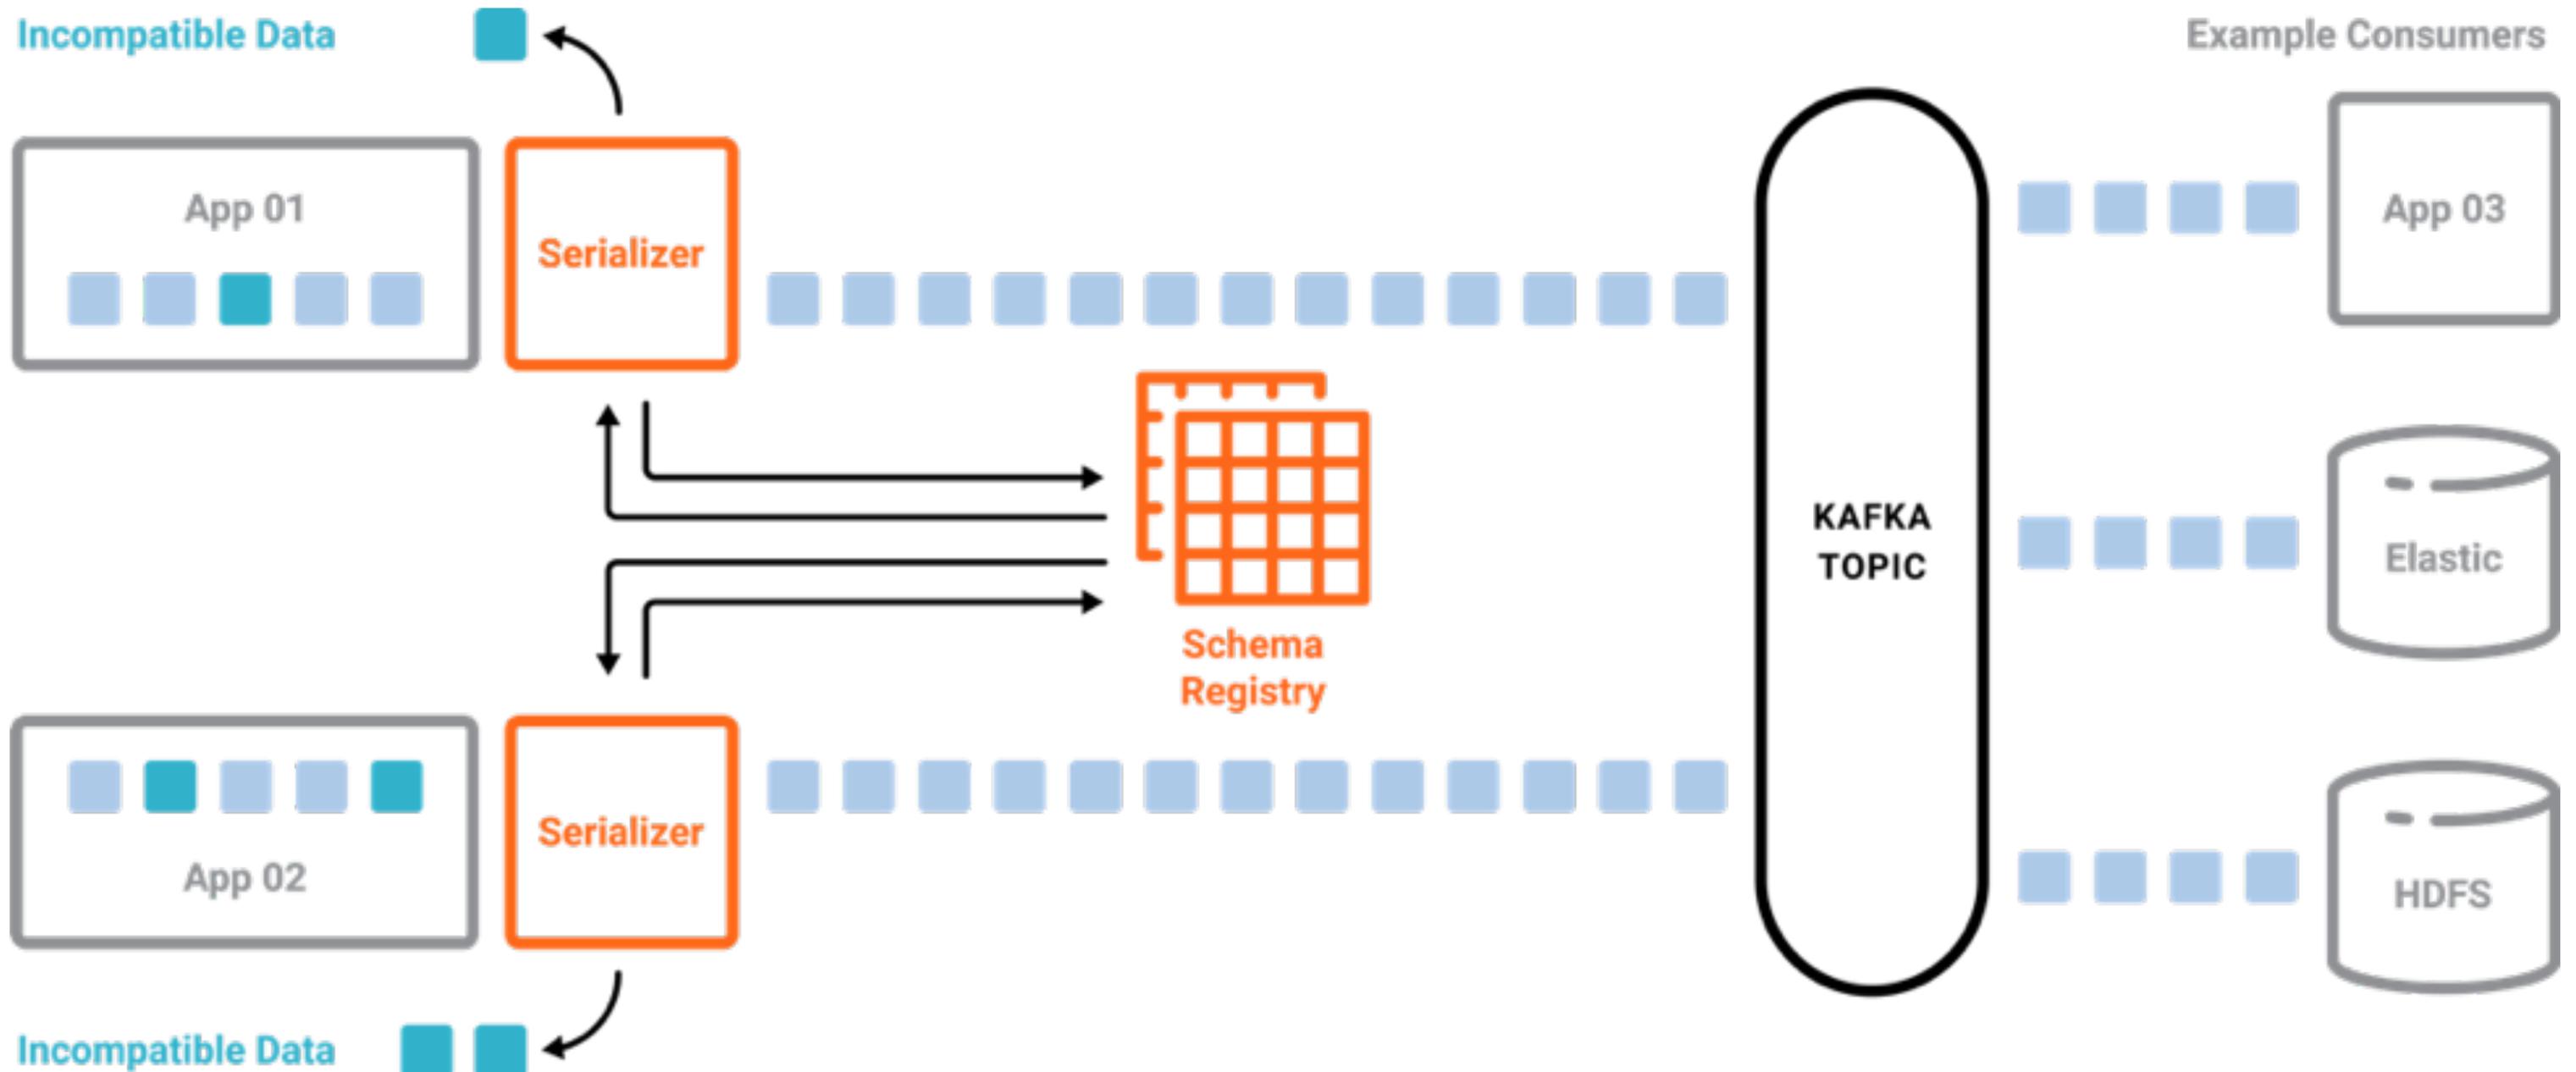


@gamussa | @jconf.dev | @confluentinc

WHAT ARE KAFKA APPS?



SCHEMA, WUT?



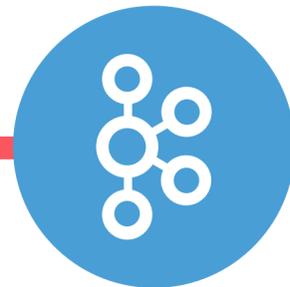
LET'S SEE SOME CODE



KAFKA STREAMS FOR STREAM PROCESSING



the **KAFKA STREAMS API** is a
JAVA API to
BUILD REAL-TIME APPLICATIONS



HOME

INTRODUCTION

QUICKSTART

USE CASES

DOCUMENTATION

Getting Started

APIs

Configuration

Design

Implementation

[DOCUMENTATION](#) » [KAFKA STREAMS](#) » [DEVELOPER GUIDE](#) »

Testing Kafka Streams

Table of Contents

- [Importing the test utilities](#)
- [Testing Streams applications](#)
- [Unit testing Processors](#)

<https://kafka.apache.org/documentation/streams/developer-guide/testing.html>

TOPOLOGY TEST DRIVER

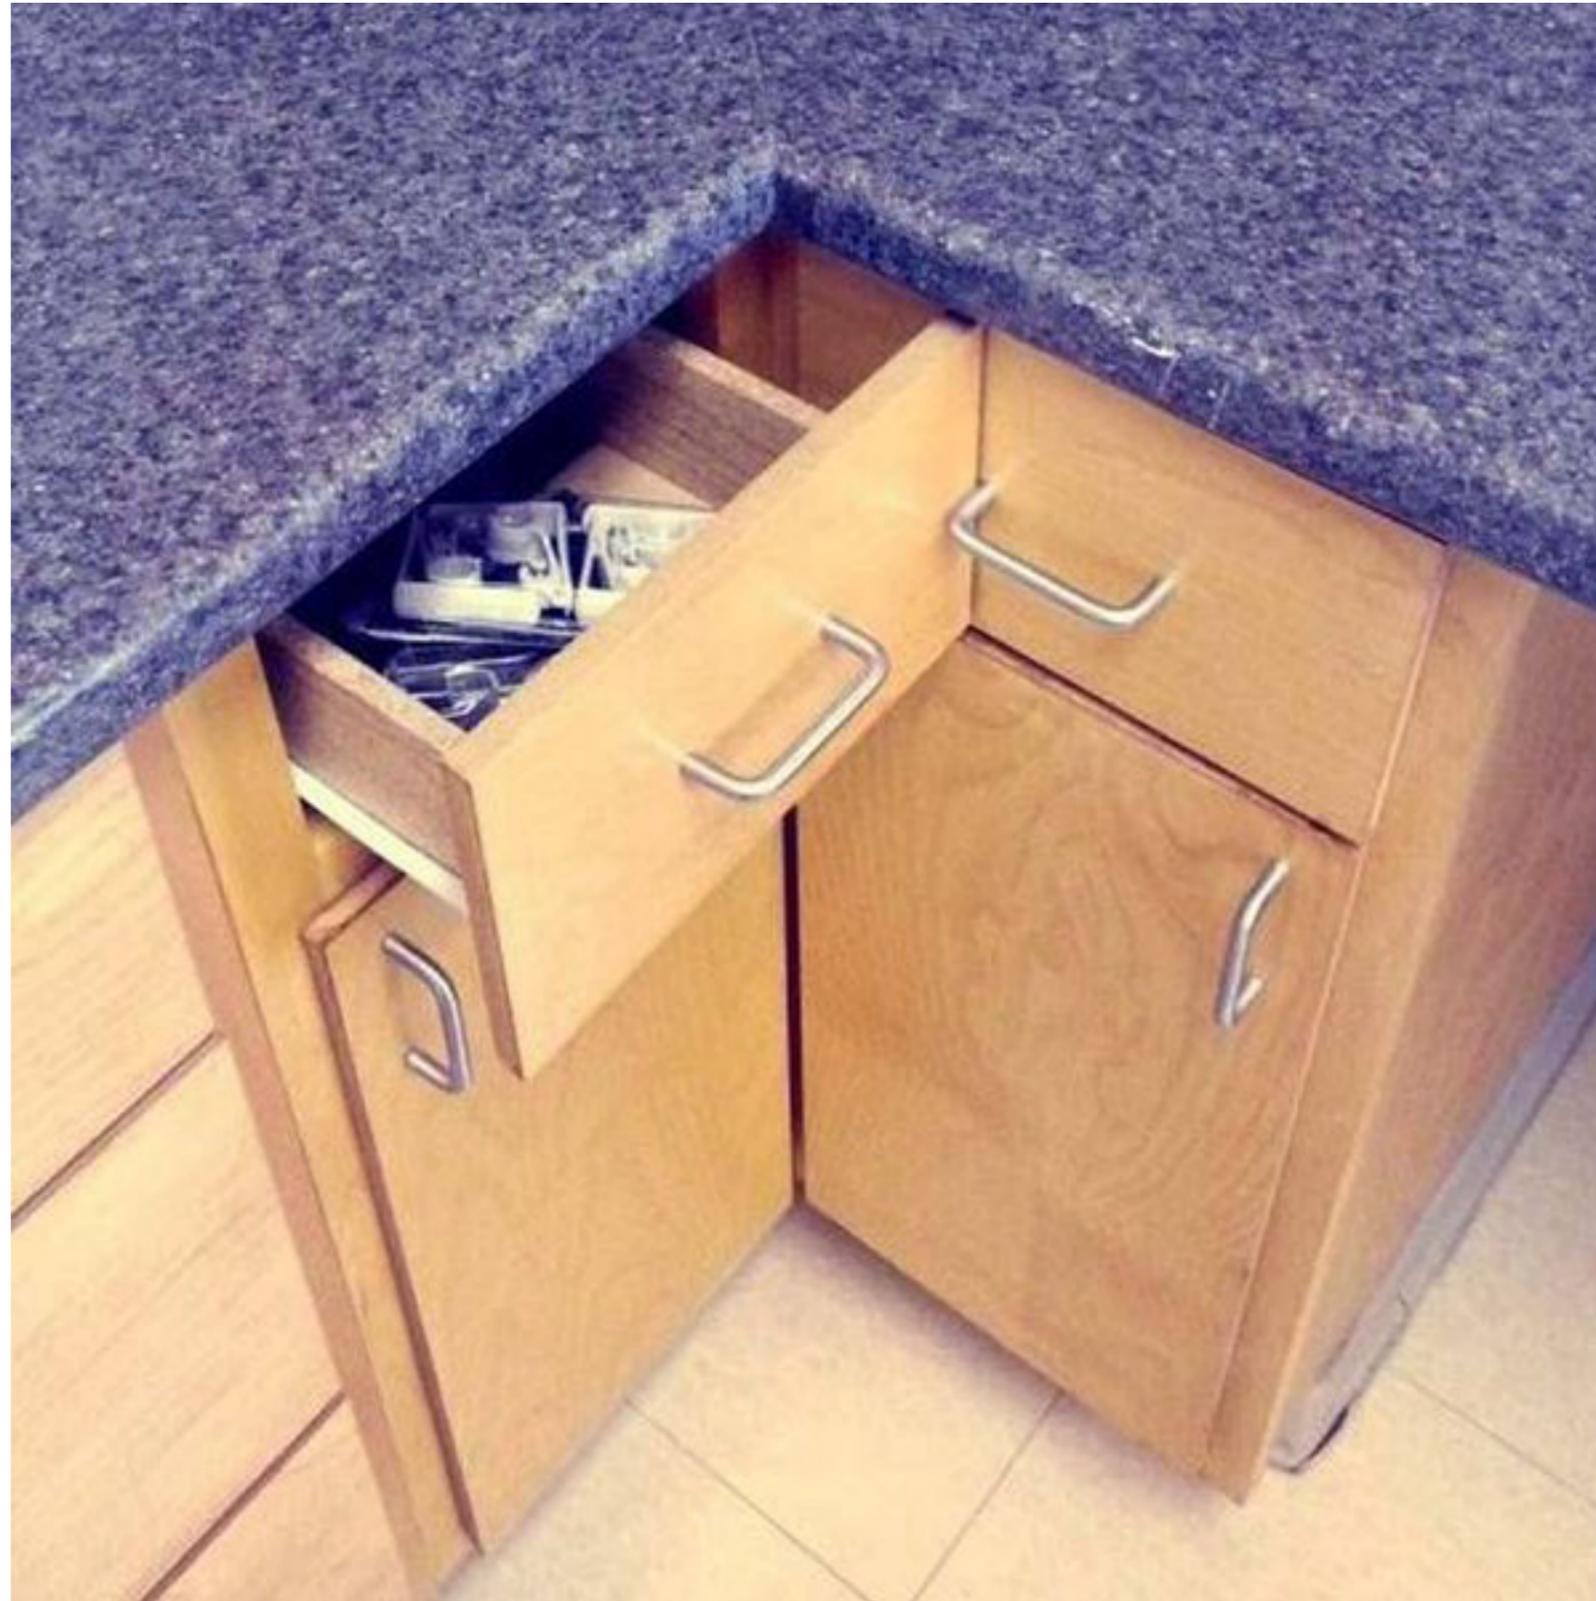


- **TopologyTestDriver**

- **Unit test** your Kafka Streams Topology
- **Doesn't Require Kafka** infra to run
- Convenience methods to **define input and output**
- Convenience methods to provide fixtures and **validate the results**

LET'S SEE SOME CODE





@gamussa | @jconf.dev | @confluentinc



@gamussa | @jconf.dev | @confluentinc

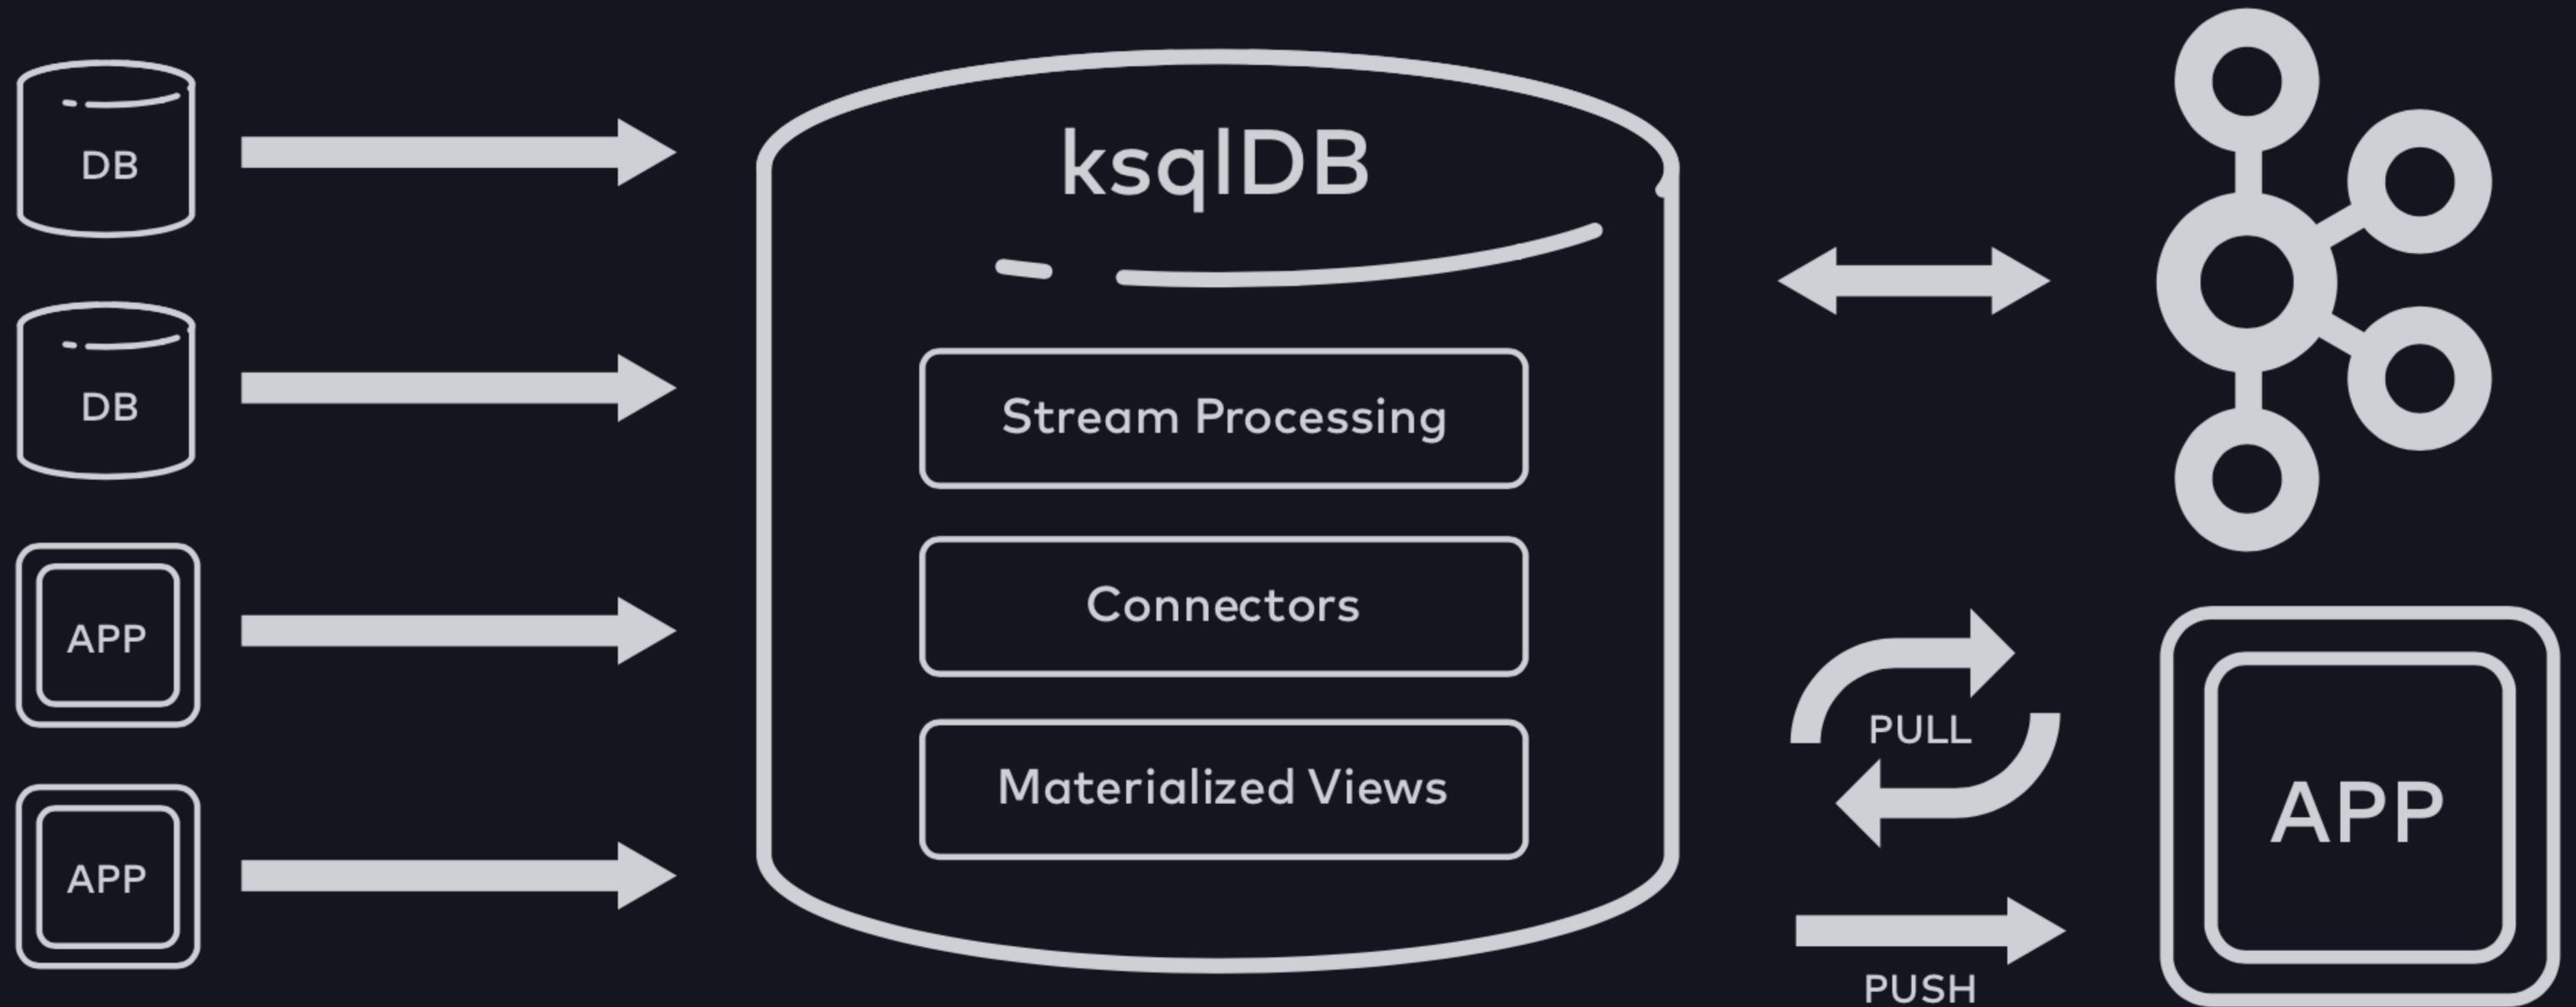
WHEN TTD DOESN'T HELP?



<https://gamov.dev/livestreams-s01e05>

@gamussa | @jconf.dev | @confluentinc

KSQldb FOR STREAM PROCESSING





```
1 NAME
2     ksql-test-runner - The KSQL testing tool
3
4 SYNOPSIS
5     ksql-test-runner {--input-file | -i} <inputFile>
6                     {--output-file | -o} <outputFile>
7                     {--sql-file | -s} <statementsFile>
8
9 OPTIONS
10    --input-file <inputFile>, -i <inputFile>
11        A JSON file containing the input records.
12
13        This option may occur a maximum of 1 time
14
15
16    --output-file <outputFile>, -o <outputFile>
17        A JSON file containing the expected output records.
18
19        This option may occur a maximum of 1 time
20
21
22    --sql-file <statementsFile>, -s <statementsFile>
23        A SQL file containing KSQL statements to be tested.
24
25        This option may occur a maximum of 1 time
```

LET'S SEE SOME CODE



CONCLUSION



What did you learn today?



- What's Kafka apps apps and how to test them
- Producer / Consumer API
- Kafka Streams
- ksqlDB
- How containers and TestContainer help a lot

A. always

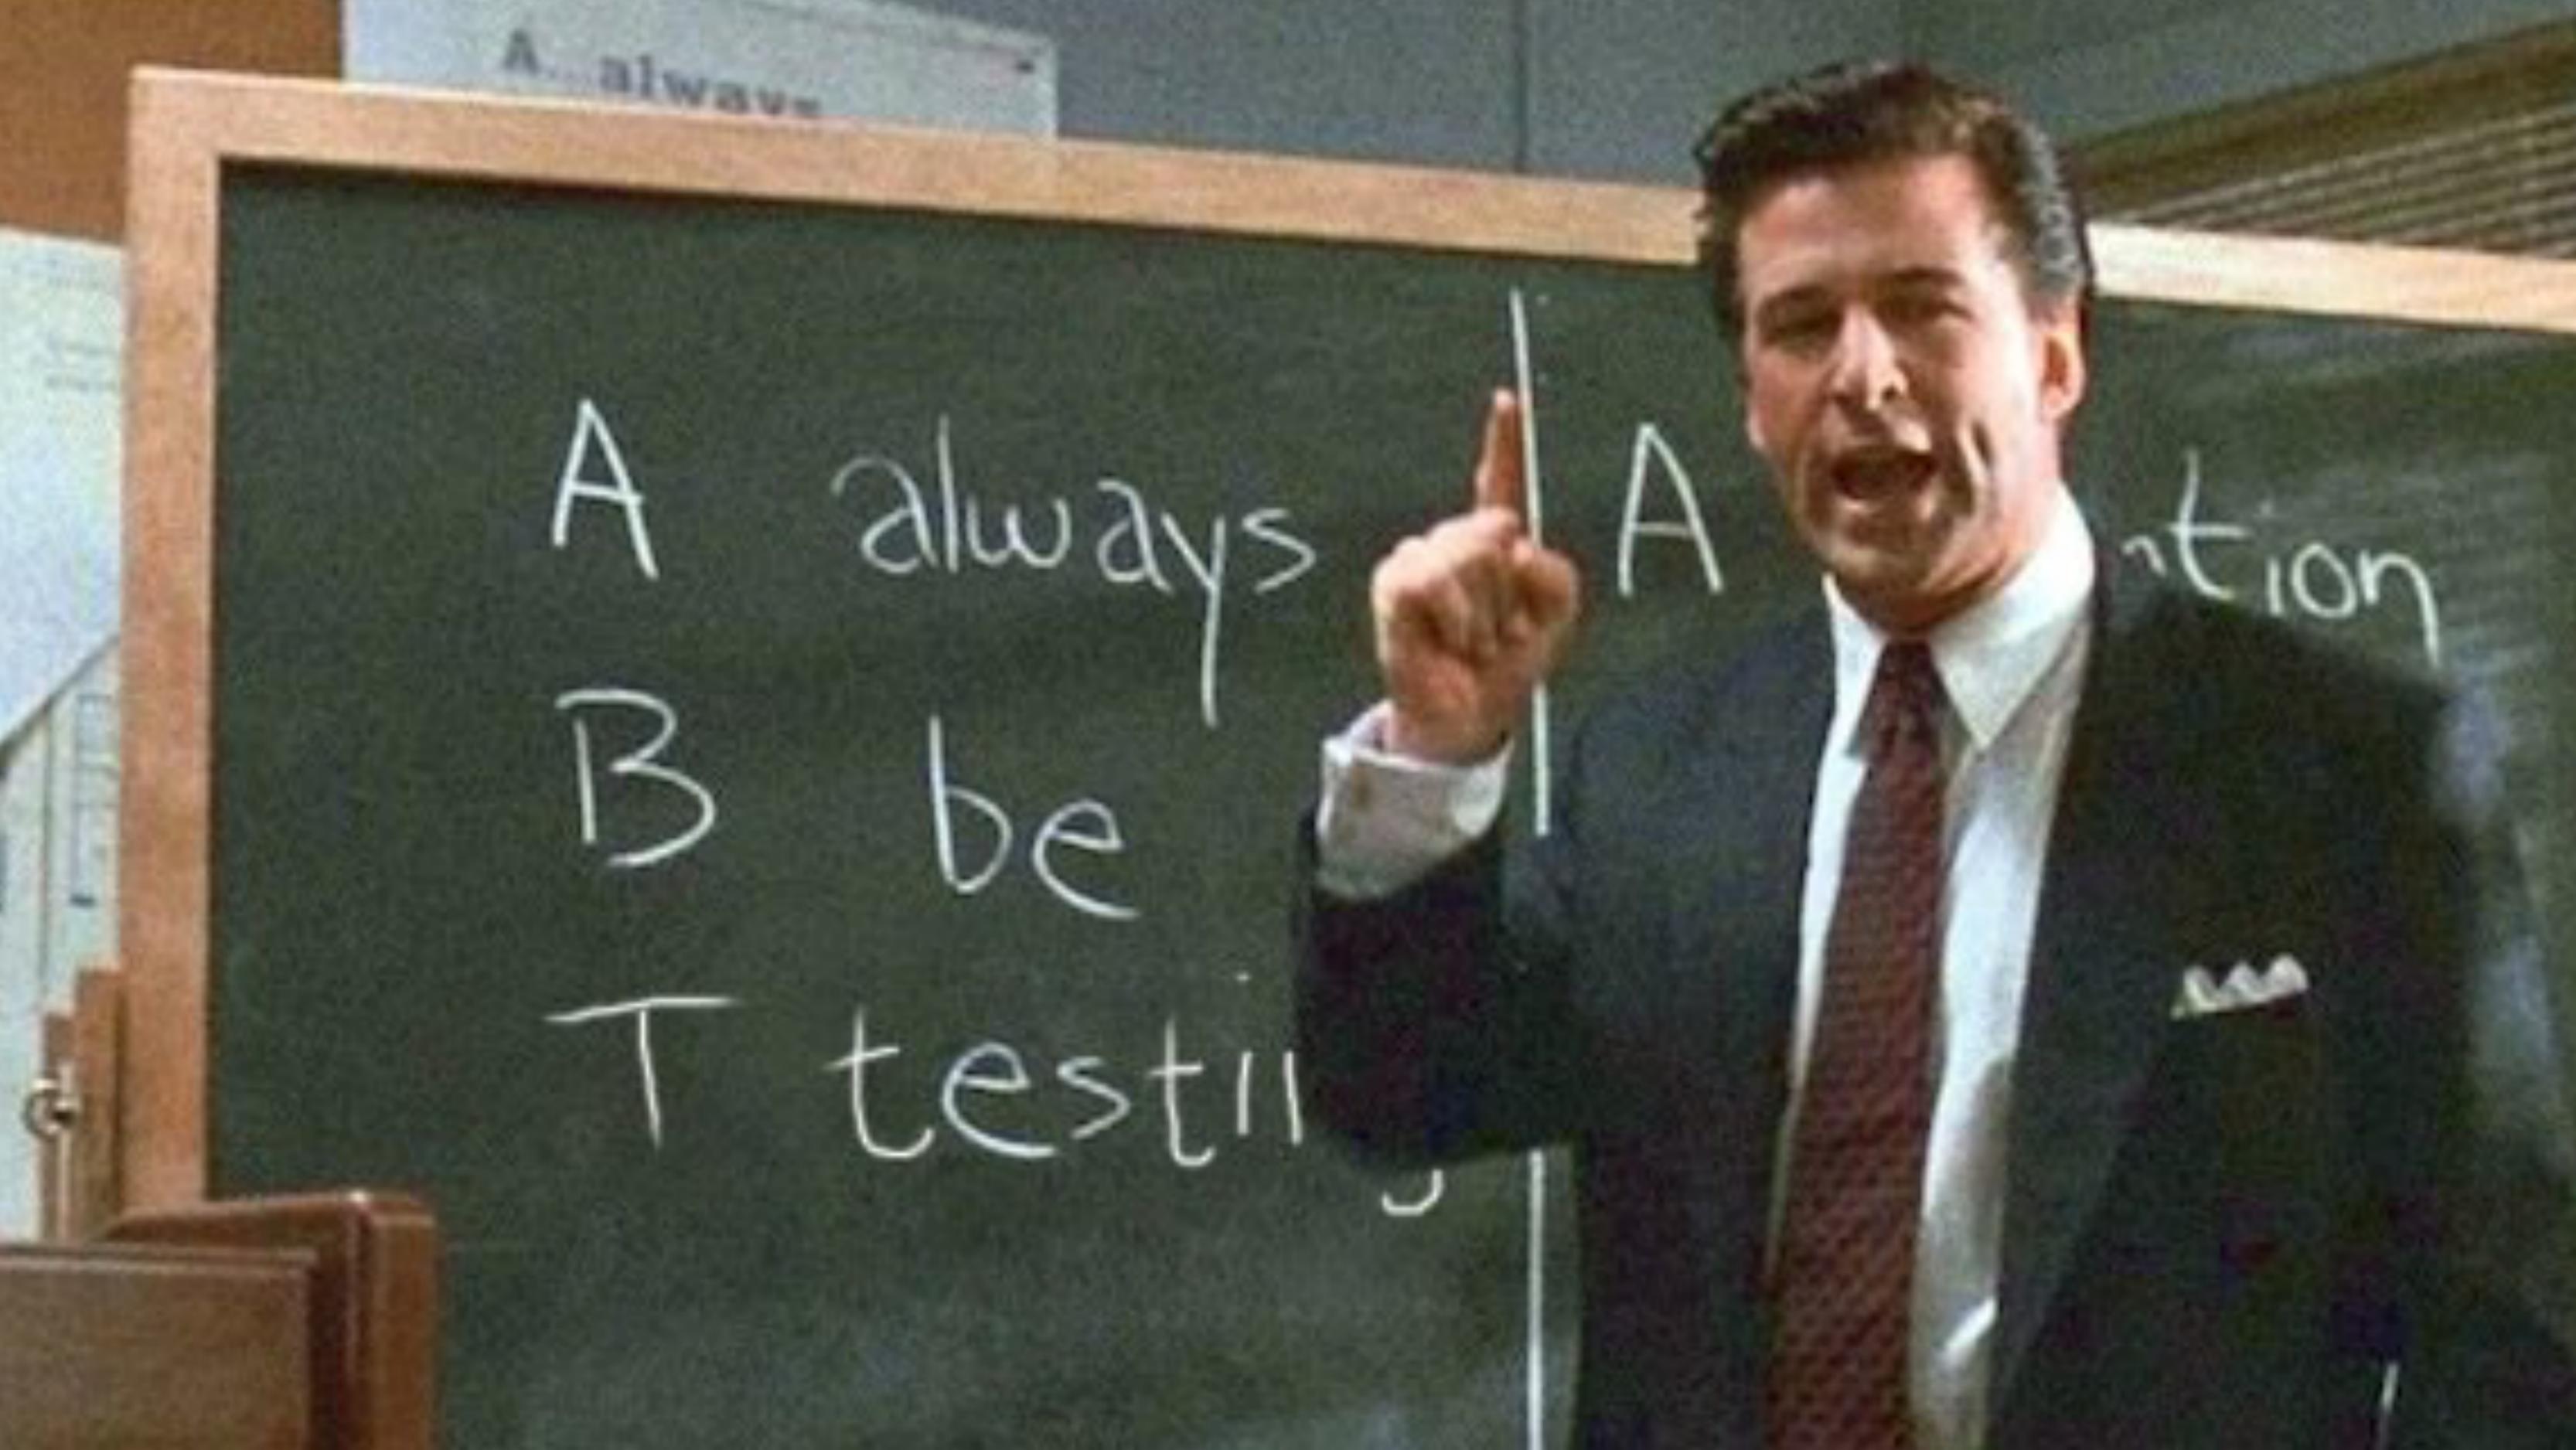
A always

B be

T testi

A

ntion





WANT TO LEARN MORE?

Learn Kafka.

Start building with
Apache Kafka at
Confluent Developer.

Watch full version
<https://gamov.dev/developer>



developer.confluent.io

<https://developer.confluent.io>



Confluent Community Slack



A vibrant community of over 16,000 members

Come along and discuss **Apache Kafka** and **Confluent Platform** on dedicated channels including `#ksqlDB`, `#connect`, `#clients`, and more

<http://cnfl.io/slack>





CONFLUENT

@gamussa

| @jconf.dev

| @confluentinc