




Client Cannot Communicate



Lessons From a Decade of Kafka Consulting





Hello!


/thoughtworks

Simple
Machines



Simon Aubury

Associate Director –
Data Platforms



Confluent
Community Catalyst



2014



"Happy" Pharrell Williams

Apache Kafka 0.8

Staying happy



It has a steep learning curve ...

Learning - onboard quickly



In mid-size / larger-sized companies don't be the blocker for upskilling others

- An **easy onboarding** experience
- A **sandbox environment** for folks to try without constraints

Learning - keep current




To stay happy **stay Kafka informed ...**

- Big surface area
- Keeps growing - need to keep learning
- Great resources



<https://github.com/pmoskovi/kafka-learning-resources>

2015



"Thinking Out Loud" Ed Sheeran

Apache Kafka 0.9

Thinking out loud



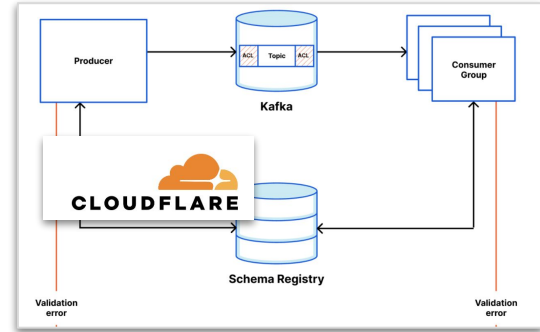
Seeing the same mistakes
across different teams

What do you do
when the traffic,
meta-data, and amount of
developers and use cases
grow?

WixEngineering

#1 Common Infra
with common
features.

Wix

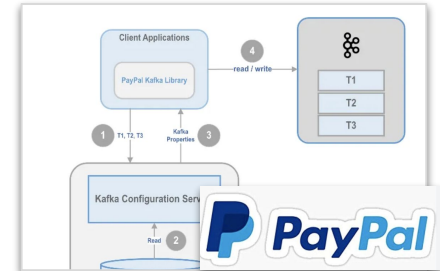


Cloudflare

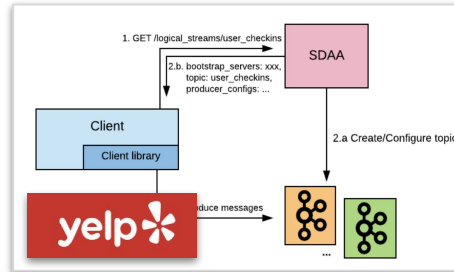
```
test-event-proto:
  owner:
    team: data-platform-team
    slack: #ask-data
  description: a test topic for proto messages
  categories:
    - events, pubsub, test, internal }
    - compact | both }
    - failover | roundrobin | replicate
```

coinbase

Coinbase



Paypal



Yelp

Learning - sensible defaults



Codifying **best practices and principles into pre-defined approaches** for common tasks

- Wrap the producer and consumer SDK
- Hide the Kafka boilerplate
- Broker location, sterilisation, key strategy, other properties

2016

JUSTIN BIEBER



"Sorry" - Justin Bieber

Apache Kafka 0.10

Sorry



Mistakes were made



Kafka is flexible enough to solve just about any problem!



Kafka is flexible enough to solve just about **any** problem?



**YOUR SCIENTISTS WERE SO PREOCCUPIED
WITH WHETHER OR NOT THEY COULD...**

Not everything
should be an
Apache Kafka
project

THEY DIDN'T STOP TO THINK IF THEY SHOULD.

Learning - when not to use Kafka



 You **don't** want to use Kafka for

- Unreliable networks
- Connectivity to 1000's of client applications
- Replacing a database or huge payloads
- CV driven development

2017



"Shape of You" Ed Sheeran

Apache Kafka 0.11

Schemas are cool



How do we managing these shapes?



Decomposition

Moving from monolith to decoupled
microservices with clear responsibilities








“Schemas are the APIs used by event-driven service ...so **a publisher and subscriber need to agree** on exactly how a message is formatted.”

(O'Reilly)

Learning to love schemas



-  Schema registry
-  Schema management - configuration repo
-  Schema ownership (and accountability)
-  Schema flexibility (and evolvability)
-  Choose the ~~right~~ appropriate compatibility mode

2018



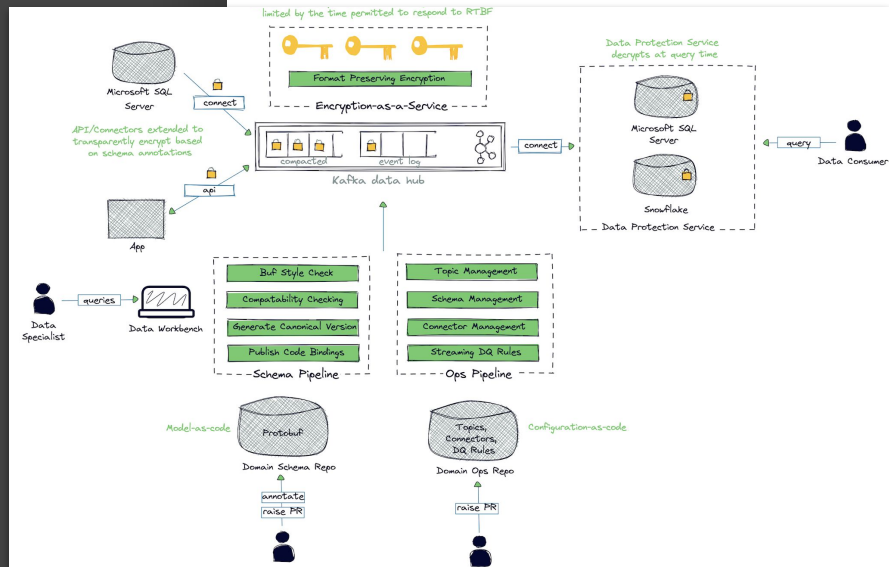
"God's Plan" Drake

Apache Kafka 1.0

Centralised Kafka



Working with the ~~god~~
"central platform team"



Introducing Apache Kafka as an enterprise backbone is not without its challenges

Saxo Bank's Best Practices for a Distributed Architecture

Learning - centralised team



Balance developers **and** platform concerns

- Establish appropriate abstractions
- Non negotiables - security & access
- Enforcing best practices by default
- Serialization, monitoring, encryption -
 - if it's important - bake it in from day one

2021

Save your tears



"Save Your Tears"

The Weeknd and Ariana Grande


Apache Kafka 2.8

Running Kafka yourself can get pretty time consuming depending how big you scale.

Learning – Building it yourself?



Getting the wrong idea from that conference talk you attended



Solving Imaginary Scaling Issues

At Scale

ONLY?

@ThePracticalDev

😱 Don't run Apache Kafka
yourself *

** unless you're very small, or very big!*

<https://onlybooks.com/books/imaginary-scaling-issues>

Take aways - greatest hits

- ✅ Codify best practices
- ❤️ Love schemas and manage like code
- ⚖️ Balance developers and platform concerns
- 📖 Keep learning





And now?

Today

Beautiful Things



The beauty of Apache Kafka lies in its ability to unify data processing, integration, and real-time analytics within a single platform.



"Beautiful Things"

by Benson Boon

Apache Kafka 3.8

Thanks!



@SimonAubury