Hops in the Cloud

•••

Jim Dowling,
CEO at Logical Clocks
@jim_dowling

BERLIN BUZZWORDS
2019 JUNE 16-18

Brief History of Hops

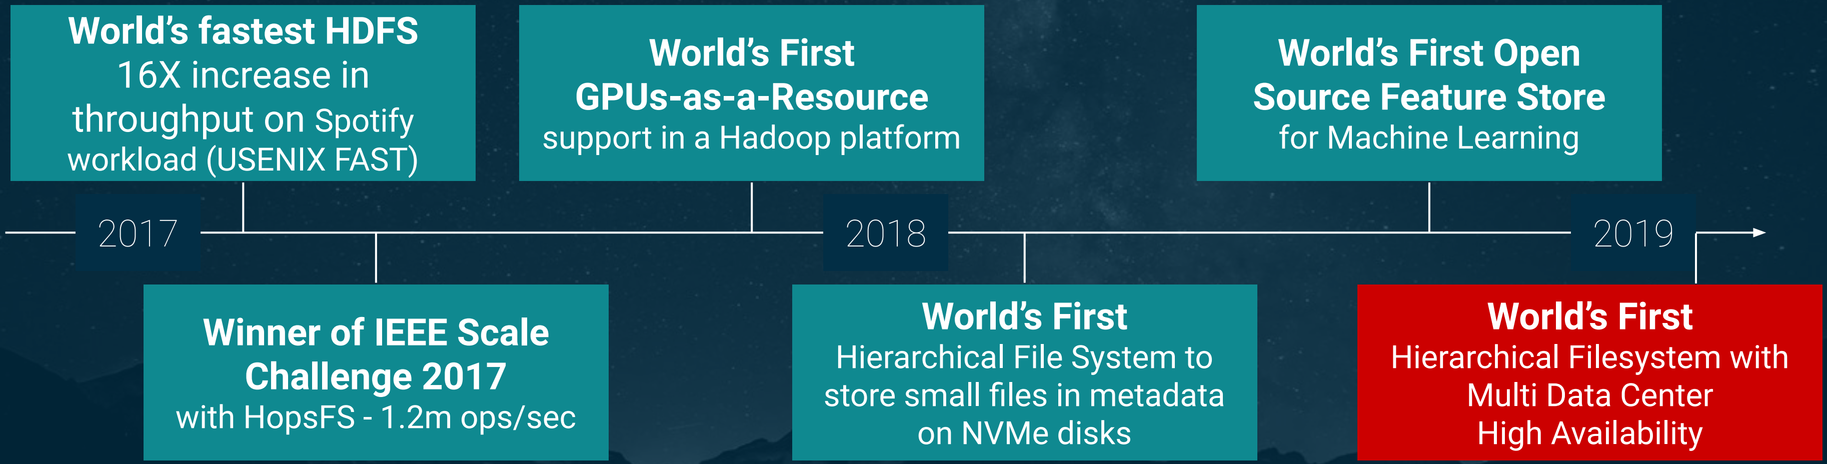

2017

World's fastest HDFS
16X increase in throughput on Spotify workload (USENIX FAST)

Winner of IEEE Scale Challenge 2017
with HopsFS - 1.2m ops/sec

World's First GPUs-as-a-Resource
support in a Hadoop platform

2018

World's First
Hierarchical File System to store small files in metadata on NVMe disks

World's First Open Source Feature Store
for Machine Learning

2019

World's First
Hierarchical Filesystem with Multi Data Center High Availability

"If you're working with big data and Hadoop, **this one paper could repay your investment** in the Morning Paper many times over.... **HopsFS is a huge win.**"

- Adrian Colyer, The Morning Paper

Quick overview of Hops/Hopsworks

The **only** open-source data platform to support:

- Project-based multi-tenancy
- On-premise resource management of GPUs (>1 server)
- Per-Project Python Dependencies with Conda
- Feature Store
- Jupyter notebooks as Jobs (Airflow)
- Free-text search for files/dirs in the filesystem
- NVMe to store small files in filesystem metadata

Example workflow in Hopsworks at Scale

1. Insert 1m images (<100kb) in seconds
2. Train a DNN classifier using 100s of GPUs
3. Run a Spark job to identify all objects in the 1m images and add the image annotations (JSON) as extended metadata to HopsFS
4. "show me the images with >3 bicycles" and get a sub-second response.

Ops folks: Remove the image directory, and elasticsearch is auto-cleaned up!
Data scientists: Do it all in Jupyter notebooks and Python (if you want)!

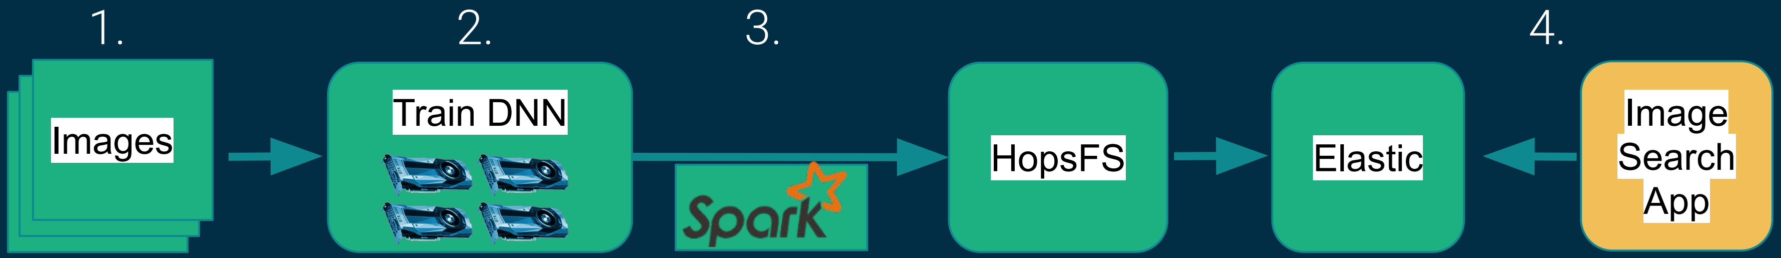

The Future is Cloud-Native...but what about the FS?

Kubernetes

Does it have to be S3?

What will the Cloud-Native Filesystem be?

A Brief History of Data

LOGICAL CLOCKS

MapReduce v0.01-alpha

IBM 082 Punch Card Sorter

Scan -> Sort -> Scan -> Sort ….

Not Fault Tolerant!

First DBMS’ and Filesystems were Disk Aware

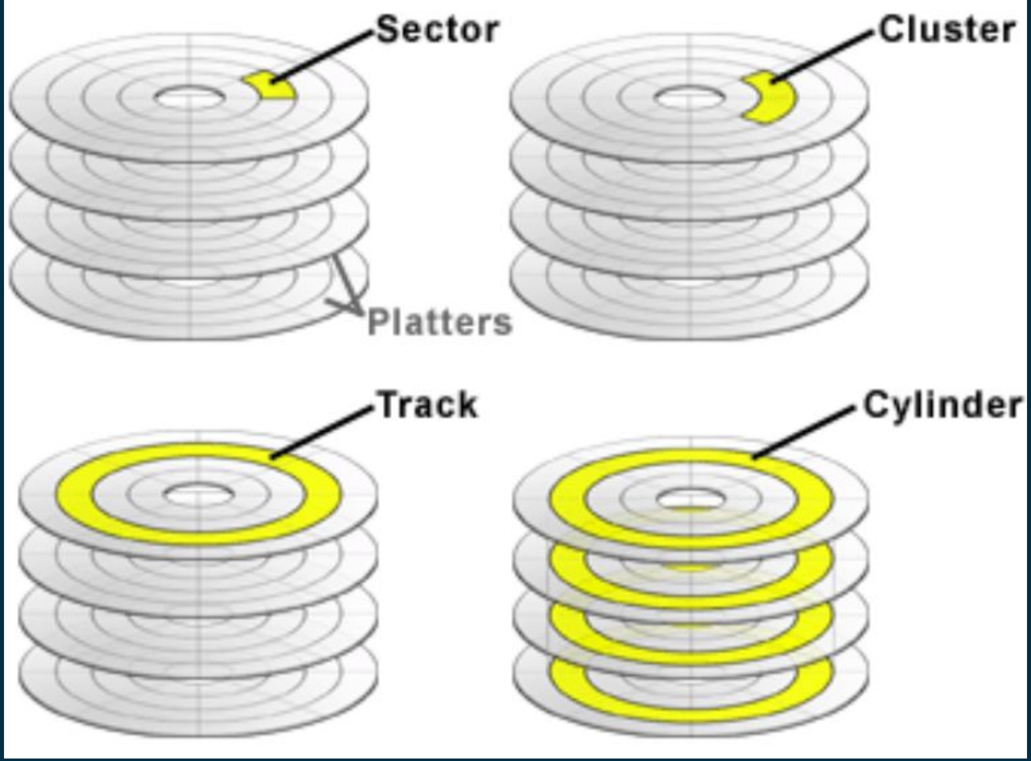

Early Filesystems’ block size was tightly coupled to the sector size of a disk

Hierarchical and Network DB Systems

You had to know what you want, and how to find it on disk.

Codd’s Relational Model and SystemR

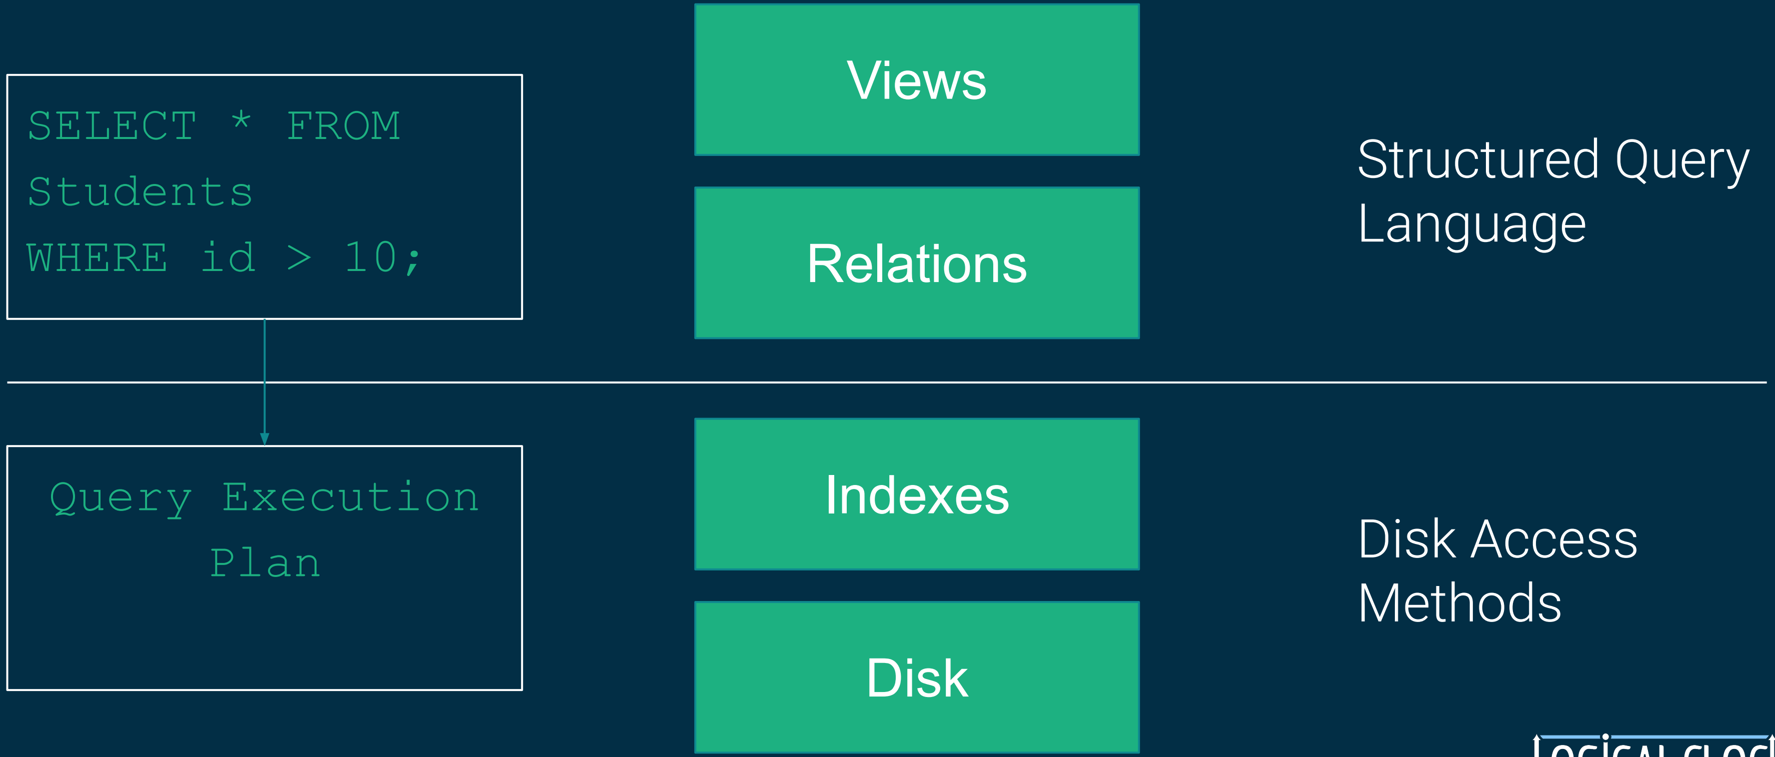

```
SELECT * FROM
Students
WHERE id > 10;
```

Query Execution Plan

Views

Relations

Structured Query Language

Indexes

Disk

Disk Access Methods

+30 years..Data Volumes outgrew Relational DBs

Data volumes got too large for single-server SQL DBs

And thus, the NoSQL movement was born...

.....only to be quickly out-evolved

Evolutionary History of SQL Datastores

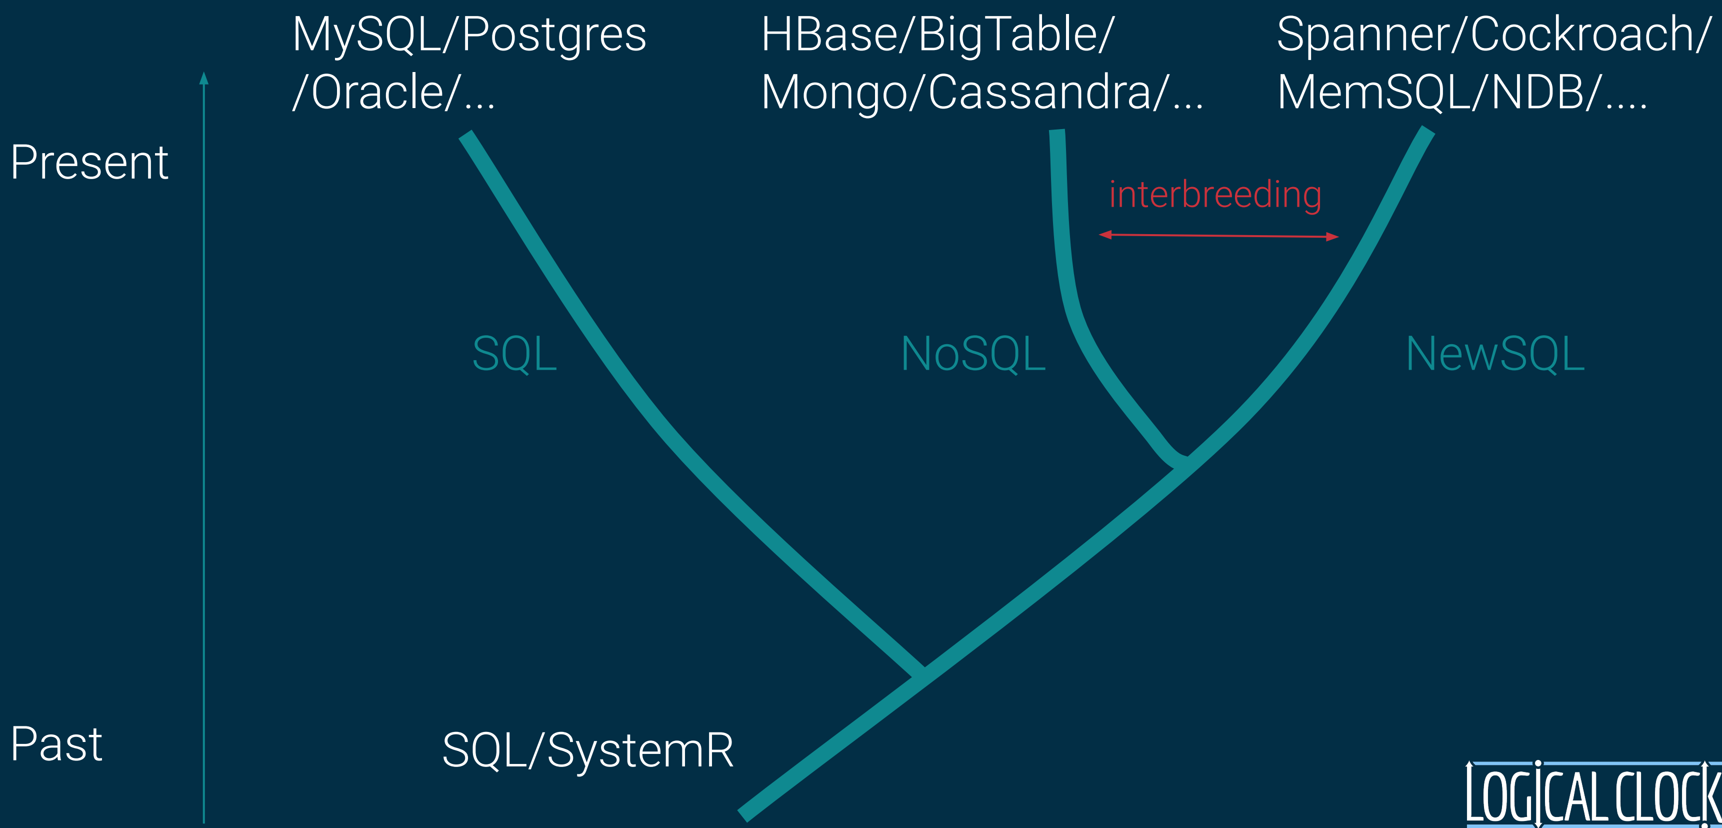

What about Filesystems?

Evolutionary History of Filesystems

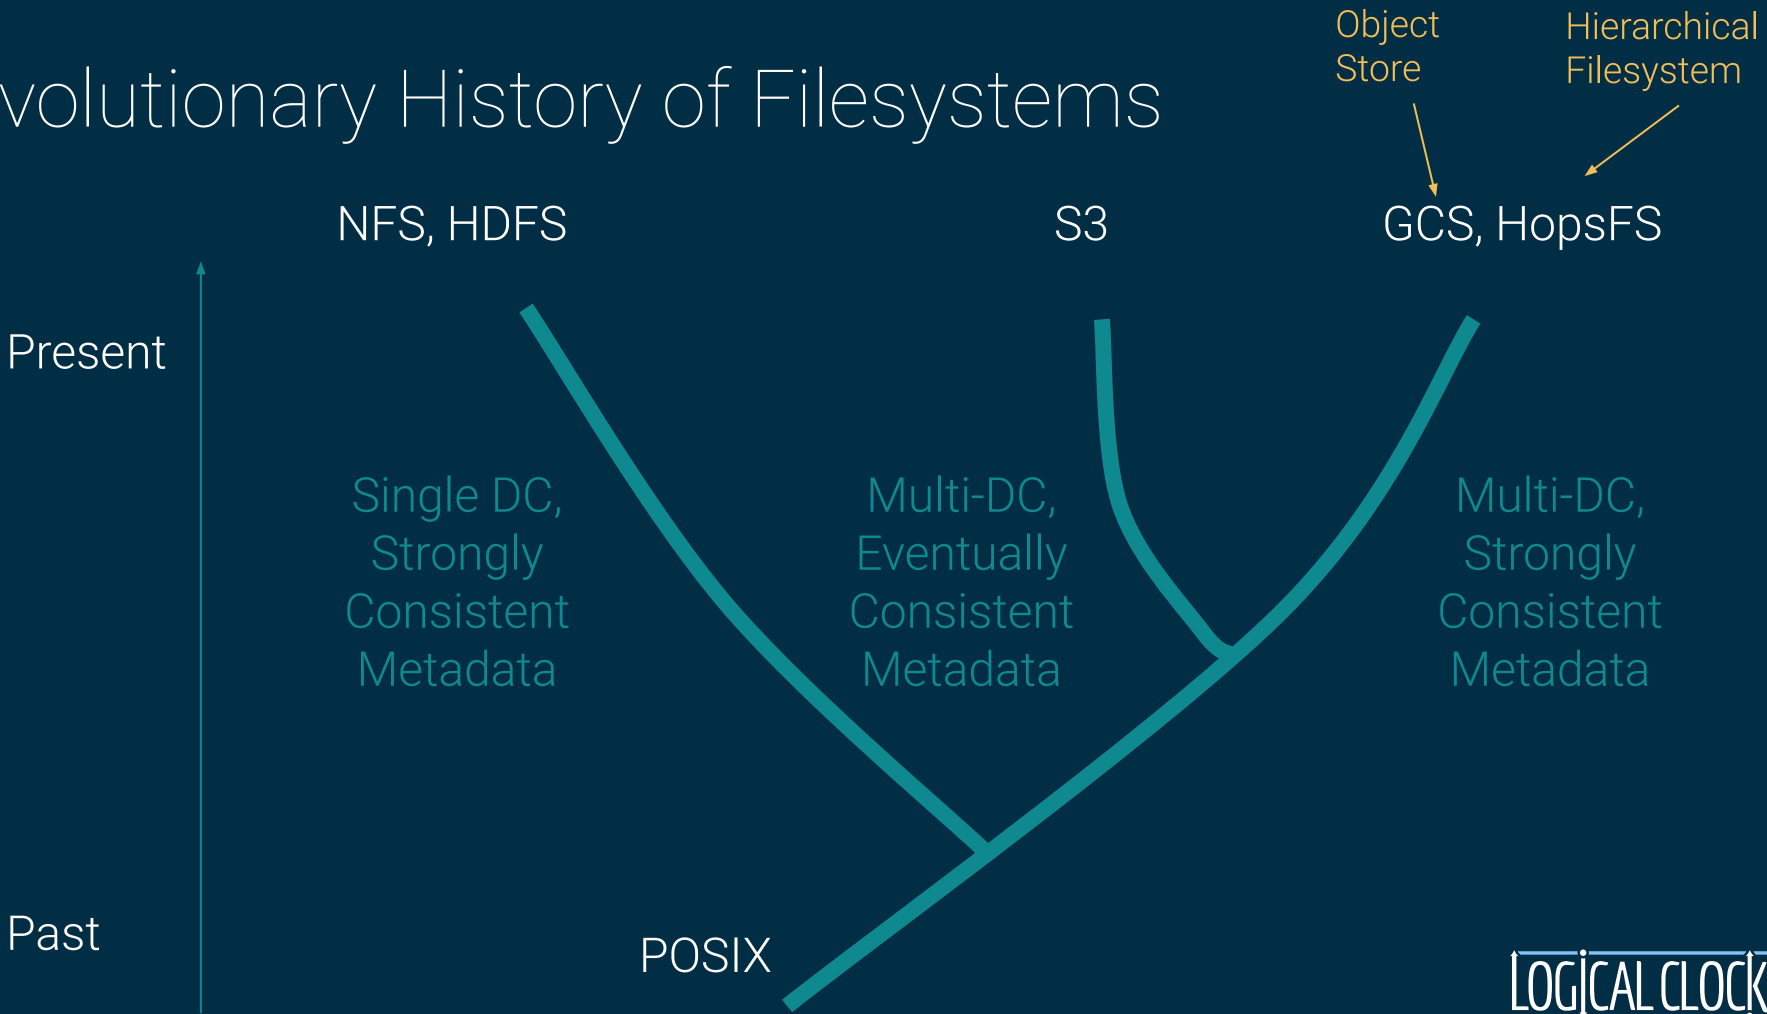

Object Store

Hierarchical Filesystem

NFS, HDFS

S3

GCS, HopsFS

Present

Single DC, Strongly Consistent Metadata

Multi-DC, Eventually Consistent Metadata

Multi-DC, Strongly Consistent Metadata

Past

POSIX

Why is Strongly Consistent Metadata important?

- POSIX-like semantics
  - Insert a file in a dir, and yes, it will be there!
- Atomic rename
  - Building block for scalable SQL systems
- Consistent change data capture (changelog)
  - Data provenance
  - Search/Index/tag the filesystem namespace

HopsFS uses NDB for Strongly Consistent Metadata

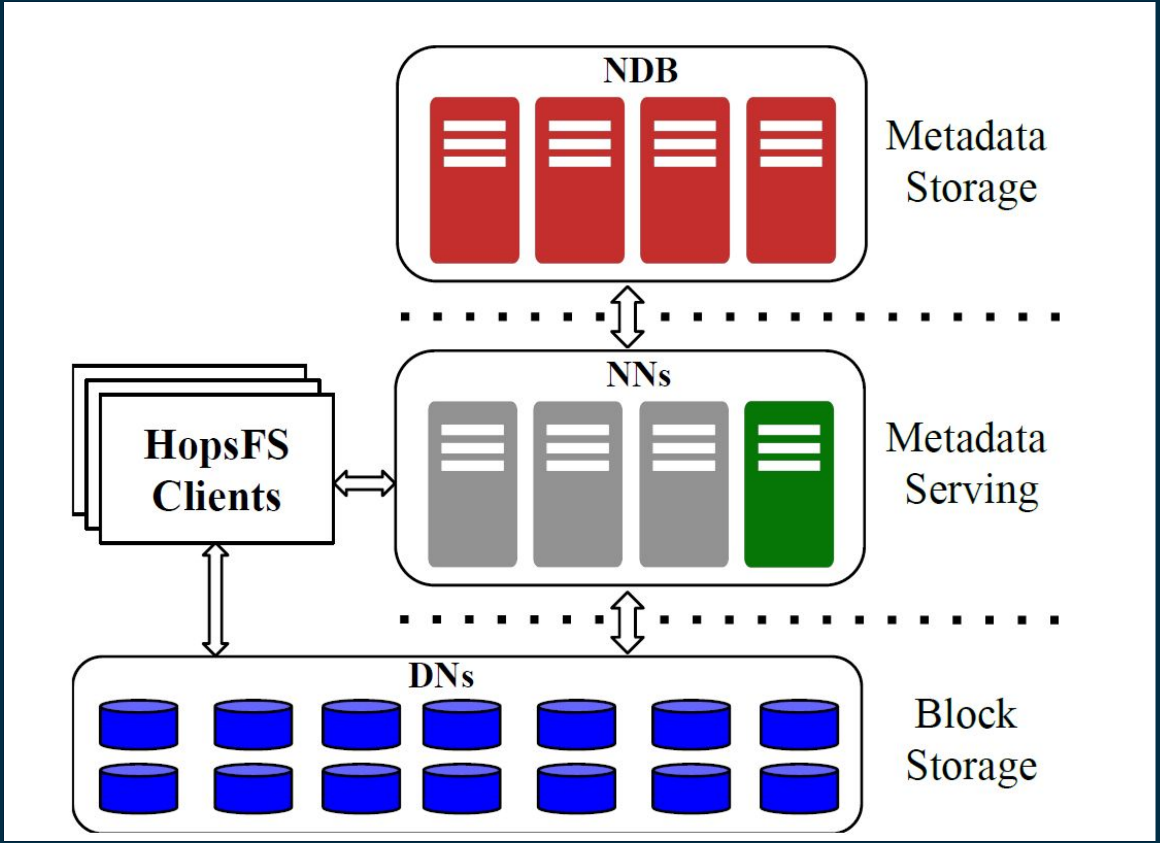

Make these Layers Data-Center HA

Multi-DC HopsFS affects every layer of the stack

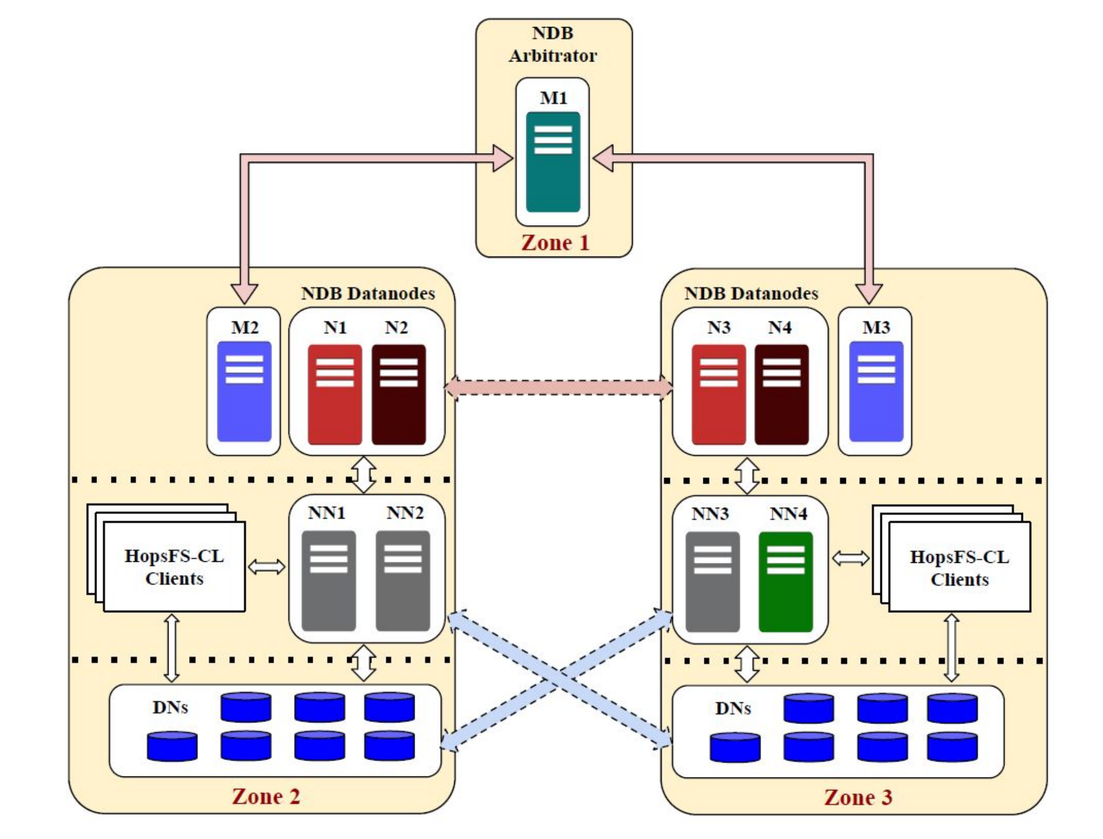

Database nodes DC-aware

Namenodes DC-aware

36% performance improvements by optimizing for DC local operations

Triple replication also possible with HopsFS

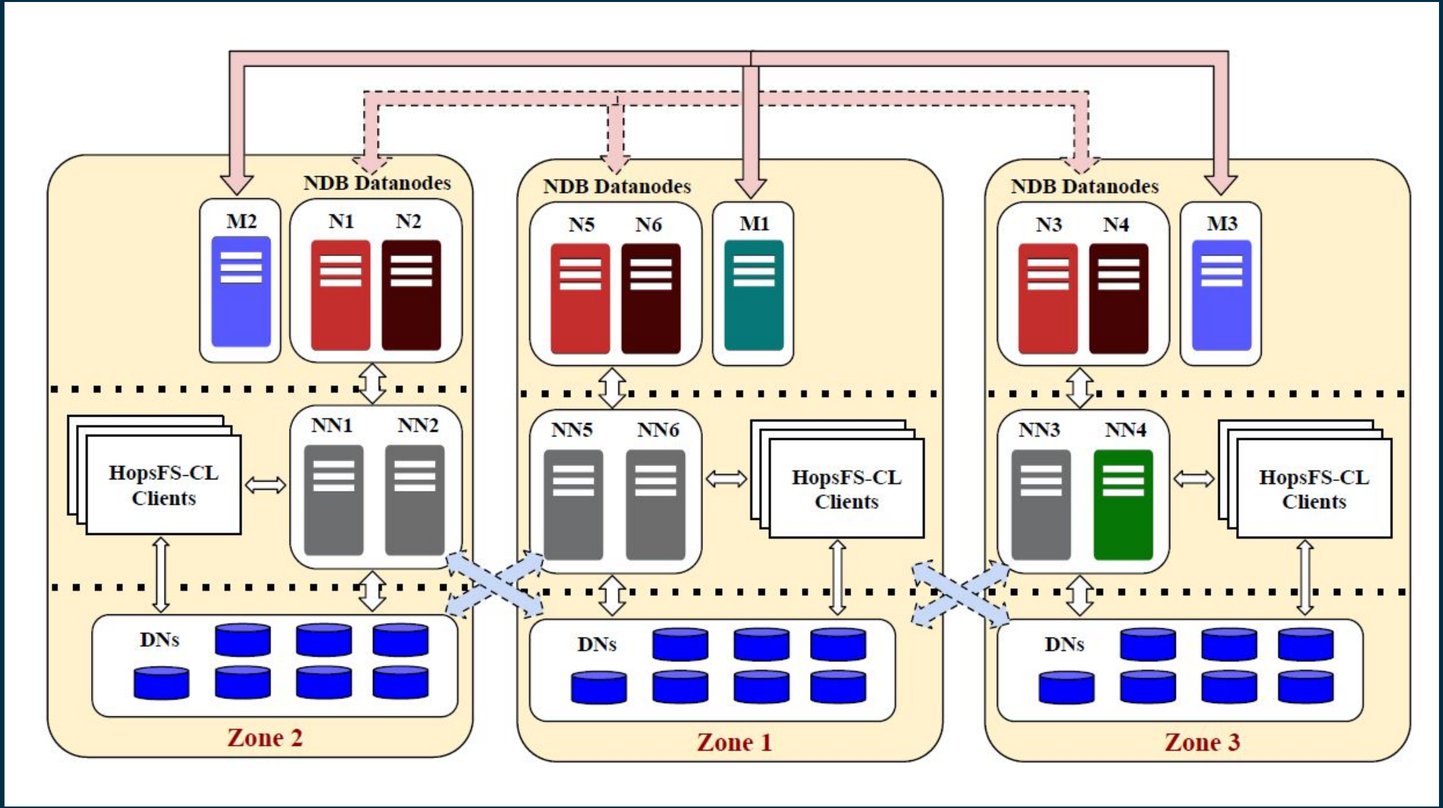

Change Data Capture for HopsFS with Epipe

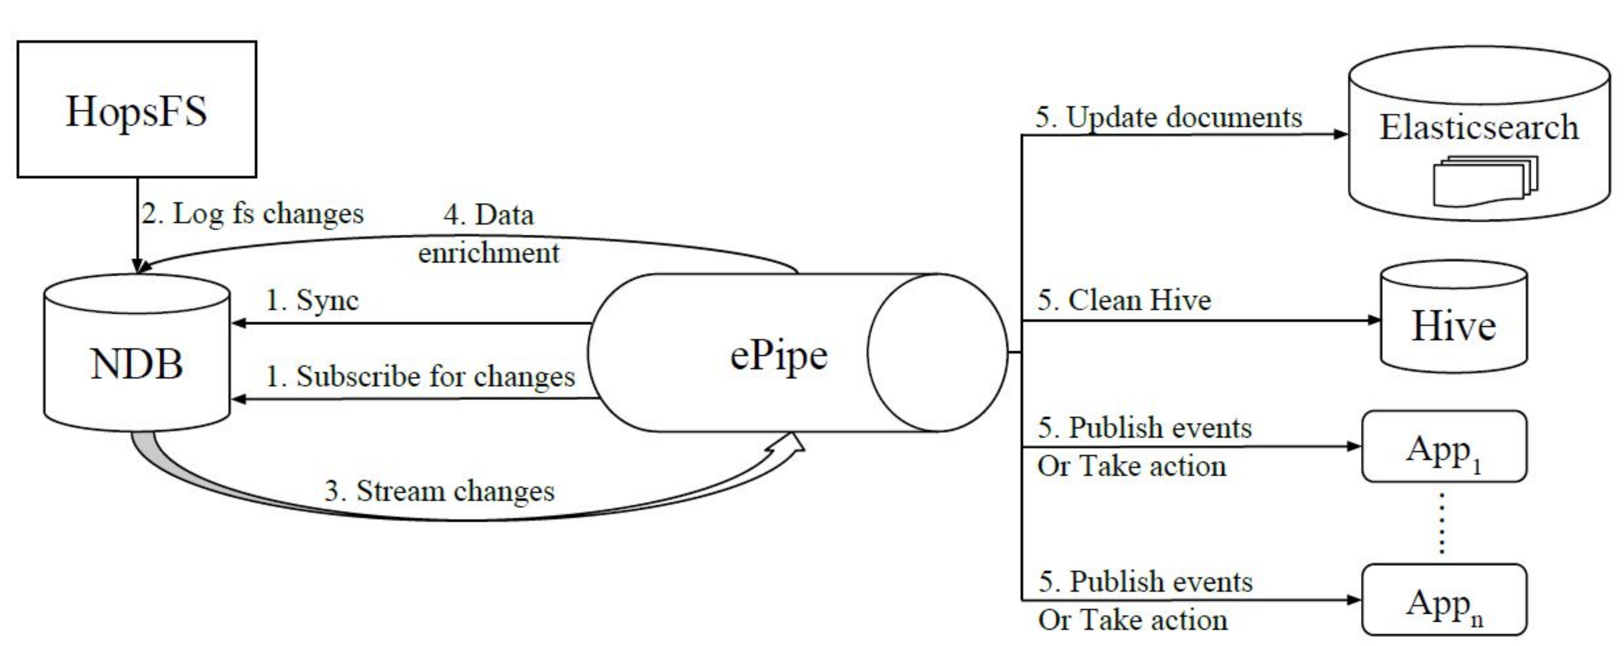

Overhead of running ePipe on the Spotify Hadoop workload: 4.77%

ePipe: Near Real-Time Polyglot Persistence of HopsFS Metadata, Ismail et al, CCGrid, 2019.

HOPS STRIKES BACK IN THE CLOUD

Availability:	Highly available across Data Centers (AZ)
Performance:	>1.6m Ops/Second on Spotify workload (GCE, 3 AZs) NVMe disks used to store small files in metadata layer
Security:	TLS-based security
HDFS API:	Native support in Spark, Flink, TensorFlow, etc.

Hopsworks - a platform for Data Intensive AI built on Hops

Hopsworks hides the Complexity of Deep Learning

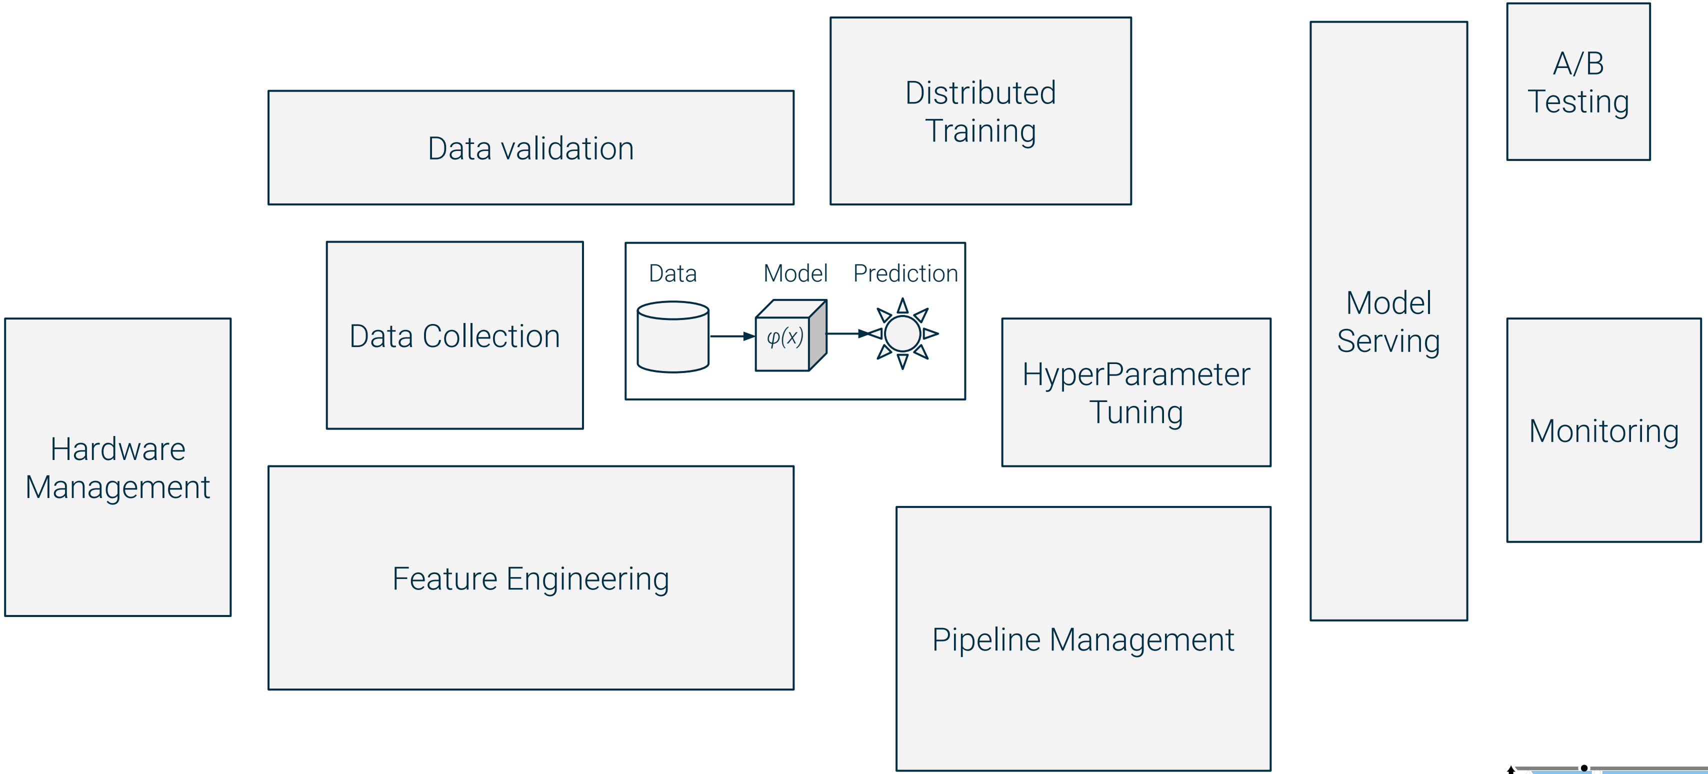

*Figure from "Technical Debt in Machine Learning Systems", Google research paper

Hopsworks hides the Complexity of Deep Learning

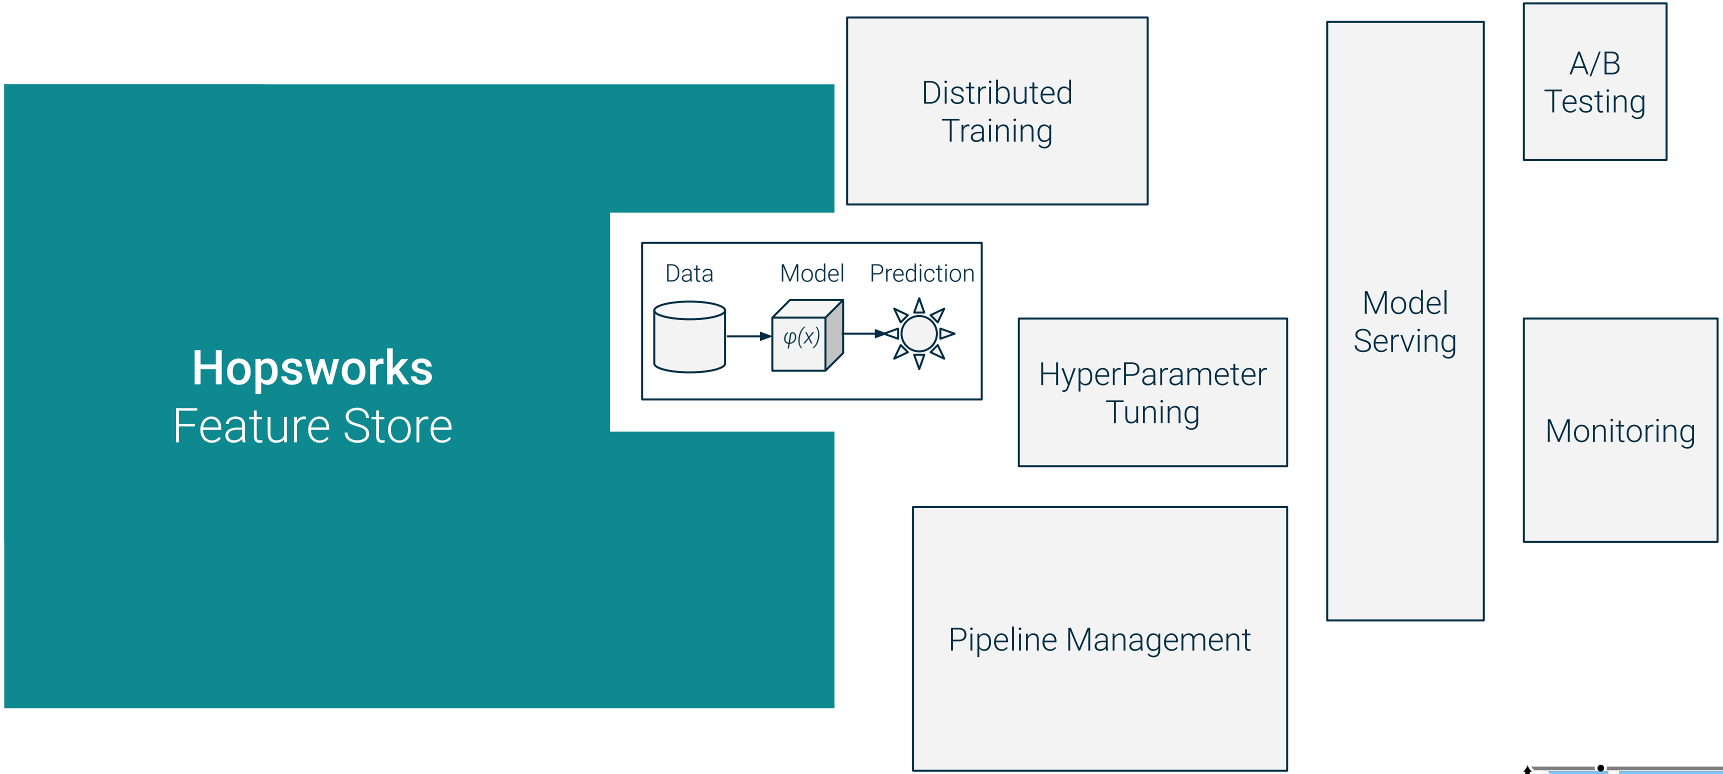

Hopsworks
Feature Store

Distributed Training

HyperParameter Tuning

Pipeline Management

Model Serving

A/B Testing

Monitoring

Hopsworks hides the Complexity of Deep Learning

Hopsworks
Feature Store

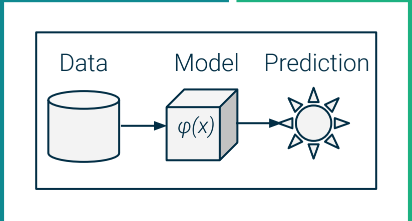

Hopsworks
REST API

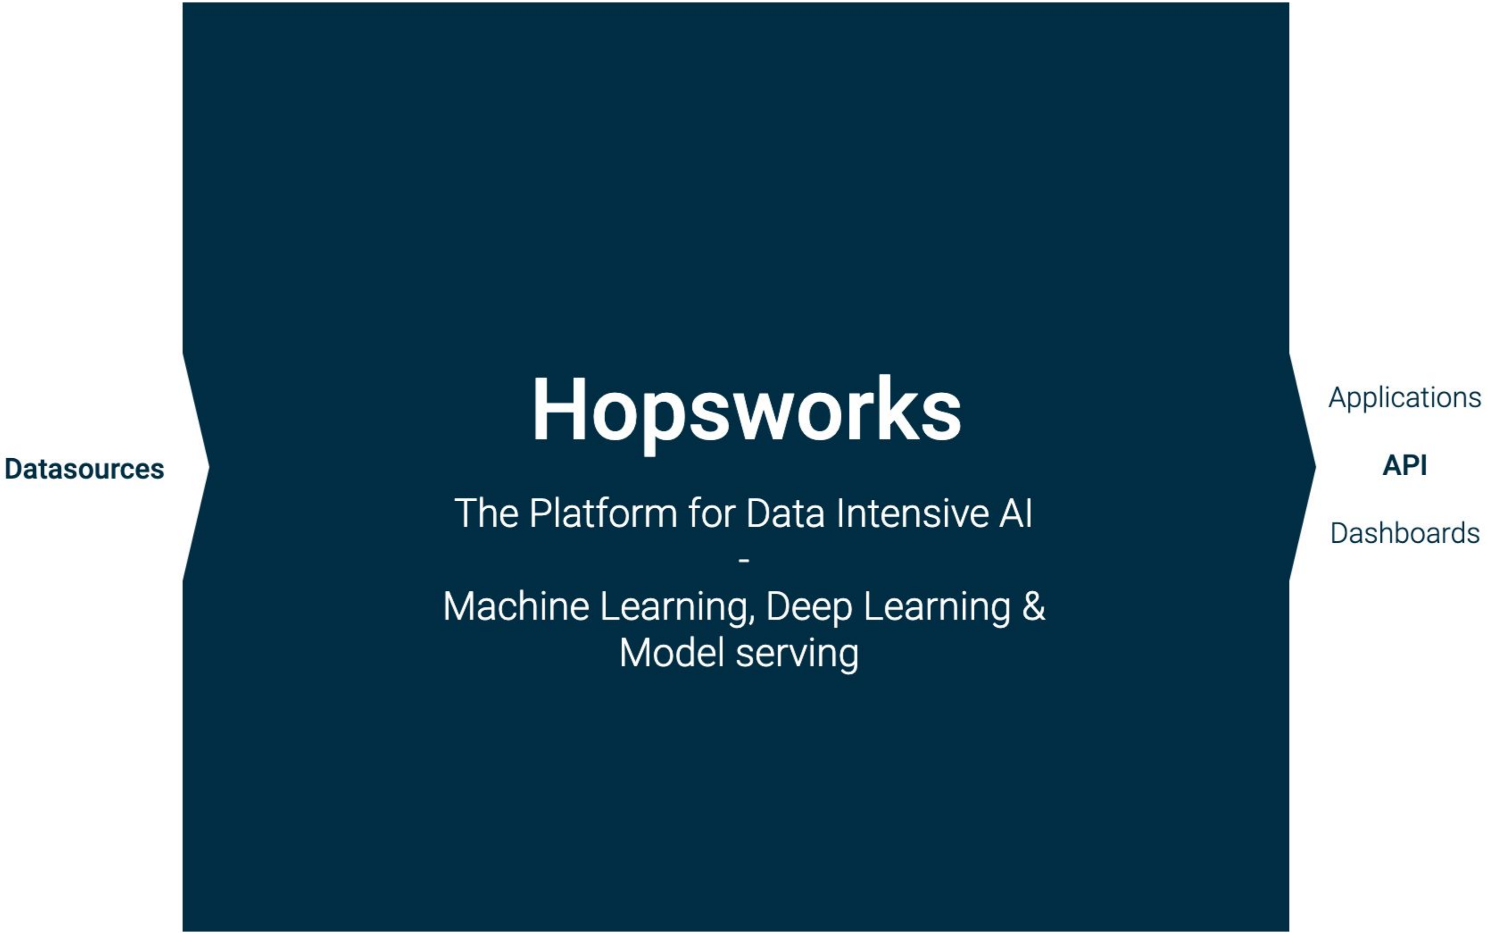

Datasources

Hopsworks

The Platform for Data Intensive AI

-

Machine Learning, Deep Learning & Model serving

Applications

API

Dashboards

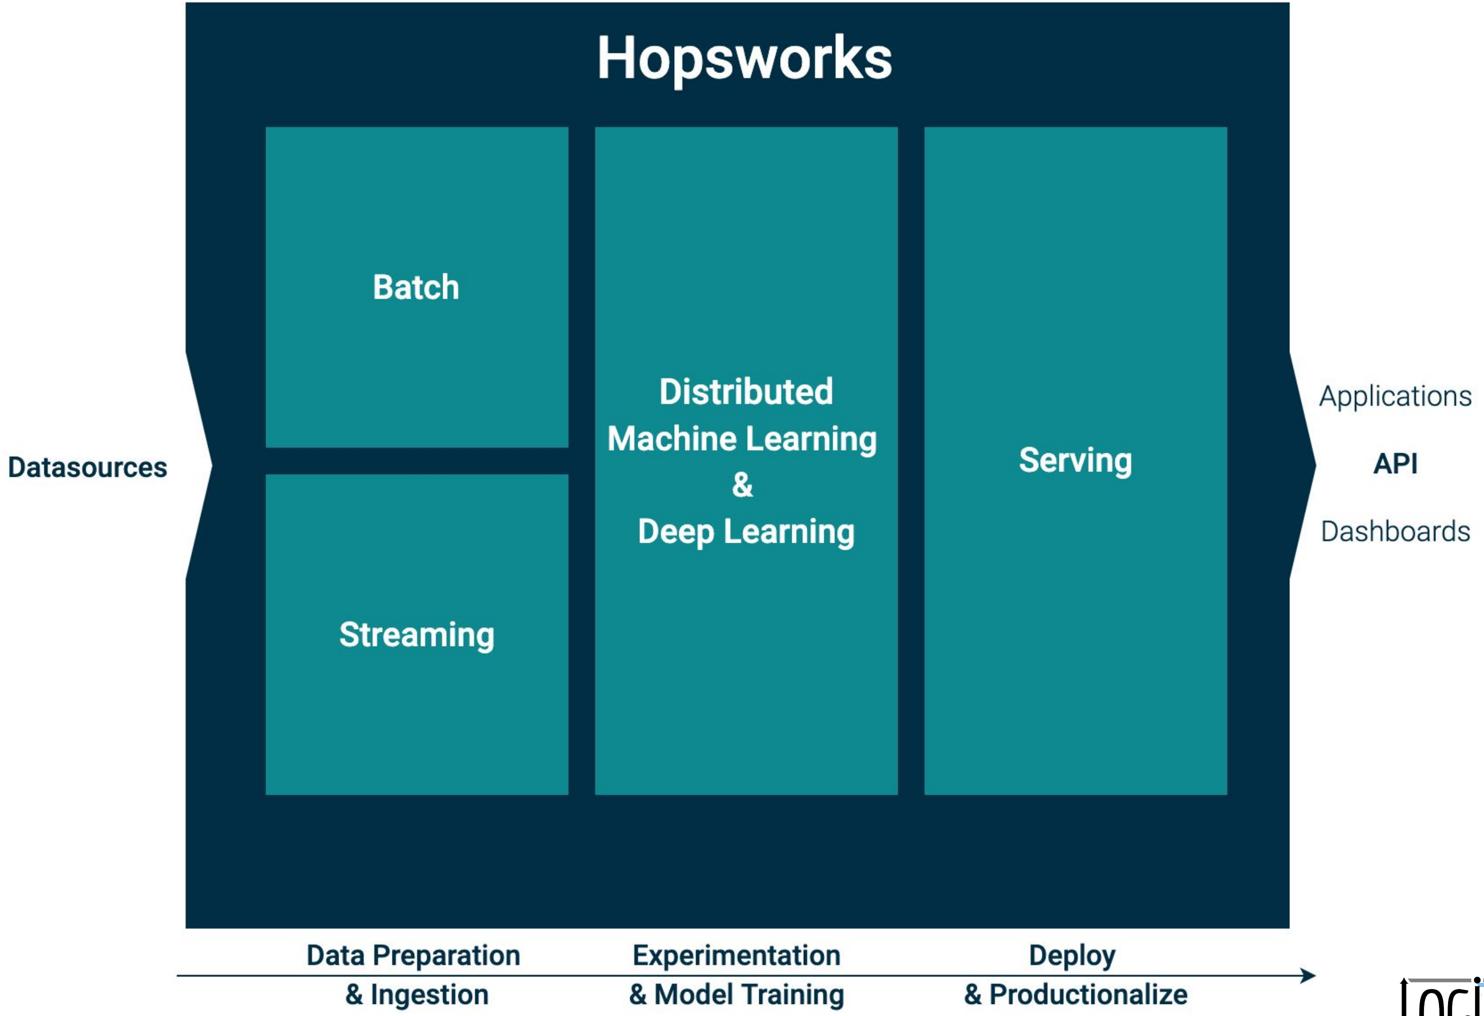

Hopsworks

Datasources

Batch

Streaming

Distributed Machine Learning & Deep Learning

Serving

Applications

API

Dashboards

Data Preparation & Ingestion

Experimentation & Model Training

Deploy & Productionalize

LOGICAL CLOCKS

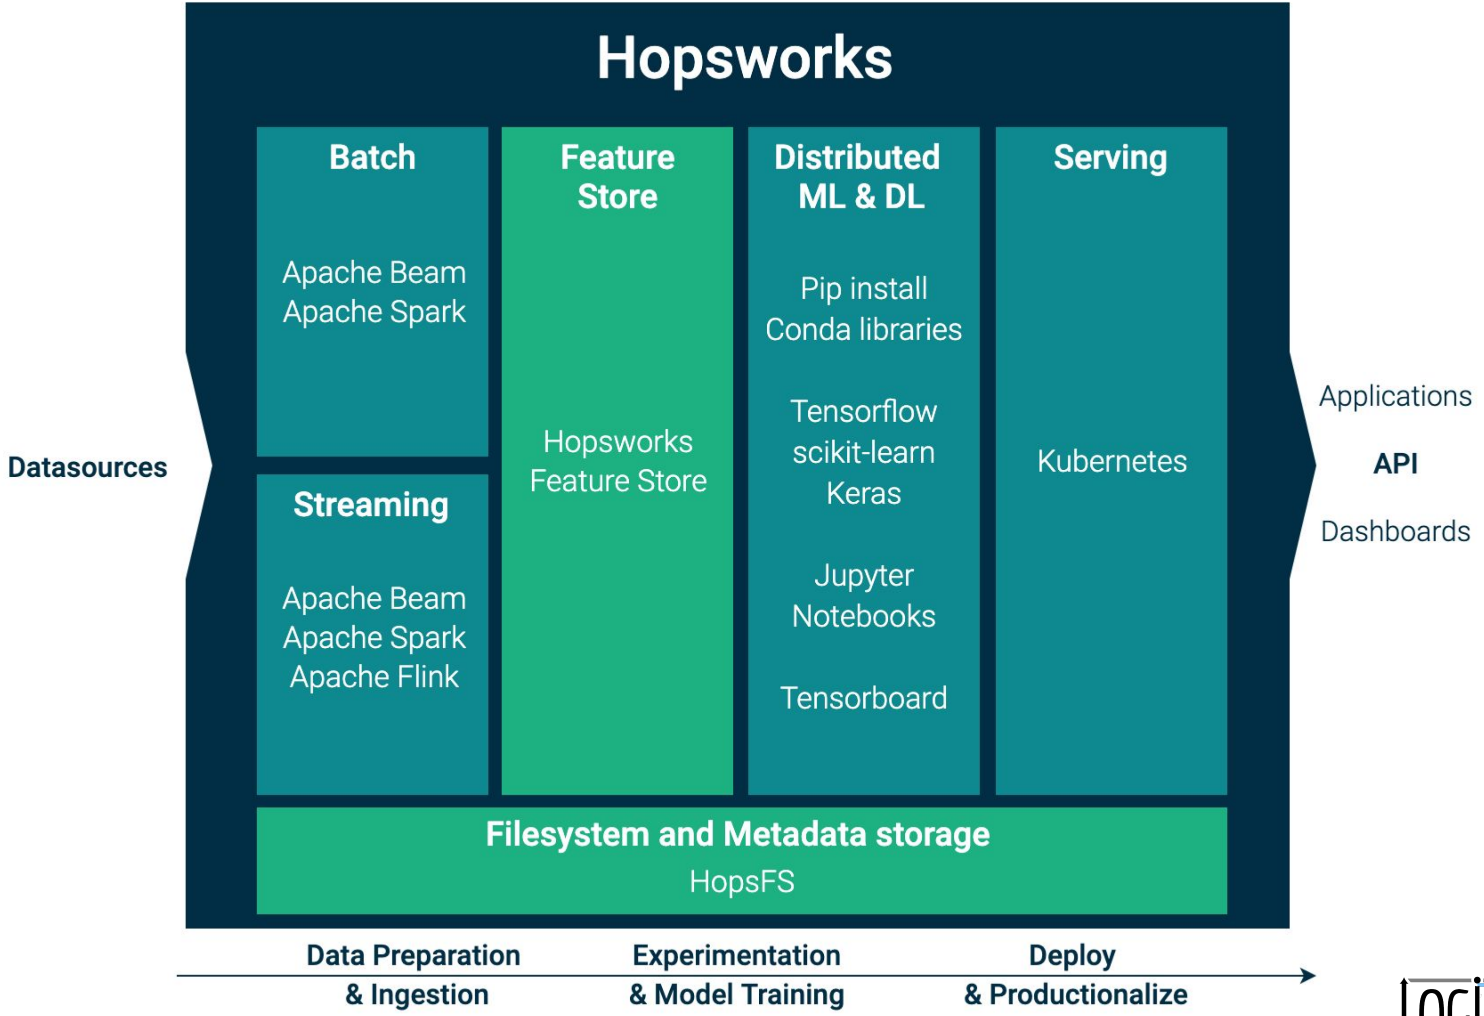

Hopsworks

Datasources

Batch

Apache Beam
Apache Spark

Streaming

Apache Beam
Apache Spark
Apache Flink

Feature Store

Hopsworks Feature Store

Distributed ML & DL

Pip install
Conda libraries

Tensorflow
scikit-learn
Keras

Jupyter
Notebooks

Tensorboard

Serving

Kubernetes

Filesystem and Metadata storage

HopsFS

Applications

API

Dashboards

Data Preparation & Ingestion

Experimentation & Model Training

Deploy & Productionalize

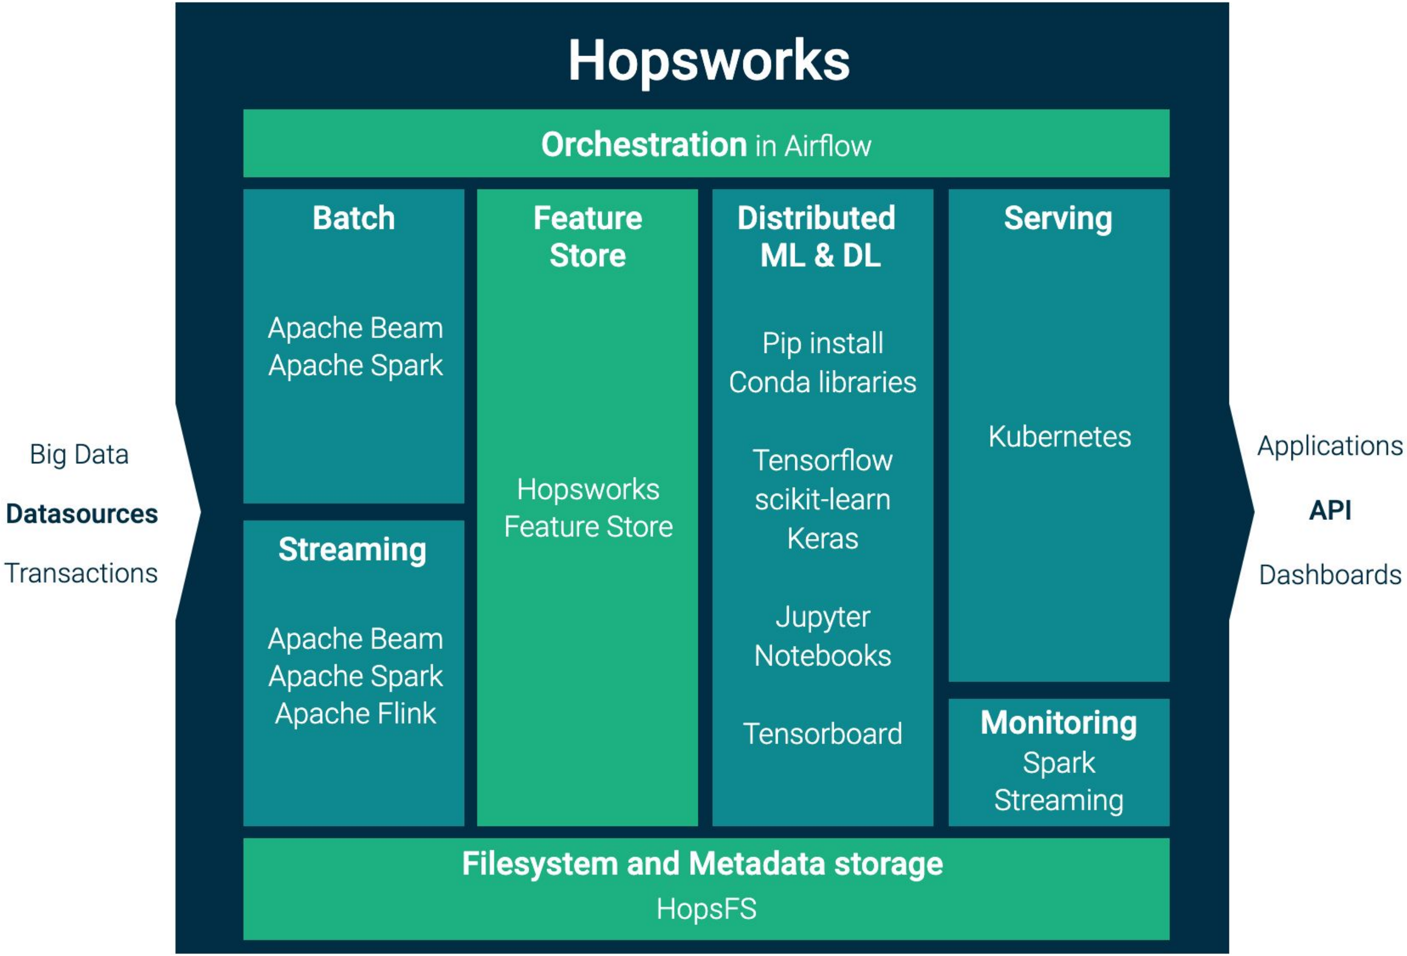

Hopsworks

Orchestration in Airflow

Batch

Apache Beam
Apache Spark

Streaming

Apache Beam
Apache Spark
Apache Flink

Feature Store

Hopsworks Feature Store

Distributed ML & DL

Pip install
Conda libraries

Tensorflow
scikit-learn
Keras

Jupyter
Notebooks

Tensorboard

Serving

Kubernetes

Monitoring

Spark Streaming

Filesystem and Metadata storage

HopsFS

Big Data
Datasources
Transactions

Applications
API
Dashboards

Encrypt everything

TSL/SSL encrypted calls between services with X.509 certificates

Secure Collaboration

Multi-Tenancy with Project-based collaboration and resource management

Data Lake Support

Integrates with your existing Data Lake or acts as your Data Lake

LOGICAL CLOCKS

What is Hopsworks?

Efficiency & Performance

Feature Store
Data warehouse for ML

Distributed Deep Learning
Faster with more GPUs

HopsFS
NVMe speed with Big Data

Horizontally Scalable
Ingestion, DataPrep,
Training, Serving

Usability & Process

Jupyter/Python Development
Notebooks in pipelines

Version Everything
Code, Infrastructure, Data

Model Serving on Kubernetes
TF Serving, MLeap, SkLearn

End-to-End ML Pipelines
Orchestrated by Airflow

Security & Governance

Secure Multi-Tenancy
Project-based restricted access

Encryption At-Rest, In-Motion
TLS/SSL everywhere

AI-Asset Governance
Models, experiments, data, GPUs

Data/Model/Feature Lineage
Discover/track dependencies

Which services require Distributed Metadata (HopsFS)?

Efficiency & Performance

Feature Store
Data warehouse for ML

Distributed Deep Learning
Faster with more GPUs

HopsFS
NVMe speed with Big Data

Horizontally Scalable
Ingestion, DataPrep,
Training, Serving

Usability & Process

Jupyter/Python Development
Notebooks in pipelines

Version Everything
Code, Infrastructure, Data

Model Serving on Kubernetes
TF Serving, MLeap, SkLearn

End-to-End ML Pipelines
Orchestrated by Airflow

Security & Governance

Secure Multi-Tenancy
Project-based restricted access

Encryption At-Rest, In-Motion
TLS/SSL everywhere

AI-Asset Governance
Models, experiments, data, GPUs

Data/Model/Feature Lineage
Discover/track dependencies

End-to-End ML Pipelines in Hopsworks

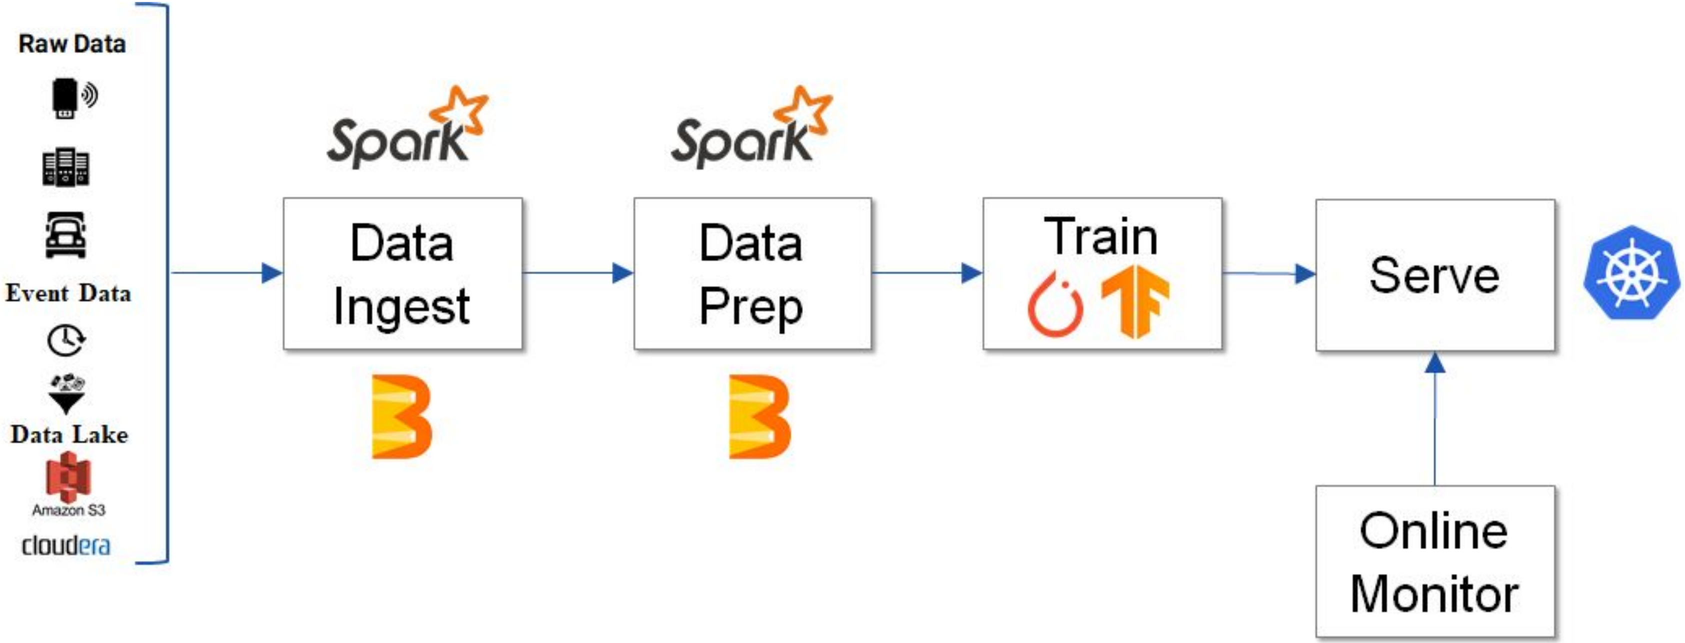

End-to-End Pipelines can be factored into stages

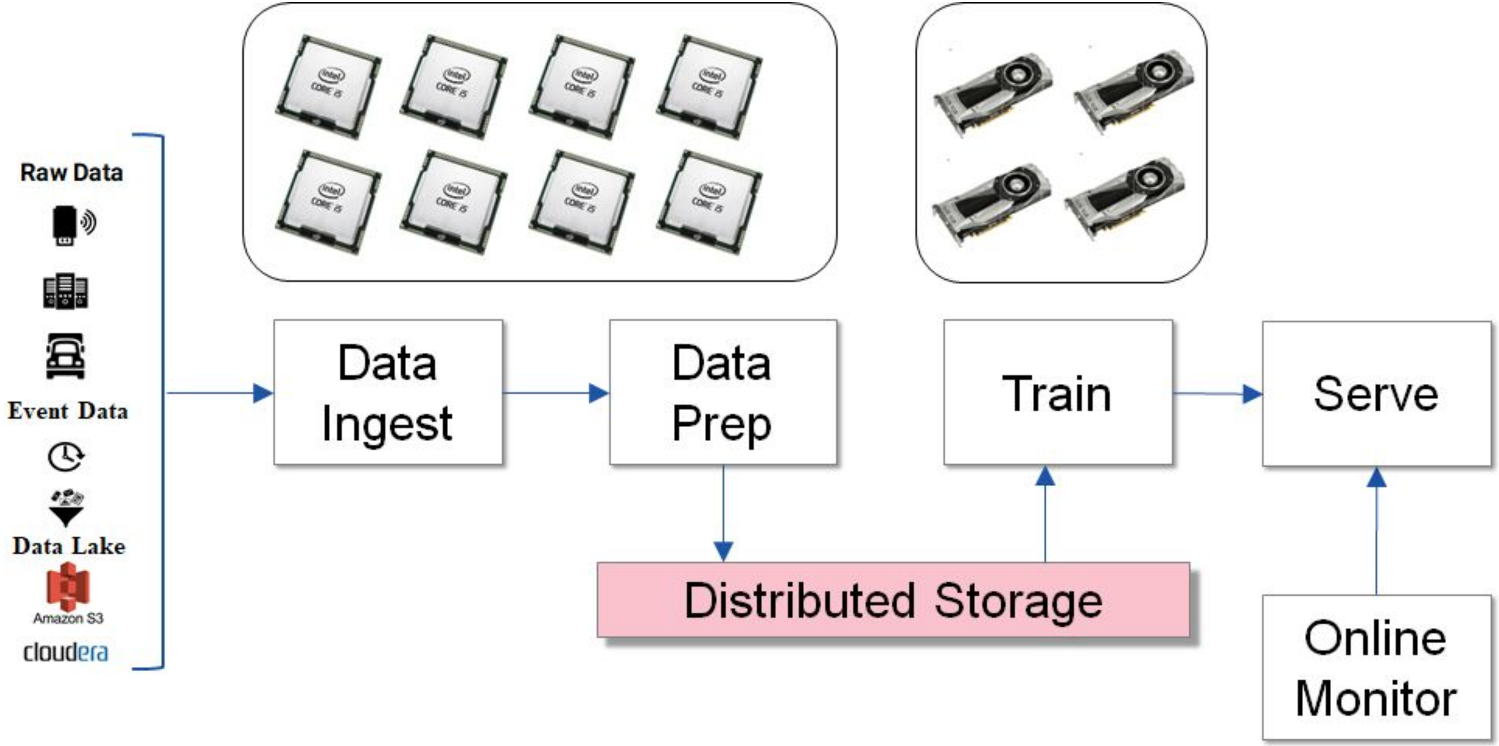

Typical Feature Store Pipelines

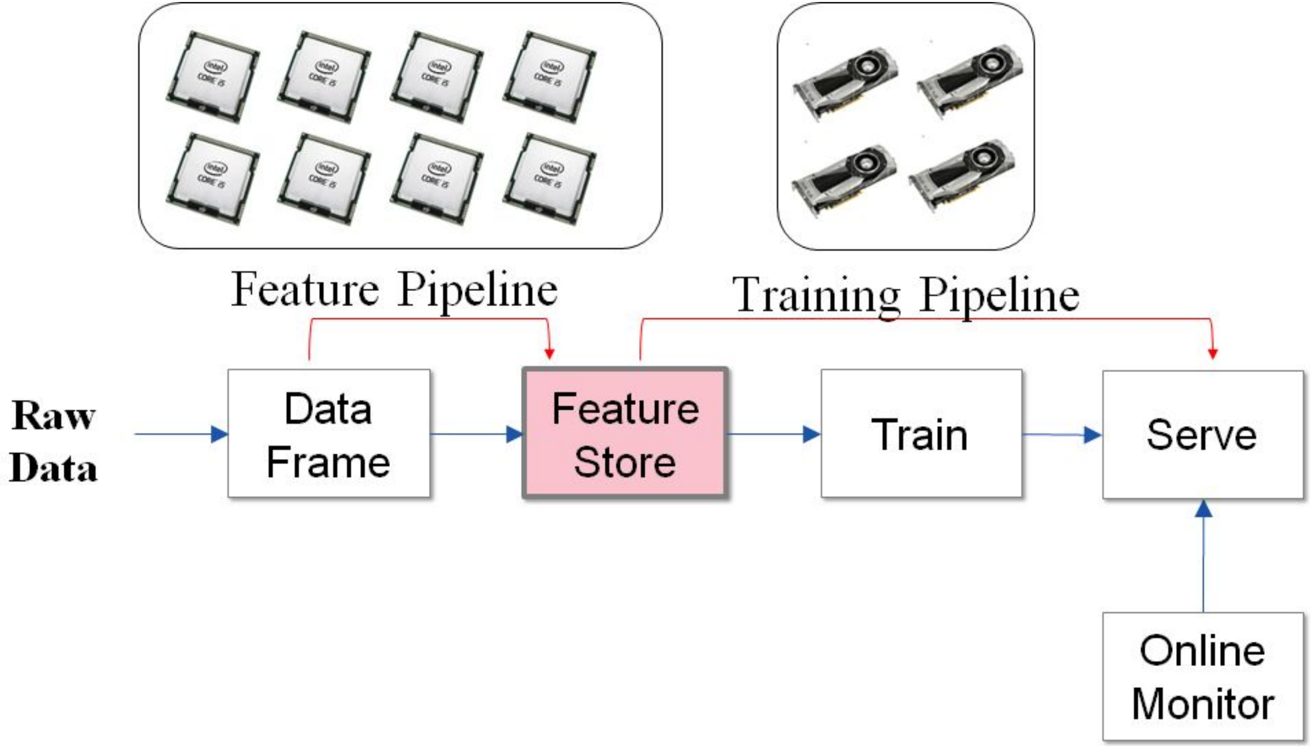

Hopsworks’ Feature Store

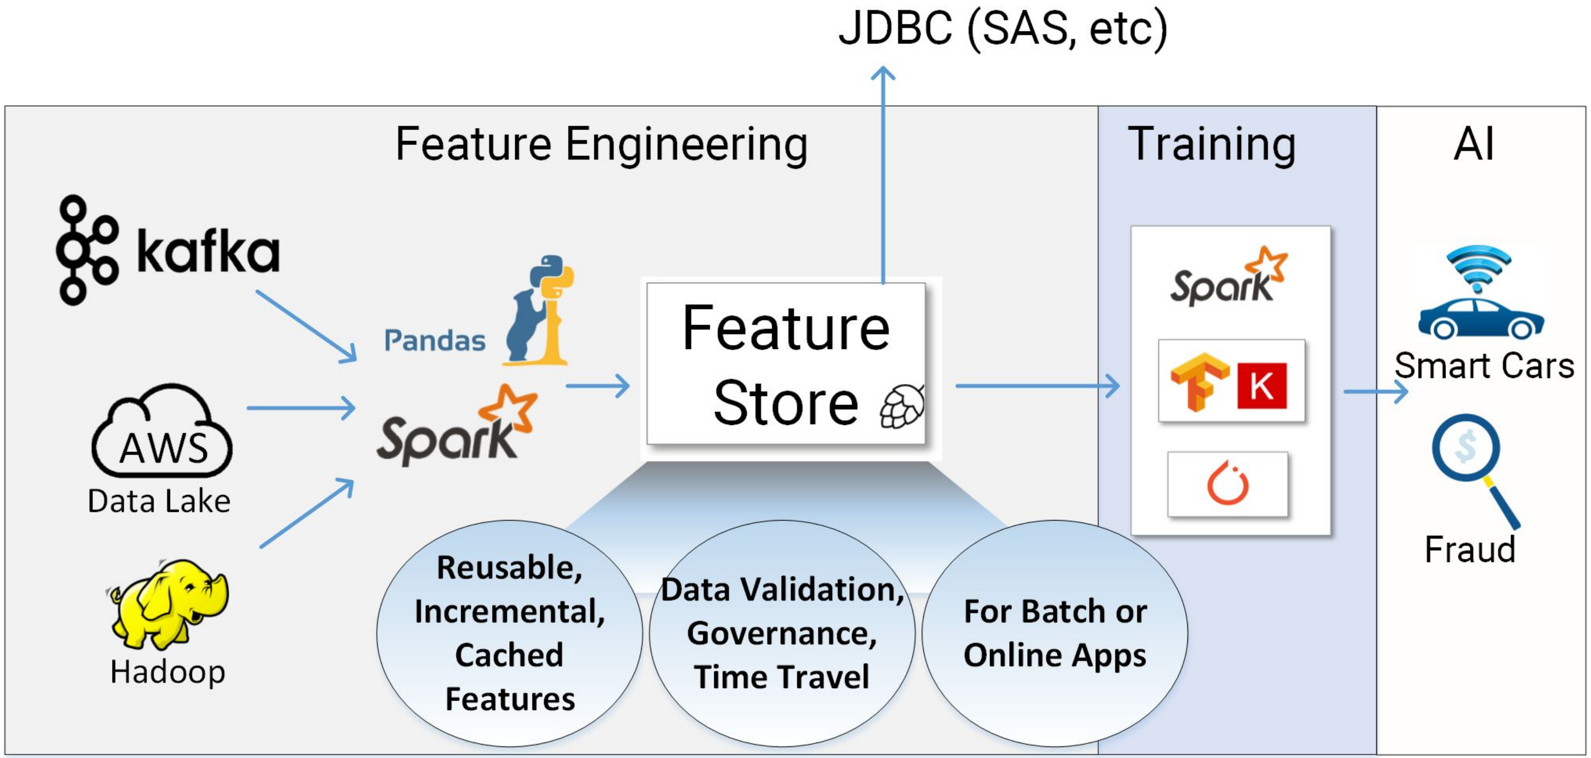

Dev View: Pipelines of Jupyter Notebooks in Airflow

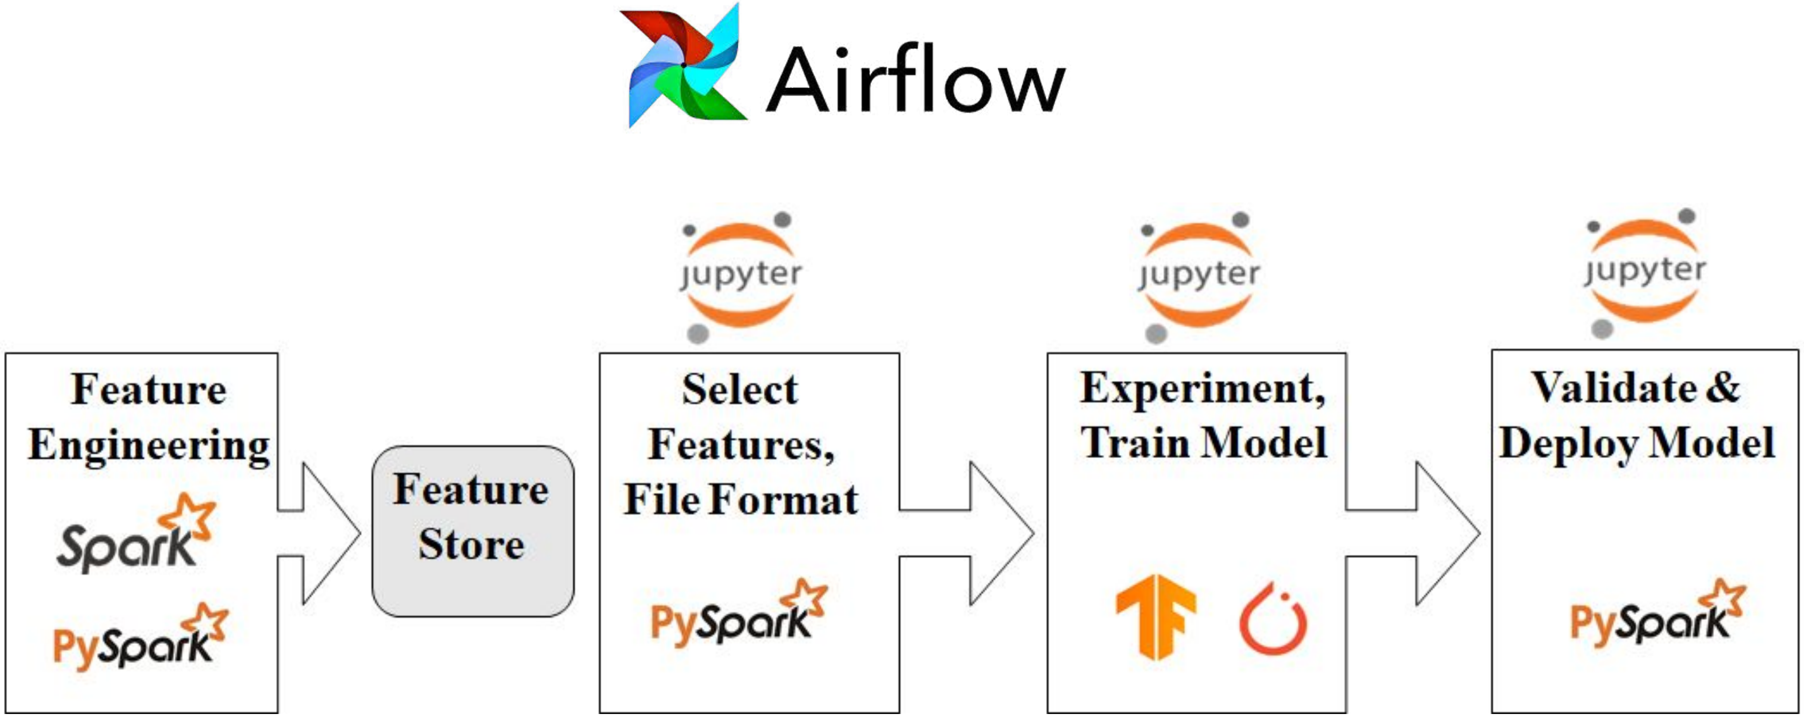

How to get started with Hopsworks?

Register for a free account at: www.hops.site
Images available for AWS, GCE, Virtualbox.

We need your support. Star us, tweet about us!

https://github.com/logicalclocks/hopsworks

@hopsworks

LOGICAL CLOCKS