



KubeCon

North America 2017

Modern Big Data Pipelines Over Kubernetes

Eliran Bivas

iguazio

@bivas

@iguazio

A Tale of 3 Engineers ...



DevOps



Application



Data



[@bivas](#)



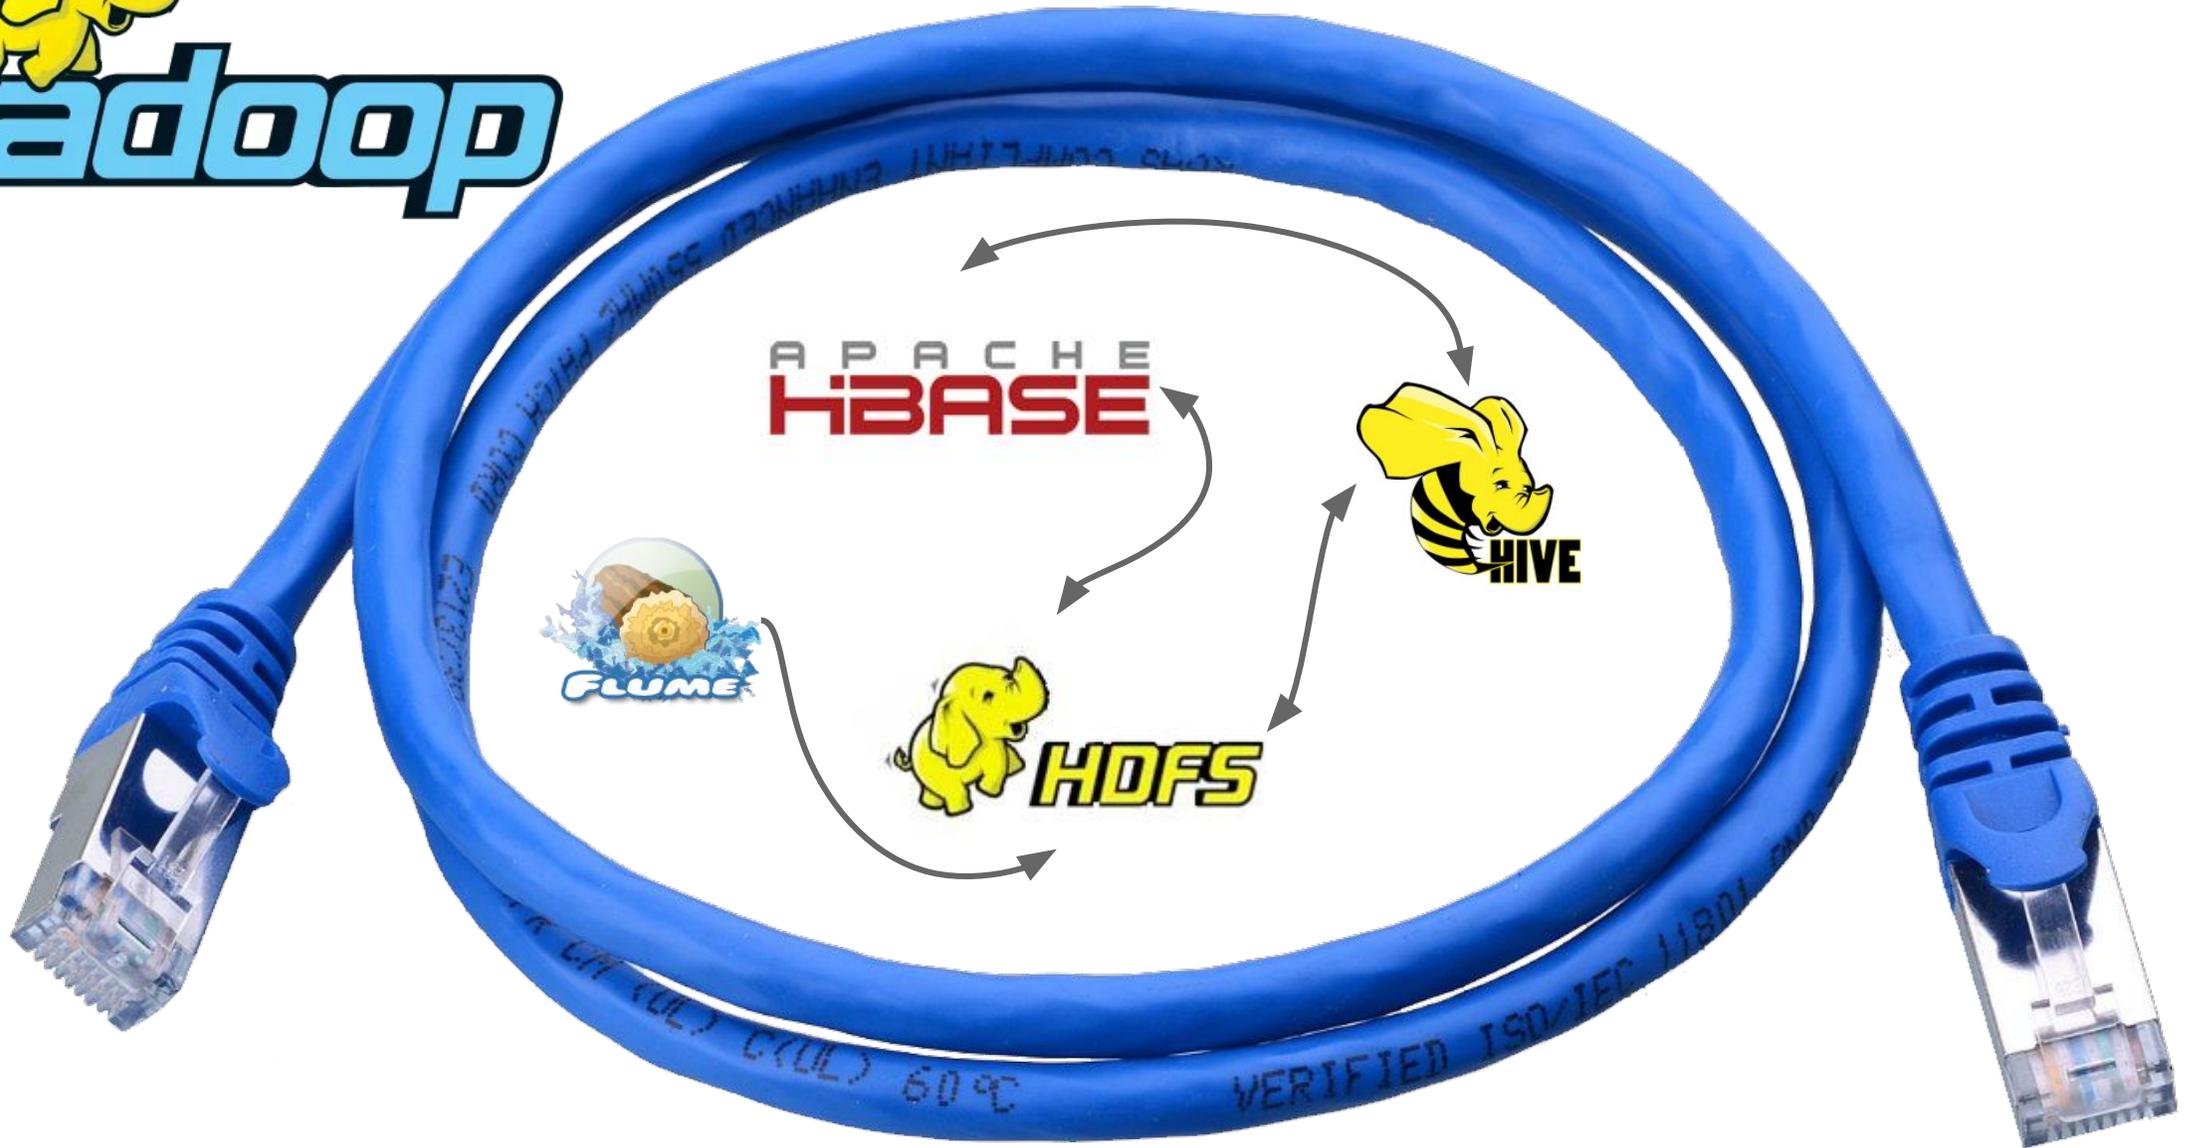
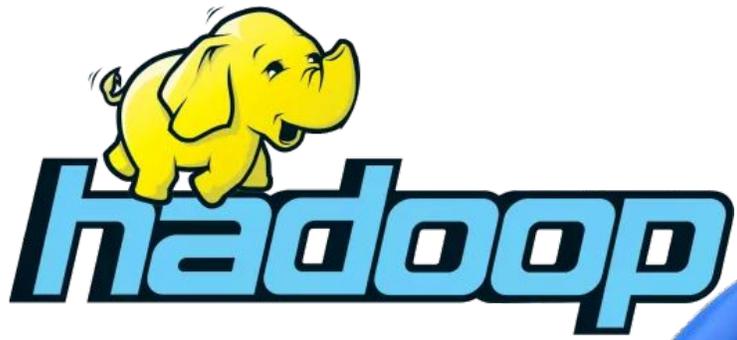
<https://github.com/bivas>

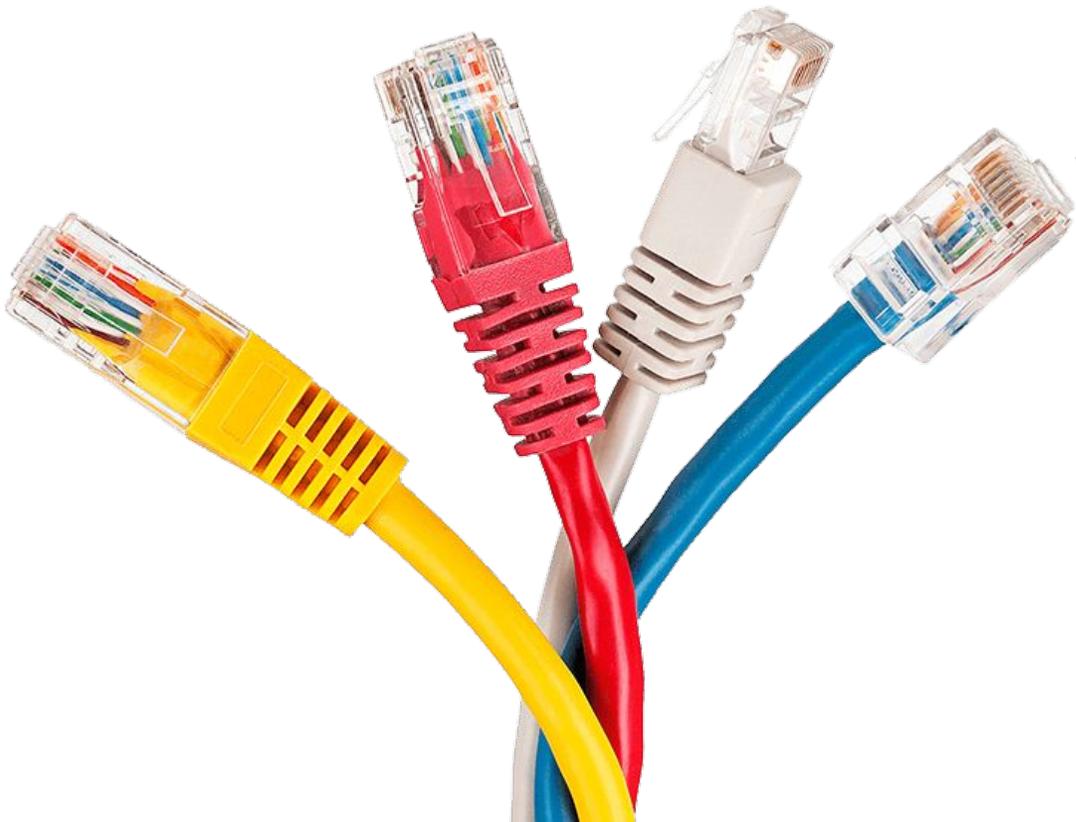


<https://linkedin.com/in/bivas>



<https://www.iguazio.com>





Spark

APACHE
HBASE



cassandra



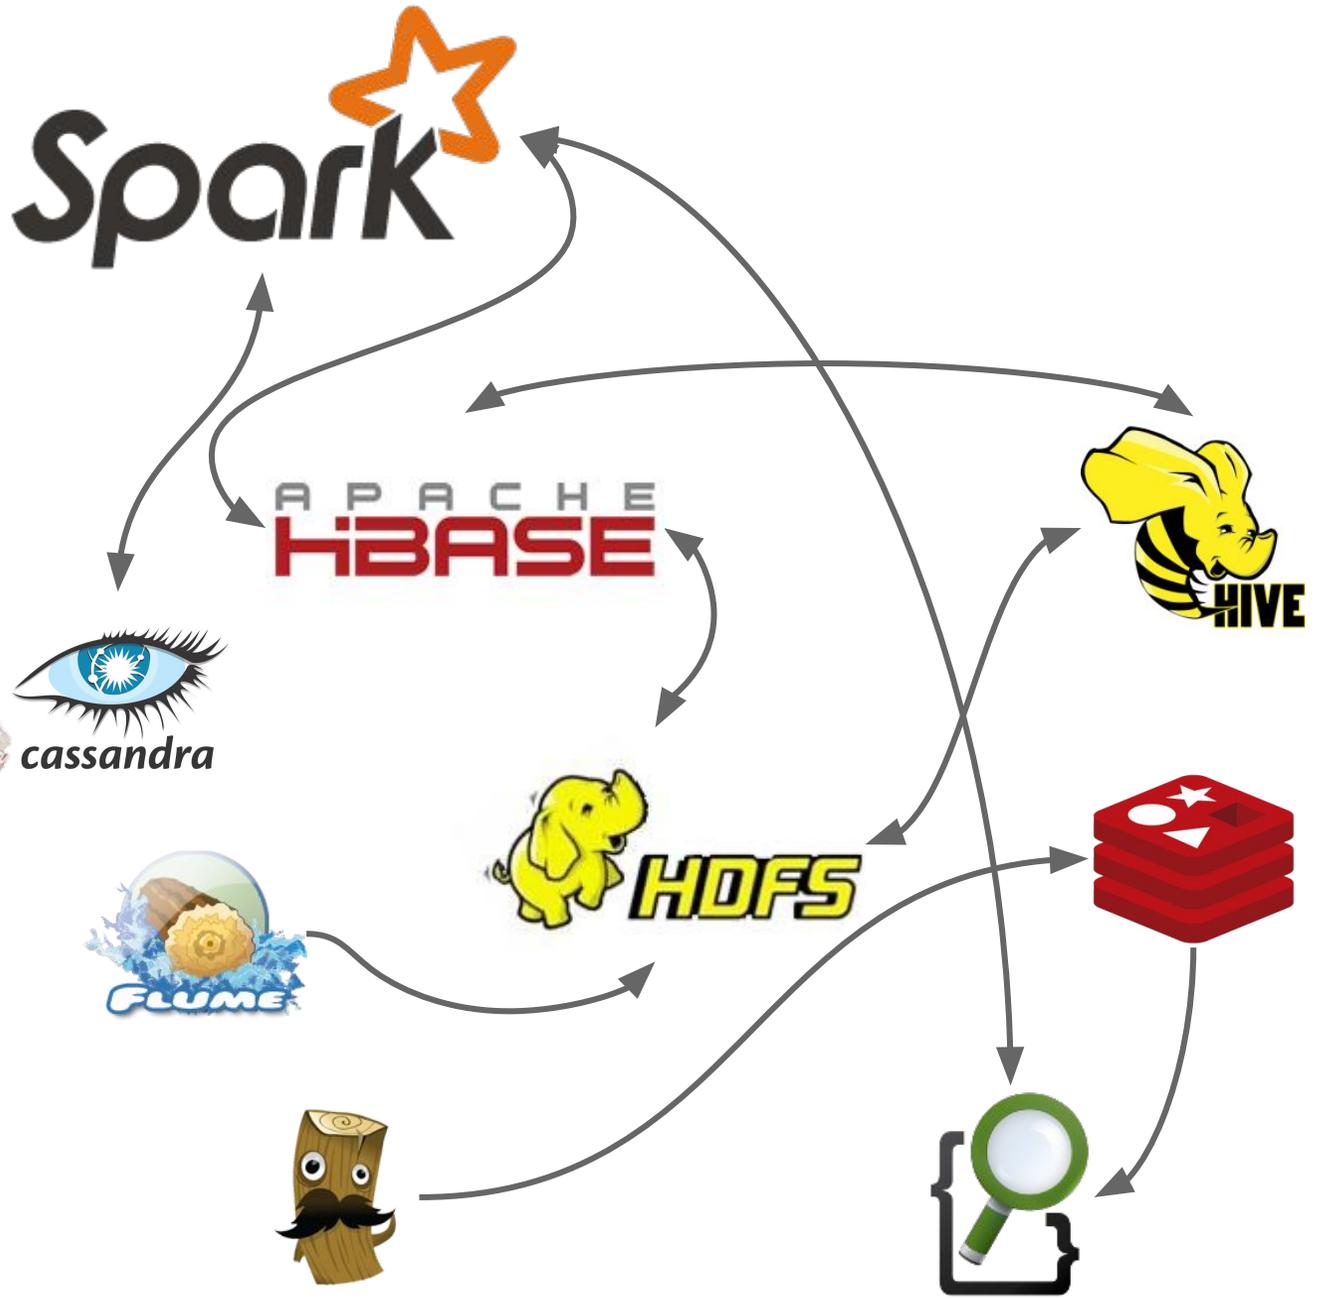
FLUME

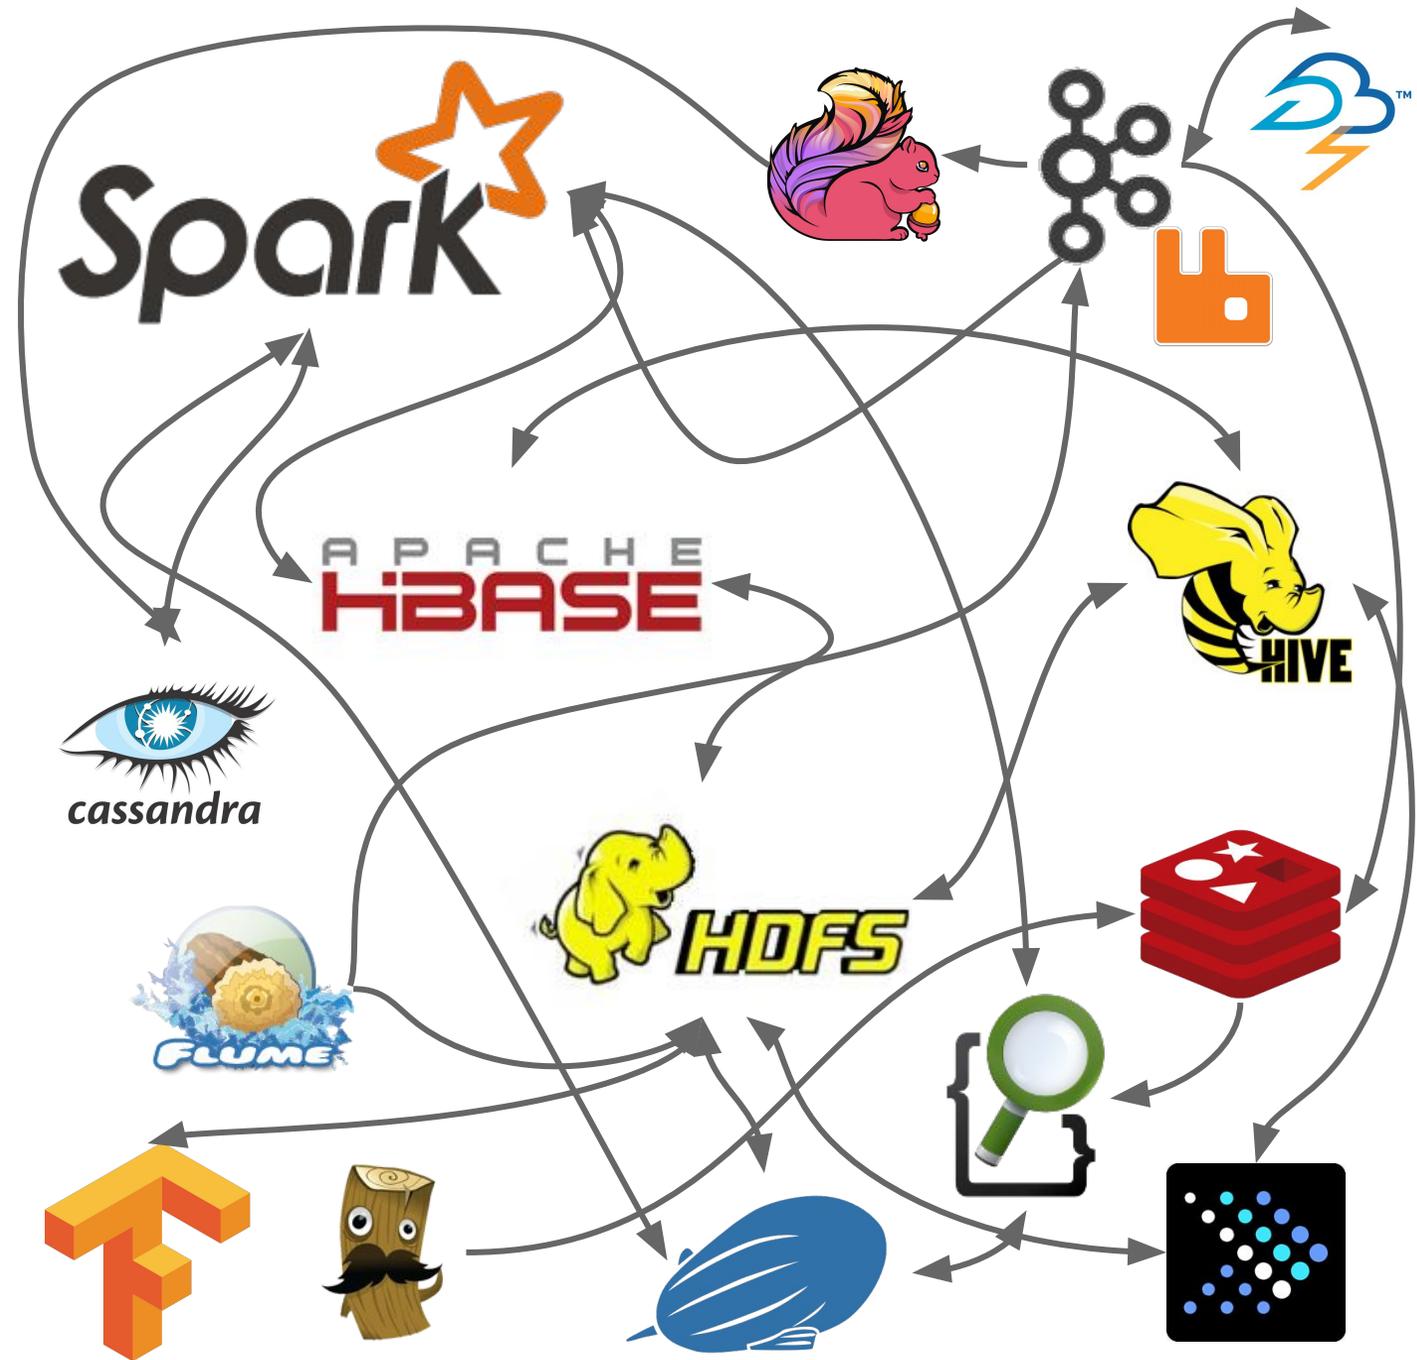


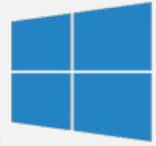
HDFS



HIVE







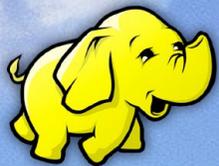
Microsoft
Azure



Google
Cloud Platform



amazon
web services

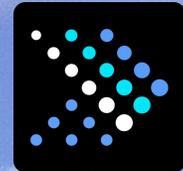


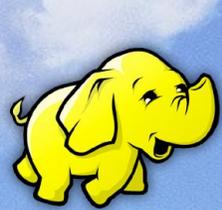
APACHE
HBASE



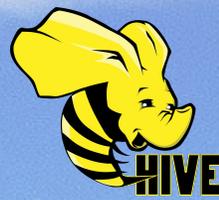
HIVE

Spark



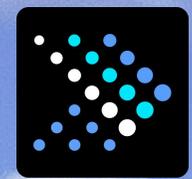


APACHE
HBASE



HIVE

Spark



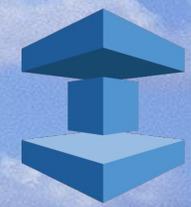
S3



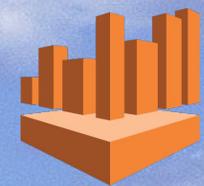
Elastic MapReduce



AI



ElastiCache



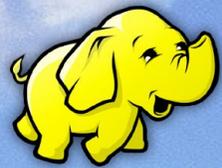
Athena



Kinesis



Lambda

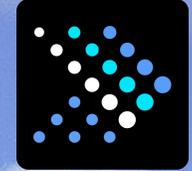


APACHE HBASE

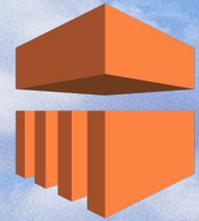


HIVE

Spark



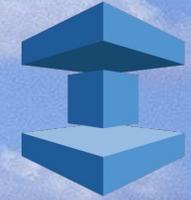
S3



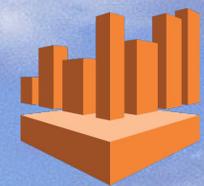
Elastic MapReduce



AI



ElastiCache



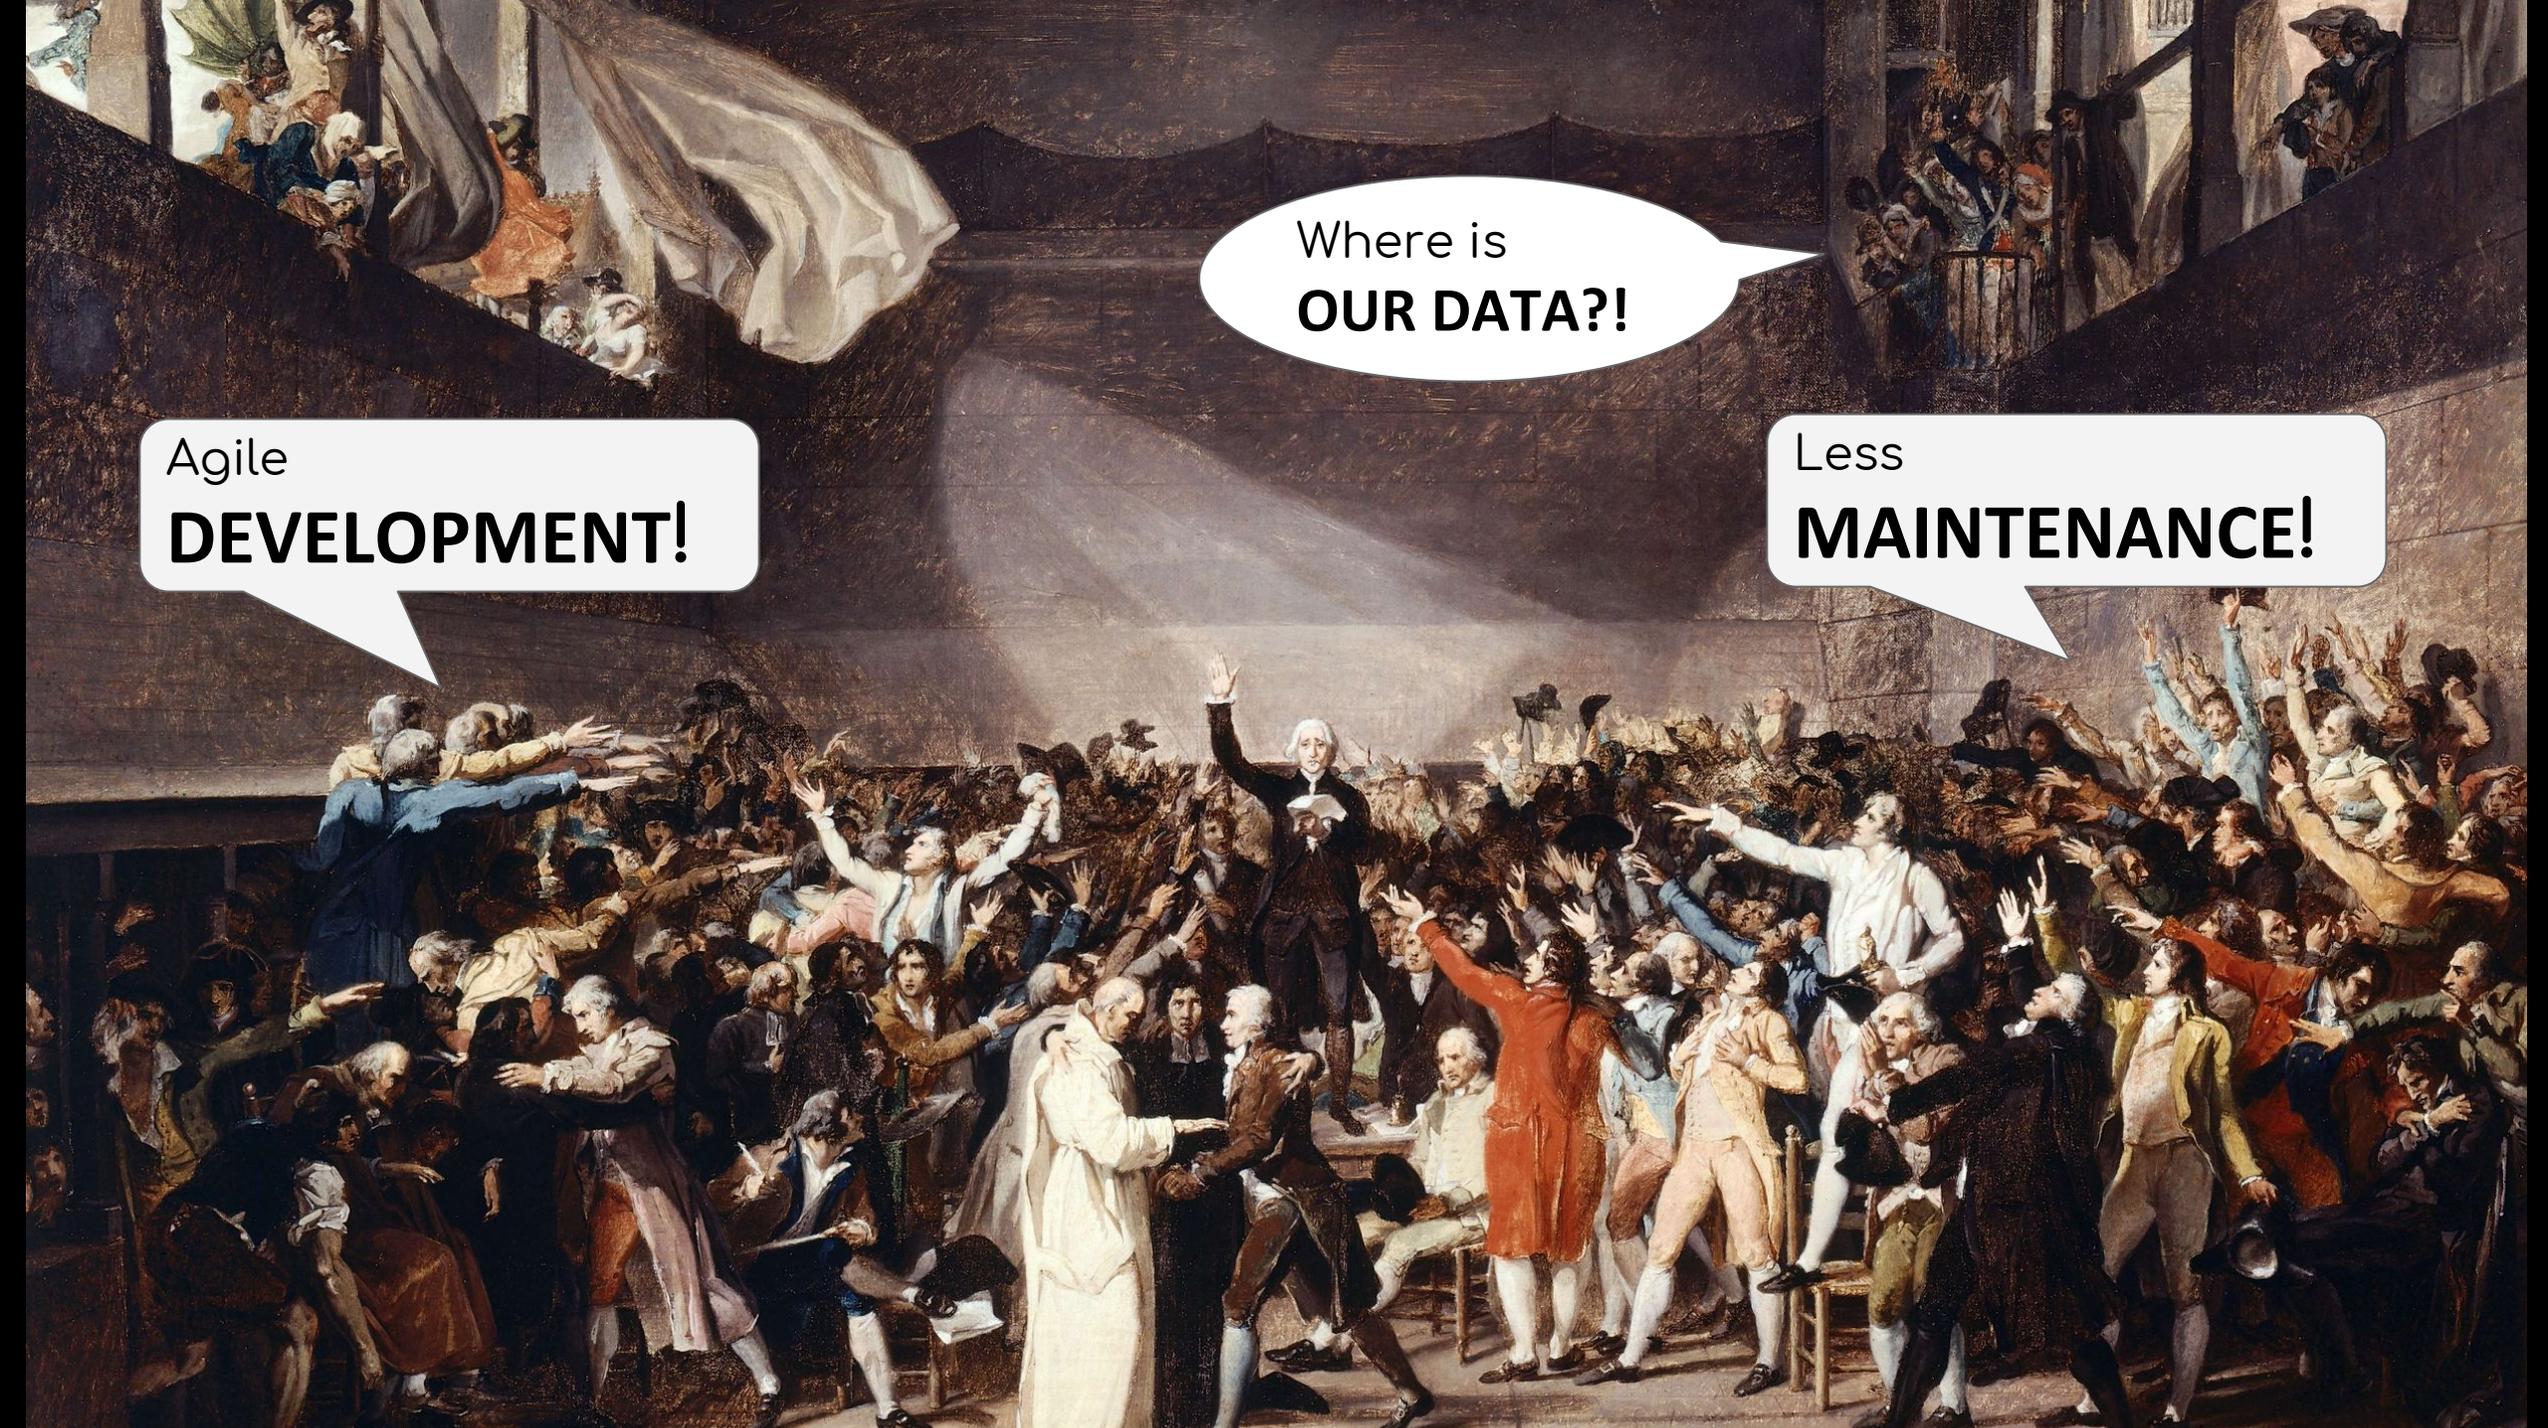
Athena



Kinesis



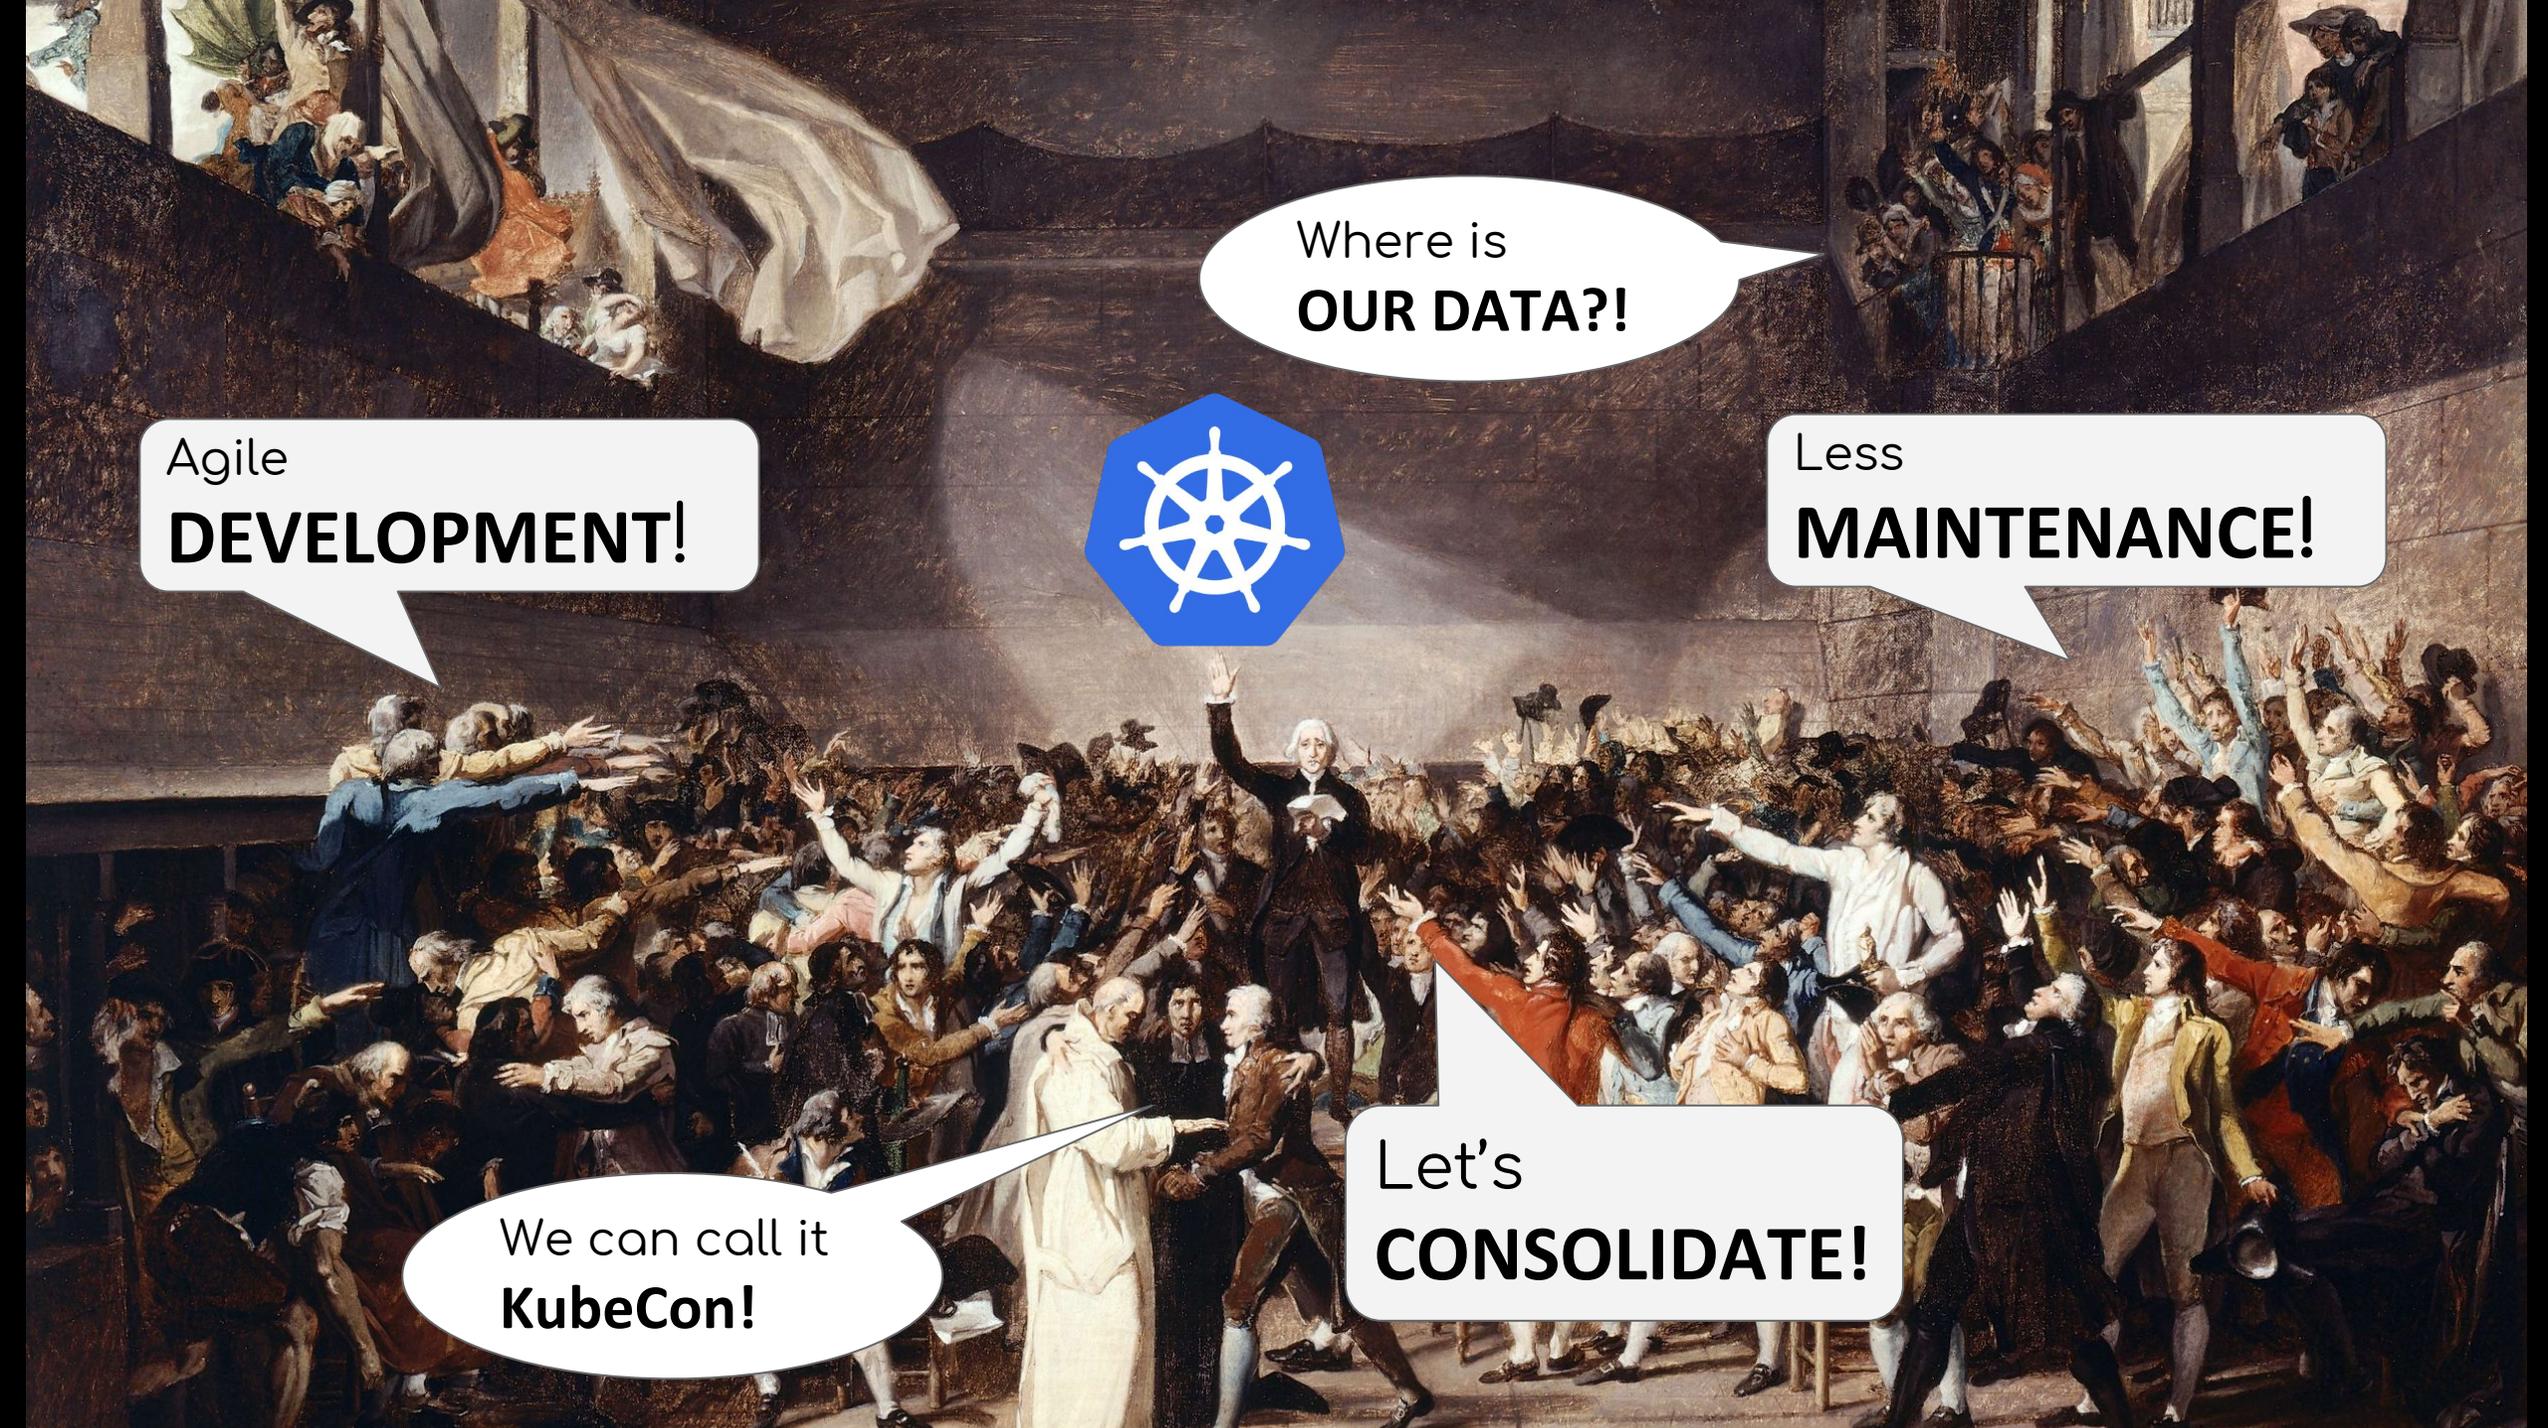
||| Pause



Where is
OUR DATA?!

Agile
DEVELOPMENT!

Less
MAINTENANCE!



Where is
OUR DATA?!

Agile
DEVELOPMENT!



Less
MAINTENANCE!

We can call it
KubeCon!

Let's
CONSOLIDATE!

Data Services



Unstructured

S3 / GlusterFS / Iguazio



Structured

DynamoDB / Iguazio / Cassandra



Stream

Kinesis / Kafka / Iguazio



**Decouple Data
Services from
Application and
Frameworks**



Data Services



Unstructured

S3 / GlusterFS / Iguazio



Structured

DynamoDB / Iguazio / Cassandra



Stream

Kinesis / Kafka / Iguazio



**Unified
Orchestration**



**Decouple Data
Services from
Application and
Frameworks**



Application



Serverless



Data Services

 Unstructured S3 / GlusterFS / Iguazio	 Structured DynamoDB / Iguazio / Cassandra	 Stream Kinesis / Kafka / Iguazio
---	--	--



Cloud Native
Application and
Frameworks

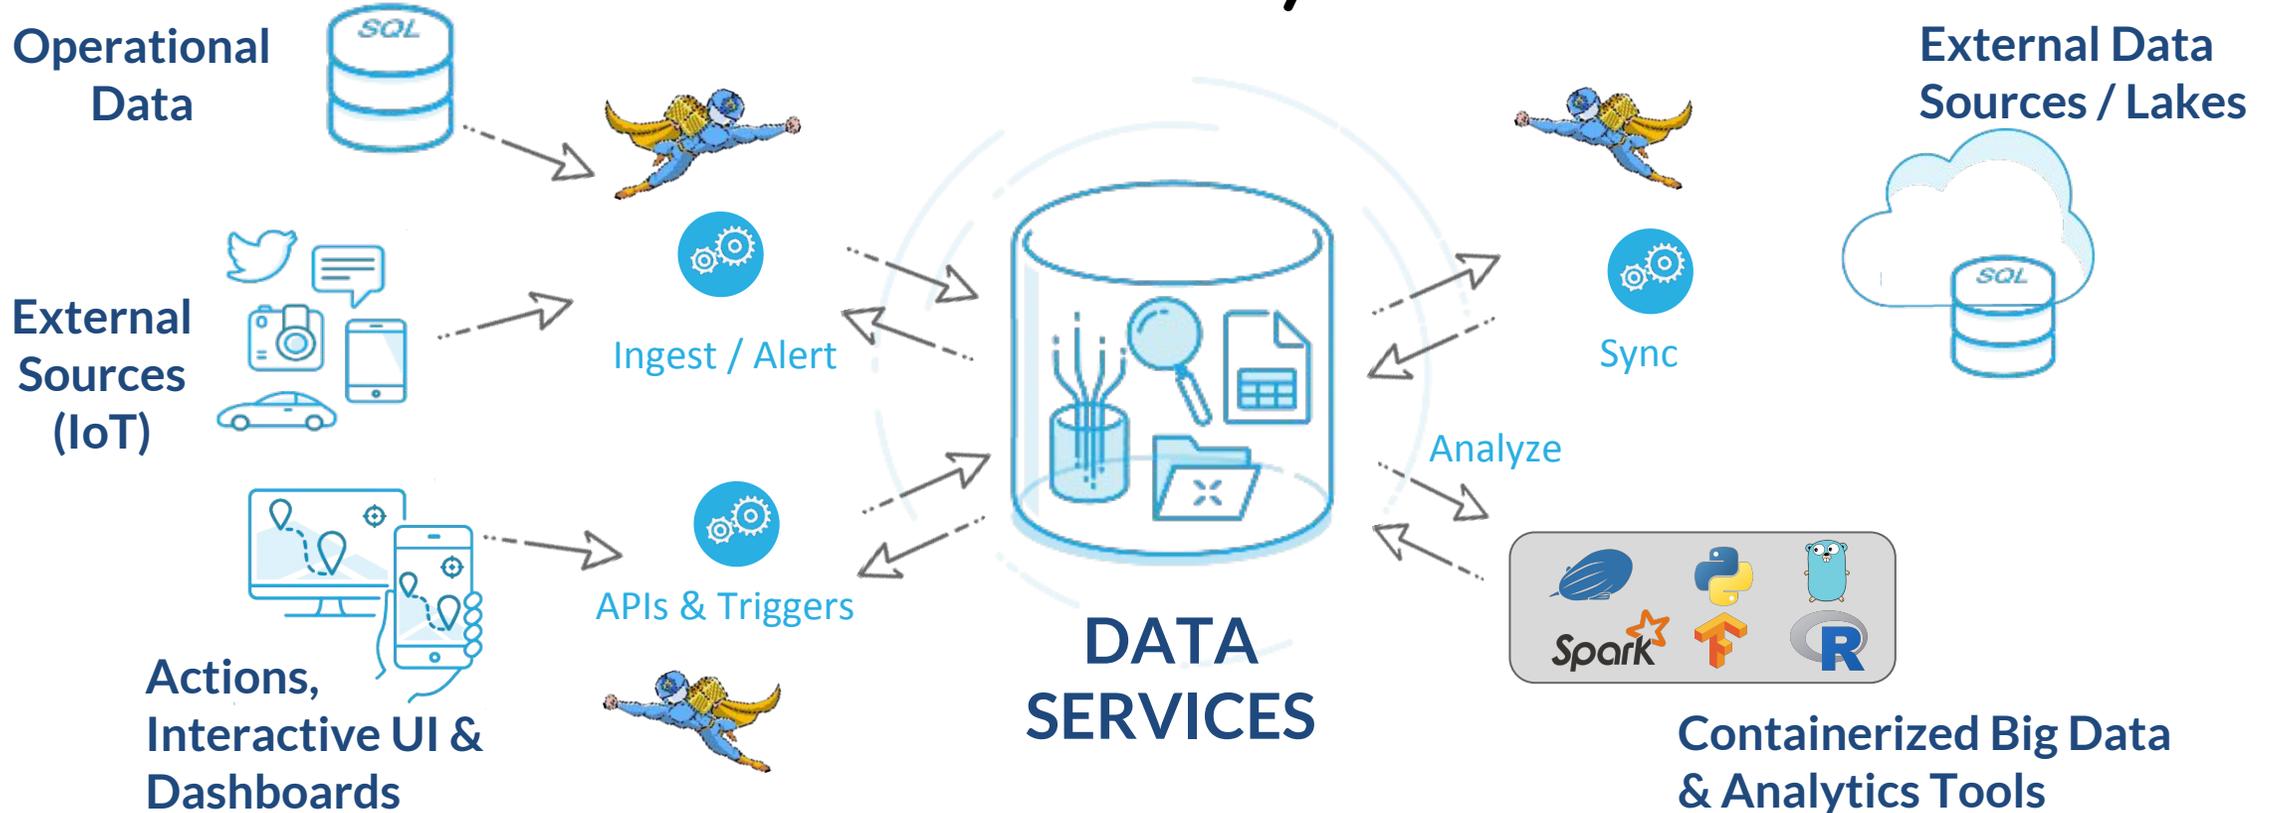


Unified
Orchestration



Decouple Data
Services from
Application and
Frameworks

Continuous Analytics Flow



Serverless

- ✓ Easy to deploy functions, no YAMLS, Dockerfile, build, etc.
- ✓ Auto-scaling and event triggers are built-in
- ✓ Use any language

BUT

- ✗ Slow performance, lack of concurrency, no stream partitioning
- ✗ Limited number of event sources (mostly HTTP, some kafka)
- ✗ Hard to debug and diagnose
- ✗ Hard to manage complex functions and dependencies

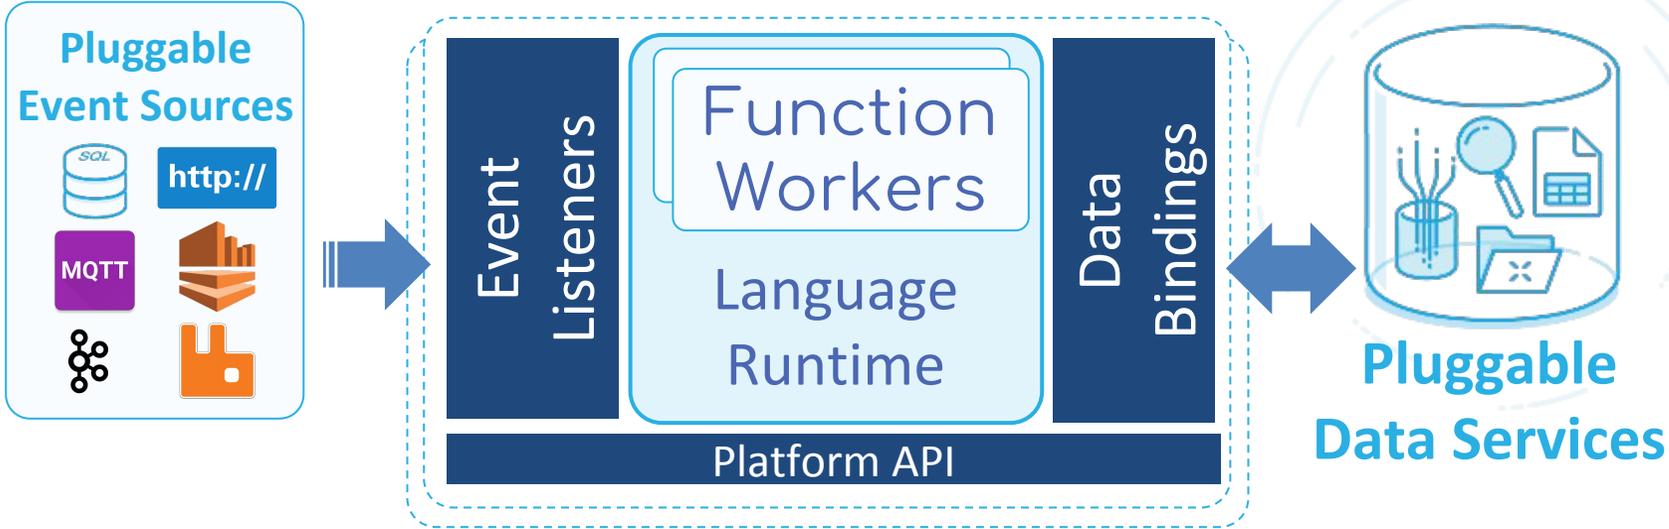


nuclio A New Real-Time Serverless Platform

Comprehensive, Open, Super Fast and runs ANYWHERE

400K Events/sec per Process

Real-time, Auto-scaling Function Processors



Common Services

- CLI
- Dealer
- Web UI
- Builder

Log & Monitor

- Log icon
- Flame icon
- elasticsearch.

Platform Specific Controller

- Raspberry Pi icon
- Docker icon
- Kubernetes icon

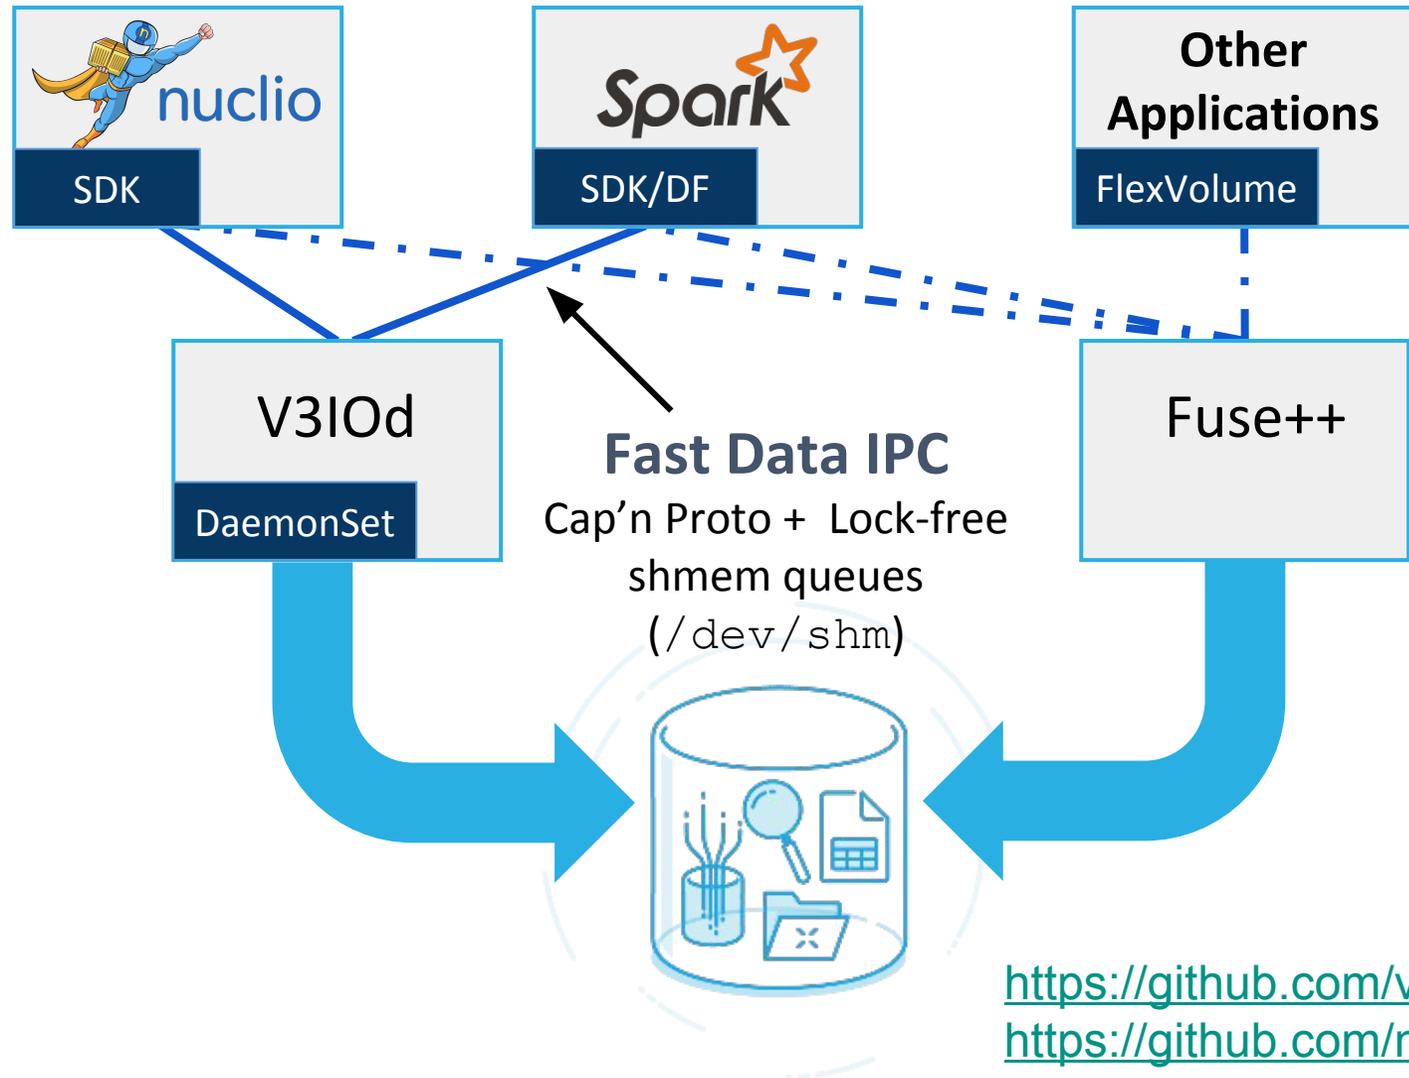


- Any Source and Any Workload
- Maximum Performance
- Dev, Test, and Run ANYWHERE
- Simple, Fast, Secure data integration

<https://nuclio.io>

Accelerate Performance Using Shared Memory

-  Fast Data Access
-  Fast Microservice Initialization
-  Share TCP/RDMA Connections
-  NO Kernel Changes!



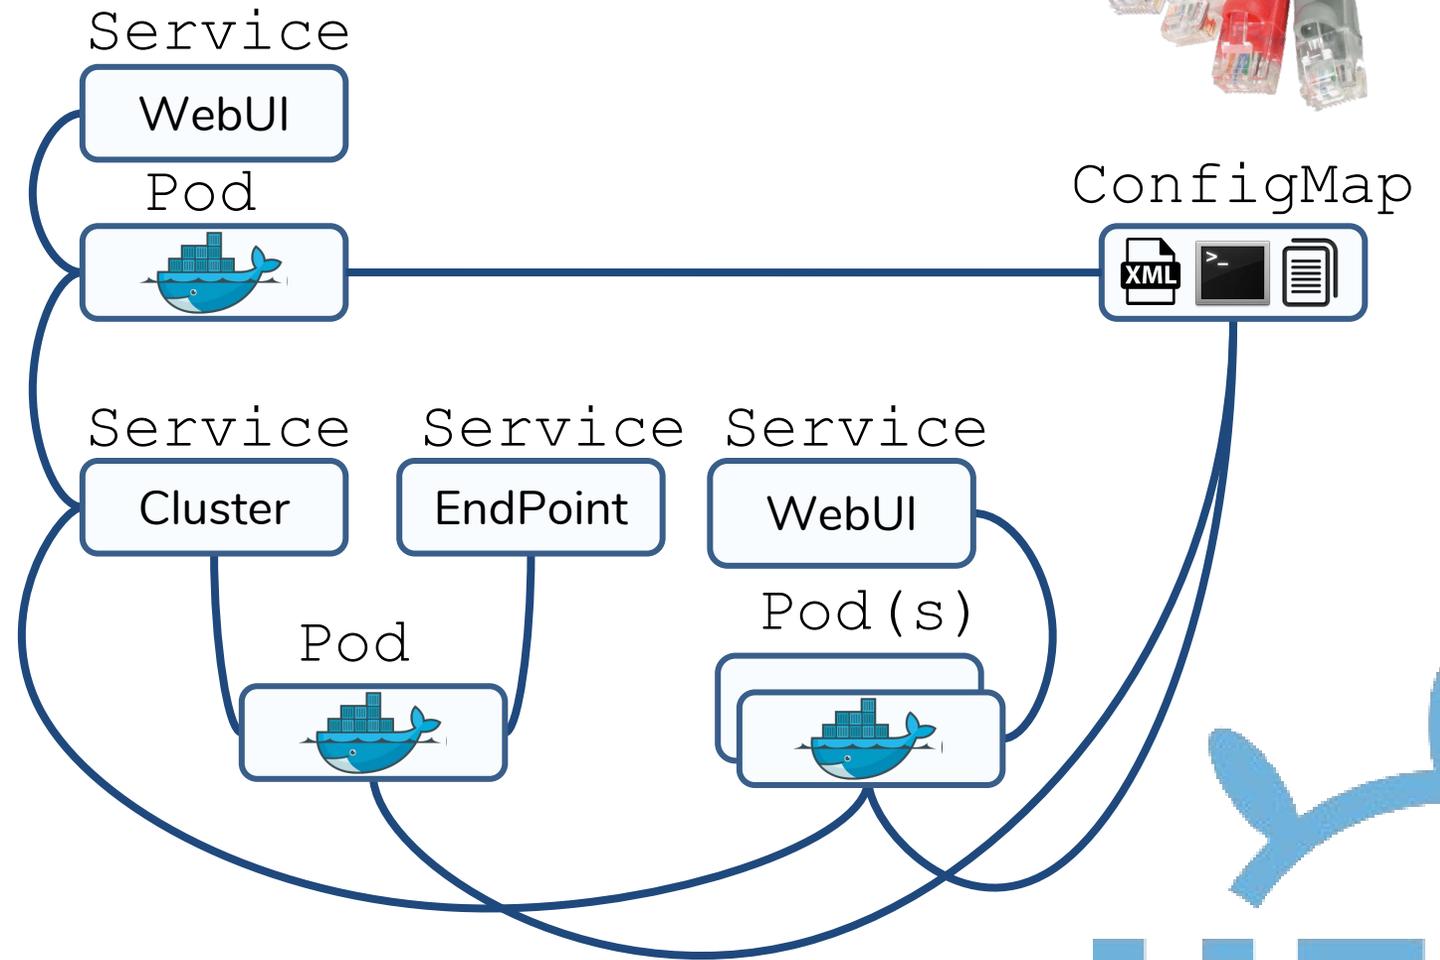
Spark + HELM USE CASE



Deployment
YAML

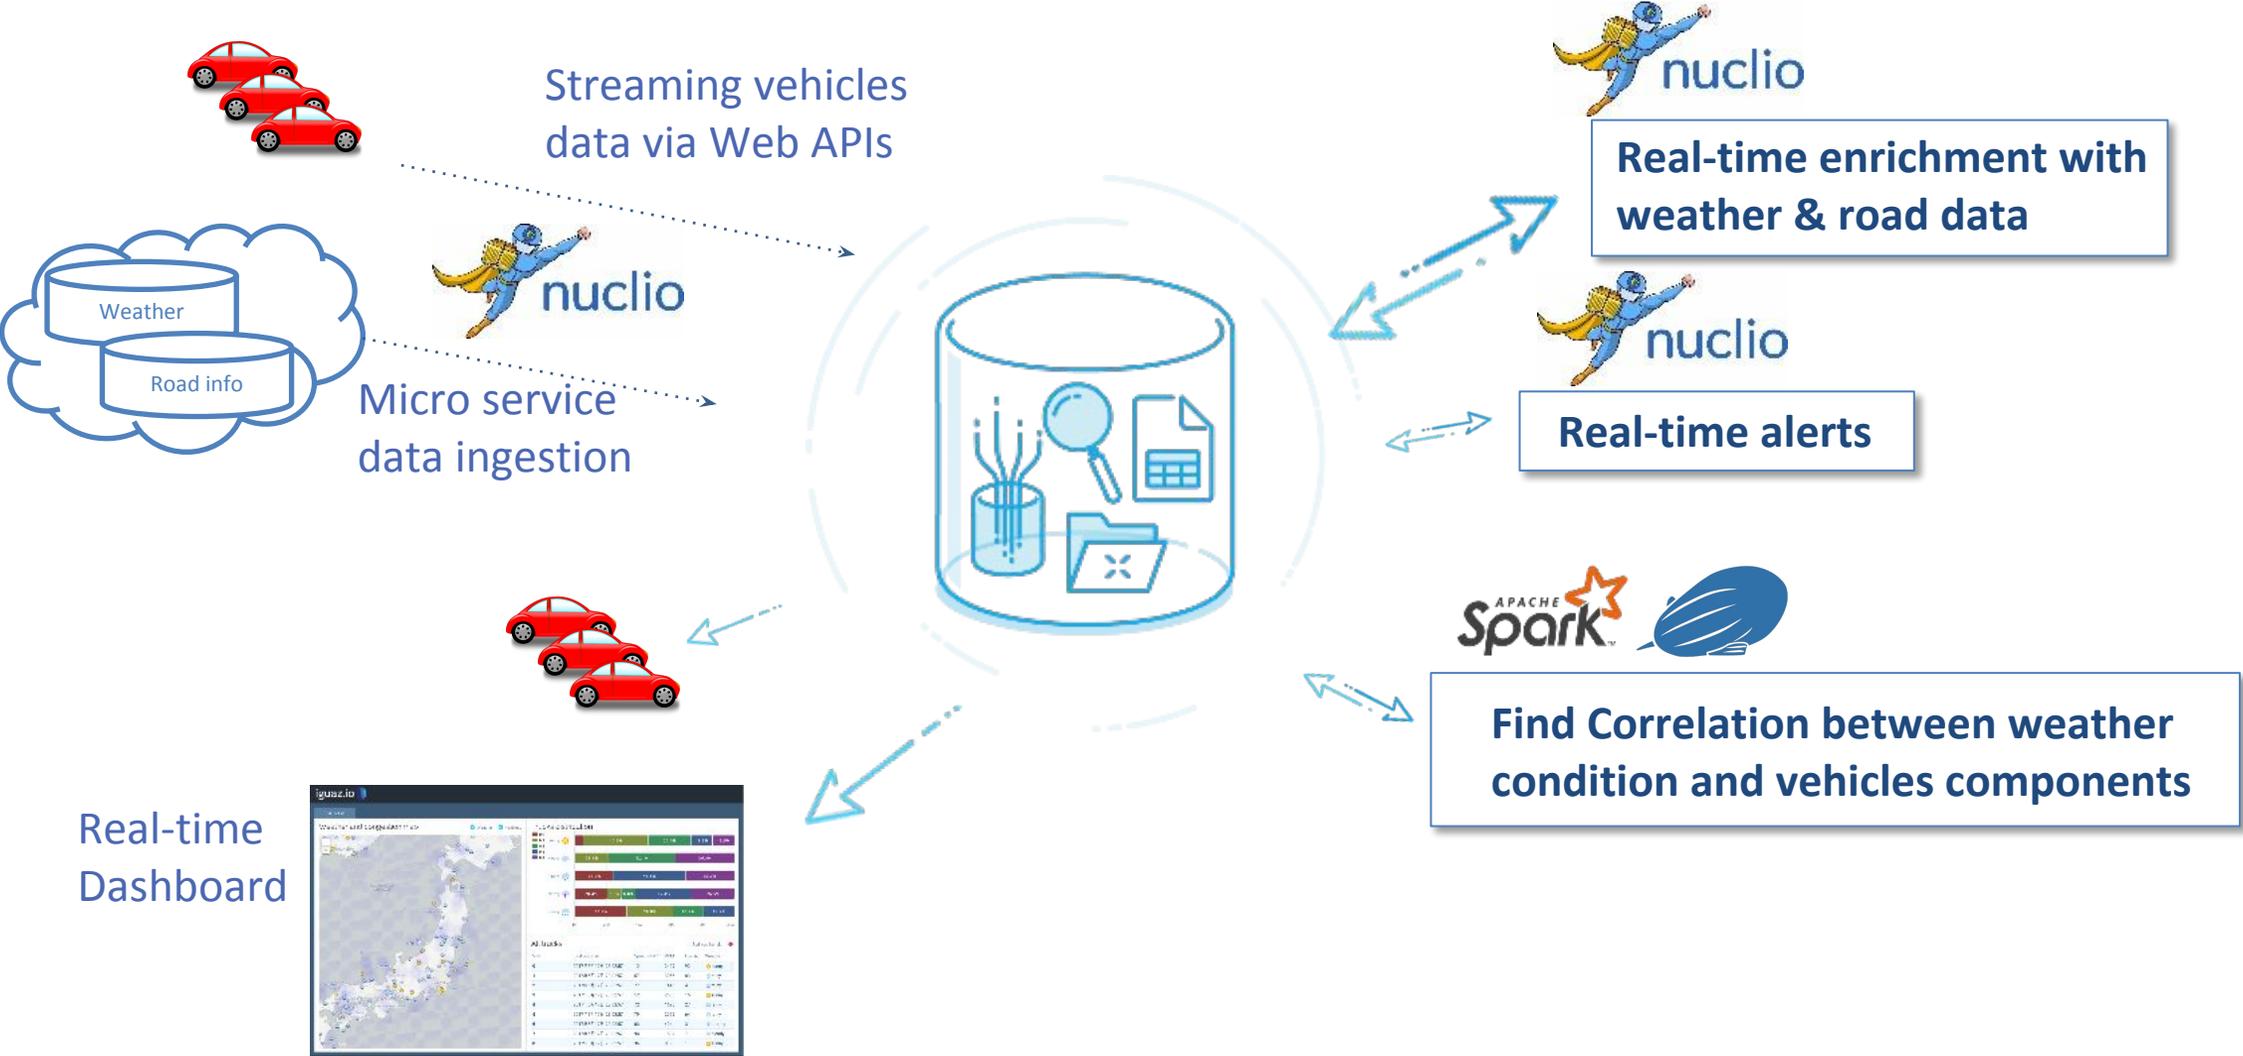


Deployment
YAML



CUSTOMER USE CASE

Real time analytics for vehicle maintenance





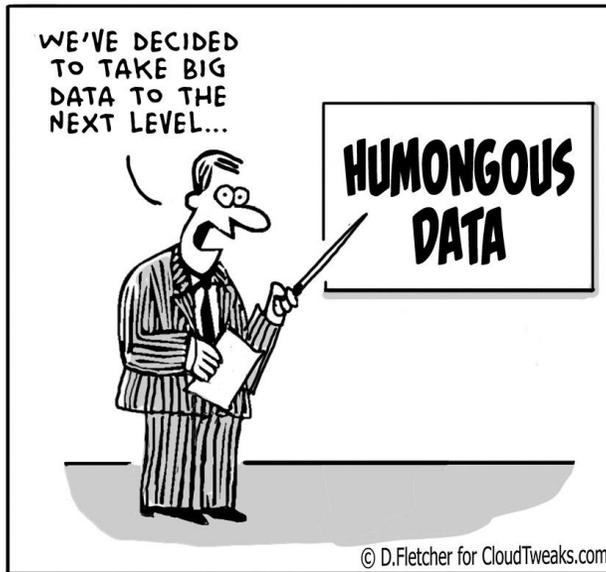
DEMO

Tips & Lessons Learned

1. RTFM!
2. Community (!/?)
3. Know `kubectl` (`port-forward...`)
4. Navigate with **Helm**
5. Do NOT overuse `NodePort`
6. Configuration goes to `ConfigMap`
7. Override `command` in `YAML` is tricky -
use `init script` and keep it in `ConfigMap`
8. Collect operational data and use it!



Thank You





KubeCon

North America 2017

