

HOPSWORKS

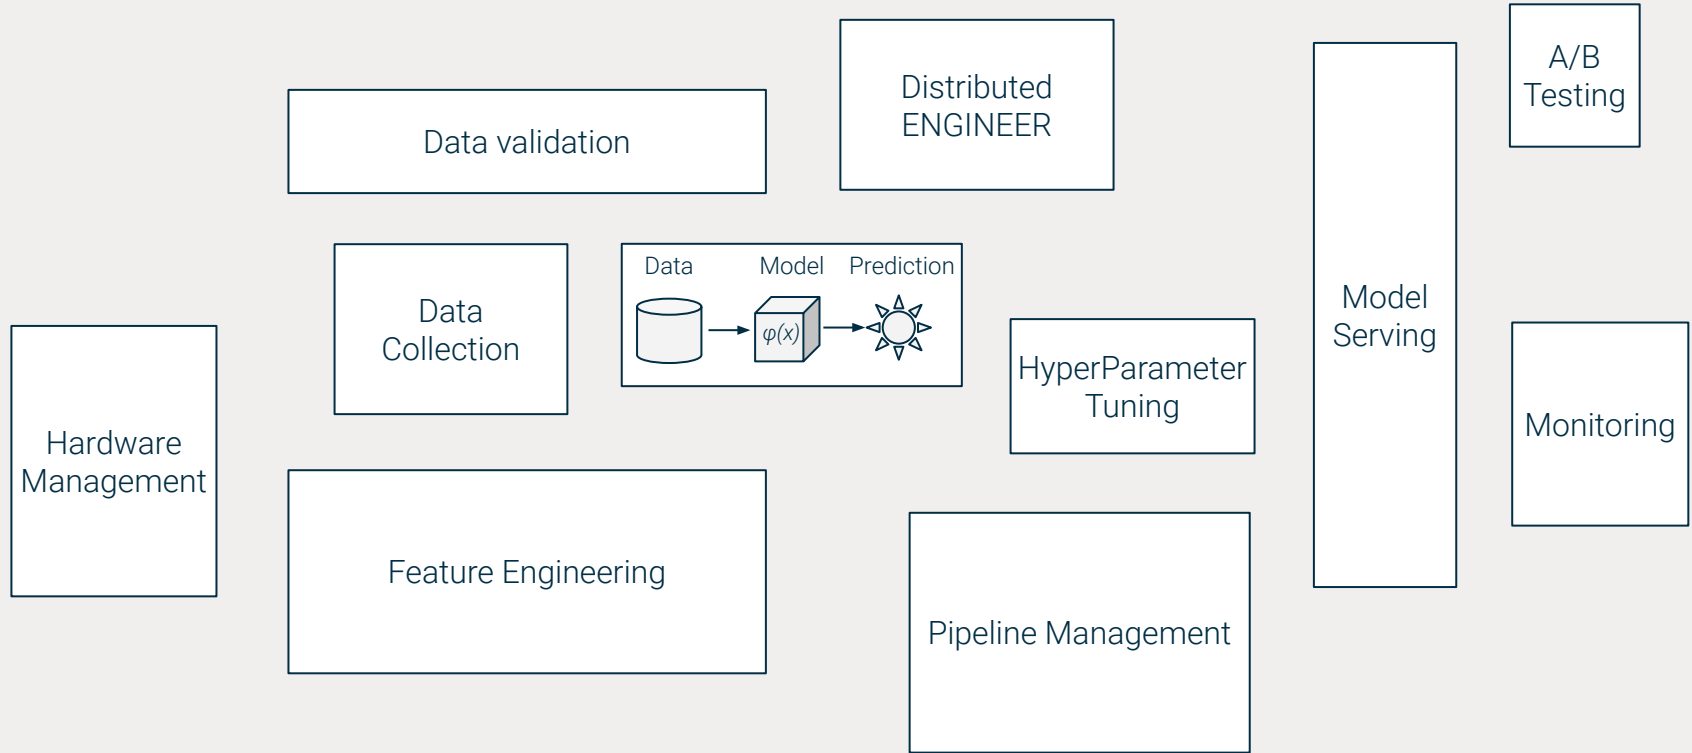
KTH Lecture for ID2223

Dr. Jim Dowling^{1,2}

Slides together with Alexandru A. Ormenisan^{1,2}, Mahmoud Ismail^{1,2}




Growing Consensus on how to manage complexity of AI



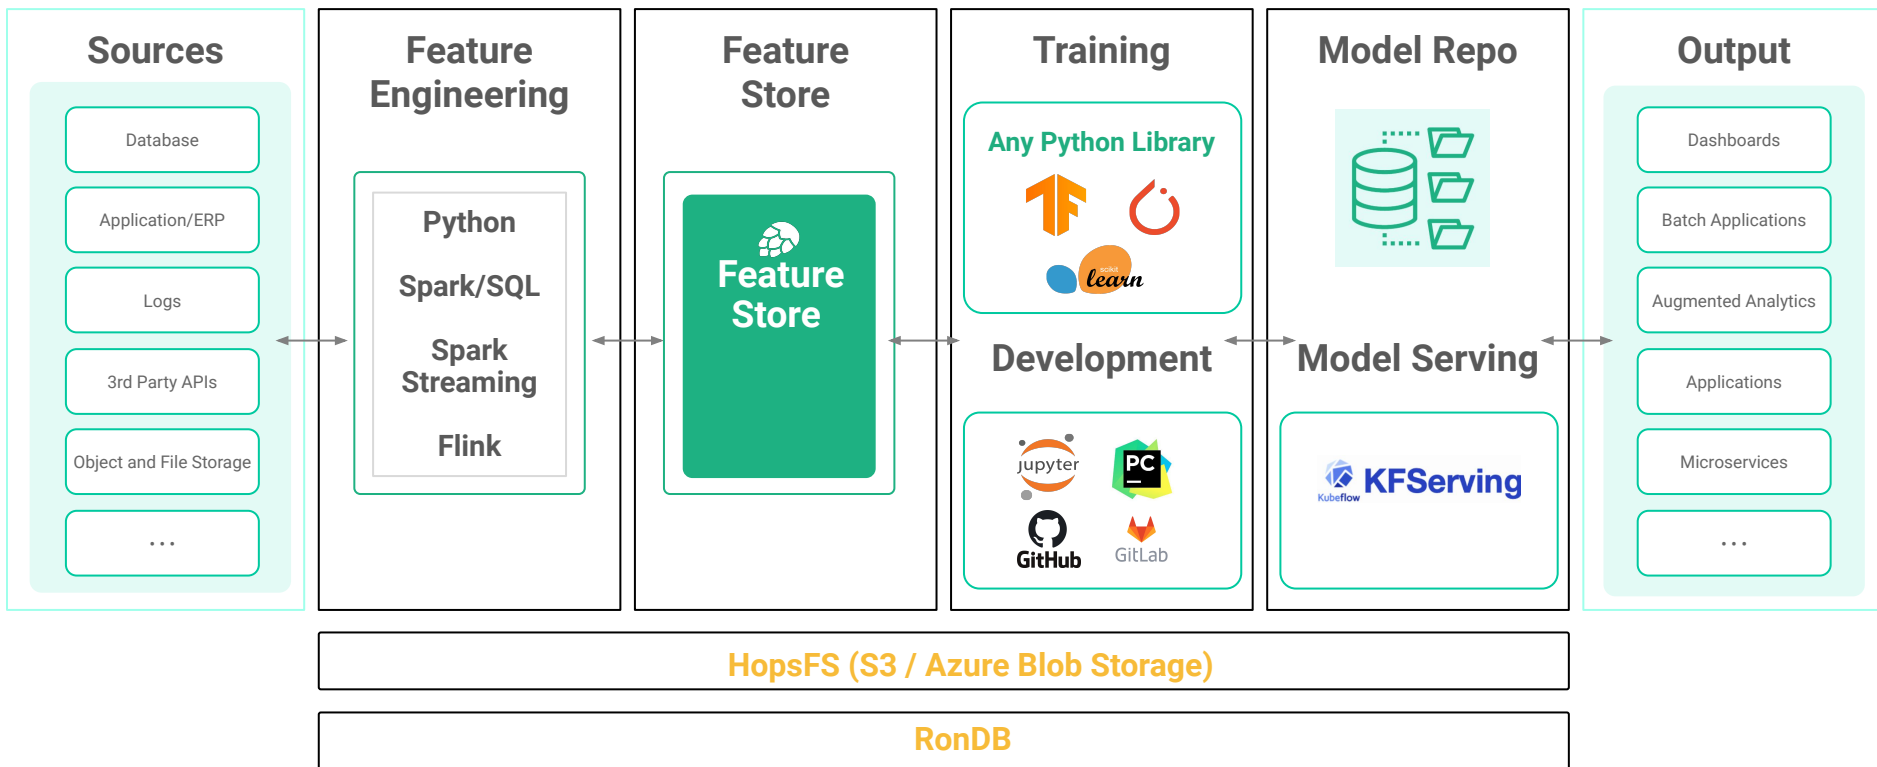


**FEATURE
STORE**



ML PLATFORM
TRAIN and SERVE

Hopworks - Design and Operate AI Applications at Scale



Hopsworks is an Open, Modular Feature Store



Teams use the tools of their choice,
integrated with the
Hopsworks Feature Store

Data Science



Azure Machine Learning



Amazon SageMaker



Google AI



databricks



DOMINO



dataiku

Model Serving



Azure Machine Learning



Amazon SageMaker



Google AI



databricks



Kubeflow



SELDON



HOPSWORKS Feature Store

Data Engineering



python™



Spark



Apache Airflow



Microsoft SQL Server



mongoDB.



snowflake®

Compliance & Regulatory



Informatica



collibra™

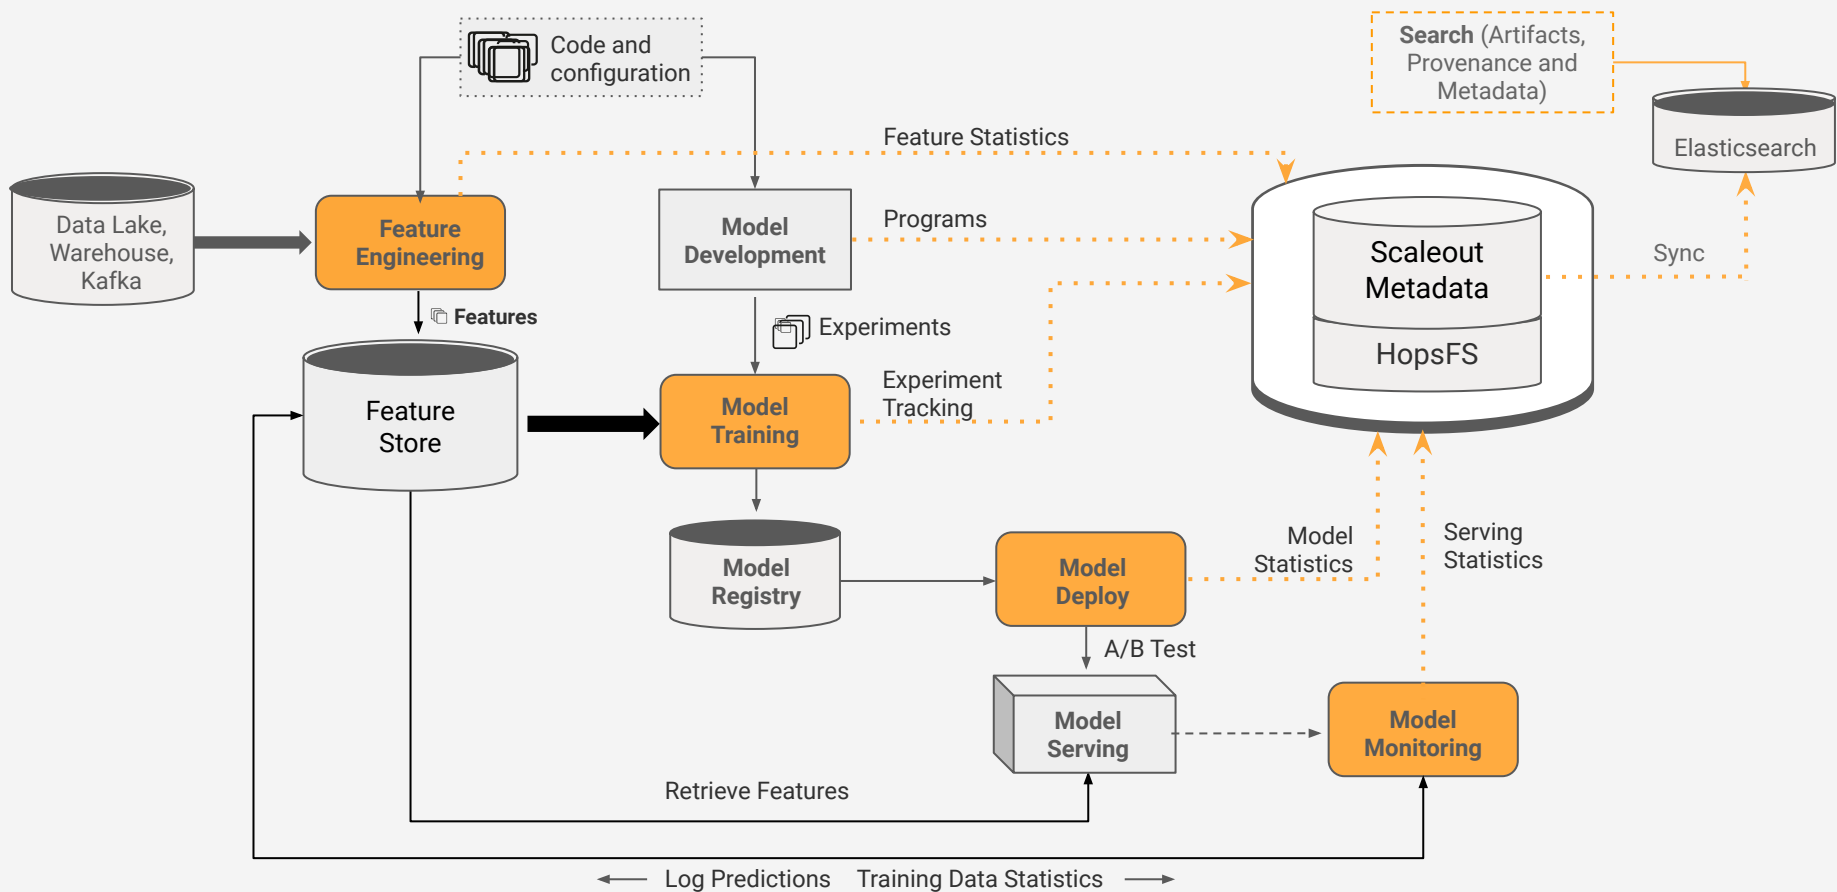


Alation

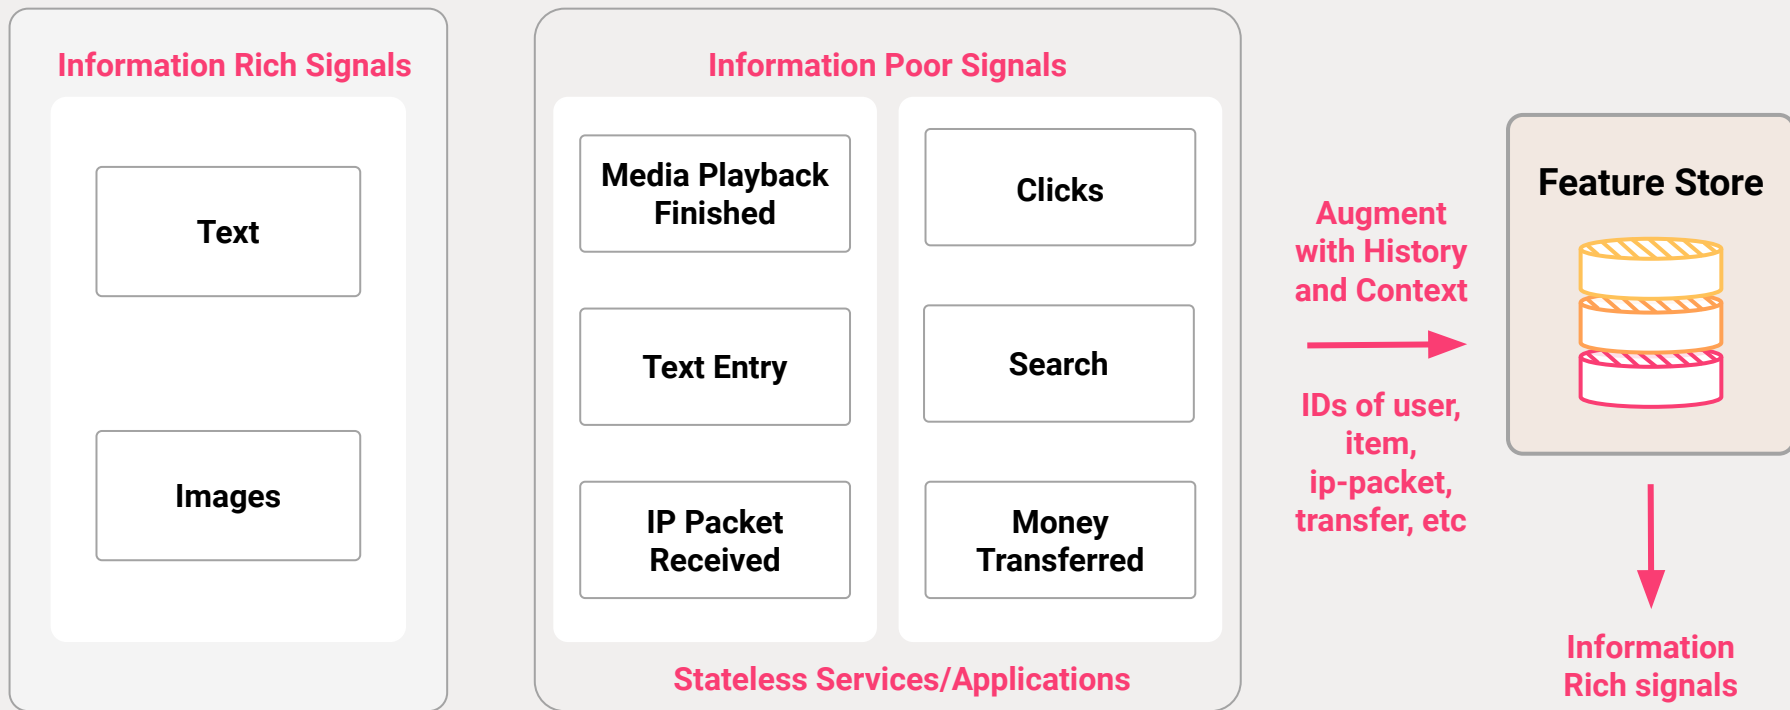


Apache Atlas

Hopworks End-to-End Machine Learning (ML) Pipelines




The Feature Store - From Information Poor Signals to Information Rich Signals



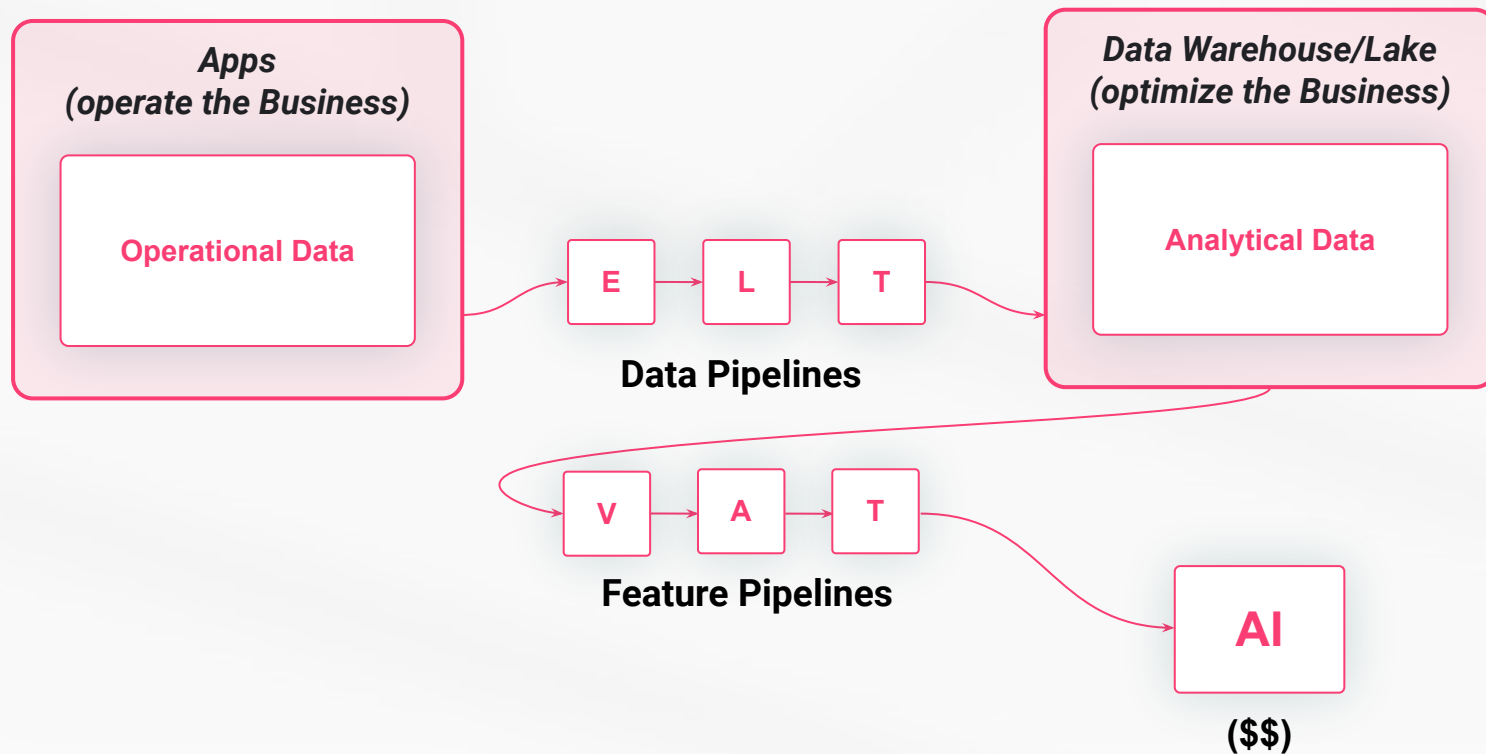
Commercial Feature Stores for Machine Learning



How the feature store fits in your data infrastructure

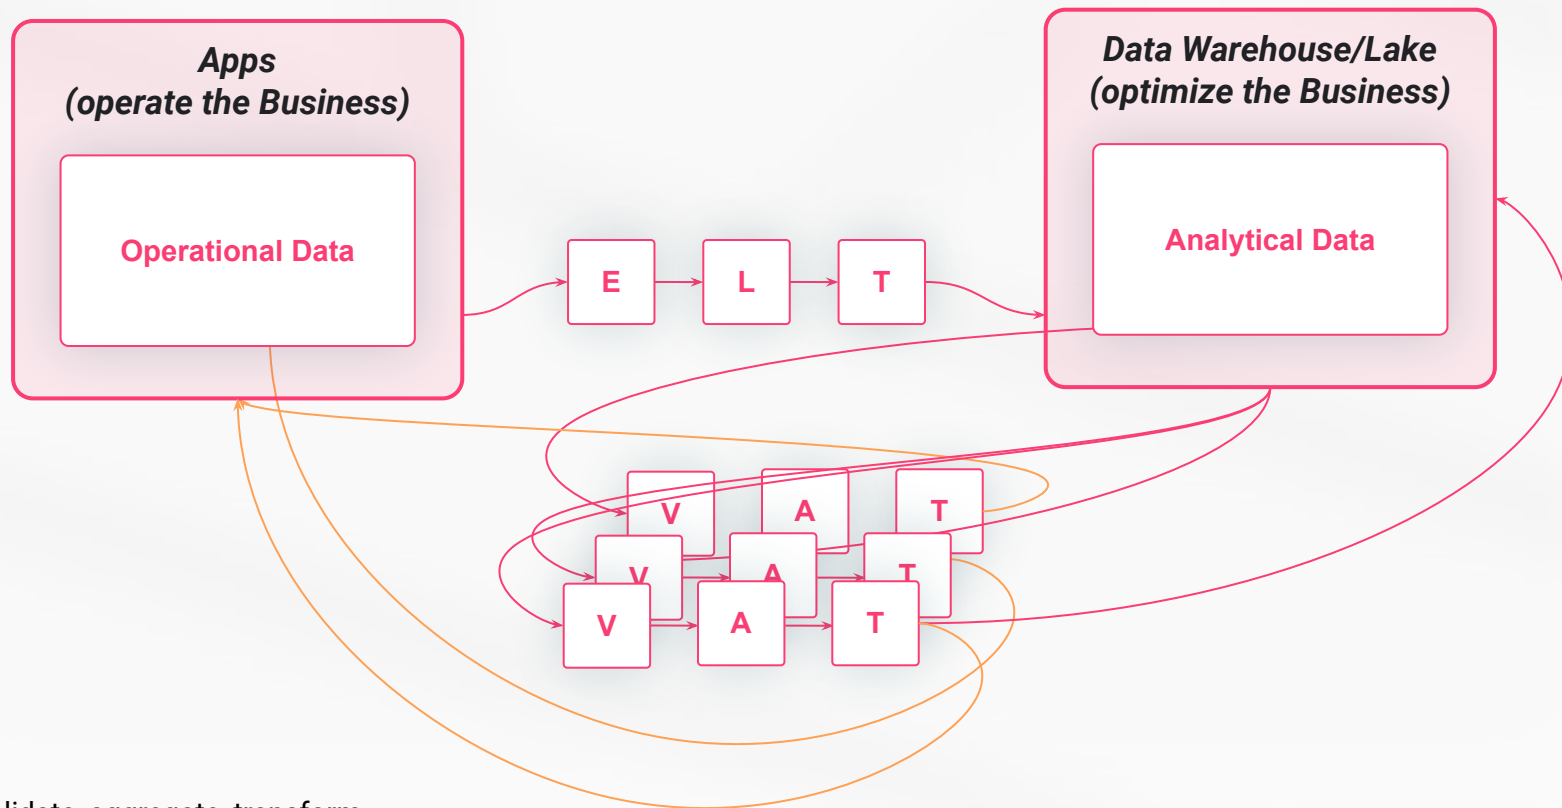


How the feature store fits in your data infrastructure



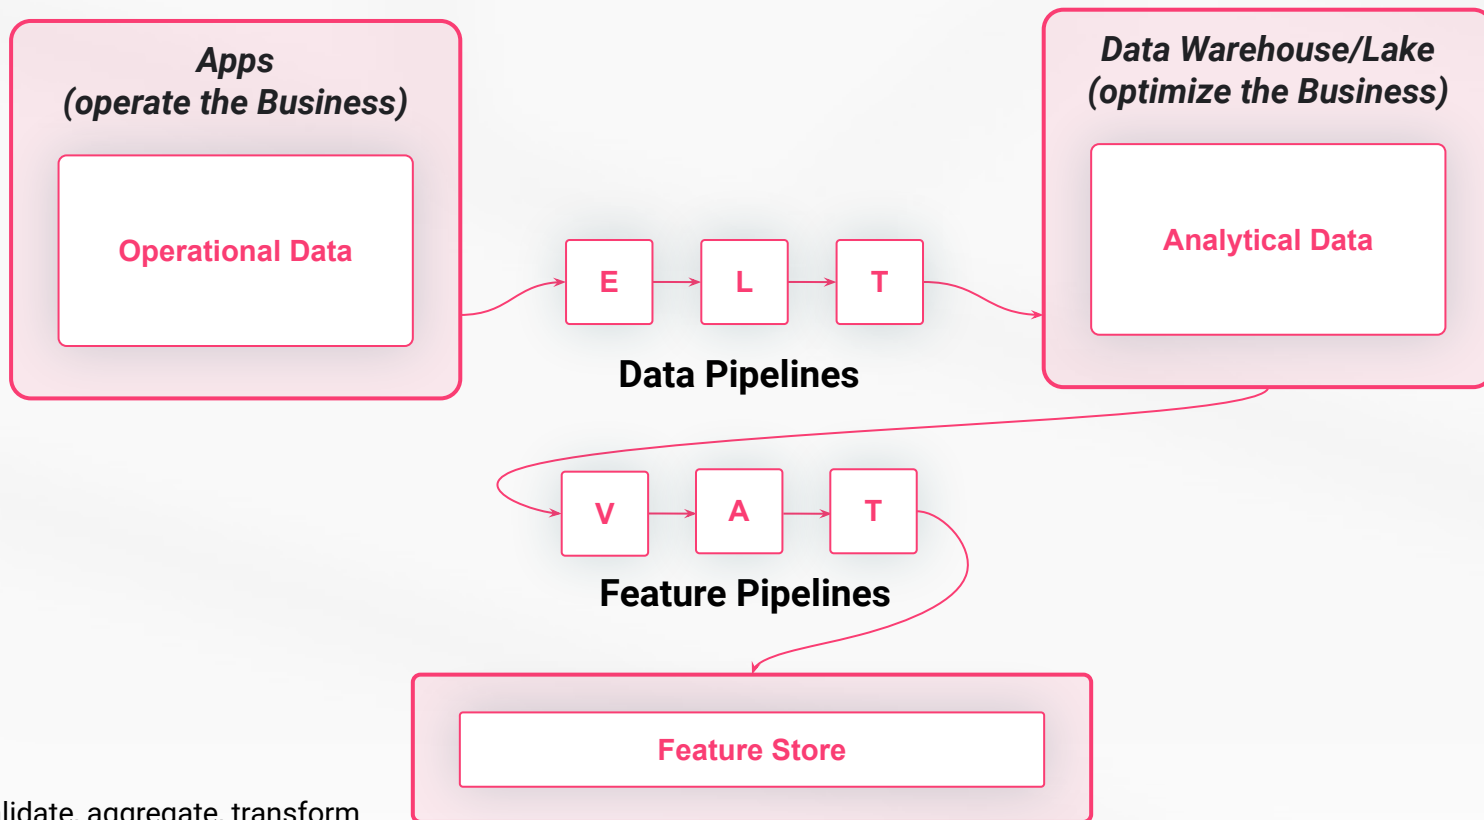
VAT = validate, aggregate, transform

How the feature store fits in your data infrastructure



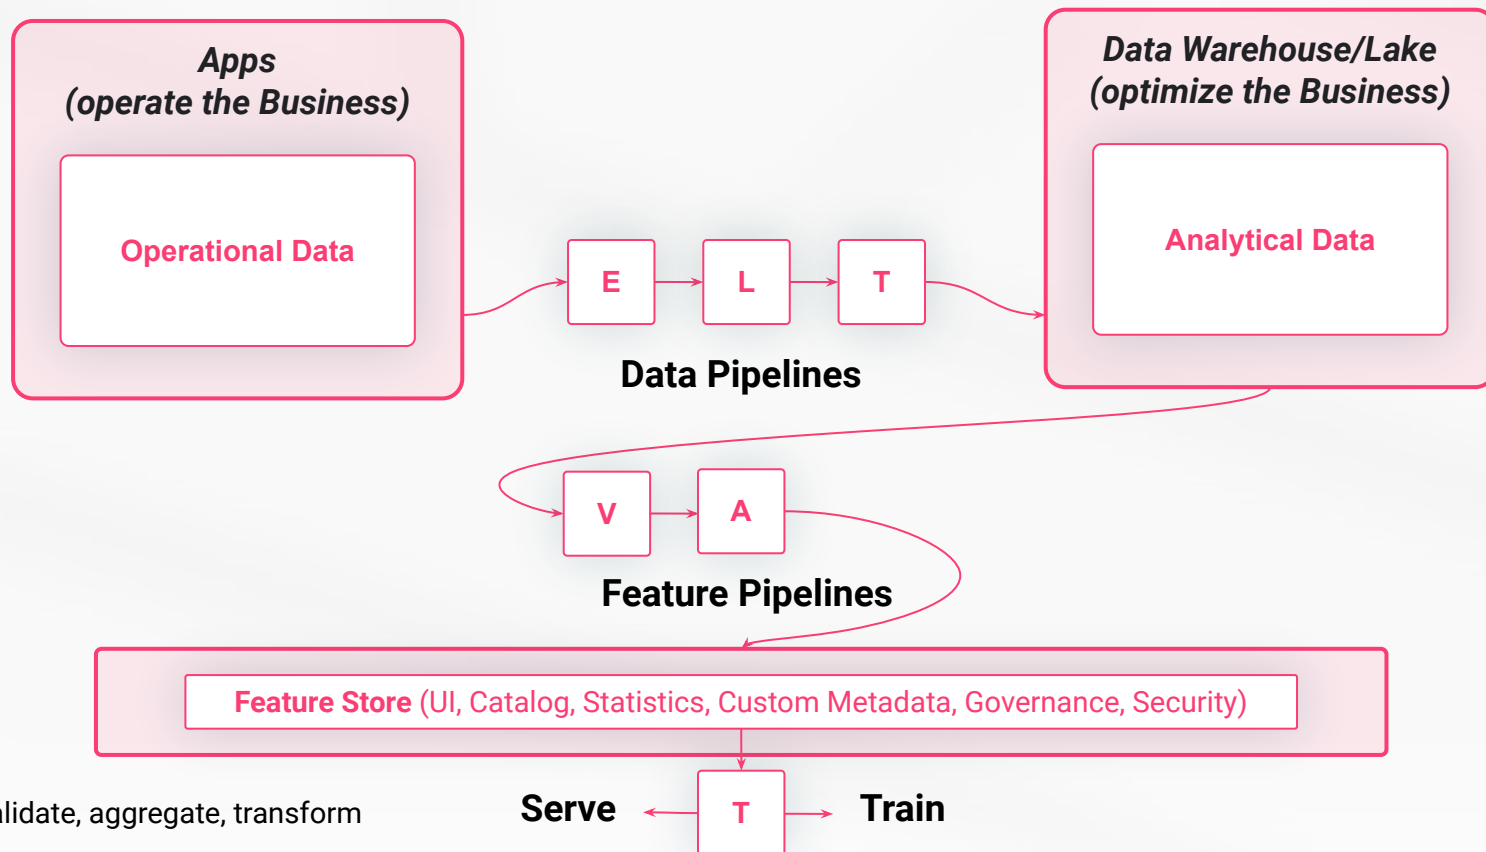
VAT = validate, aggregate, transform

How the feature store fits in your data infrastructure



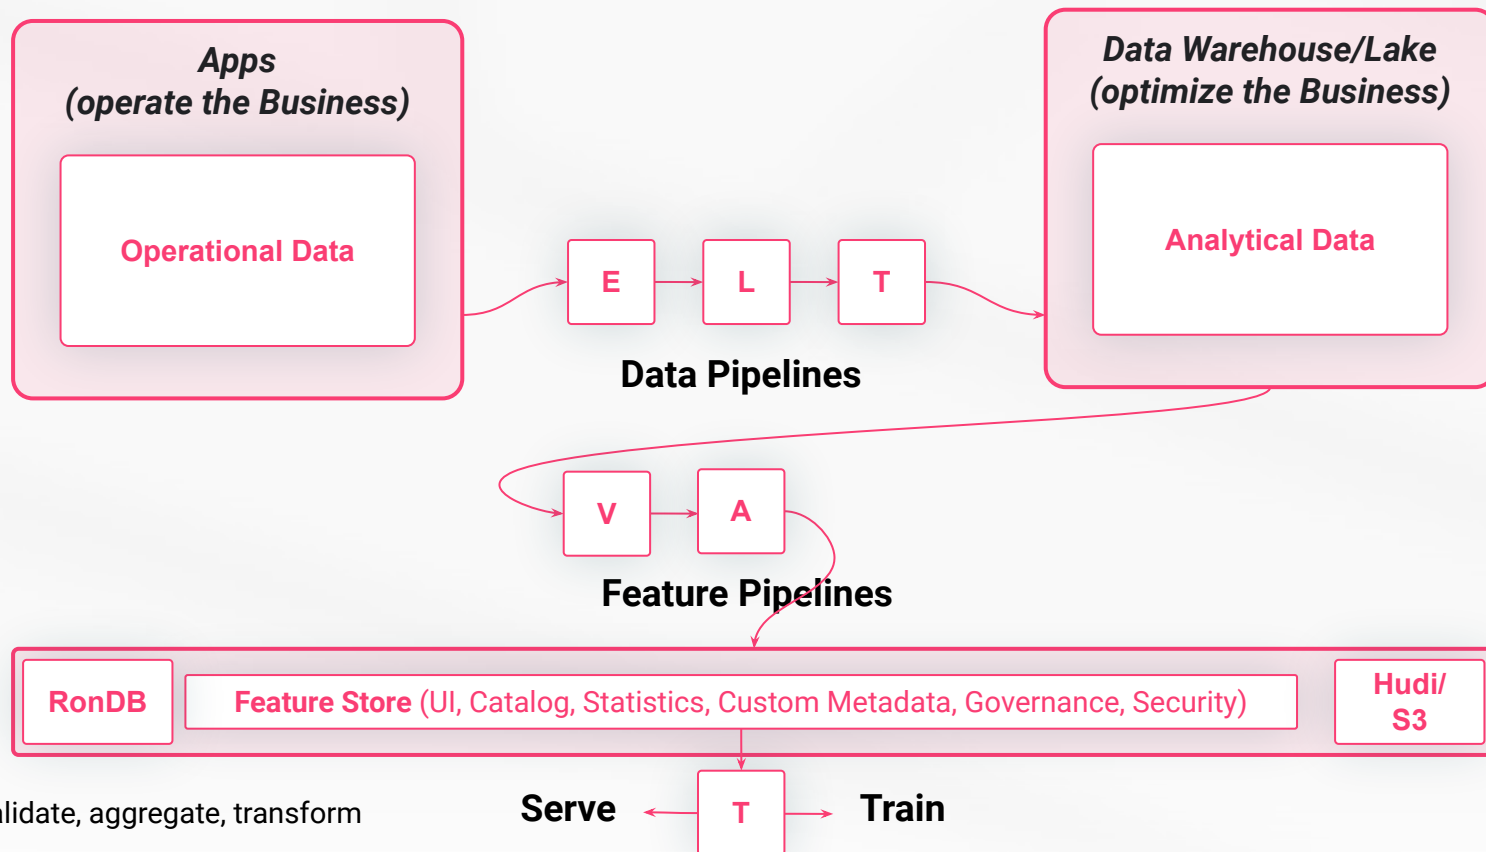
VAT = validate, aggregate, transform

How the feature store fits in your data infrastructure



VAT = validate, aggregate, transform

How the feature store fits in your data infrastructure



VAT = validate, aggregate, transform

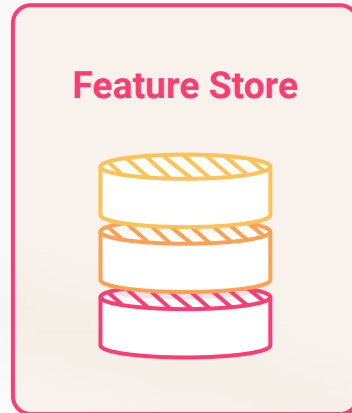
Enables Collaboration between folks who speak different languages



Data Scientist
Python

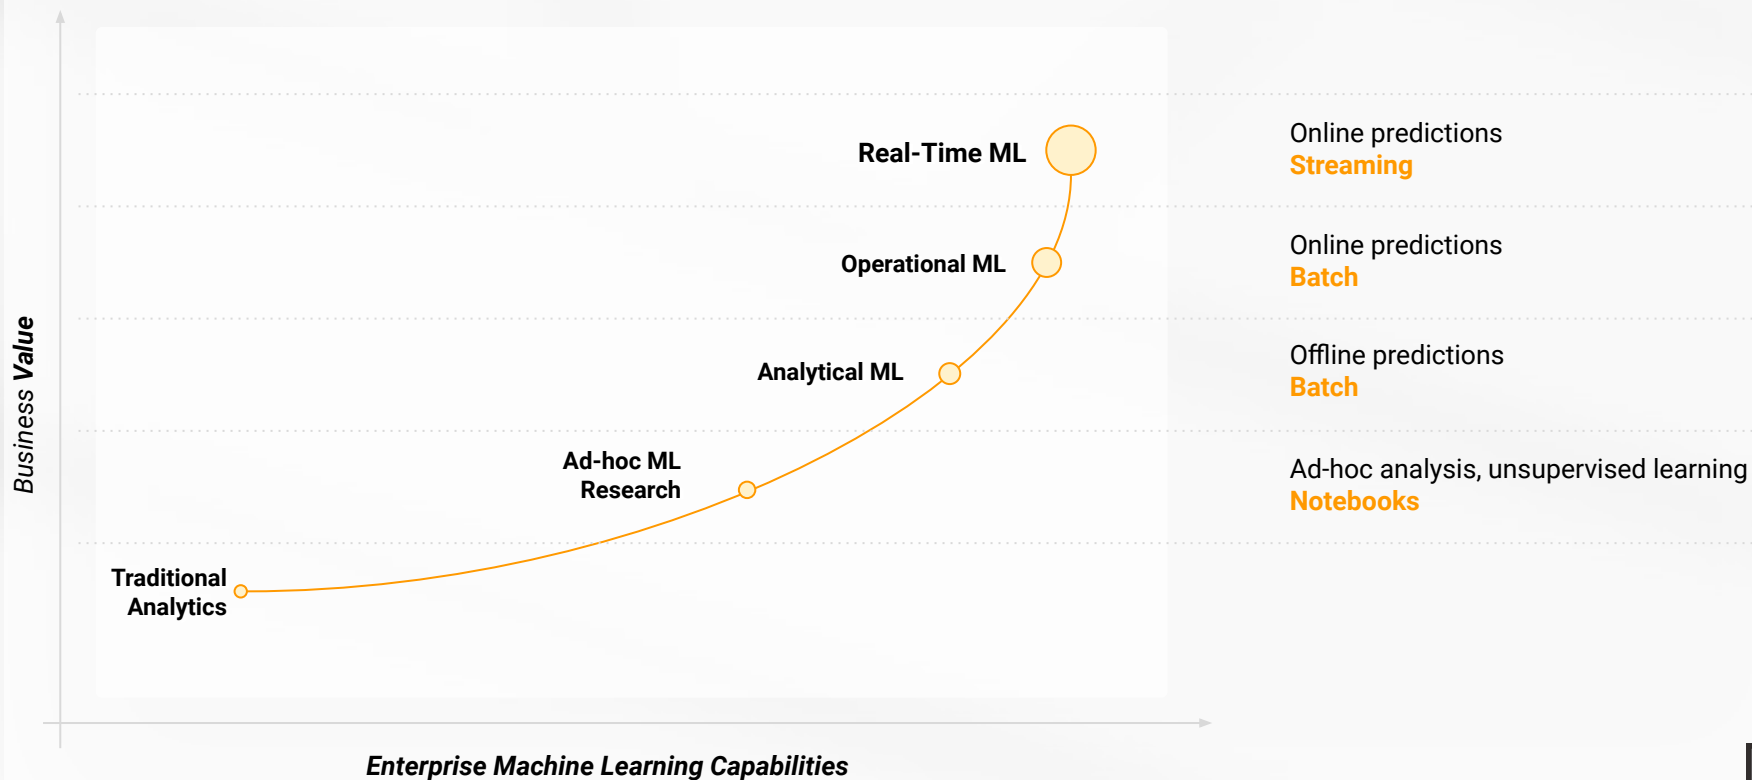


Data Engineer
SQL, Spark, Flink, Python

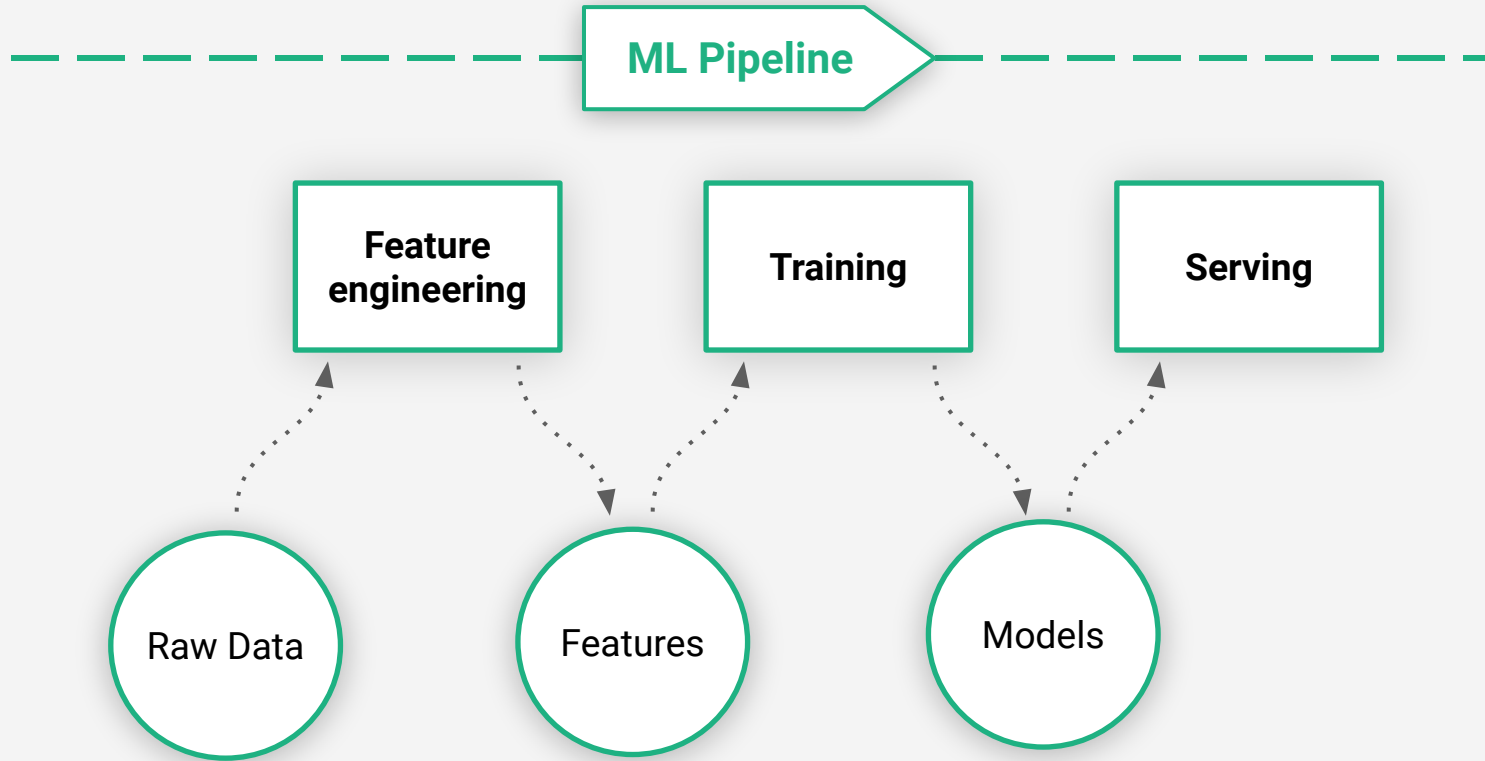


ML Engineer
Kubernetes, Serverless

Road To AI Value



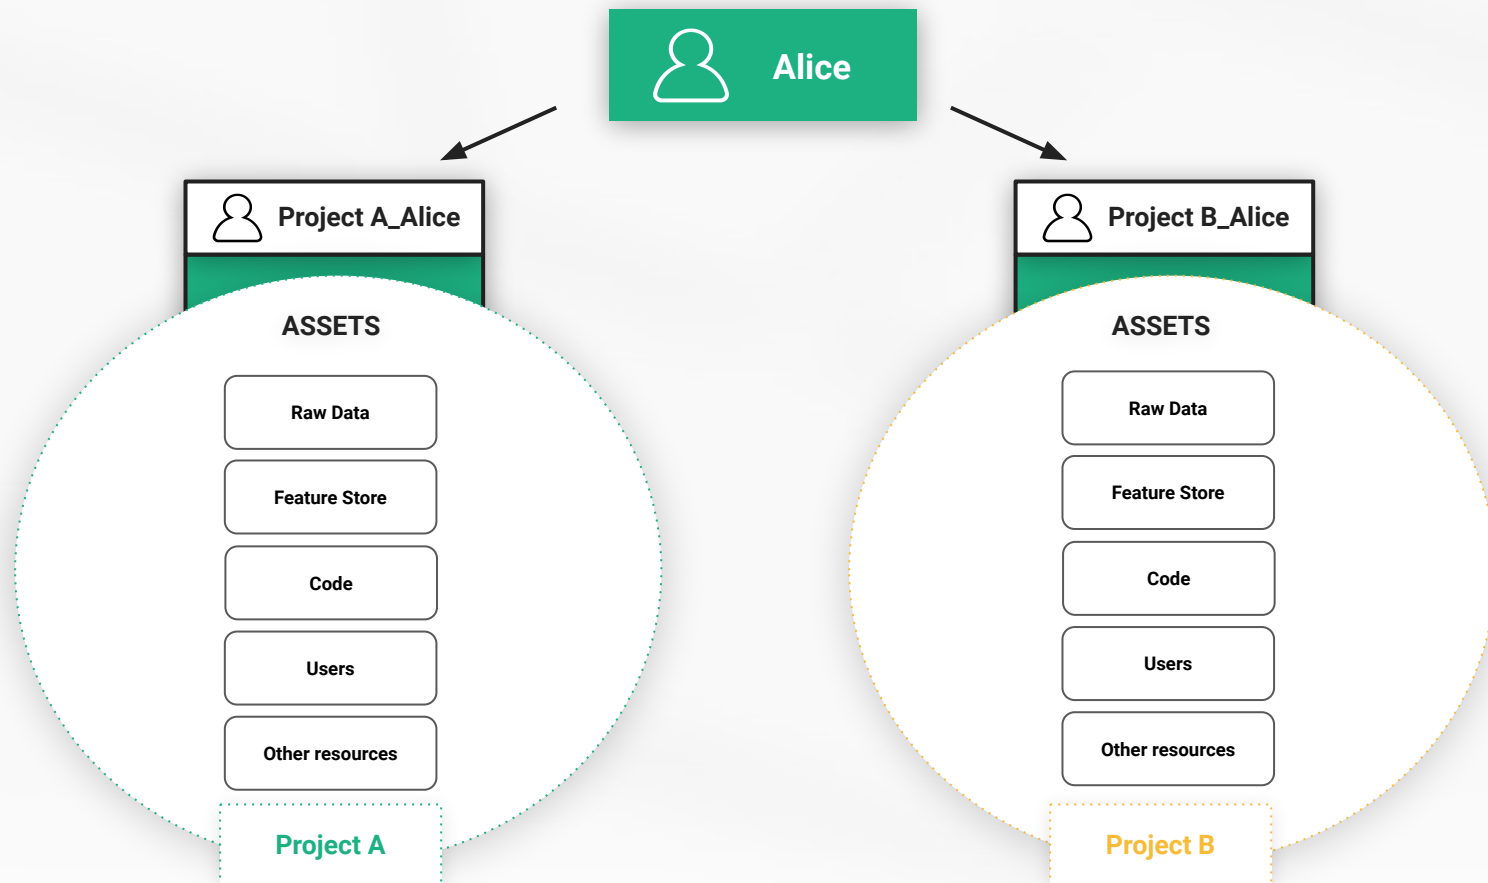
What are ML Pipelines?



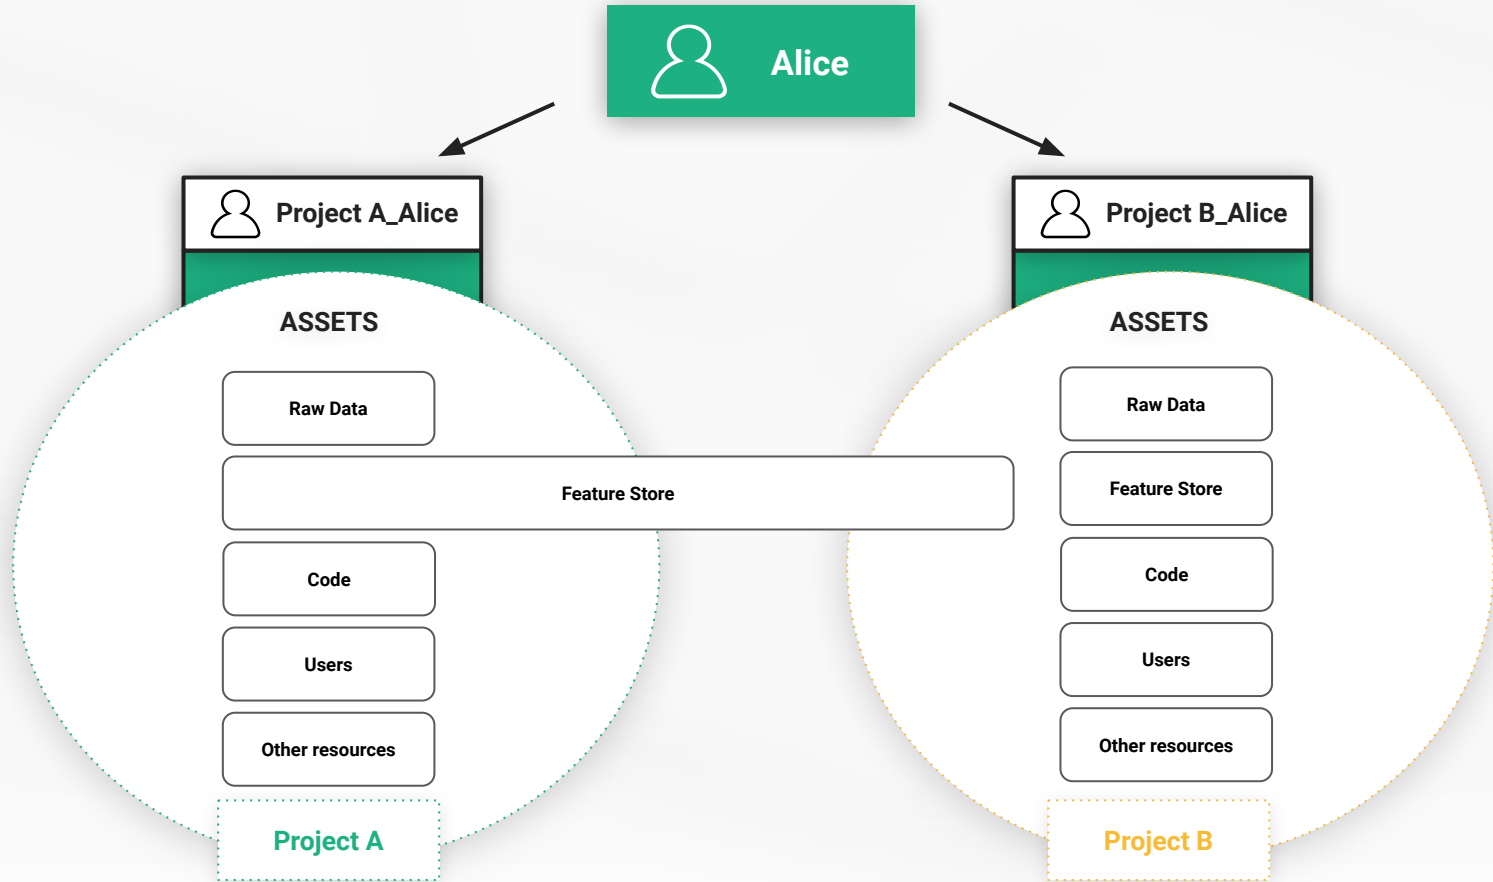


How it Started

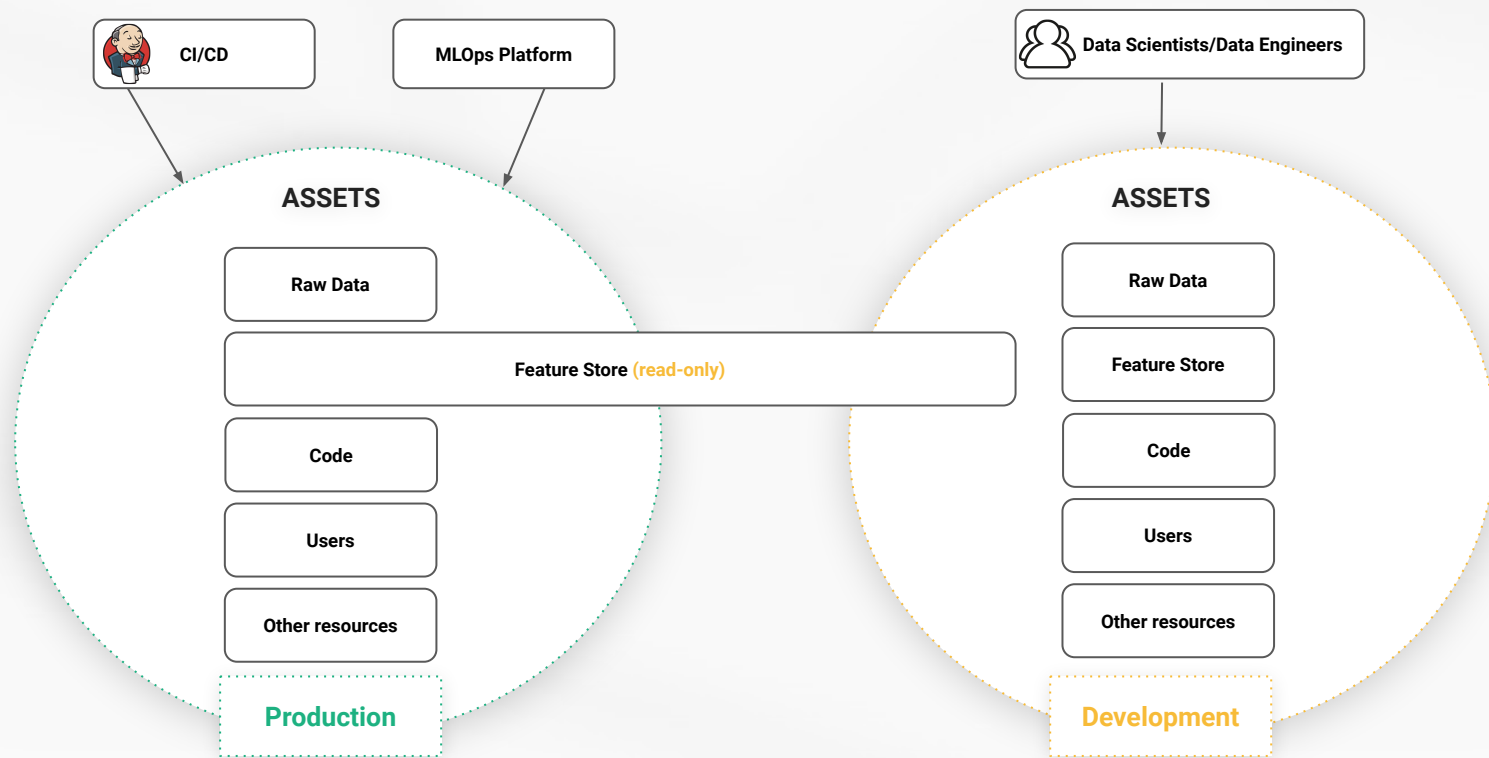
Project Based Multi Tenancy



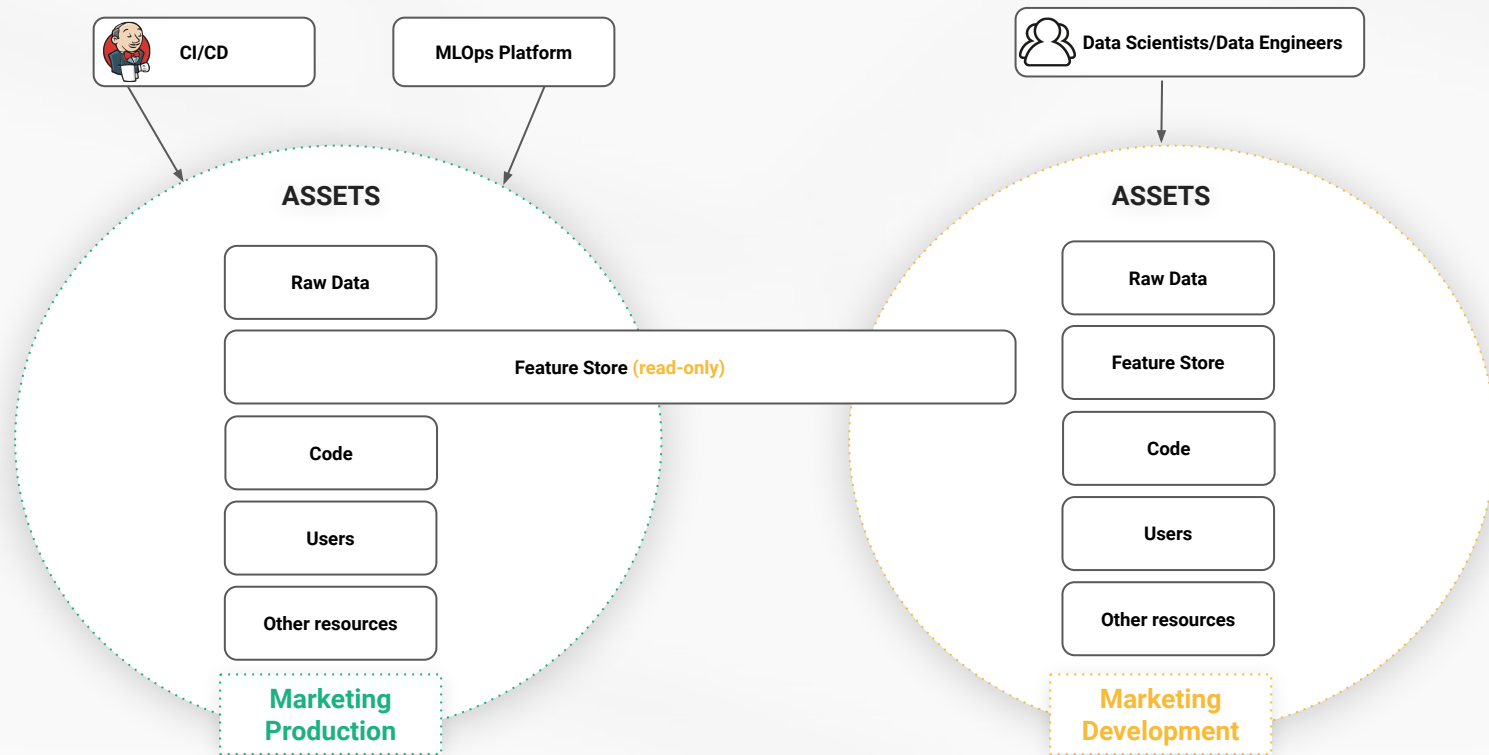
Project Based Multi Tenancy - Shared Data



Project Based Multi Tenancy - Production/Development



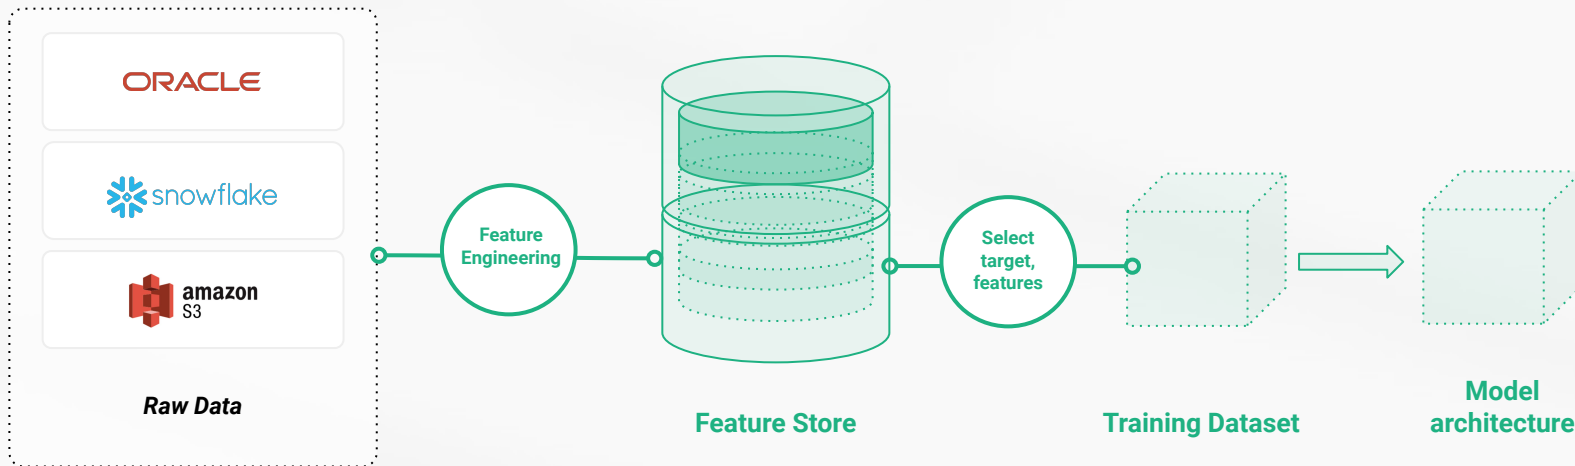
Project Based Multi Tenancy - Mix Structure



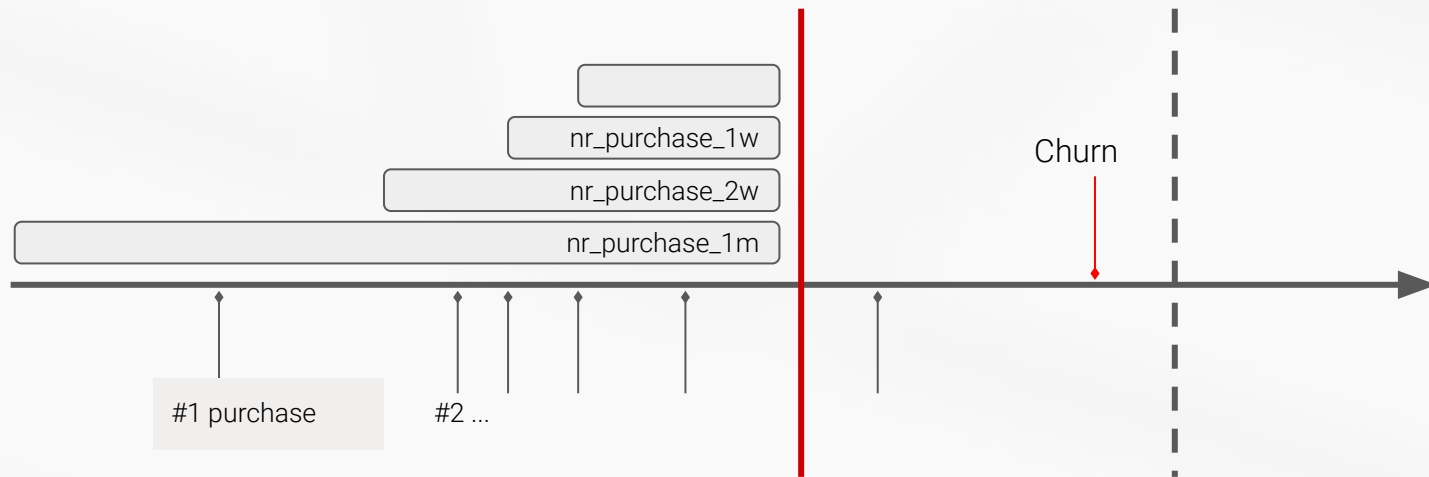


Feature Engineering

From Data to Features to Training Data to Models




Feature engineering - Aggregation example



Transaction				Features	
<i>transaction_time</i>	<i>id</i>		<i>Entity</i>	<i>1w</i>	<i>1m</i>
2021/02/10	23	<i>Feature Aggregation</i> →	1	1	2
2021/01/24	42				
2021/01/20	47				

Feature engineering




Existing feature groups


Dataframe = (Python/PySpark/Spark/Flink based feature engineering)

```
fg = fs.create_feature_group("churn",  
    version=1,  
    description="Customer information about activity of contract",  
    online_enabled=True,  
    primary_key=["customer_id", "contract_id"],  
    event_time="ts")
```


```
fg.save(dataframe)
```



Online feature store



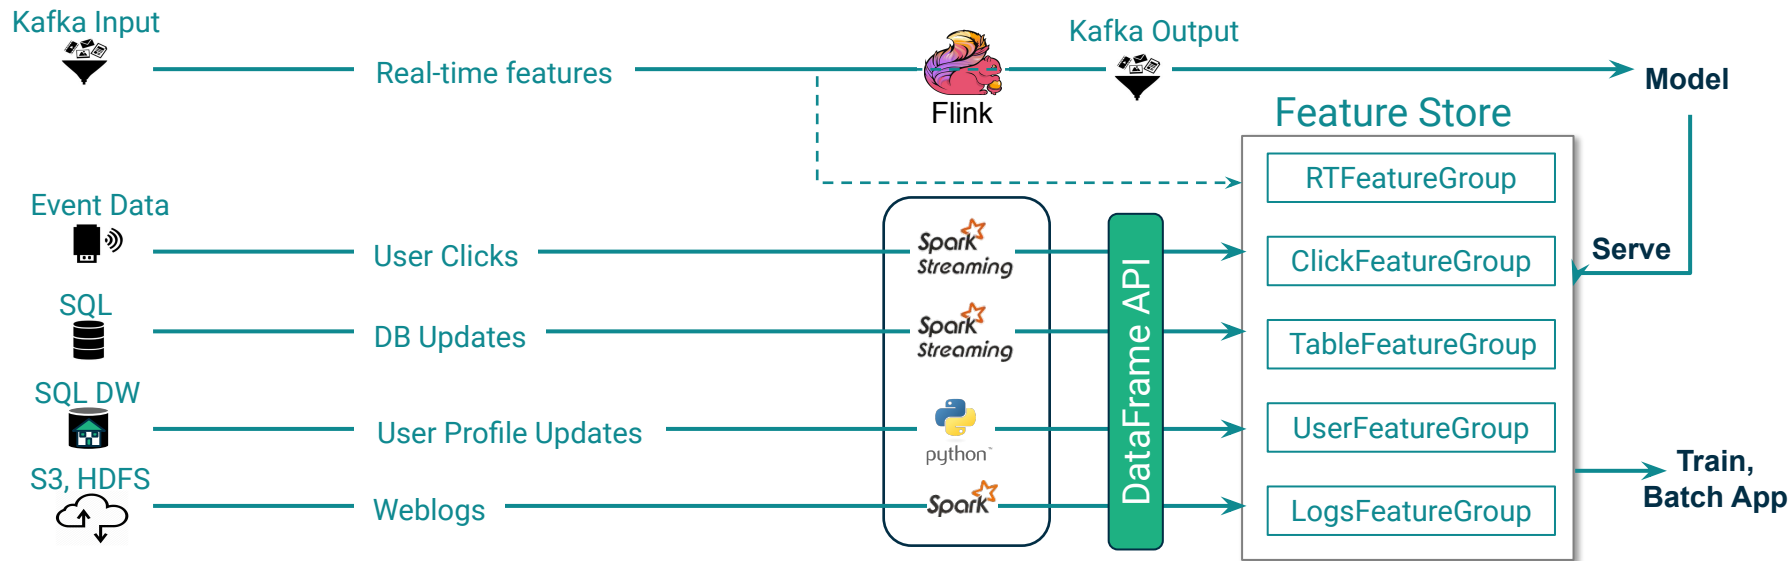
Offline feature store



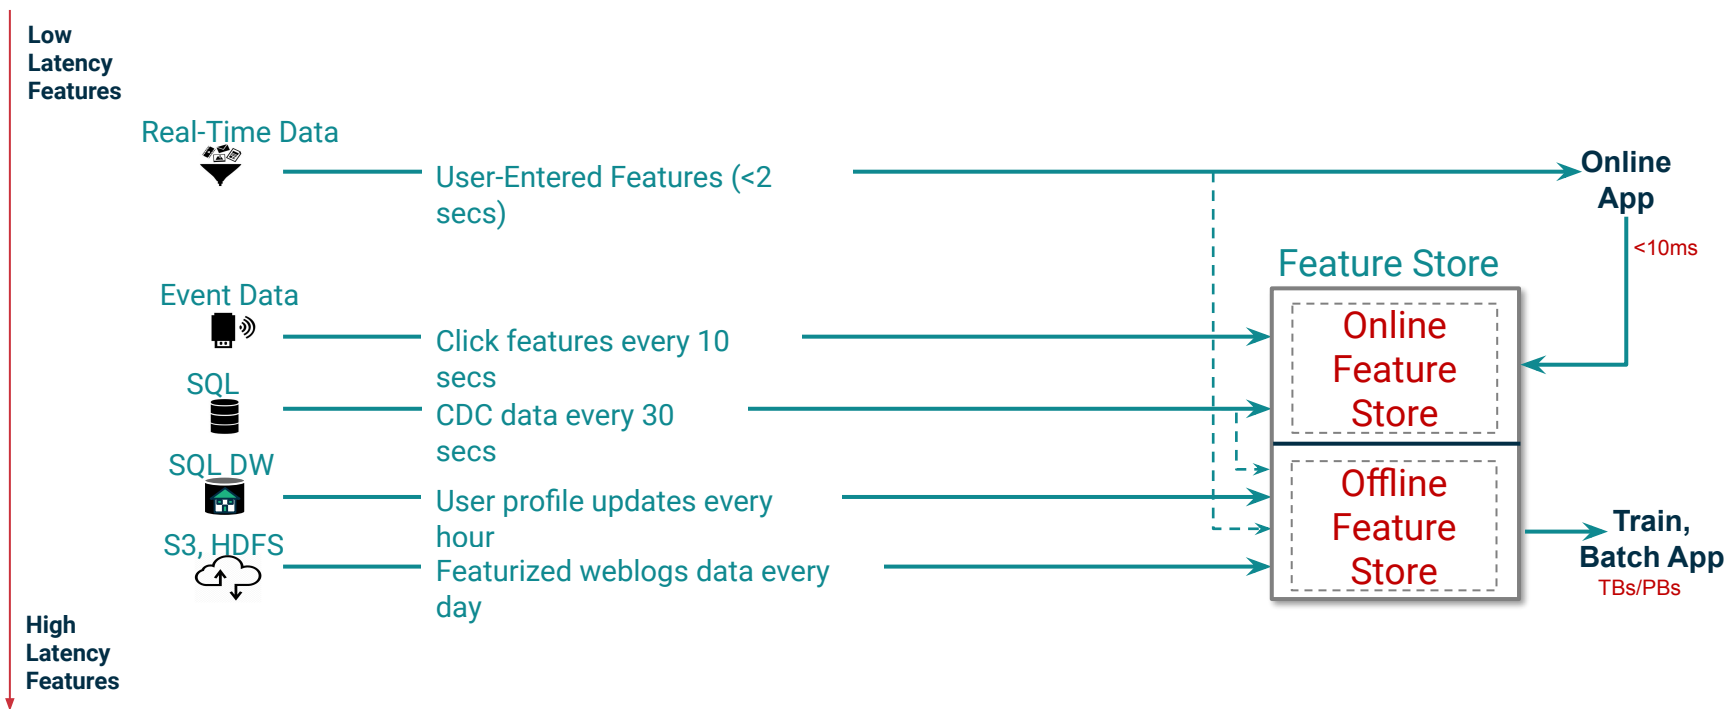
Other Data sources



Reusable Features are stored in FeatureGroups

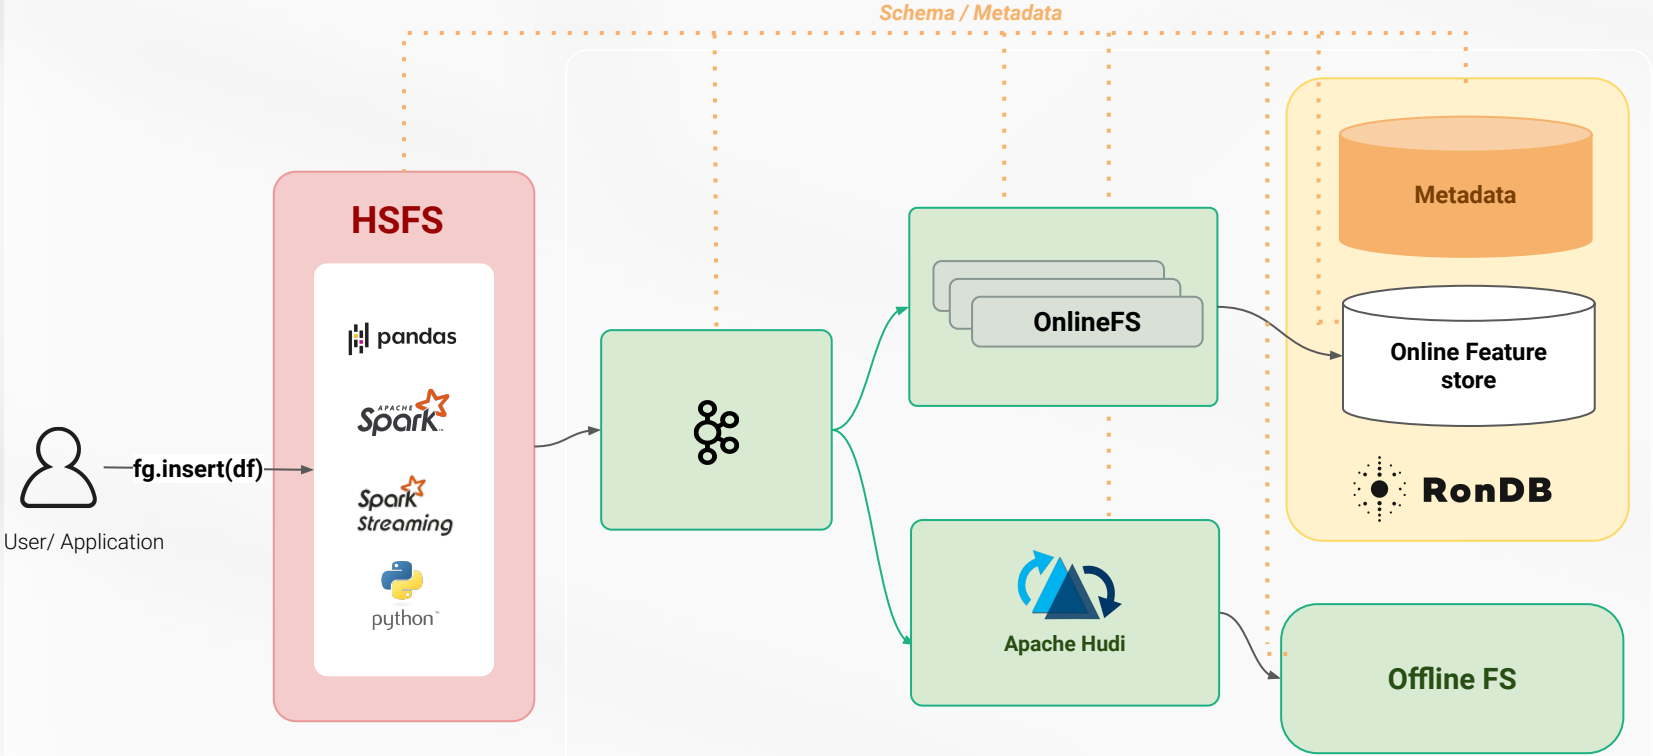


Feature Pipelines update the Feature Store (2 Databases!) with data from backend Platforms



No existing database is both scalable (PBs) and low latency (<10ms). Hence, online + offline Feature Stores.

Streaming Applications can write Fresh Features






Feature Groups


Cached feature groups




- Feature groups stored on Hopworks
- Can be available both offline and online.
- Documentation:
https://docs.hopsworks.ai/feature-store-api/latest/generated/feature_group/
- Example:
https://examples.hopsworks.ai/featurestore/hsfs/basics/feature_engineering/



Online Feature Store



Offline Feature Store





Schema versioning

- Each feature group has a version number
- Version numbers allow users to identify breaking changes to the schema (feature dropped, change in the way a feature is being computed)
- Appending feature to a feature group is not considered a breaking change

Data versioning

- Calling `insert()/save()` on a feature group generates a new data commit.
- Data commits allow users to track how the data changed during the lifetime of a feature group.
- Users can navigate the commit history using the [Activity UI](#)
- Using the `as_of` method, users can retrieve features from a feature group, at a specific point in time.



List actions performed on a feature group:

- Feature group creation
- Data ingestion
- Statistics computation
- Data validation

The screenshot shows a web interface for viewing activity on a feature group. The breadcrumb path is 'cc_fraud > card_transactions #21'. The activity log is filtered for 'events all' and shows a list of 'Data ingestion' events for '12 Oct. 2021'. Each event includes a commit ID and row counts for new, updated, and deleted rows.

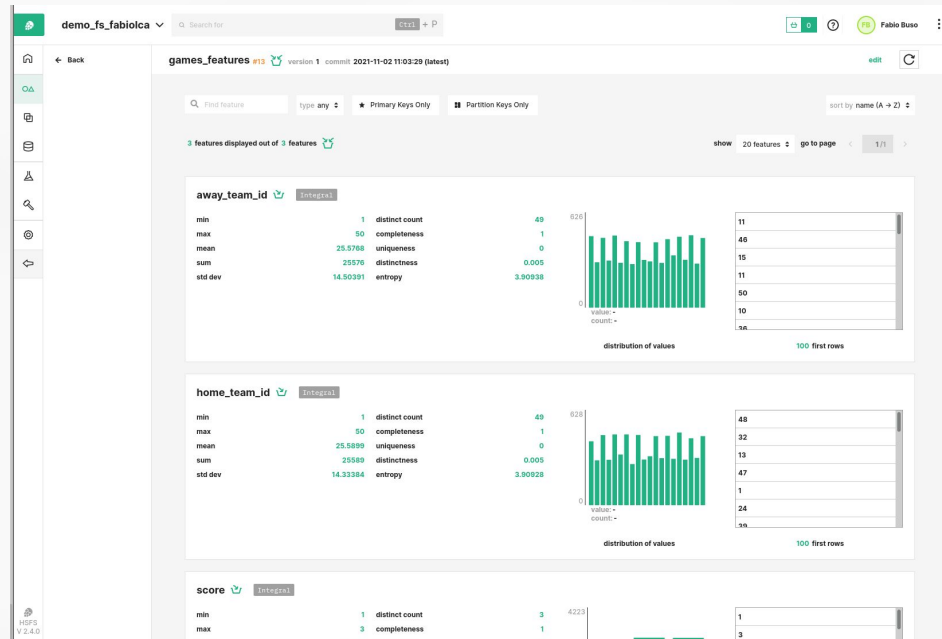
Date	Action	Status	Commit ID	Row Counts
12 Oct. 2021	Data ingestion	Warning	commit 2021-10-12 18:49:33	0 new rows, 43k updated rows, 0 deleted rows
12 Oct. 2021	Data ingestion	Warning	commit 2021-10-12 18:46:32	0 new rows, 6 updated rows, 0 deleted rows
12 Oct. 2021	Data ingestion	Warning	commit 2021-10-12 18:44:26	0 new rows, 43k updated rows, 0 deleted rows
12 Oct. 2021	Data ingestion	Warning	commit 2021-10-12 18:41:15	0 new rows, 1 updated rows, 0 deleted rows
12 Oct. 2021	Data ingestion	Warning	commit 2021-10-12 16:01:12	0 new rows, 43k updated rows, 0 deleted rows
12 Oct. 2021	Data ingestion	Warning	commit 2021-10-12 15:57:58	0 new rows, 7 updated rows, 0 deleted rows
12 Oct. 2021	Data ingestion	Warning	commit 2021-10-12 15:27:21	57 new rows, 43k updated rows, 0 deleted rows
12 Oct. 2021	Data ingestion	Success	commit 2021-10-12 15:21:36	87k new rows, 0 updated rows, 0 deleted rows



Statistics are computed at feature group level for each [data commit](#)

Hopworks computes automatically: descriptive statistics, histograms and correlations between features

Statistics can be explored from the UI.





Tags allow users to specify arbitrary metadata and make it searchable throughout the feature store.

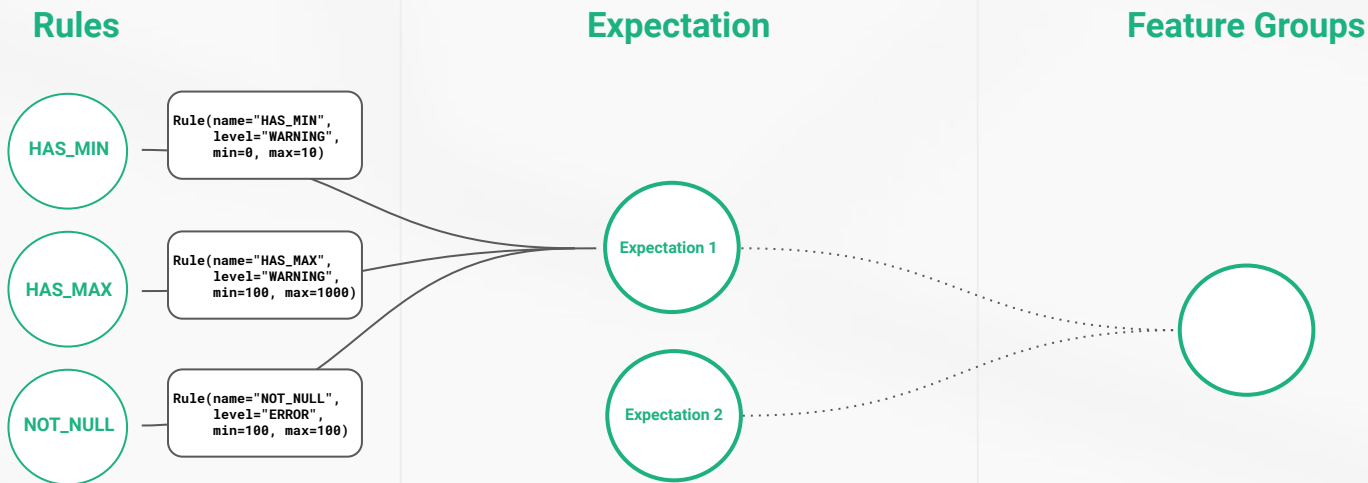
Tags require a schema that can be defined and enforced at platform level.

Tags can be manipulated through the UI or through the [APIs](#)

Tags		edit
pii		edit properties
pii	true	
privacy-officer	Fabio Buso	



Data Validation



These are the rules supported by the platform:
https://docs.hopworks.ai/feature-store-api/latest/generated/feature_validation/#rule-definitions

An expectation is a set of rules with specific values and rule severity

Expectations can be applied to specific features on a feature group, or to the entire set of features of a feature group

https://docs.hopworks.ai/feature-store-api/latest/generated/feature_validation/#rule-definitions

Validation types



Validation Type	Success	Warning	Failure
Strict	Insertion	Reject	Reject
Warning	Insertion	Insertion	Reject
All	Insertion	Insertion	Insertion
None	No data validation performed		

An expectation is a set of rules with specific values and rule severity severity

Expectations can be applied to specific features on a feature group, or to the entire set of features of a feature group

https://docs.hopsworks.ai/feature-store-api/latest/generated/feature_validation/#rule-definitions

```
Dataframe = (Python/PySpark/Spark/Flink based feature engineering)
fg = fs.create_feature_group("churn",
                             version=1,
                             description="Customer/contract information about activity of
contract",
                             validation_type="STRICT",
                             primary_key=["customer_id", "contract_id"])
fg.save(dataframe)
```


Data Validation




```
Dataframe = (Python/PySpark/Spark/Flink based feature engineering)

fg = fs.create_feature_group("churn",
    version=1,
    description="Customer/contract information about activity of
contract",
    online_enabled=True,
    primary_key=["customer_id", "contract_id"],
    event_time="ts")

fg.save(dataframe)
```



Data validation

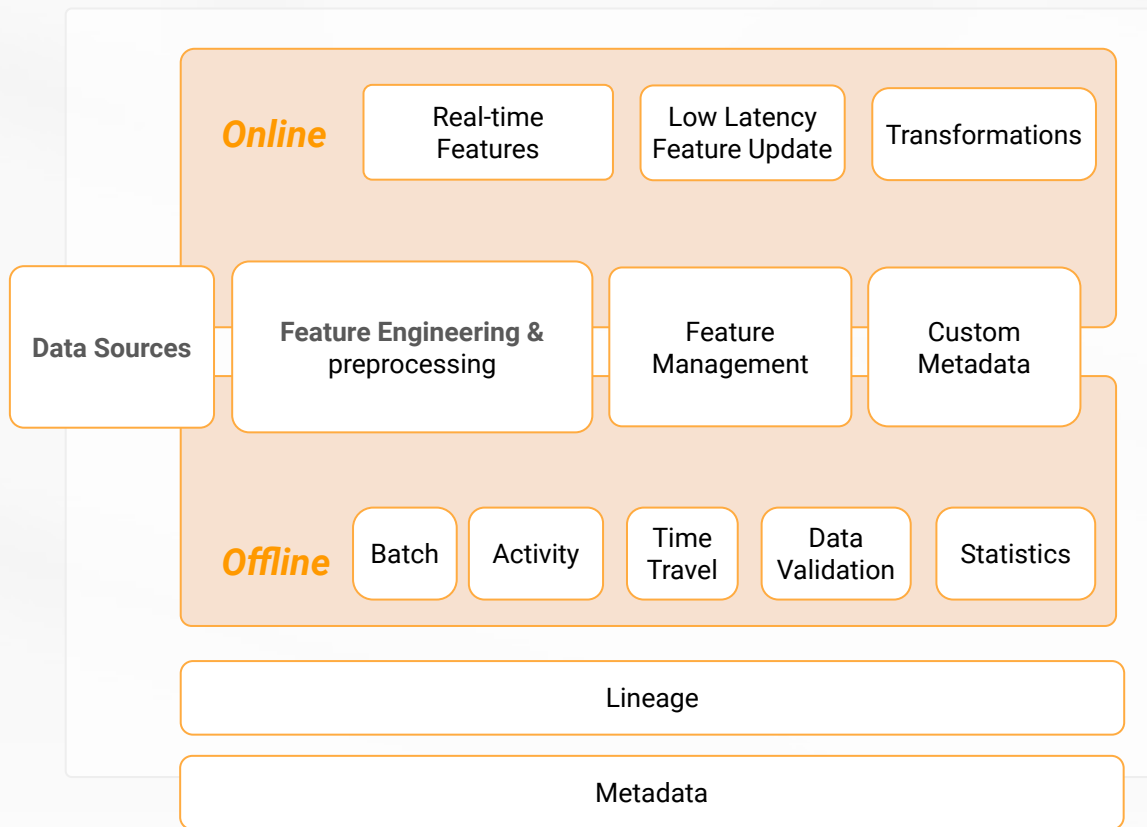


Feature store



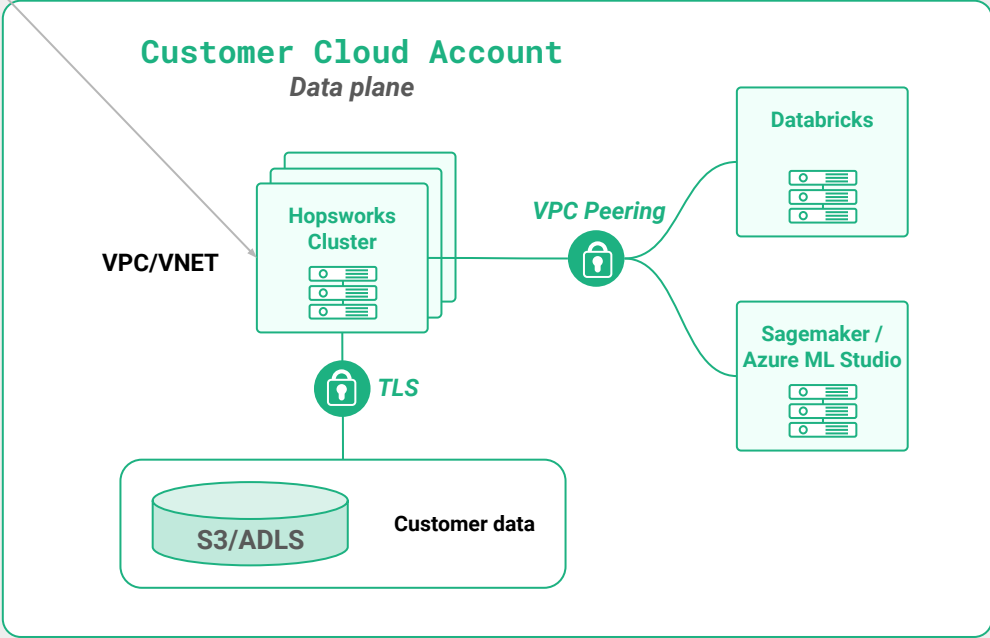
 slack

Hopsworks Feature Store

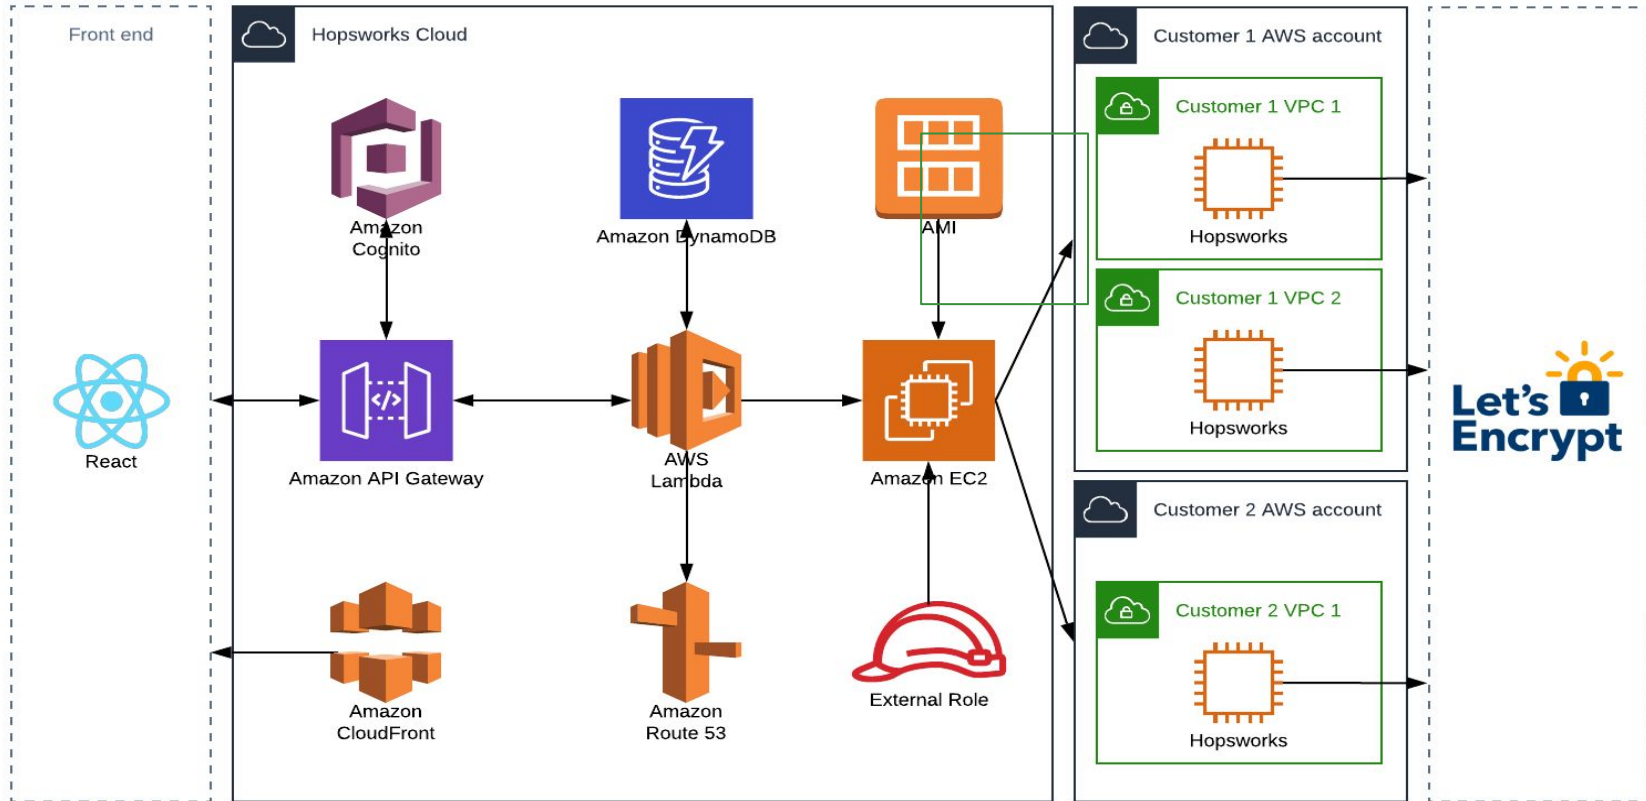


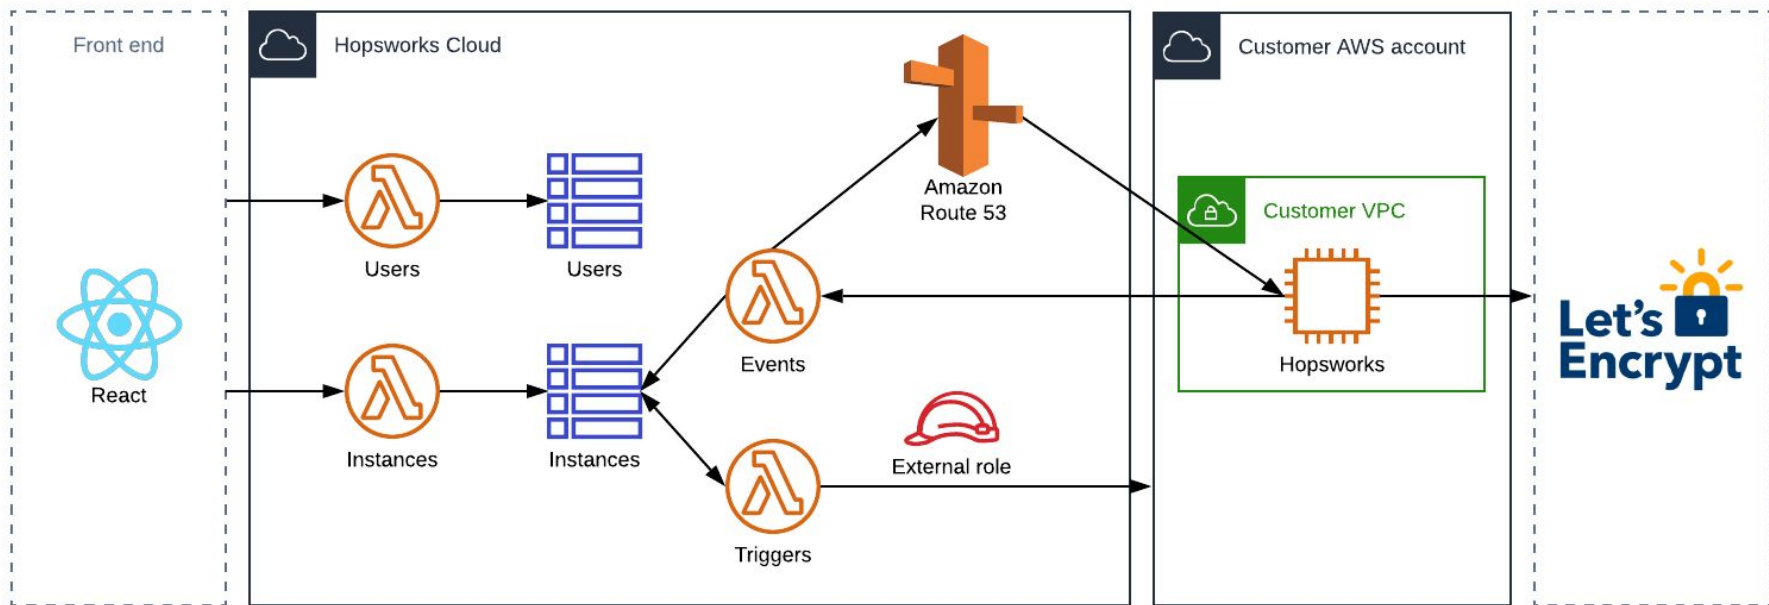


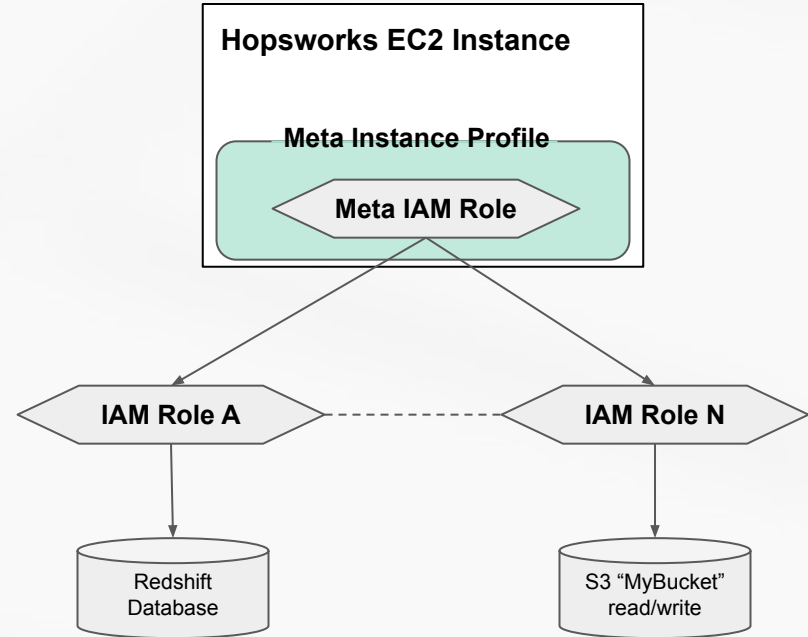
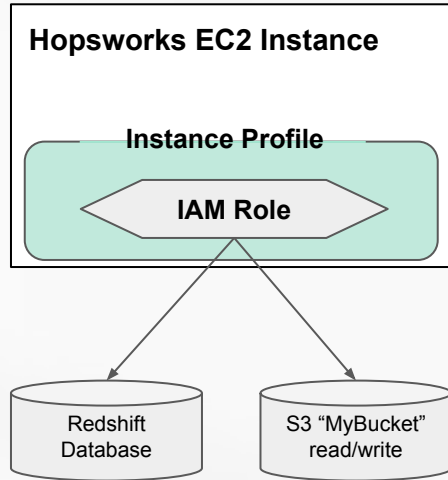
Hopsworks.ai
Deployment / Security




Serverless Platform on AWS - Amplify, Cognito, CloudFront, Lambdas, Route 53, DynamoDB








Configuring AWS IAM Role Chaining



Example: Only allow admins of Project 'RawFeatures' to read from Redshift



Assuming AWS IAM Roles

The screenshot shows the Hopsworks interface for a project named 'project1'. The 'Cloud Roles' tab is active, displaying a list of roles accessible from the project. The roles are:

- Role id: 1 (Default), Project Role: ALL, Cloud Role: arn:aws:iam:xxxxxxxxxxxx:role/s3-role
- Role id: 1026, Project Role: ALL, Cloud Role: arn:aws:iam:xxxxxxxxxxxx:role/s3-role1
- Role id: 1029, Project Role: Data scientist, Cloud Role: arn:aws:iam:xxxxxxxxxxxx:role/s3-role2

```
from hops.credentials_provider import get_role, assume_role
credentials = assume_role(role_arn=get_role(1))
spark.read.csv("s3a://resource/test.csv").show()
```

```
import io.hops.util.CredentialsProvider
val creds = CredentialsProvider.assumeRole(CredentialsProvider.getRole(1))
spark.read.csv("s3a://resource/test.csv").show()
```



Metadata is data that describes other data.



Metastore (Database)

Provenance queries

- SQL or Free-Text or Graph?
- Update Throughput?
- Latency of queries?
- Size of Metadata?

File System (S3, HopsFS, etc)

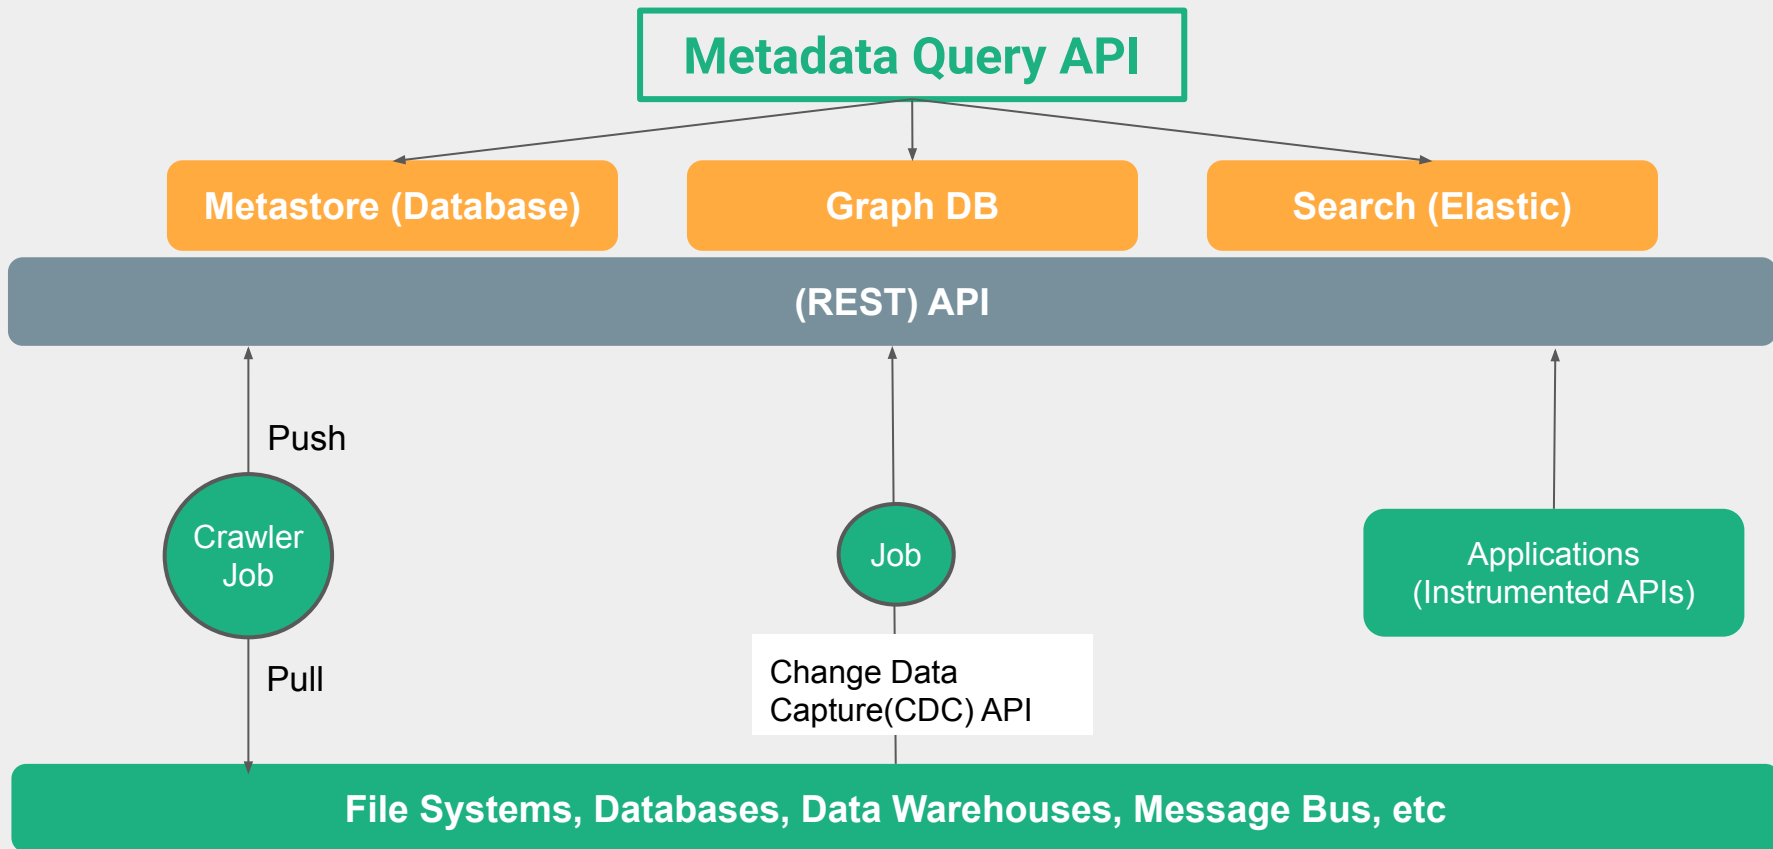
Metadata Cataloging Systems - a whole industry



<https://www.dataplatformschool.com/blog/w0y8g0-the-data-governance-zoo>

miro

3 Mechanisms for Metadata Collection. Polyglot Metadata Storage for Efficient Querying.





Metastore



Consistency issues
Synchronization



File System (S3, HopsFS, etc)



Metadata is data that describes other data.

Unspoken Assumption:

Why are Data and Metadata always separate stores?

Artifacts and Metadata in End-to-End ML Pipelines



Metastore (Database)

Governance
(Privileges, Audit,
Retention, etc)

Feature Stats
(Min, Max, Std,
Mean, Distrib.)

Experiments
(HParams, Env,
Results, Graphs)

Model Desc
(Privileges, Perf,
Provenance, etc)

Jobs
(Executions,
Lineage, Prov.)

Metadata

Artifacts

Raw Data

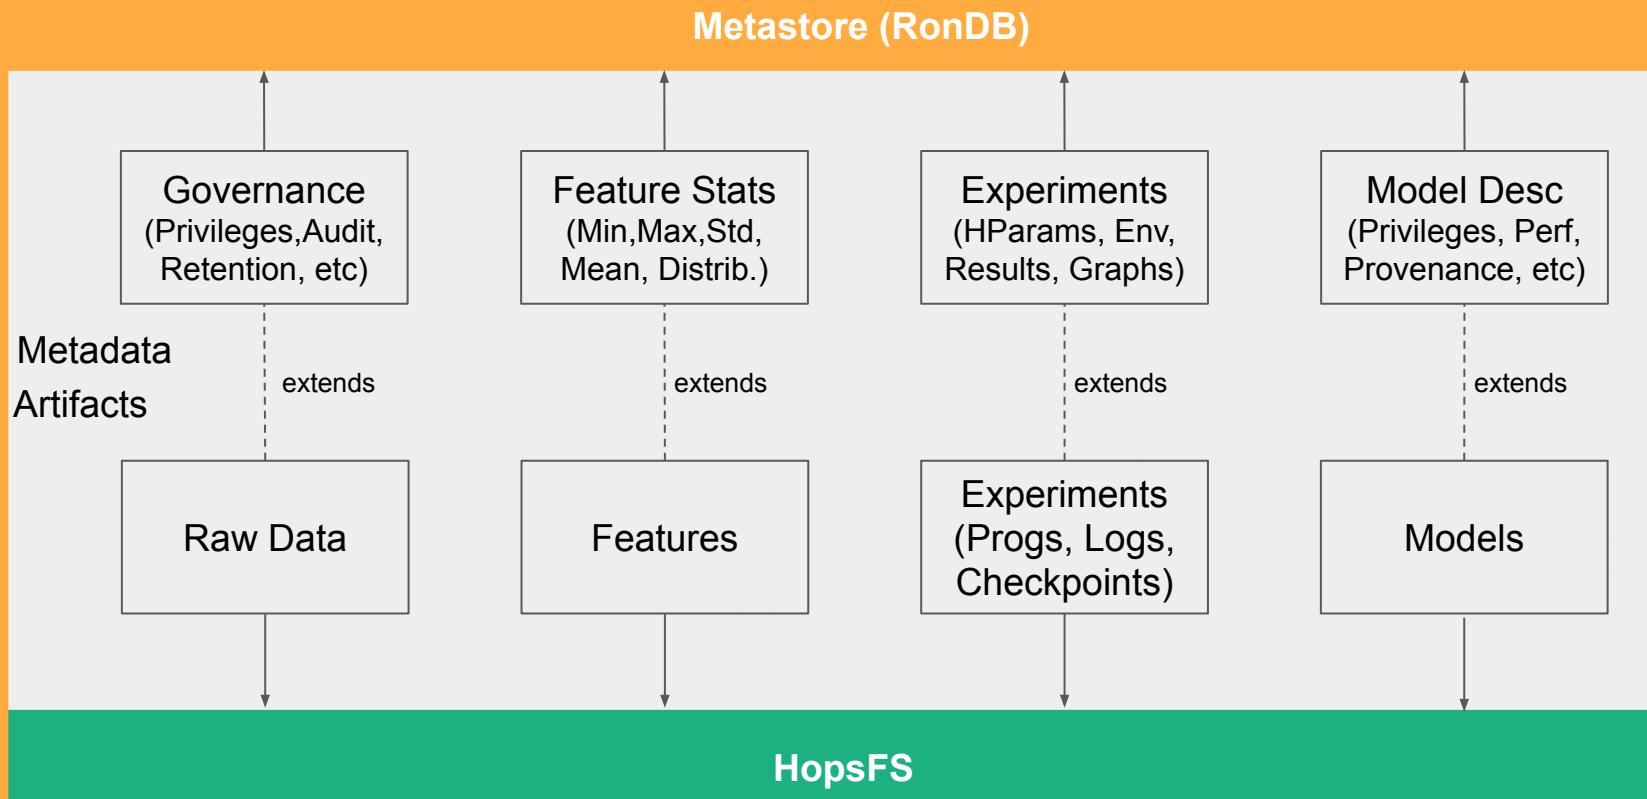
Features

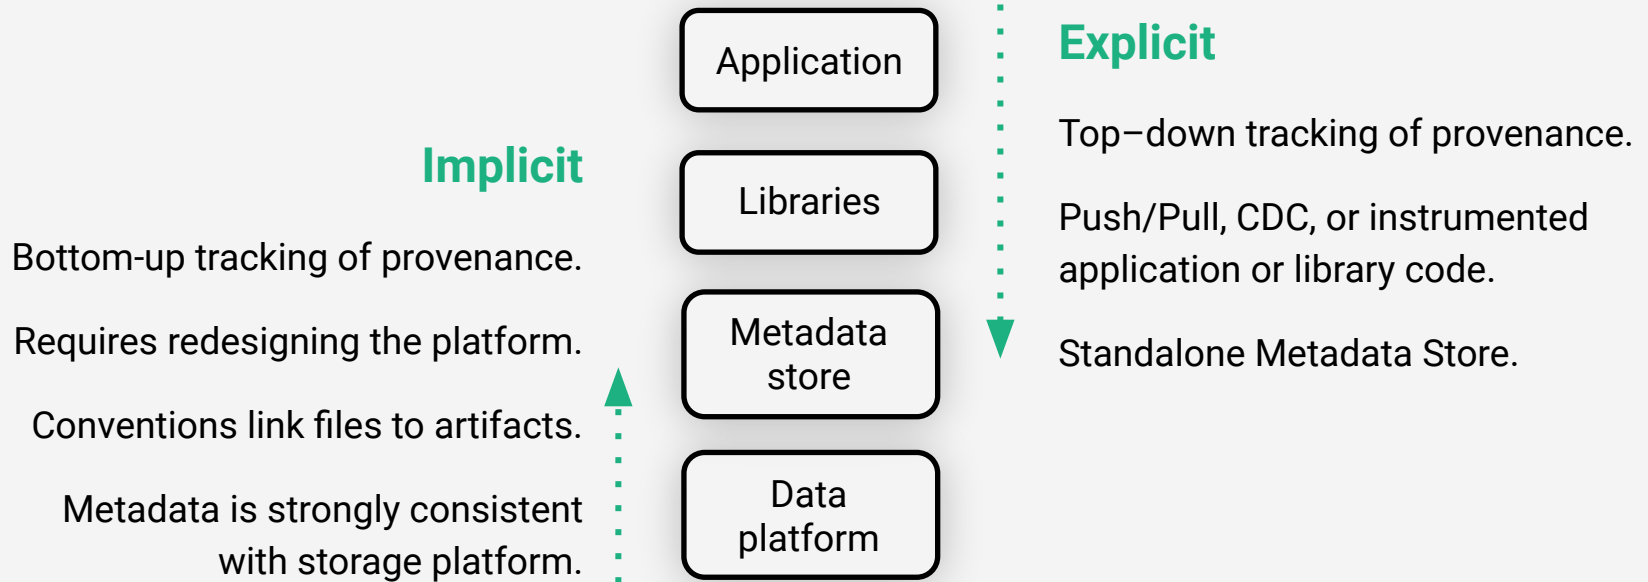
Experiments
(Src Code, Logs,
Checkpoints)

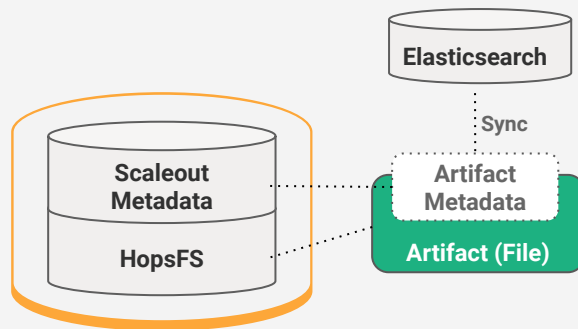
Models

File System (S3, HopsFS, etc)

Mechanism 4: Artifacts and Metadata in the same system - a Unified Metadata Layer (Hopsworks)







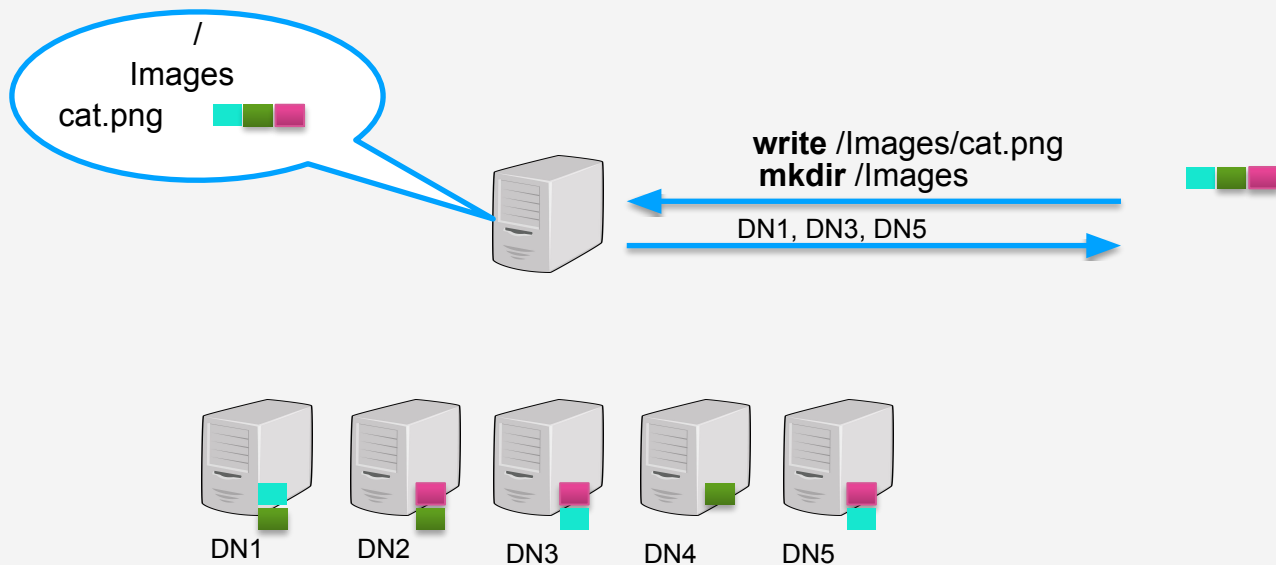
ePipe: Near Real-Time Polyglot Persistence of HopsFS Metadata

Mahmoud Ismail¹, Mikael Ronström², Seif Haridi¹, Jim Dowling¹

¹KTH - Royal Institute of Technology ²Oracle




- Highly scalable next-generation distribution of **HDFS**



What is HopsFS?



Metadata storage



inodeID	name	parentID	
1	/	0	
2	Images	1	
3	cat.png	2	



write /images/cat.png



Block storage (Datenodes)






- Drop-in replacement distribution of HDFS
- **16X - 37X** the throughput of HDFS
- **37** larger clusters than HDFS
- **10** times lower latency




ePipe: Near Real-Time Polyglot Persistence of HopsFS Metadata



Full-text search is not supported by RonDB

Polyglot Persistence - Replicating Metadata to External Systems for Efficient Querying



ePipe: Near Real-Time ~~Polyglot Persistence~~ of HopsFS Metadata




- ePipe is a databus that provides replicated metadata as a service for HopsFS
- ePipe internally
 - creates a consistent and correctly ordered change stream for HopsFS metadata
 - and eventually delivers the change stream with low latency (sub second) (**Near Real-time**) to consumers



- Extend HopsFS with a logging table to log file system changes
- Leverage the RonDB event API to stream changes on the logging table to ePipe
- ePipe enriches the file system events with appropriate data and publish the enriched events to the consumers

Inodes table and logging table updated in the same Transaction to ensure Consistency/Integrity



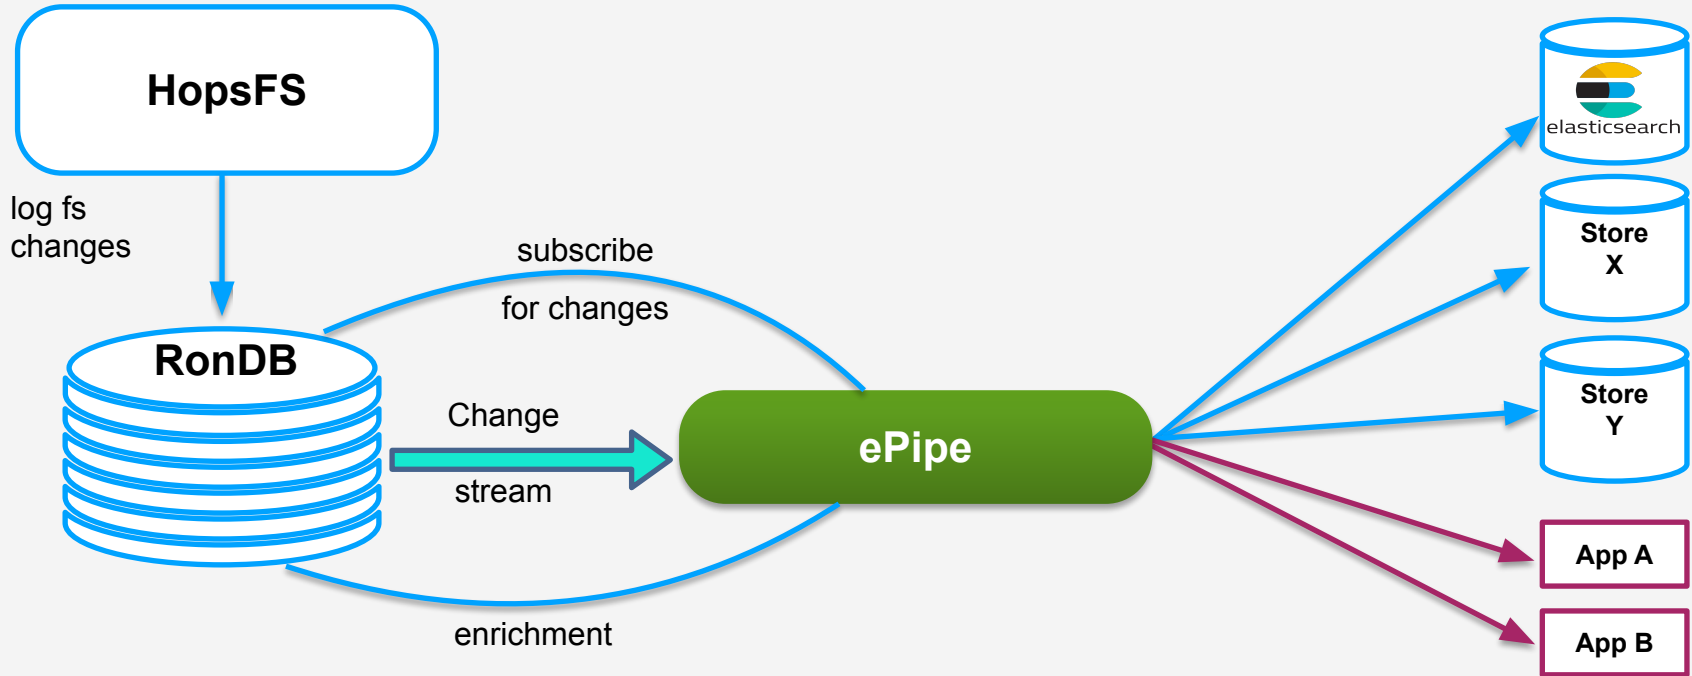
RonDB

Inodes table

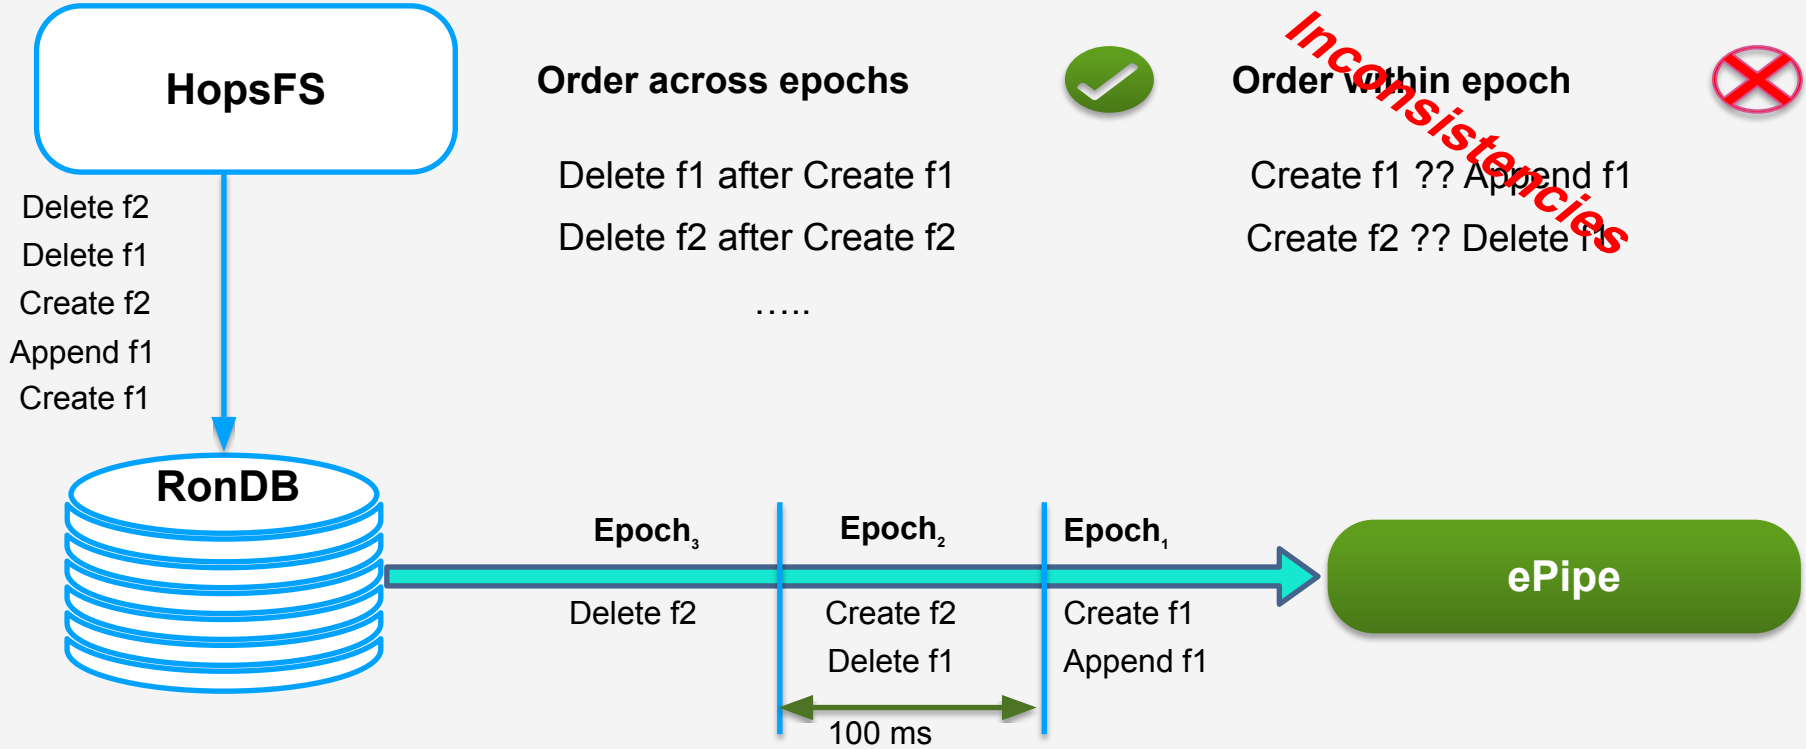
inodeID	name	parentID	
1	/	0	
2	f1	1	
3	f2	1	

logging table

name	operation
f1	CREATE
f2	CREATE
f2	DELETE
f1	DELETE



Ordering of Log Entries





- **Property 1:** The epochs are totally ordered.
- **Property 2:** The changes within the same transaction happen in the same epoch.
- **Property 3:** The changes on files are ordered only if they are in different epochs, that is, no ordering is guaranteed within the same epoch.

Strengthening NDB Ordering Properties



HopsFS


Delete f2 ,2
Delete f1 ,3
Create f2 ,1
Append f1 ,2
Create f1 ,1

We introduced a version number per inode which we will increment whenever a change occurs to an inode.

Append f1 after Create f1
Create f2 ?? Delete f1



RonDB

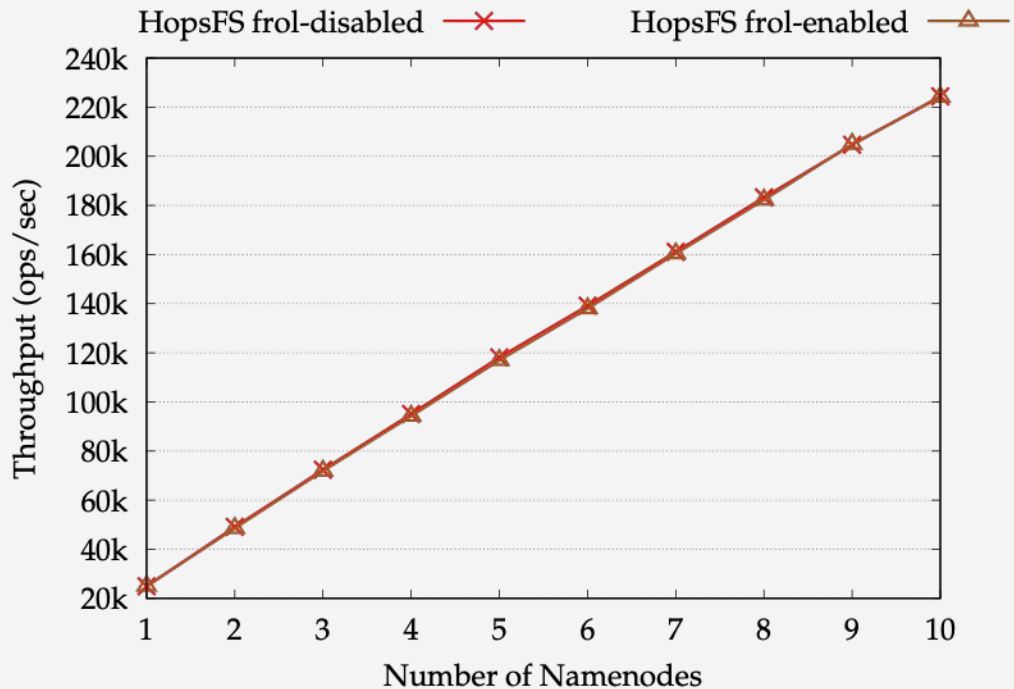


ePipe

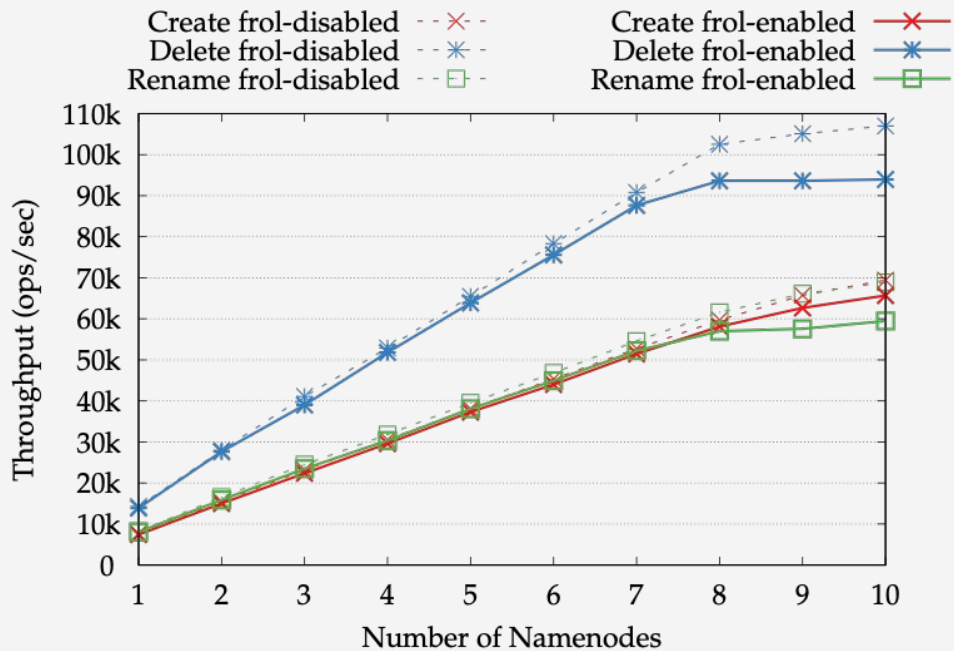


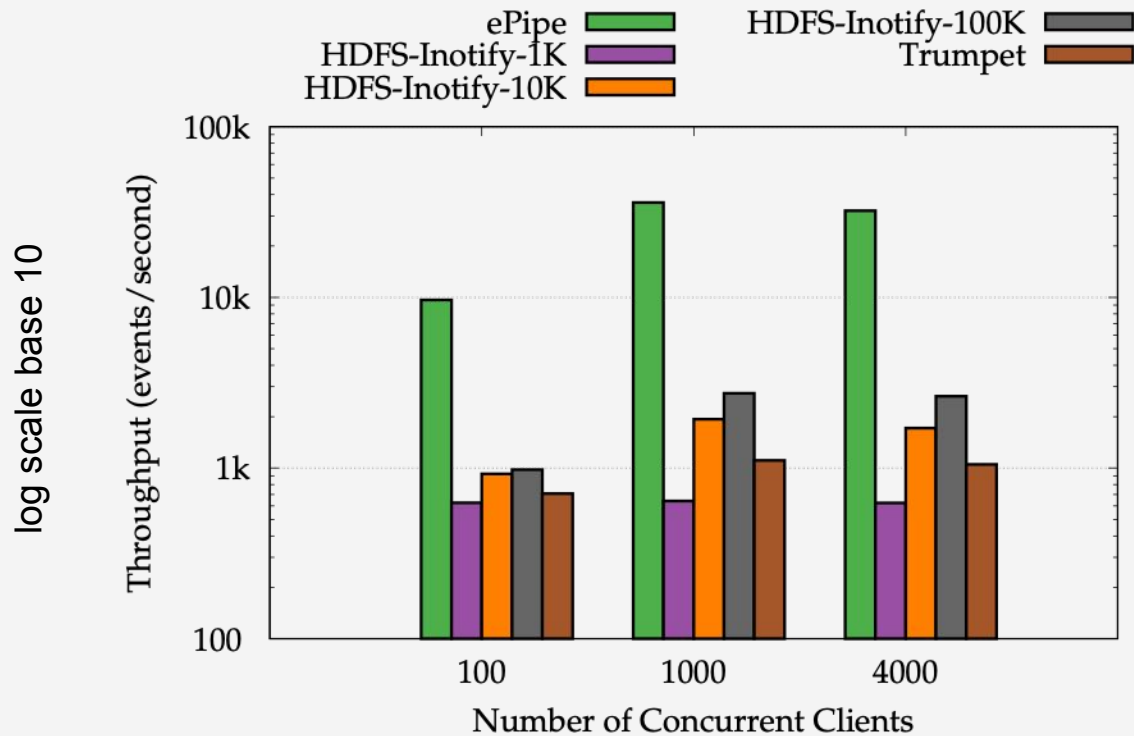
- **Property 1 & 2 & 3**
- **Property 4 & 5:** The version number ensures the serializability of the changes on the same file/directory within epochs.
- **Property 6:** The order of changes for different files/directories within the same epoch doesn't matter.

Logging overhead on HopsFS



Logging overhead on HopsFS






Latency: average Lag Time



log scale base 10



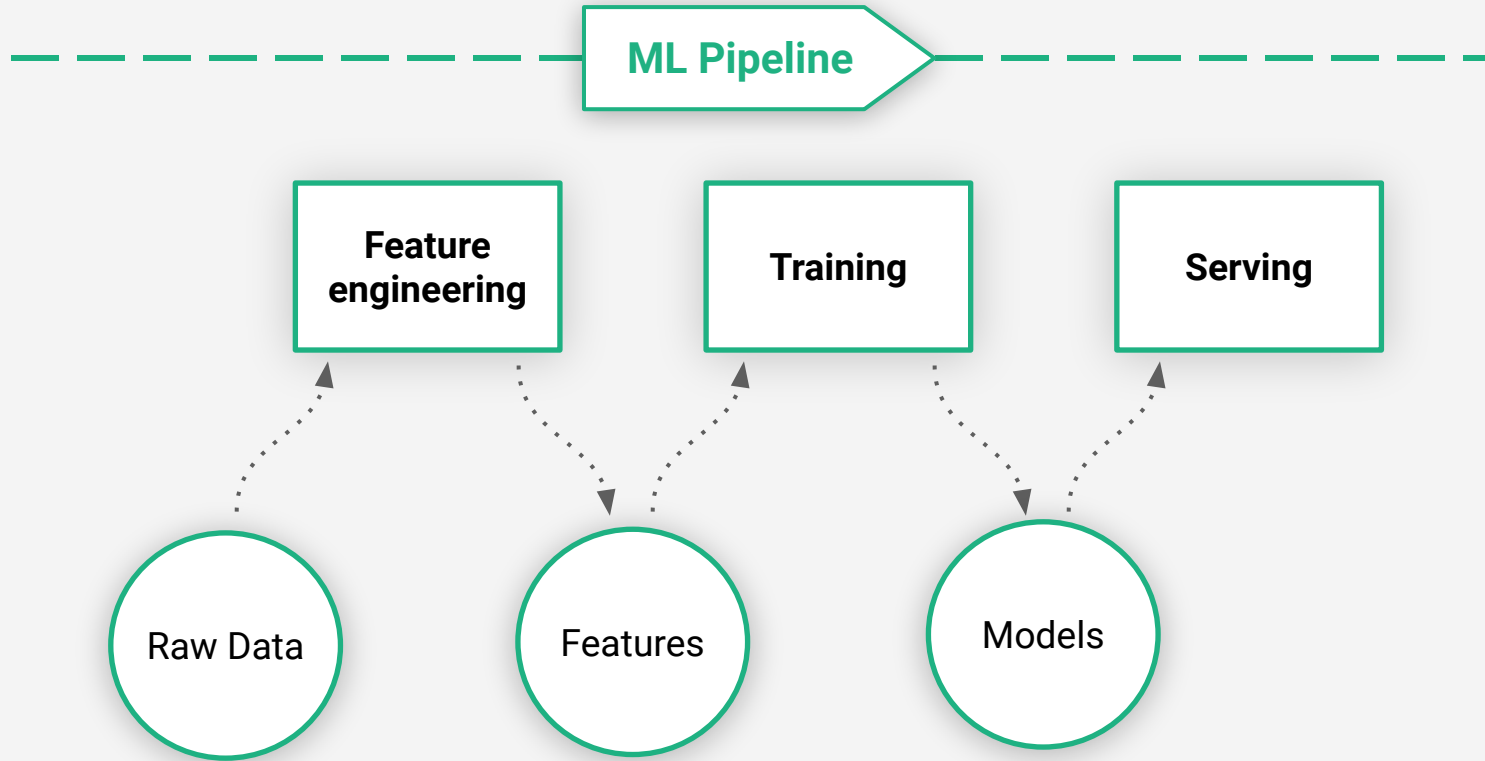


- Supports failure recovery thanks to the persistent logging table
 - The log entries are deleted only once the associated events are successfully replicated to the downstream consumers.
 - At least once delivery semantics.
- Pluggable architecture
 - For example, filter events based on file name or any other attribute.
- Not Limited to HopsFS
 - Can be extended to watch for other logging tables for different purposes.



- A databus that provides replicated metadata as a service for HopsFS
- Low overhead on HopsFS
- Low replication lag (sub-second)
- High throughput
- Pluggable architecture

What is provenance - ML Pipeline





- Provenance improves understanding of complex ML Pipelines.
- Provenance should not change the core ML pipeline code.
- Provenance facilitates Debugging, Analyzing, Automating and Cleaning of ML Pipelines.
- Provenance and Time Travel facilitate reproducibility of experiments.
- In Hopsworks, we introduced a new mechanism for provenance based on embedded metadata in a scale-out consistent metadata layer.



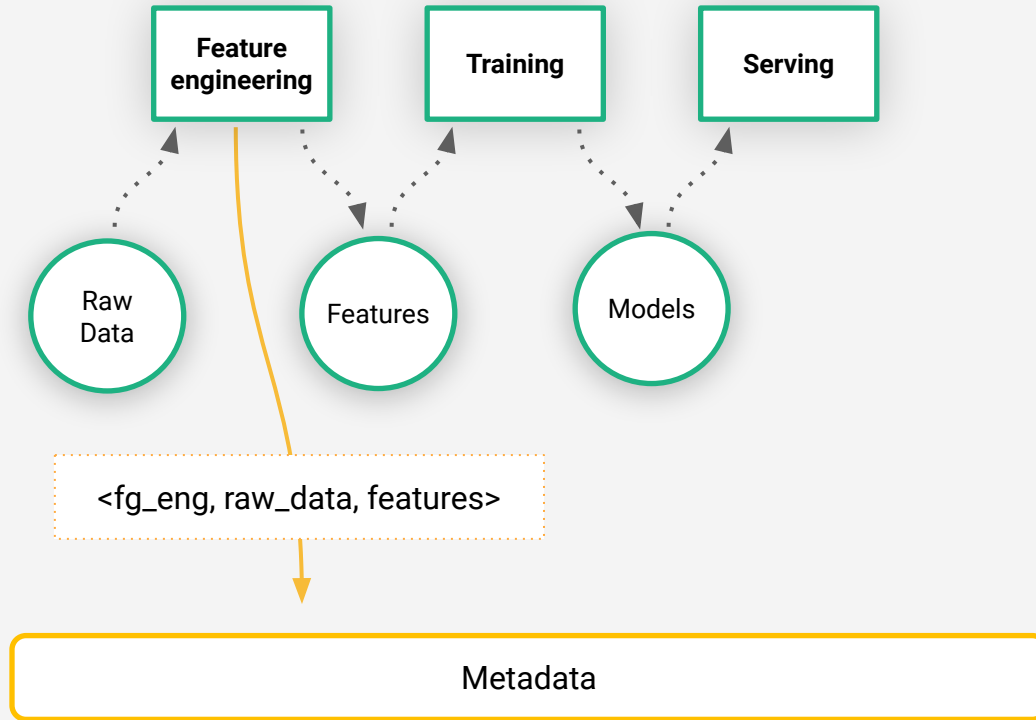
```
def train(data_path, max_depth, min_child_weight, estimators, model_name):
    X_train, X_test, y_train, y_test = build_data(..)
    mlflow.set_tracking_uri("jdbc:mysql://username:password@host:3306/database")
    mlflow.set_experiment("My Experiment")
    with mlflow.start_run() as run:
        ...
        mlflow.log_param("max_depth", max_depth)
        mlflow.log_param("min_child_weight", min_child_weight)
        mlflow.log_param("estimators", estimators)
        with open("test.txt", "w") as f:
            f.write("hello world!")
        mlflow.log_artifacts("/full/path/to/test.txt")
        ...
        model.fit(X_train, y_train) # auto-logging
        ...
        mlflow.tensorflow.log_model(model, "tensorflow-model",
            registered_model_name=model_name)
```



```
def train(data_path, max_depth, min_child_weight, estimators):
    X_train, X_test, y_train, y_test = build_data(..)
    ...
    print("hello world") # monkeypatched - prints in notebook
    ...
    model.fit(X_train, y_train) # auto-logging
    ...
    #Saves model to "hopsfs://Projects/myProj/models/.."
    hops.export_model(model, "tensorflow",...,model_name)
    ...
    # maggy makes an API call to track this dict
    return {'accuracy': accuracy, 'loss': loss, 'diagram': 'diagram.png'}

from maggy import experiment
experiment.lagom(train, name="My Experiment", ...)
```


What is provenance - Metadata

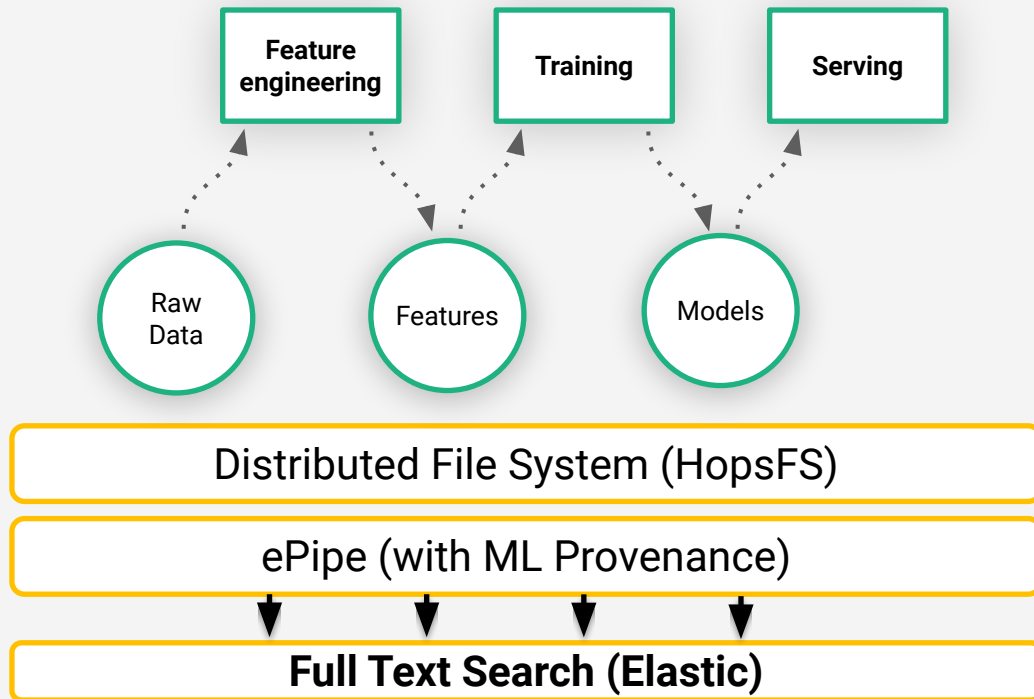


Pipeline code

In []:

```
add(fg_eng, raw_data, features)  
...  
add(training, features, model)
```

Let the platform manage the metadata!






ML Artifacts



Features, Feature Metadata,
Train/Test Datasets
Models, Model Metadata



Possibly thousands of files

Distributed File System



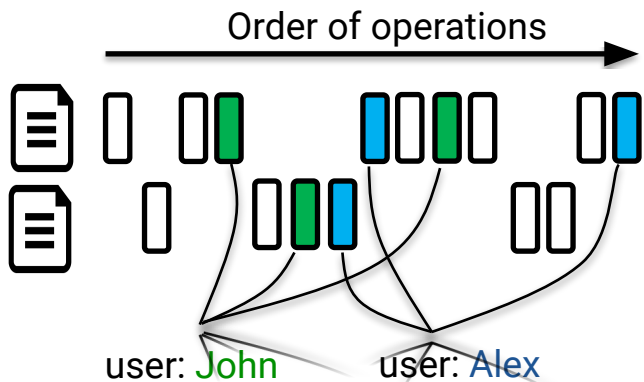
Generate thousands of operations

Change Data Capture (CDC)

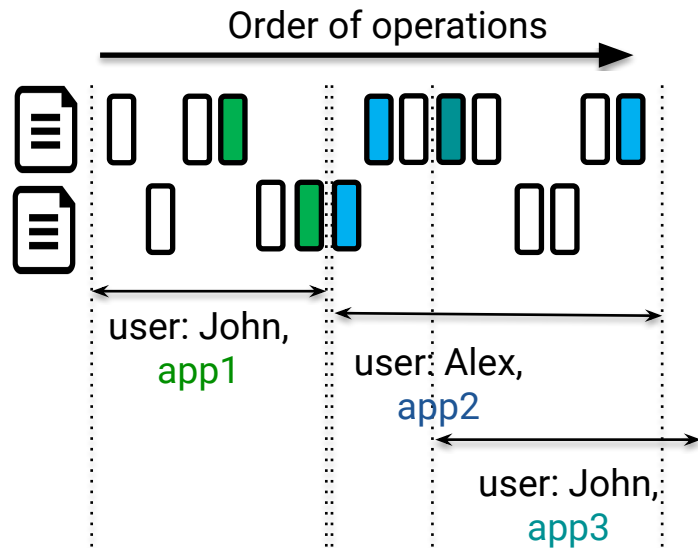
Capture only relevant operations



More context for file system operations?



Are any of these operations related?




Certificates (with AppId) enabled FS Operation



Distributed File System

Read/Write/Create/Delete/XAttr/Metadata



Additional Context

`<file, op, user_id, app_id, job_id, pipeline_id>`

Resource Manager - Yarn (Application Context)

Link input/output files via Apps

Job Manager - Hopsworks (Job Context)


Different Executions of the same Job

Workflow Manager - Airflow (Pipeline Context)

Jobs as Stages of the same Pipeline



Hopsworks Conventions






- **Path based filtering**

Example

Project — /featurestore
 — /training_datasets
 — /models





Path based filtering


- **Tag based filtering**

Example:

Custom metadata based on HDFS XAttr.

Tag: <tutorial>, <debug>

Tags can enable logging of all operations, if path based filtering is not easy to set





Path based filtering

Tag based filtering

- **Coalesce FS Operations**

Example:

Read file₁


Read file₂

...


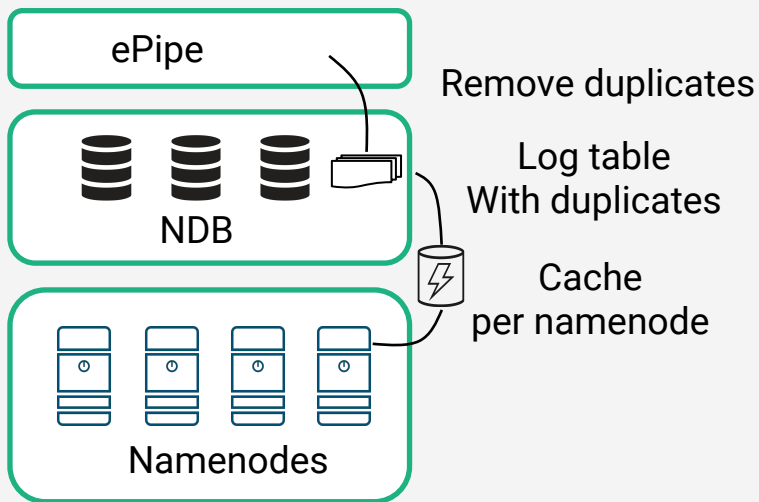
Read file_n



Access₁
Training Dataset



Optimization - FS Operation Coalesce



In []:

```
hops.load_training_dataset(  
    "/Projects/LC/Training_Datasets/ImageNet")  
...  
hops.save_model("/Projects/LC/Models/ResNet")
```

Parent Create	Artifact Create
Parent Delete	Artifact Delete
Children Read	Artifact Access
Children Create/Delete/Append/Truncate	Artifact Mutation



Path based filtering

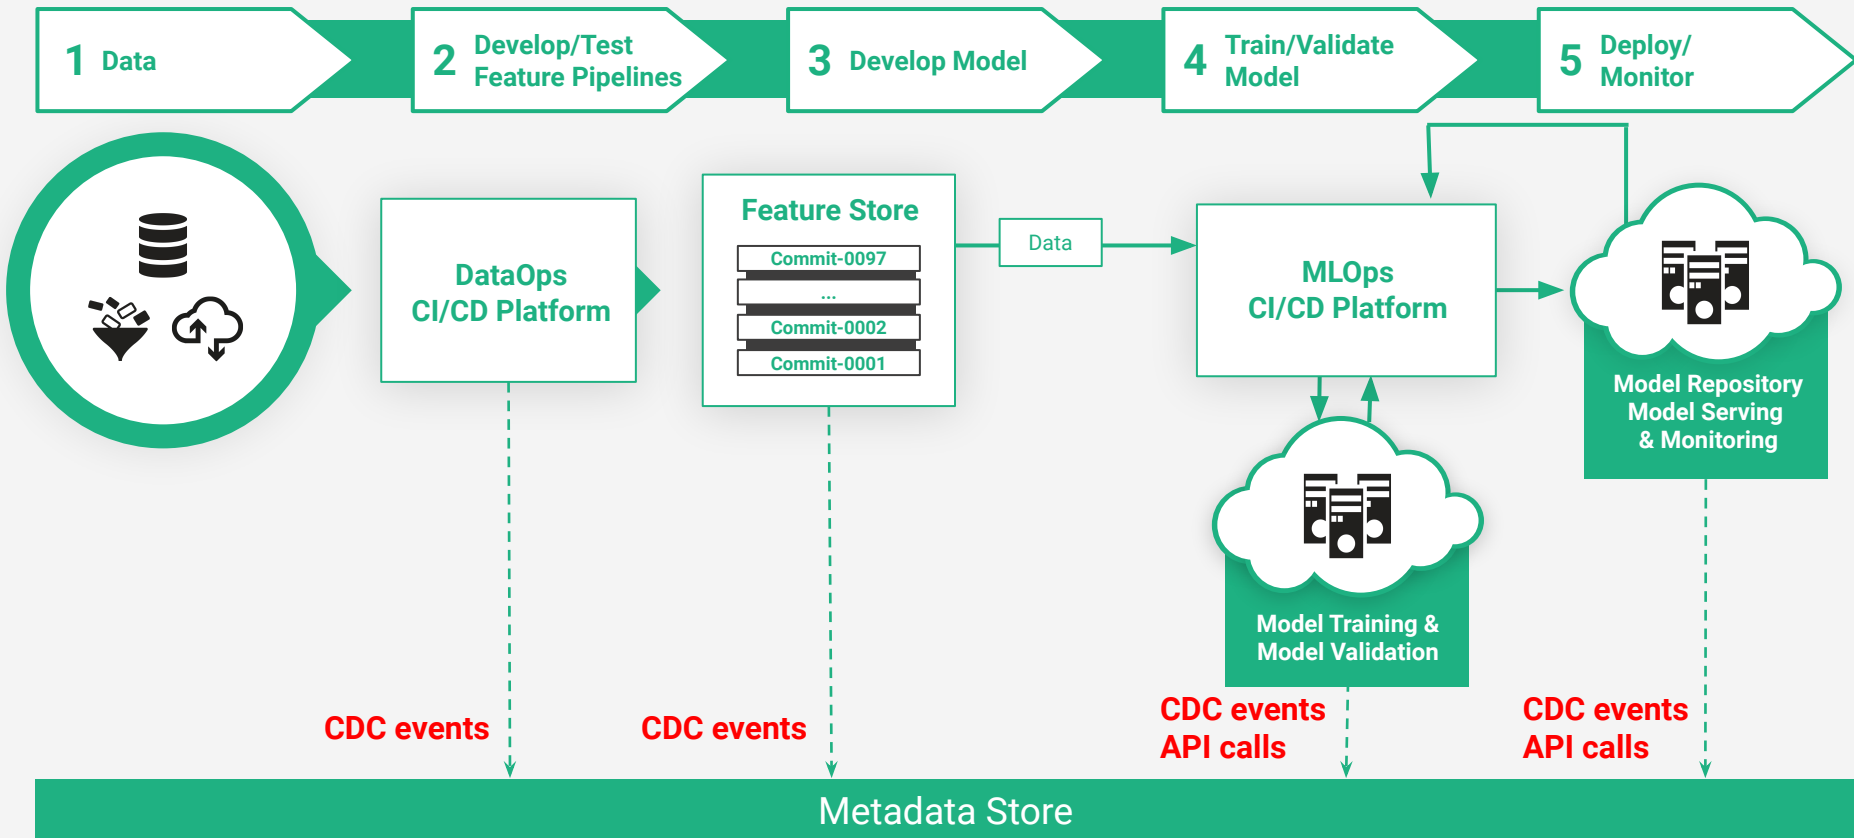
Tag based filtering

Coalesce FS Operations

- **Filtered Operations**

Filesystem Op	Metadata Stored
Create/Delete	Artifact existence
XAttr	Add metadata to artifact
Read	Artifact used by ..
Children Files Create/Delete	Artifact mutation
Append/Truncate	Artifact mutation
Permissions/ACL	Artifact metadata mutation

Hopsworks ML Pipelines





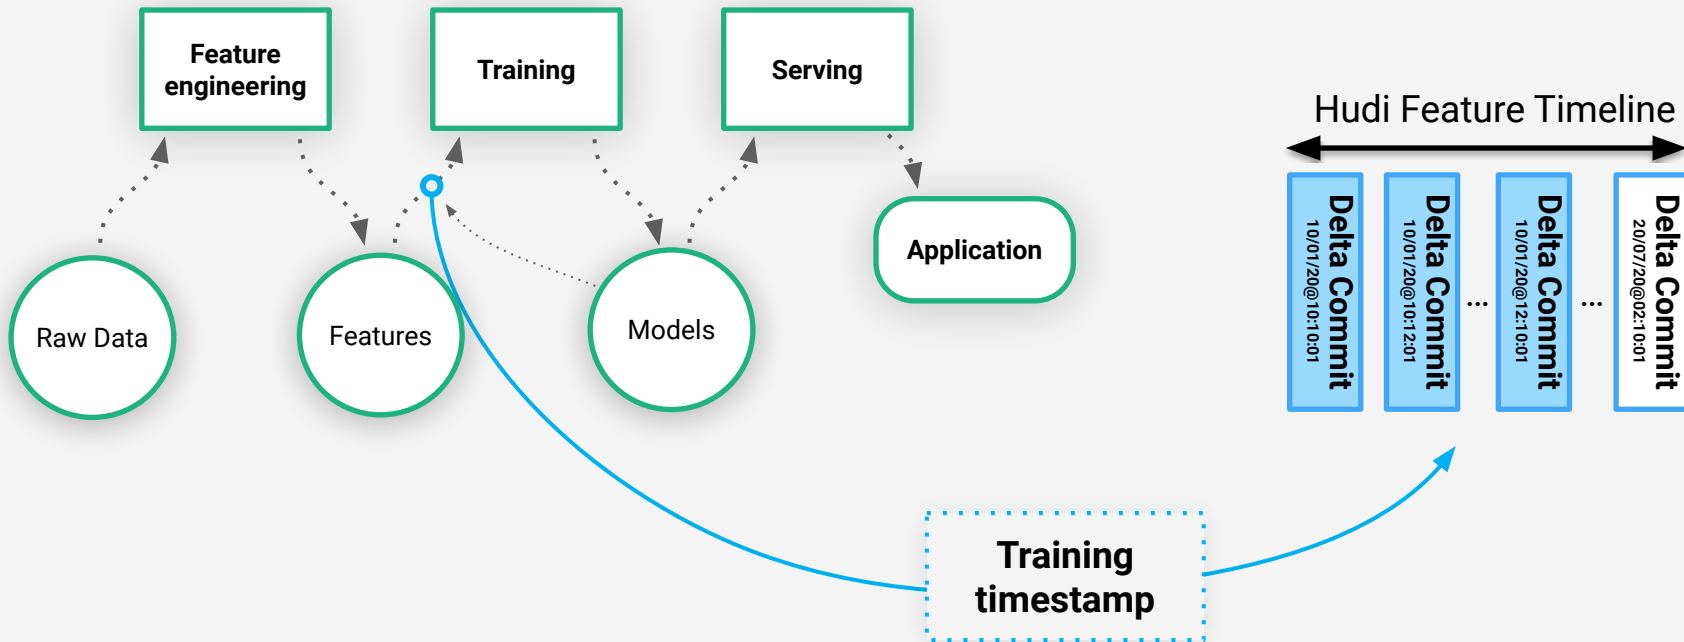
Bias Detected !

What do I do ?



Claim of Model Bias!

Can we determine the exact features used?





HOPSWORKS.ai

BY LOGICAL CLOCKS



@hopsworks



<http://github.com/logicalclocks/hopsworks>



- [Ormenisan et al, Time-travel and Provenance for ML Pipelines, Usenix OpML 2020](#)
- [Niazi et al, HopsFS, Usenix Fast 2017](#)
- [Ismail et al, ePipe, CCGrid 2019](#)
- [Small Files in HopsFS, ACM Middleware 2018](#)
- [Ismail et al, HopsFS-S3, ACM Middleware 2020](#)
- [Meister et al, Oblivious Training Functions, 2020](#)
- [Hopsworks](#)