





# OLX Data Hub

Jakub Orłowski, Krzysztof Antończak, Facundo Guerrero

Presto Summit 2019, New York City

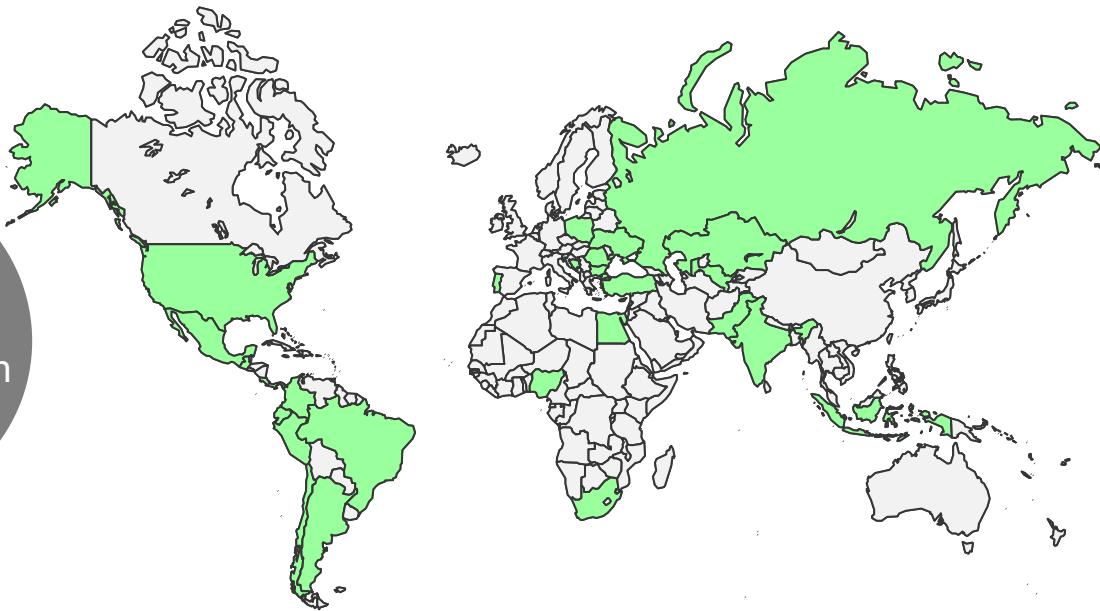


Meet OLX, the biggest  
Web company you've  
never heard of

# Within classified ads, OLX Group is the largest global player

>300m MAUs

Present in  
30 markets,  
Leading position  
in 27



... with a strong local presence



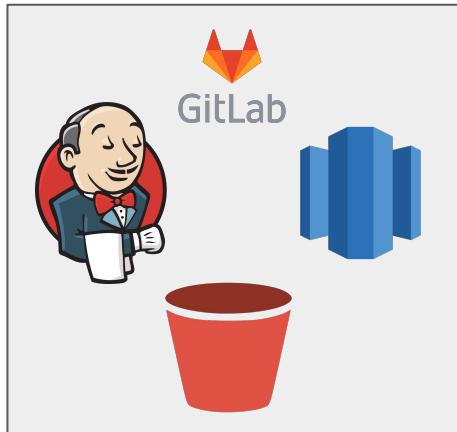
+ 30 offices  
globally



+ 5,500 dedicated  
employees

# Anatomy of a typical “BI Stack”

## Typical Data Stack

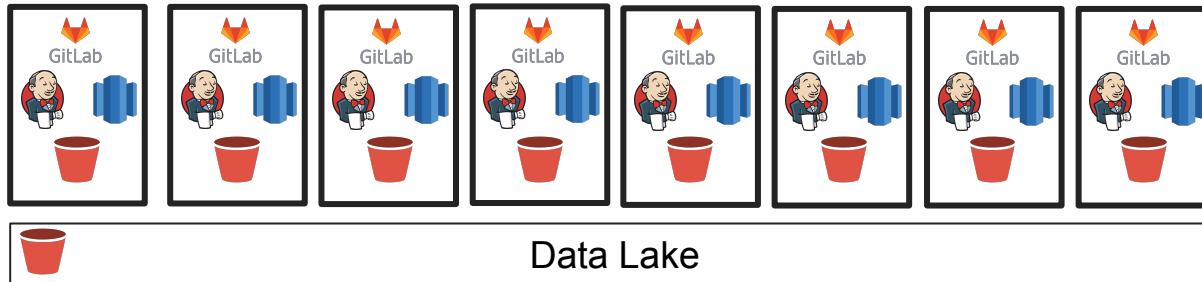


## S3, Redshift, GitLab, Jenkins

- Tight coupling between compute nodes and storage
- Data is stored on the compute nodes
- Low usage of S3 (Spectrum adoption is slower than expected)
- Limited dependency management
- No scheduling standards (random low quality python scripts)

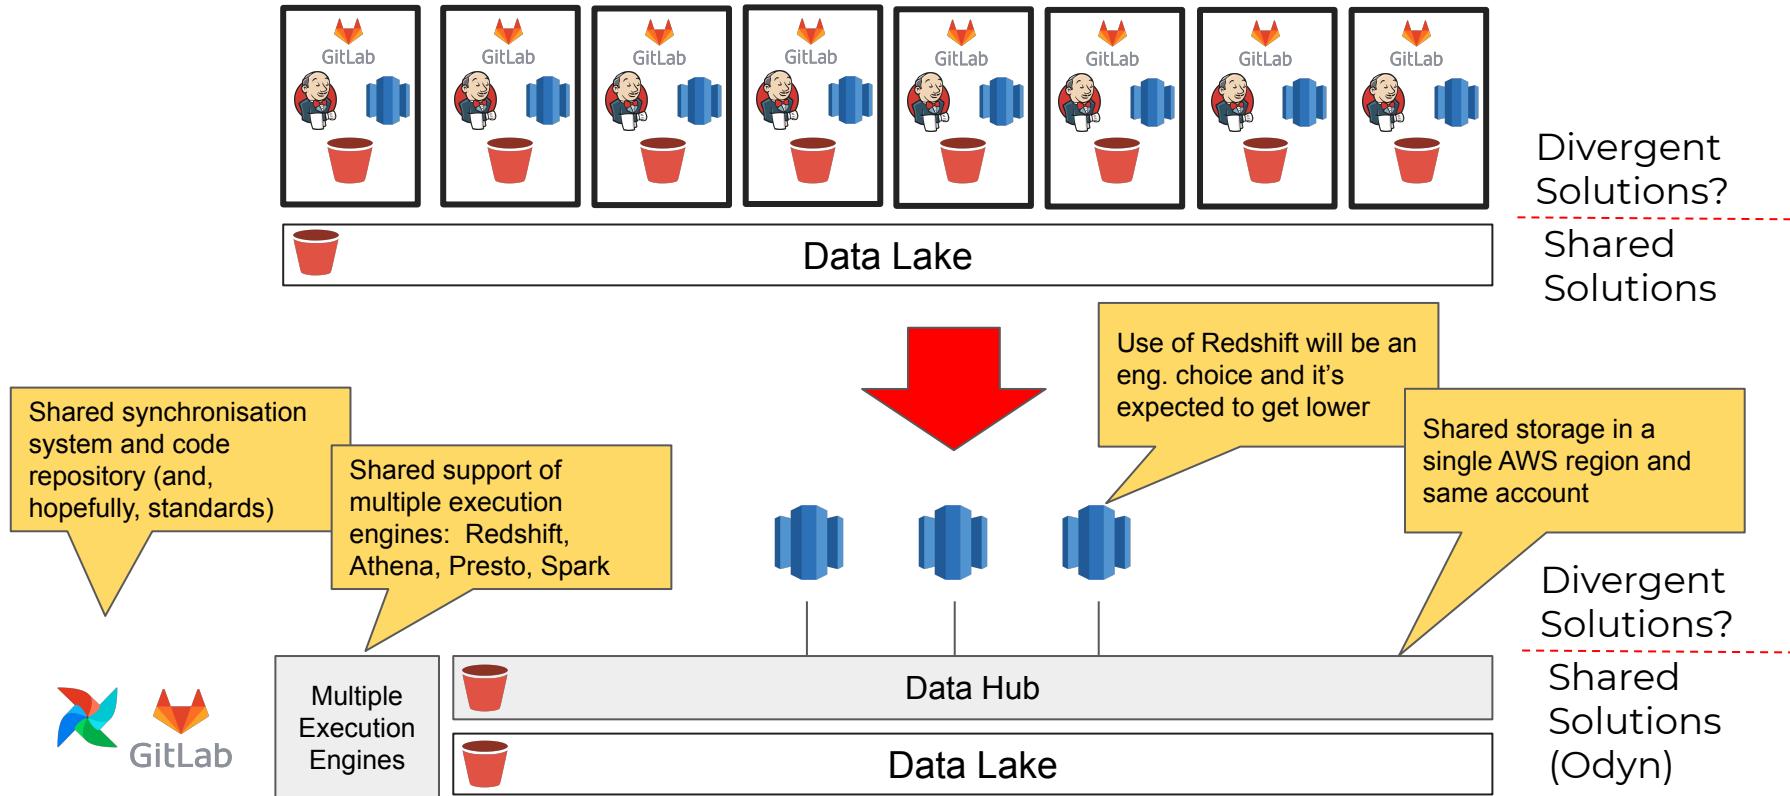
# What are the problems we aim to solve?

- Complex cross-stack **synchronisation** mechanism
- “Reservoir” design discourages **building on each other**
- Use of multiple AWS regions makes **sharing** difficult and increase costs
- Separated ETL scheduling **standards**

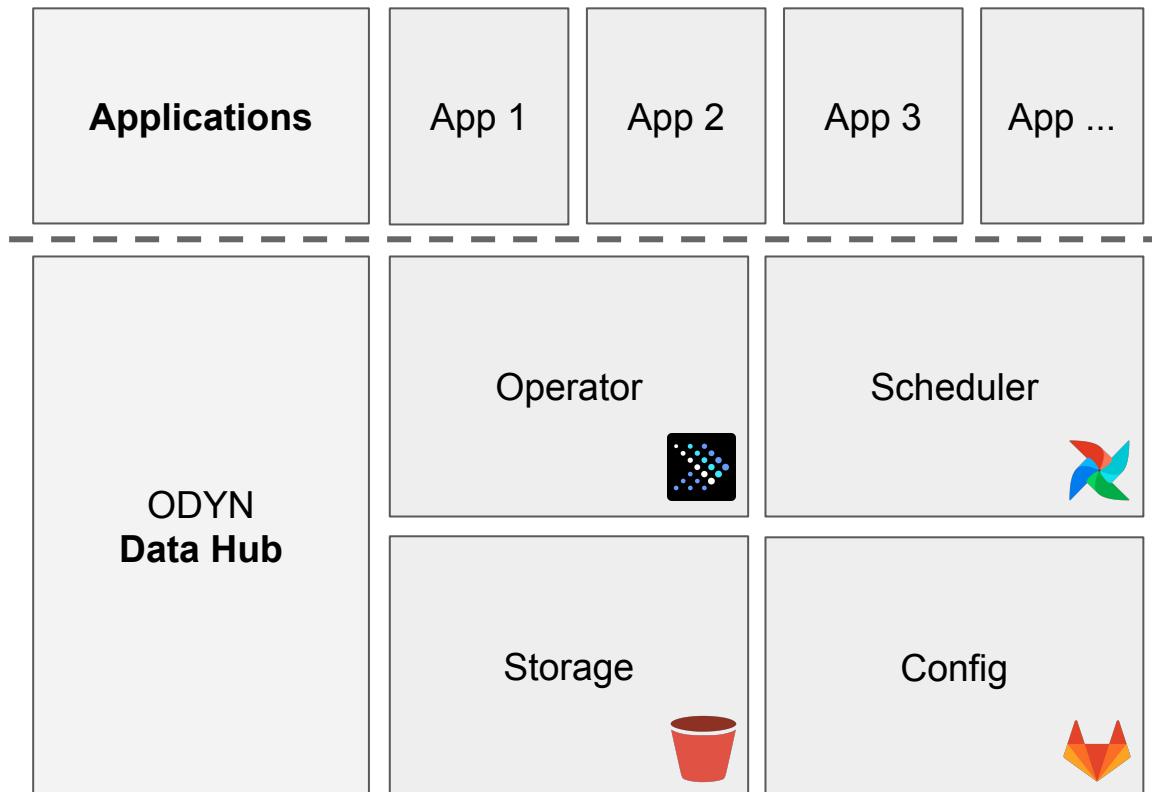


Divergent  
Solutions?  
Shared  
Solutions

# ...and what if?



# OLX Data Hub (“Odyn”) high level architecture overview



# Actual OLX Data Hub (“Odyn”) task configuration example

experiment\_events\_app.odyn 3.64 KB

```
1 type: presto
2 flow: experiment_events_app
3 presto:
4   schema: odyn_experimentation
5 schedule:
6   start_time: '2019-01-01'
7   cron: '30 04 * * *'
8 template:
9   dubizzle_dubizzle_android:
10     region: 'dubizzle'
11     brand: 'dubizzle'
12     channel: 'android'
13     platform: 'android'
14     table: horizontals_dubizzle_hydra_android
15   dubizzle_dubizzle_ios:
16     region: 'dubizzle'
17     brand: 'dubizzle'
18     channel: 'ios'
19     platform: 'ios'
20     table: horizontals_dubizzle_hydra_ios
```

experiment\_events\_app.sql 1.07 KB

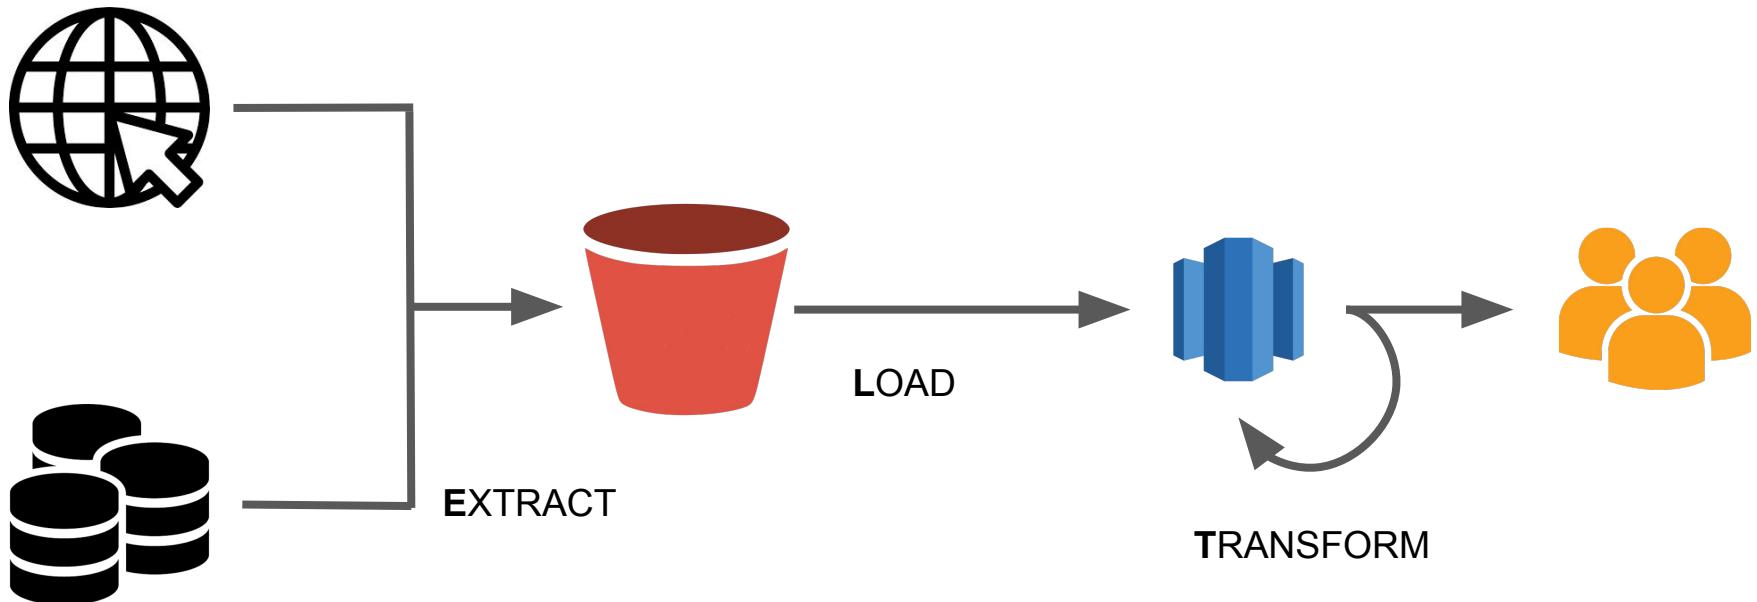
```
1 DELETE FROM awsdatacatalog.odyn_experimentation.experiment_events
2 WHERE region = '${vars:region}'
3   AND brand = '${vars:brand}'
4   AND channel = '${vars:channel}'
5   AND year = '${vars:year}'
6   AND month = '${vars:month}'
7   AND day = '${vars:day}';
8
9 INSERT INTO awsdatacatalog.odyn_experimentation.experiment_events
10 SELECT CAST(meta_date AS INTEGER) AS meta_date,
11        LOWER(params_cc) AS country,
12        '${vars:platform}' AS platform,
13        meta_session_long AS session_long,
14        meta_session AS session,
15        params_en AS event_name,
16        UPPER(params_test_name) AS test_name,
17        params_test_variation AS test_variation,
18        params_test_definition AS test_definition,
19        '${vars:region}' AS region,
20        '${vars:brand}' AS brand,
21        '${vars:channel}' AS channel,
```

# Migrating to Presto

*Why we decided to move out of the Redshift comfort zone*

# Typical data workflow of a “BI stack”

---

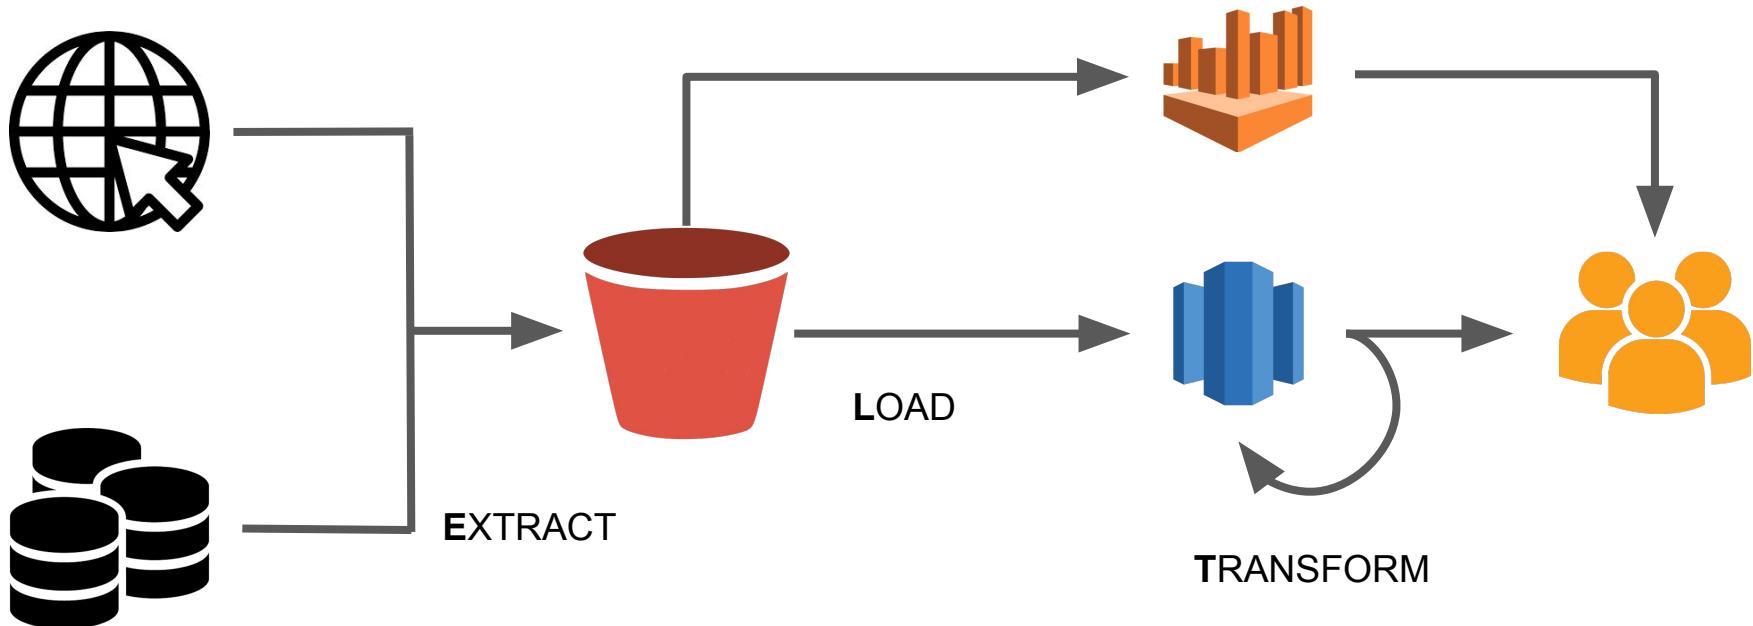


“If you were entering Hadoop ecosystem 8-10 years ago, there was this mantra: bring compute to your storage, tie them together; shipping data is so expensive.

That is no longer true. All modern architectures right now separate storage from compute. Grow your data without limit, scale your compute power whenever you need.”

Kamil Bajda-Pawlikowski, Data Council NY, Nov 7-8, 2018

# Introduced Athena for querying raw data



# Athena adoption failed :-(

- Query exhausted resources
- The query timeout is 30 minutes
- Generic raw data not so friendly for queries
- CTaS usage increase



**Amazon Athena**

Your query has the following error(s):

Query exhausted resources at this scale factor

This query ran against the "olxgroup\_reservoir\_user\_moderation" database, unless qu

New query 1

⚠ New query 2 ✖

New query 3 ✖

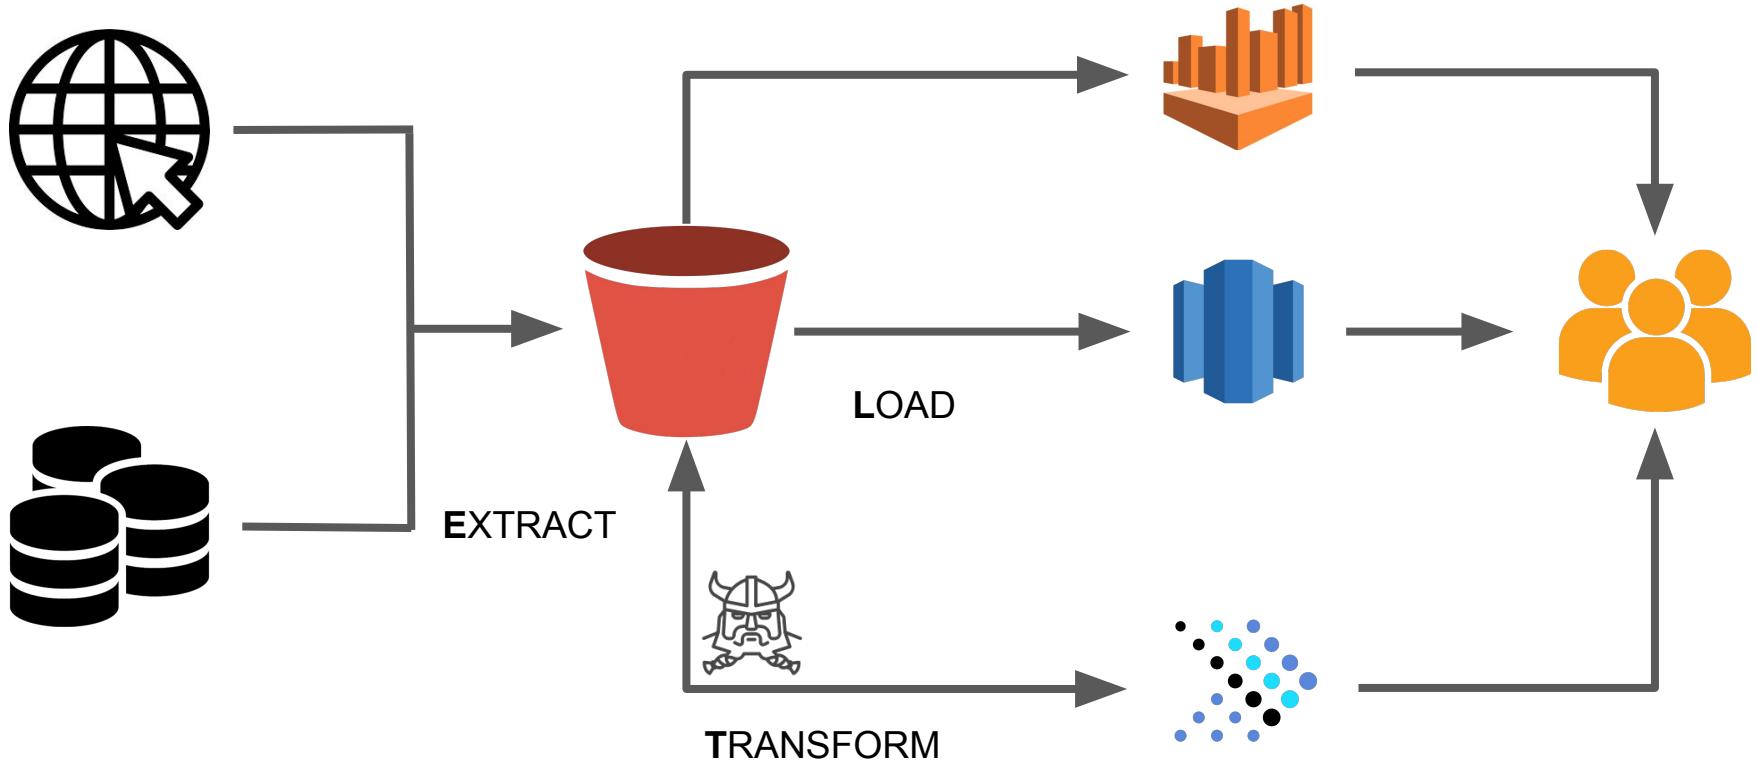
+

Looking for the best query execution engine for our needs

---



# Introduced Presto for processing data



# Presto in production at OLX

- 30+ nodes in AWS (r5.8xlarge)
- 20K+ queries daily
- 100+ users in 20 teams over 5 countries
- 1PB+ data on S3 (Parquet, ORC, JSON)





## RUNNING QUERIES

18

## ACTIVE WORKERS

25

## ROWS/SEC

6.26M

## QUEUED QUERIES

0

## RUNNABLE DRIVERS

28.5K

## BYTES/SEC

1.68G

## BLOCKED QUERIES

0

## RESERVED MEMORY (B)

1.53T

## WORKER PARALLELISM

27.5

## QUERY DETAILS

User, source, query ID, resource group, or query text

State:

 Running Queued Finished Failed

Sort ▾

Reorder Interval ▾

Show ▾

20191202\_200525\_02240\_csbtg

9:05pm

PLANNING

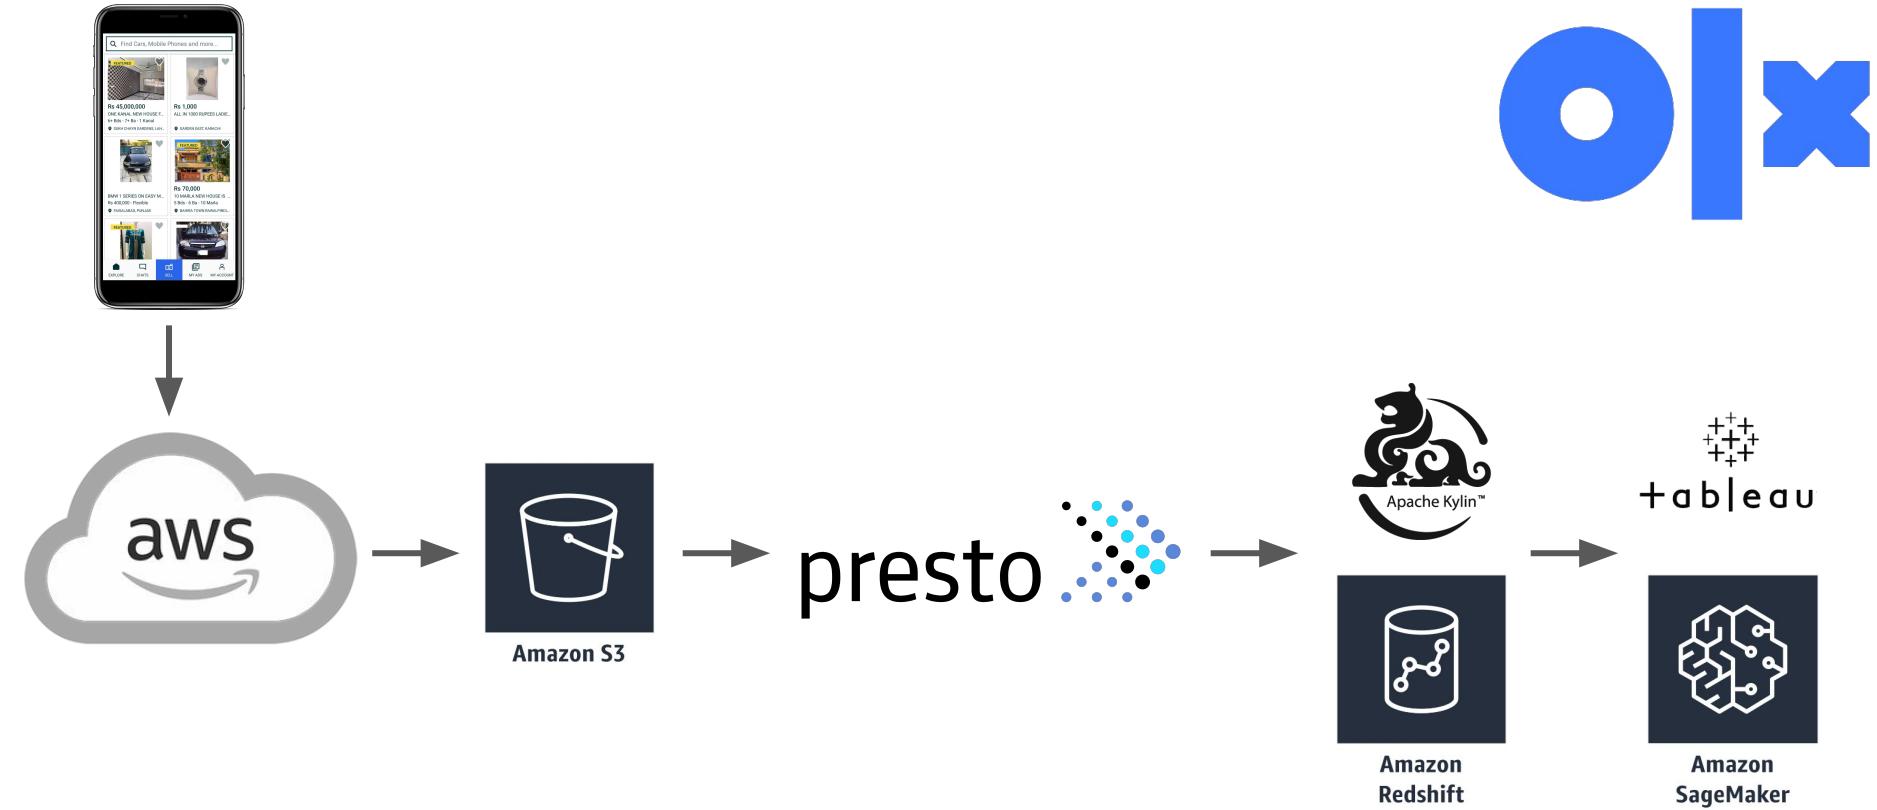
airflow

odyn-airflow data-preparation:ODYN\_...

global

```
-- DELETE
-- FROM odyn_global.fact_user_hydra_browsing_activity_per_hour
-- WHERE site_sk in ('services|eupl')
```

# OLX Data Platform



# Presto Infrastructure

*Where and how we run Presto*

# Where Presto is Running?

- Kubernetes cluster
  - AWS EKS in Ireland
  - Staging and Production
  - Single Amazon availability zone
- We move Presto from EMR to Kubernetes (EKS) using a mix of spot and on-demand instances
- Store metrics in Prometheus and show them in Grafana

Sizes:

- Production = 25 \* r5.8xlarge
- Staging = 16 \* r4.2xlarge



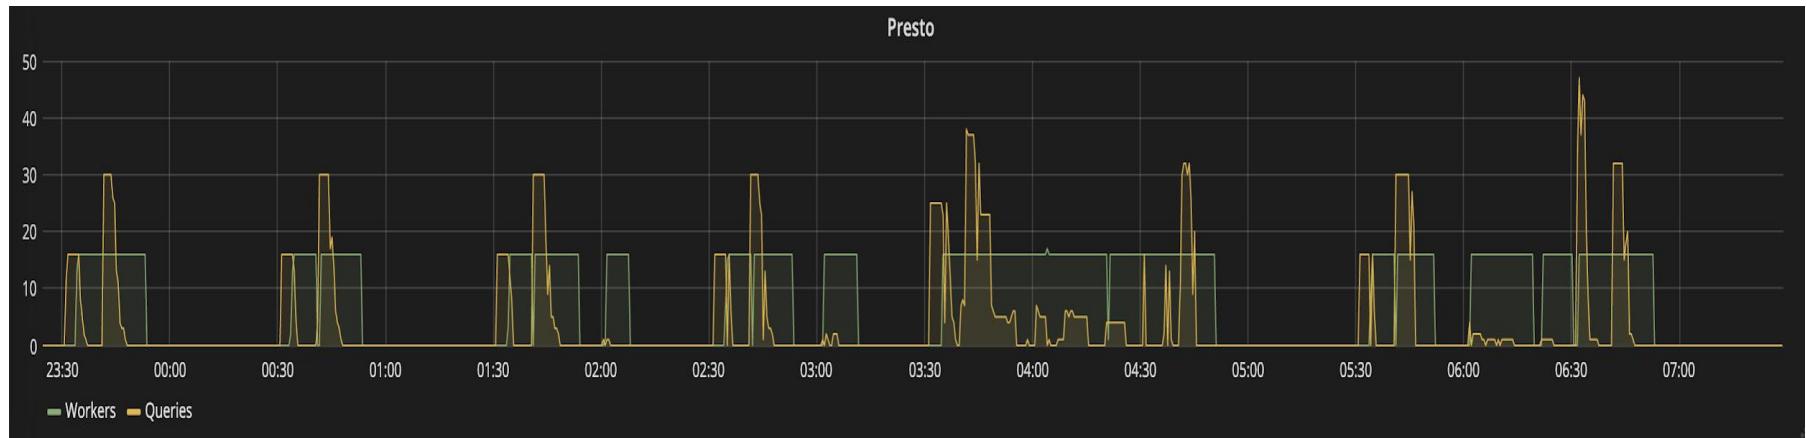
# Challenge

Presto has a static size for the cluster even where there is nothing to do, we need to have the workers nodes up



# Presto “AutoScaling”

We developed our own “auto-scaling” solution for presto workers, allowing us to reduce the cost of the cluster when no queries are running on it



# Next challenges

---

Presto still not 100% integrated in our current ecosystem.

- Cluster for analysts login using our Single Sign (OKTA) on system
- Use different IAM roles depending on user / catalog / table (GDPR).
- Cost-Based Optimizer (using Hive Metastore)



## What's your passion?

**6****Business Intelligence**[LEARN MORE](#)**25****Data Science and Analytics**[LEARN MORE](#)**58****Engineering**[LEARN MORE](#)**8****Finance**[LEARN MORE](#)**5****Human Resources**[LEARN MORE](#)**1****Legal**[LEARN MORE](#)**1****Logistics and Support Staff**[LEARN MORE](#)**2****Management**[LEARN MORE](#)**12****Marketing**[LEARN MORE](#)**8****Operations**[LEARN MORE](#)**17****Product**[LEARN MORE](#)**1****Risk Management**[LEARN MORE](#)**50****Sales**[LEARN MORE](#)



**OLX GROUP**