SPARK IN PRODUCTION: POWERFUL BUT PERILOUS

Aleš Kotnik, 25th of April 2024

WHAT IS APACHE SPARK?

Apache Spark™ is a multi-language engine for executing data engineering, data science, and machine learning on single-node machines or clusters.

HISTORY

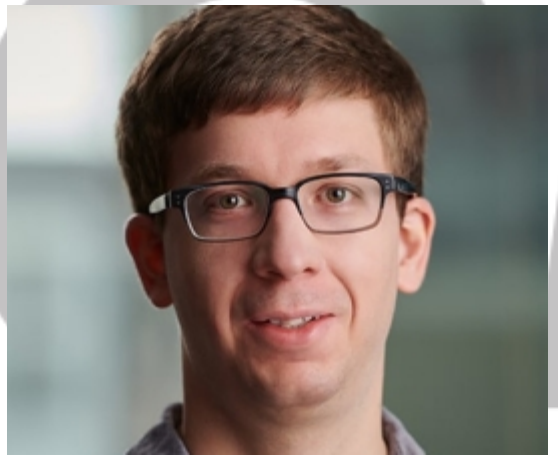

HISTORY

Developed by Matei Zaharia in 2009.

HISTORY

Developed by Matei Zaharia in 2009.

Open-sourced in 2010.

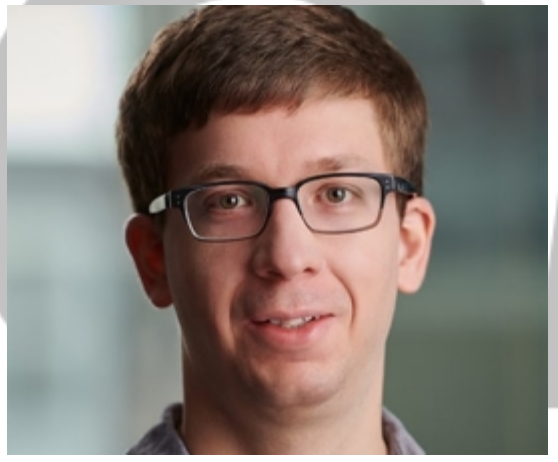

HISTORY

Developed by Matei Zaharia in 2009.

Open-sourced in 2010.

Matei is now Chief Technologist at Databricks

HISTORY

Developed by Matei Zaharia in 2009.

Open-sourced in 2010.

Matei is now Chief Technologist at Databricks

Spark has more than 2,000 contributors

POPULAR DEPLOYMENTS OPTIONS

POPULAR DEPLOYMENTS OPTIONS

- Databricks

POPULAR DEPLOYMENTS OPTIONS

- Databricks
- AWS EMR

POPULAR DEPLOYMENTS OPTIONS

- Databricks
- AWS EMR
- Spark operator for Kubernetes

SPARK OPERATOR

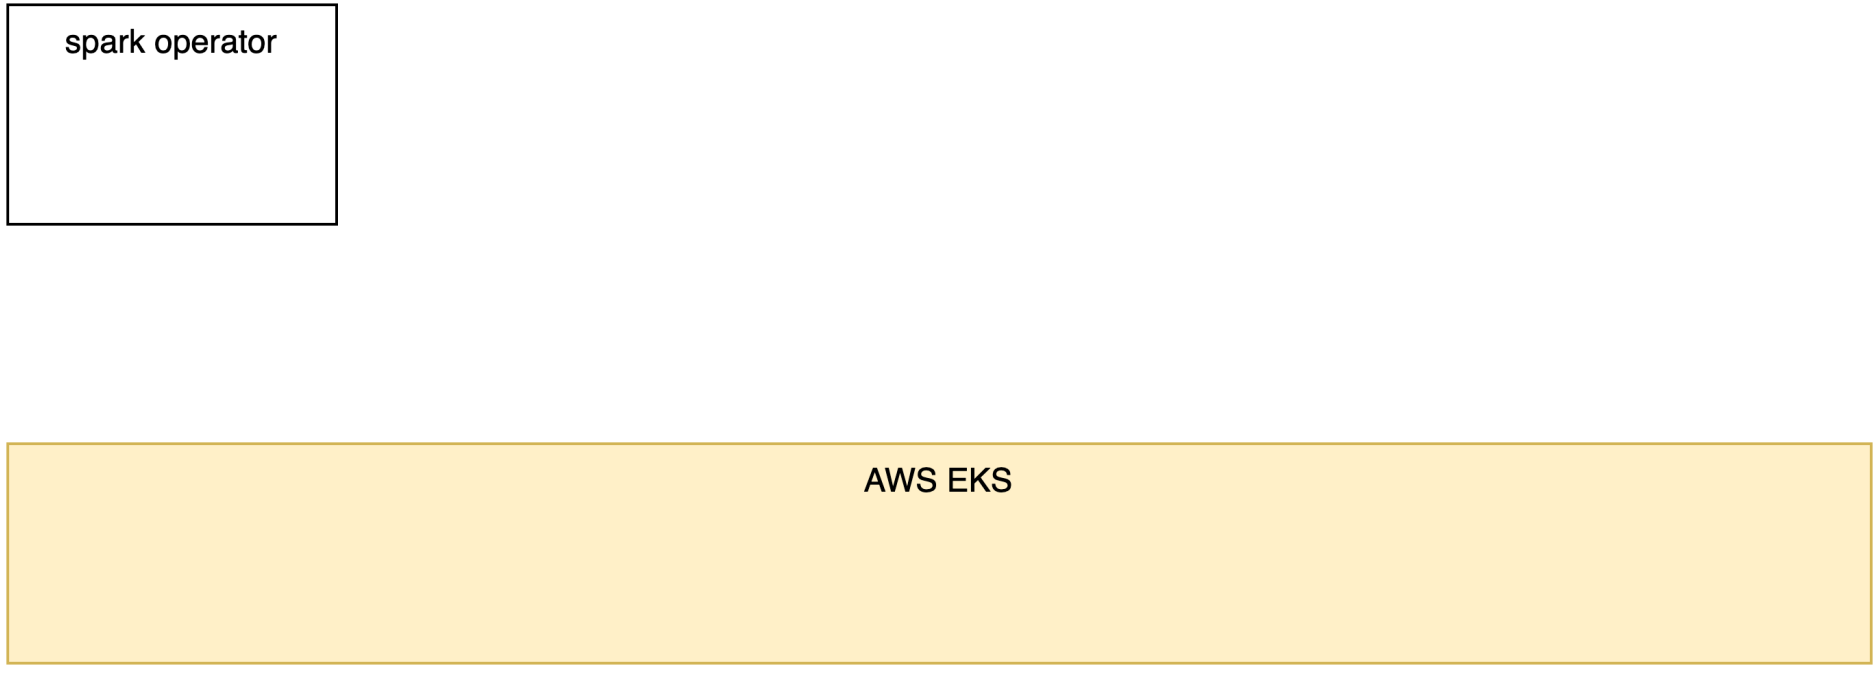

SPARK APPLICATION

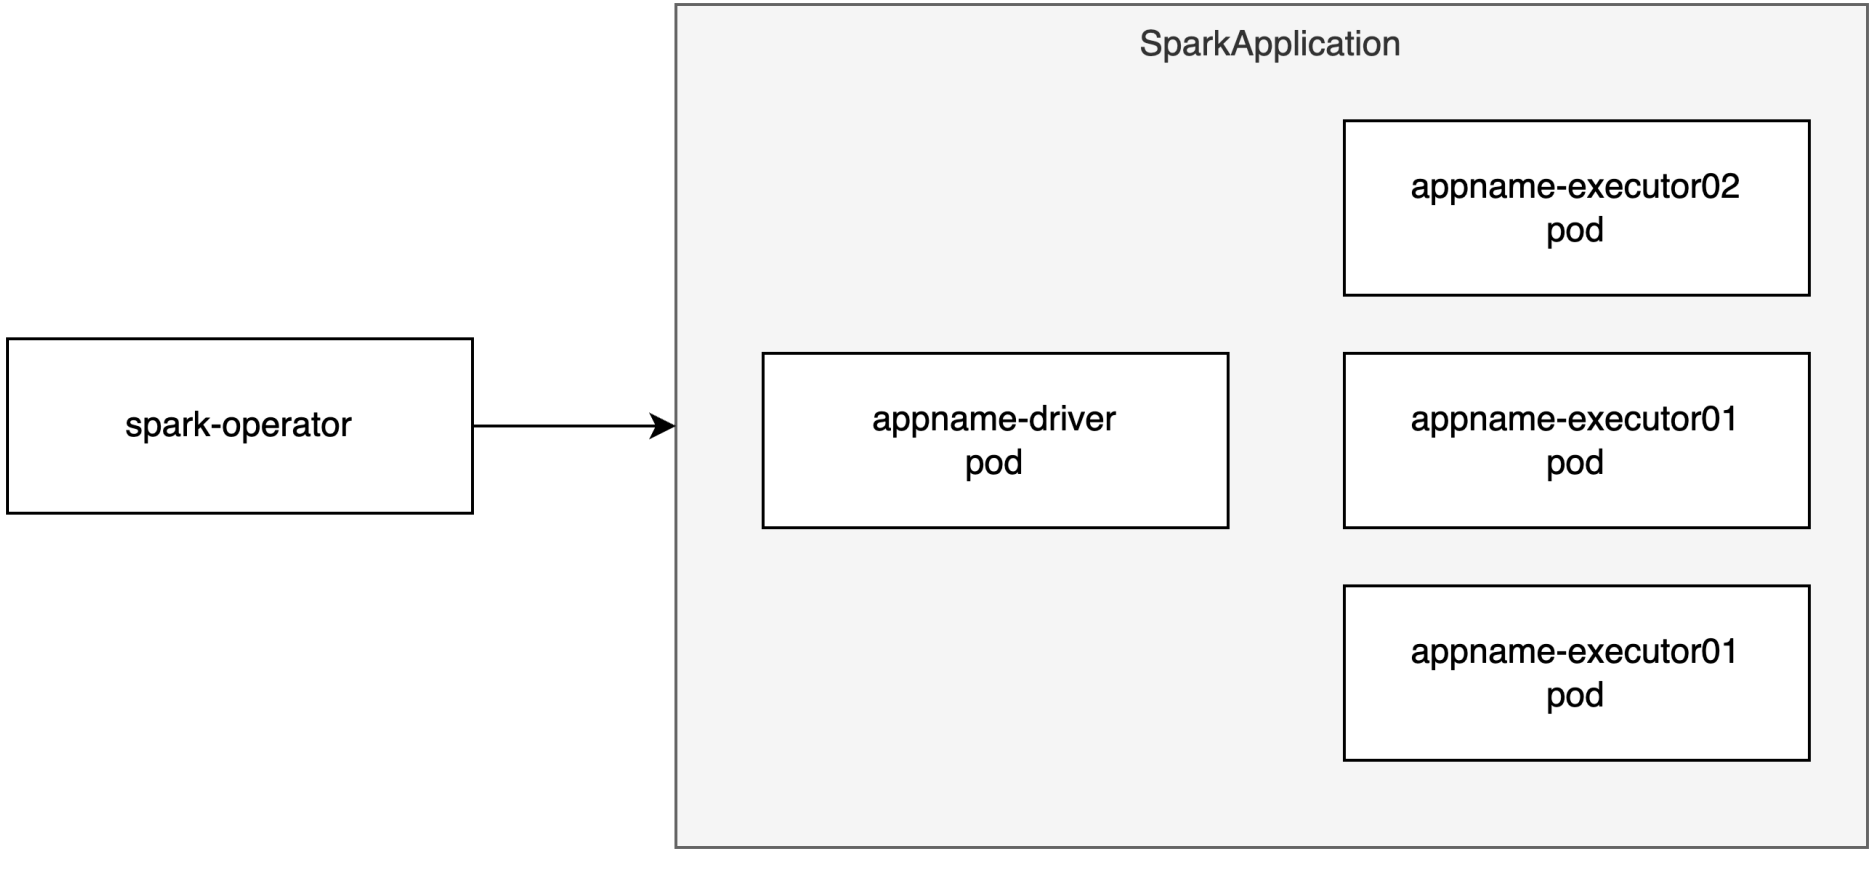

APACHE ICEBERG

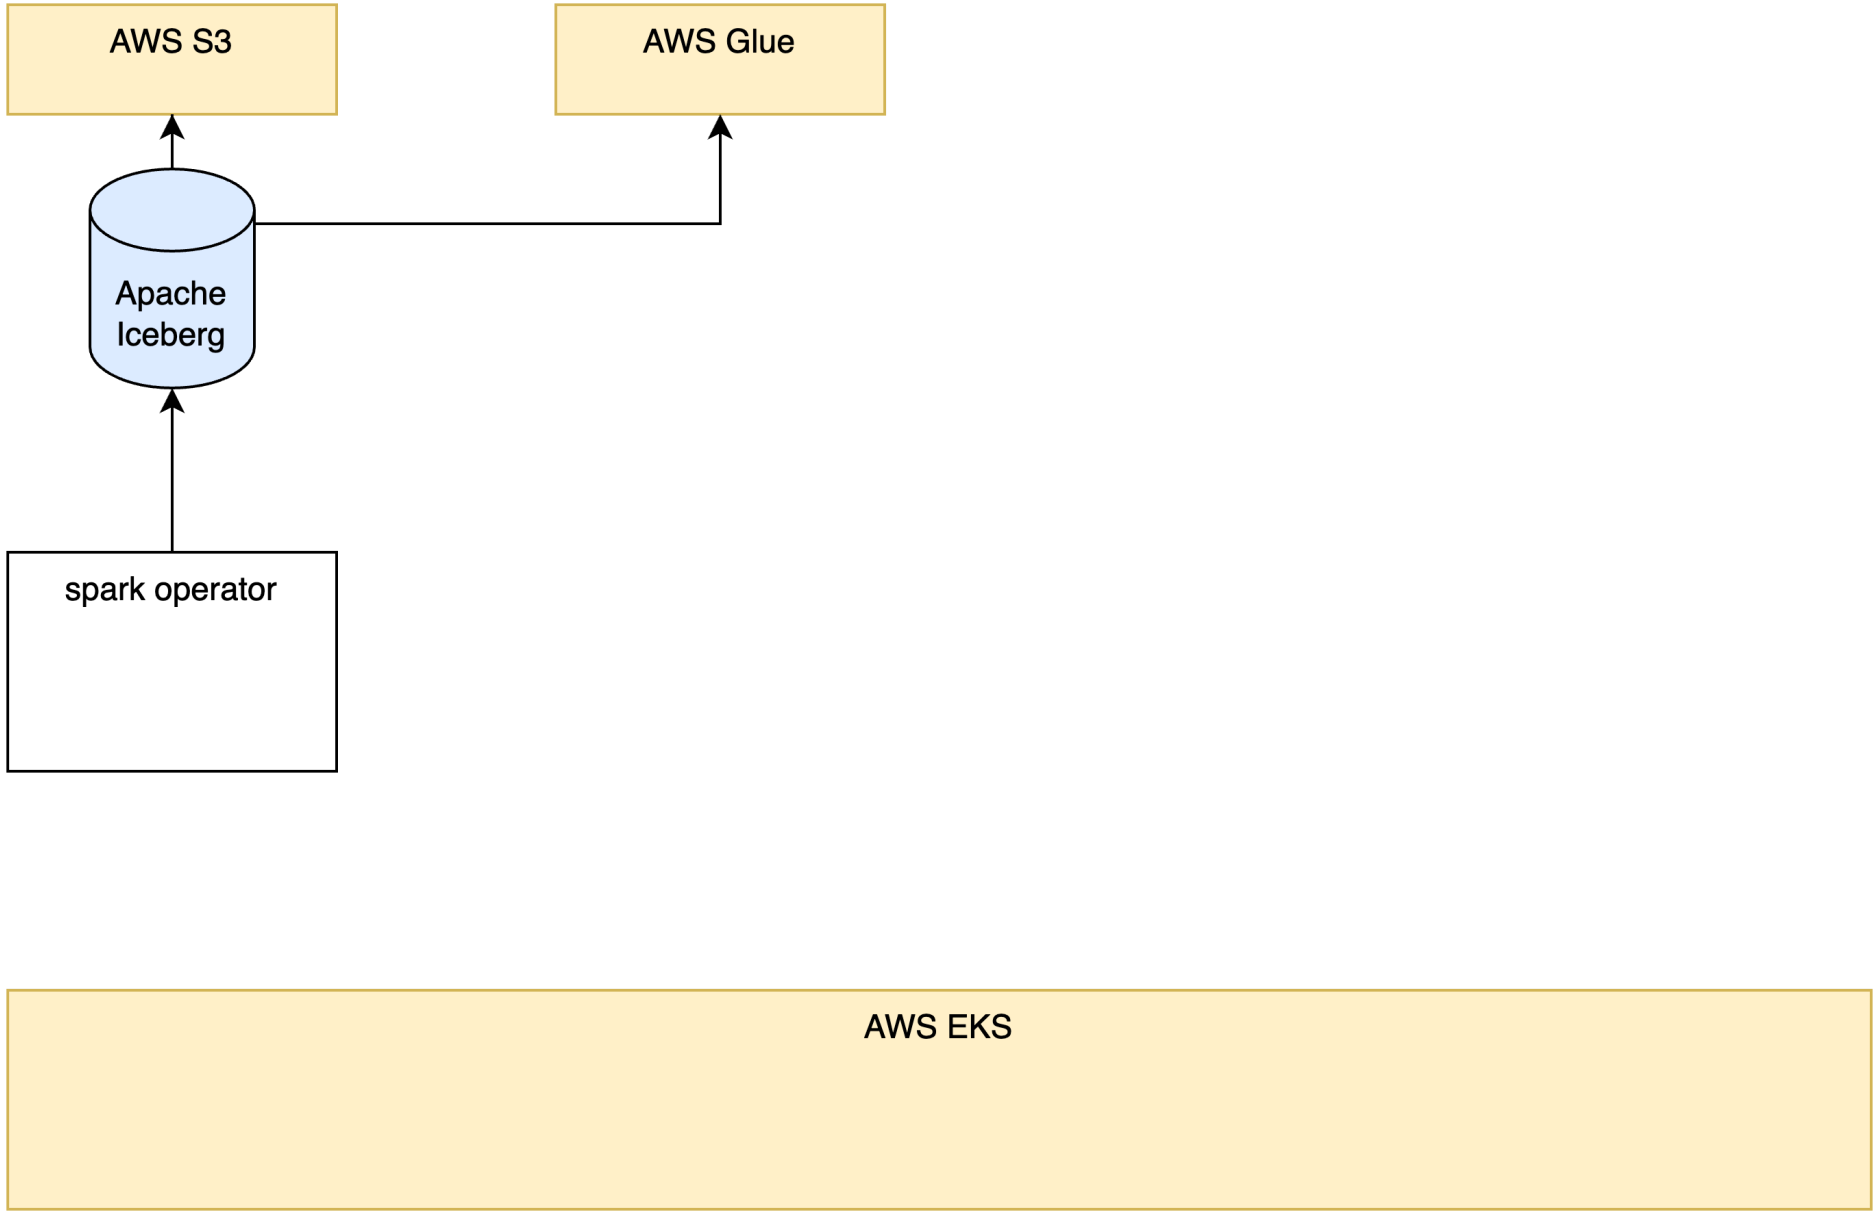
AWS S3
AWS Glue
Apache
Iceberg
spark operator
AWS EKS

APACHE ICEBERG

is a high-performance format for huge analytic tables

APACHE ICEBERG

APACHE ICEBERG

Decouples storage (uses Parquet for data and stores metadata on S3).

APACHE ICEBERG

Decouples storage (uses Parquet for data and stores metadata on S3).

Integrates with Spark, Presto (Athena), Clickhouse, ...

APACHE ICEBERG

Decouples storage (uses Parquet for data and stores metadata on S3).

Integrates with Spark, Presto (Athena), Clickhouse, ...

Schema evolution

APACHE ICEBERG

Decouples storage (uses Parquet for data and stores metadata on S3).

Integrates with Spark, Presto (Athena), Clickhouse, ...

Schema evolution

Time travel

JUPYTER NOTEBOOK

jupyter Untitled (unsaved changes)

Logout

File Edit View Insert Cell Kernel Help Trusted Python 3 (ipykernel)

Run Code

In [2]:
```
# Import Libraries
from pyspark.sql.types import StructType, StructField, FloatType, BooleanType
from pyspark.sql.types import DoubleType, IntegerType, StringType
import pyspark
from pyspark import SQLContext

# Setup the Configuration
conf = pyspark.SparkConf()

spark_context = SparkSession.builder.config(conf=conf).getOrCreate()
sqlcontext = SQLContext(sc)

# Setup the Schema
schema = StructType([
            StructField("User ID", IntegerType(),True),
            StructField("Username", StringType(),True),
            StructField("Browser", StringType(),True),
            StructField("OS", StringType(),True),
        ])

# Add Data
data = ([(1580, "Barry", "FireFox", "Windows" ),
         (5820, "Sam", "MS Edge", "Linux"),
         (2340, "Harry", "Vivaldi", "Windows"),
         (7860, "Albert", "Chrome", "Windows"),
         (1123, "May", "Safari", "macOS")
        ])

# Setup the Data Frame
user_data_df = sqlcontext.createDataFrame(data,schema=schema)
```

In [4]:
```
user_data_df.show()
```

```
+-------+--------+-------+-------+
|User ID|Username|Browser|     OS|
+-------+--------+-------+-------+
|   1580|   Barry|FireFox|Windows|
|   5820|     Sam|MS Edge|  Linux|
|   2340|   Harry|Vivaldi|Windows|
|   7860|  Albert| Chrome|Windows|
|   1123|     May| Safari|  macOS|
+-------+--------+-------+-------+
```

In [ ]:

JUPYTER NOTEBOOK

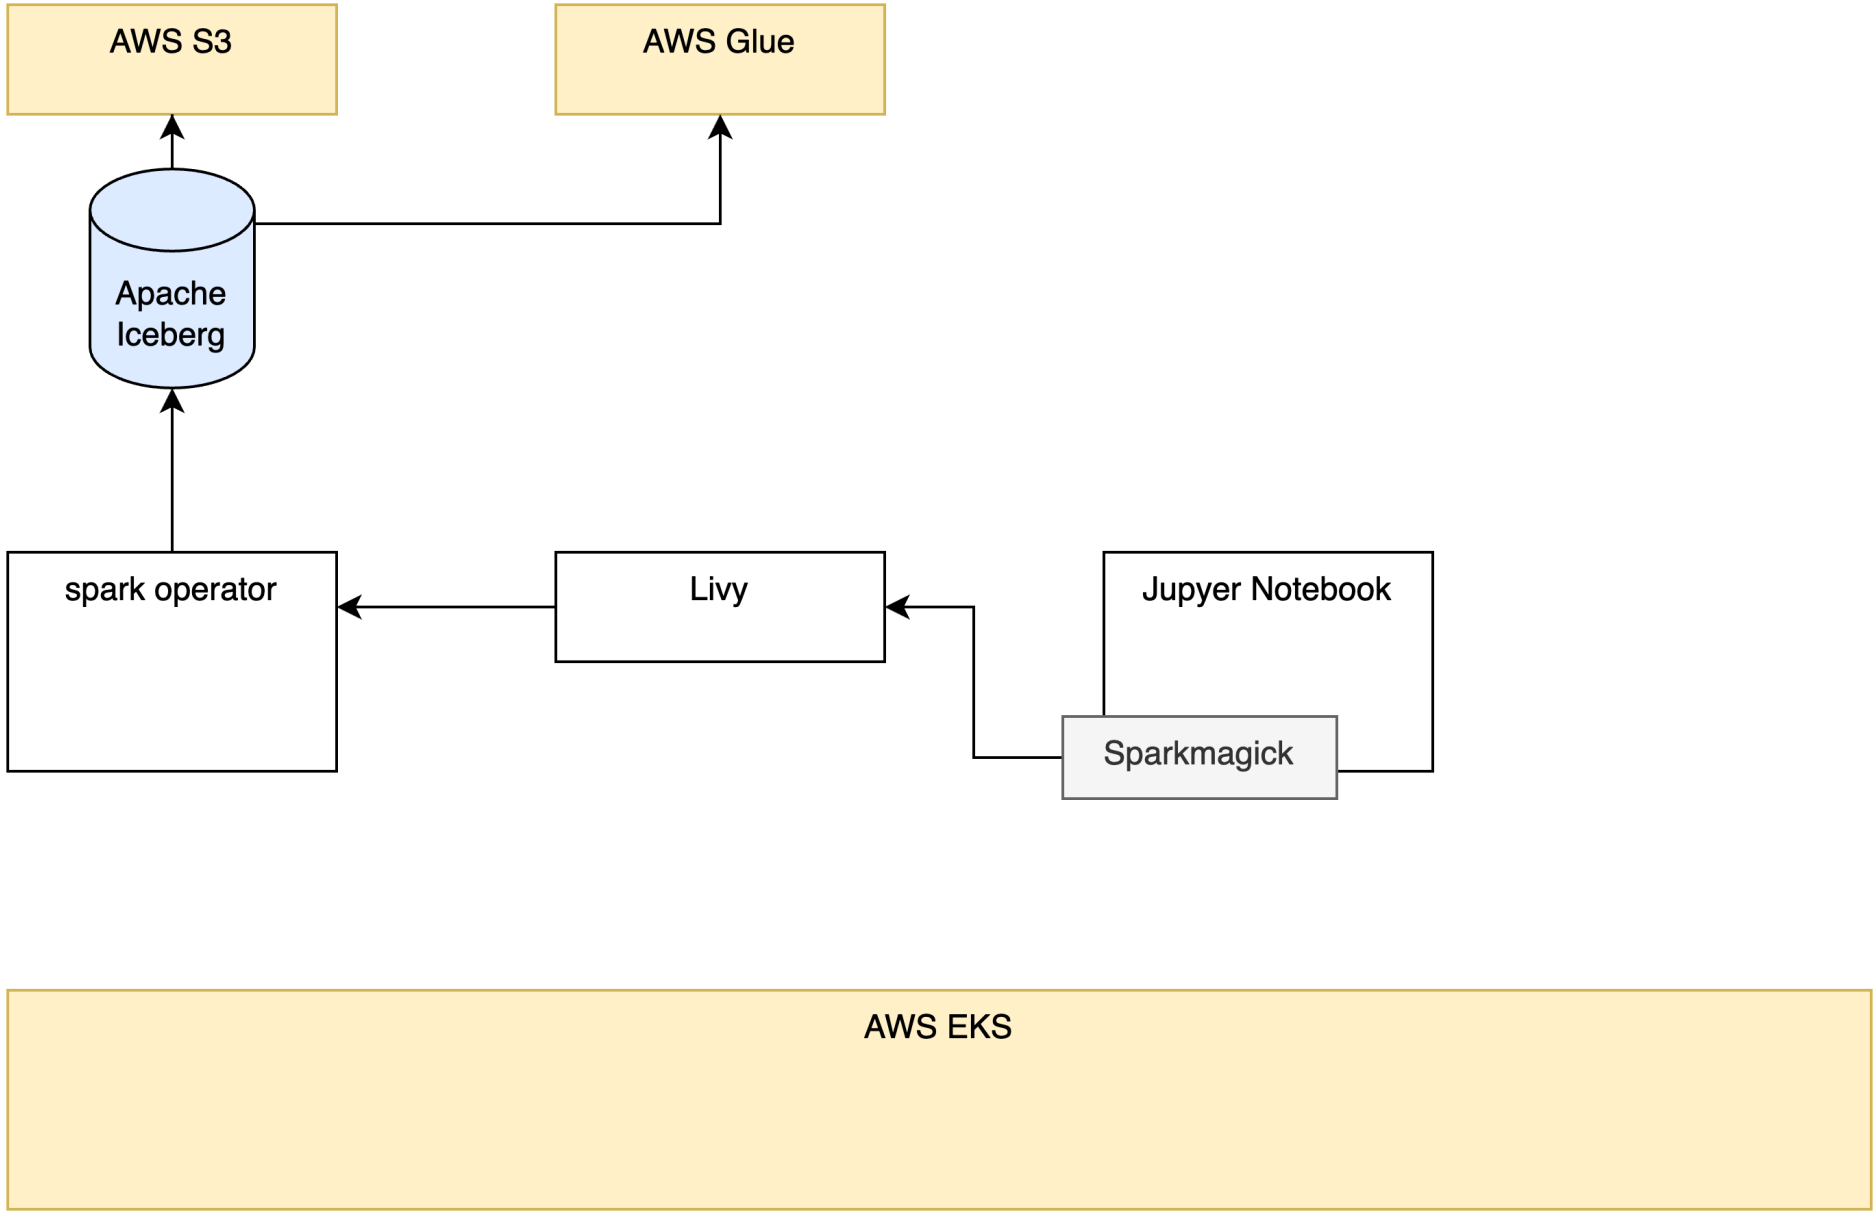
AWS S3
AWS Glue
Apache
Iceberg
spark operator
Livy
Jupyer Notebook
Sparkmagick
AWS EKS

JUPYTER NOTEBOOK

JUPYTER NOTEBOOK

Web-based interface (no local installation needed)

JUPYTER NOTEBOOK

Web-based interface (no local installation needed)

Supports multiple languages, including Python, R and Scala

JUPYTER NOTEBOOK

Web-based interface (no local installation needed)

Supports multiple languages, including Python, R and Scala

Interactive data visualization capabilities

JUPYTER NOTEBOOK

Web-based interface (no local installation needed)

Supports multiple languages, including Python, R and Scala

Interactive data visualization capabilities

Provides a Spark connectivity via SparkMagic and Livy

OTHER CLIENTS

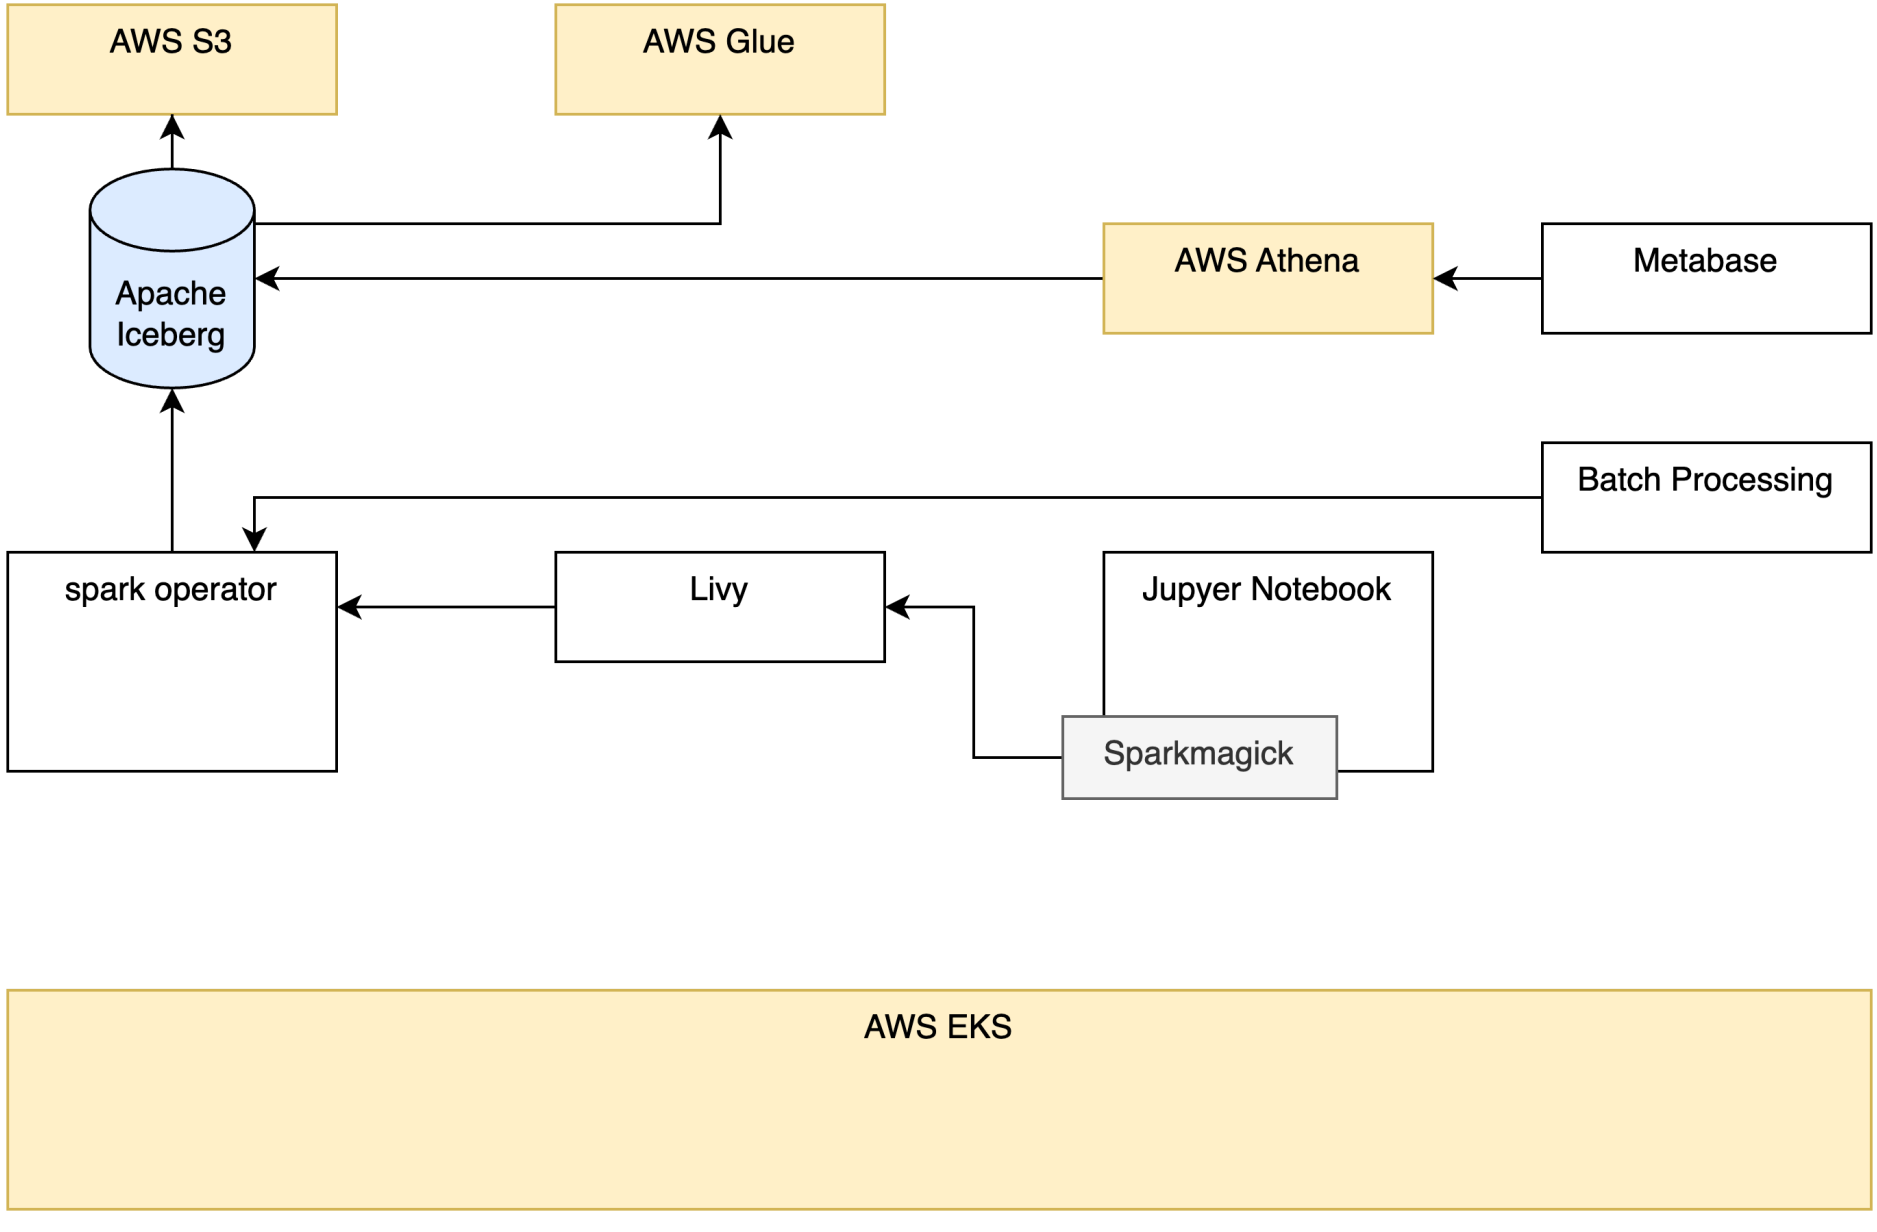
AWS S3
AWS Glue
Apache Iceberg
AWS Athena
Metabase
Batch Processing
spark operator
Livy
Jupyer Notebook
Sparkmagick
AWS EKS

WHAT CAN GO WRONG?

SPARK-OP CAN BRING K8S CLUSTER DOWN

SPARK-OP CAN BRING K8S CLUSTER DOWN

Failure generating cert for Admission Webhook

SPARK-OP CAN BRING K8S CLUSTER DOWN

Failure generating cert for Admission Webhook

Executors lack toleration for spot instance

SPARK-OP CAN BRING K8S CLUSTER DOWN

Failure generating cert for Admission Webhook

Executors lack toleration for spot instance

Starts vicious cycle

SPARK-OP CAN BRING K8S CLUSTER DOWN

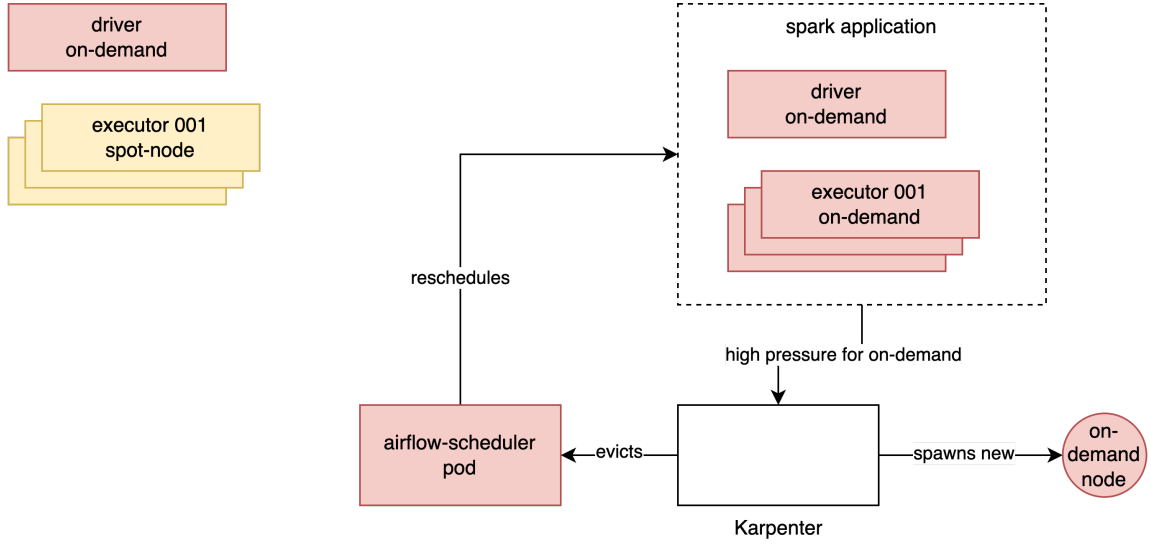

DEVELOPER UNFRIENDLY

DEVELOPER UNFRIENDLY

Spark + Iceberg -> requires cluster

DEVELOPER UNFRIENDLY

Spark + Iceberg -> requires cluster

Need to Build a Docker image to test any code change

DEVELOPER UNFRIENDLY

Spark + Iceberg -> requires cluster

Need to Build a Docker image to test any code change

A lot of time is spent in complex Spark UI

DEVELOPER UNFRIENDLY

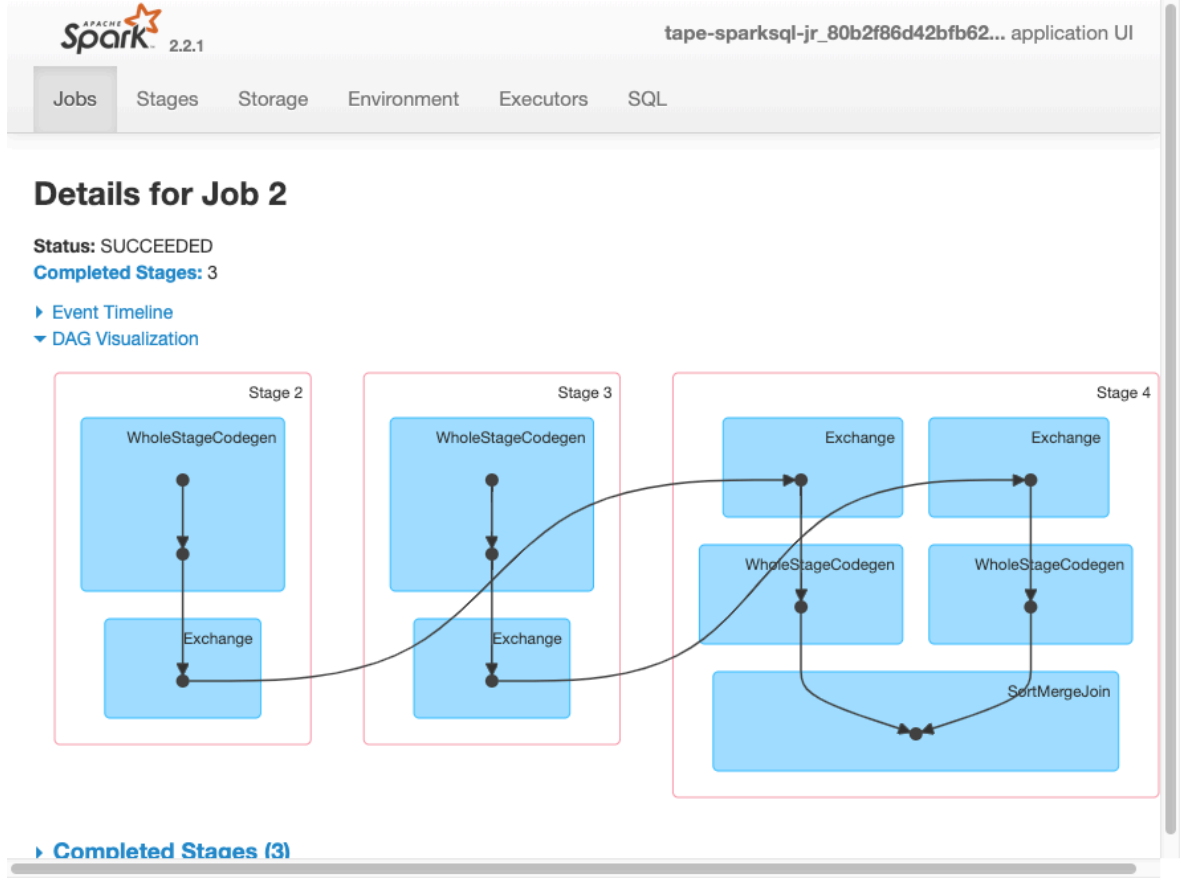

PEOPLE PROBLEM

PEOPLE PROBLEM

Spark is challenging to use

PEOPLE PROBLEM

Spark is challenging to use

When distressed, the clueless begin to randomly tweak parameters (guilty of that, too)

PEOPLE PROBLEM

Spark is challenging to use

When distressed, the clueless begin to randomly tweak parameters (guilty of that, too)

Can burn lots of $$$ and still fail

PEOPLE PROBLEM

Spark is challenging to use

When distressed, the clueless begin to randomly tweak parameters (guilty of that, too)

Can burn lots of $$$ and still fail

Data scientist often feel at home with pandas

BECOMES A HAMMER

BECOMES A HAMMER

How can do X with Spark?

BECOMES A HAMMER

How can do X with Spark?

Raises complexity

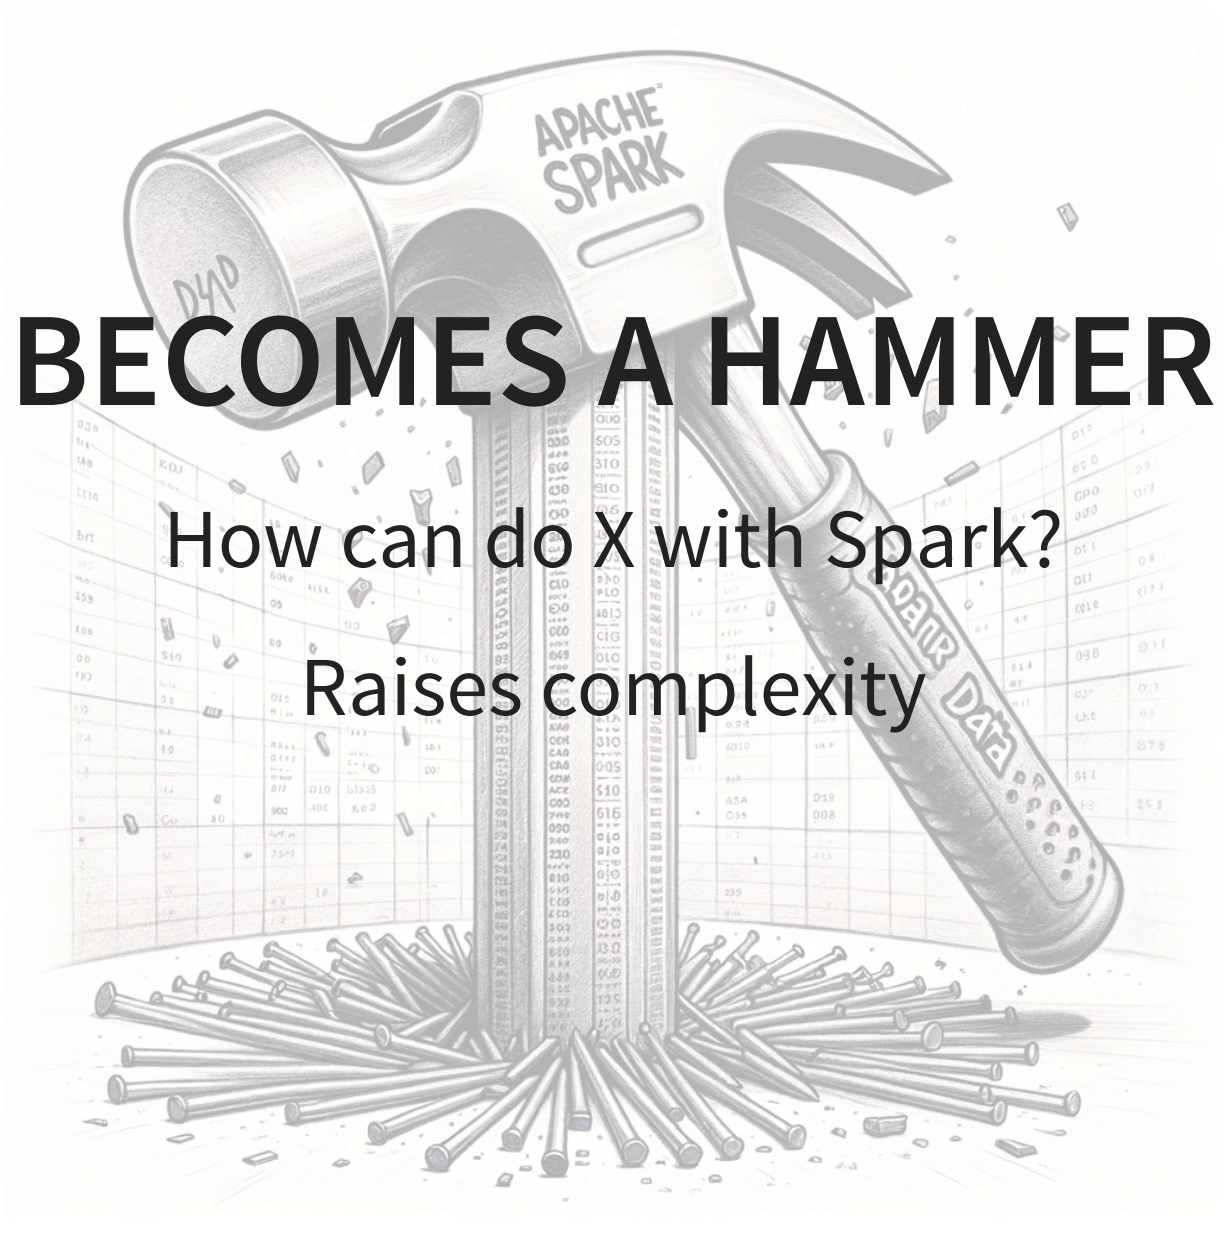

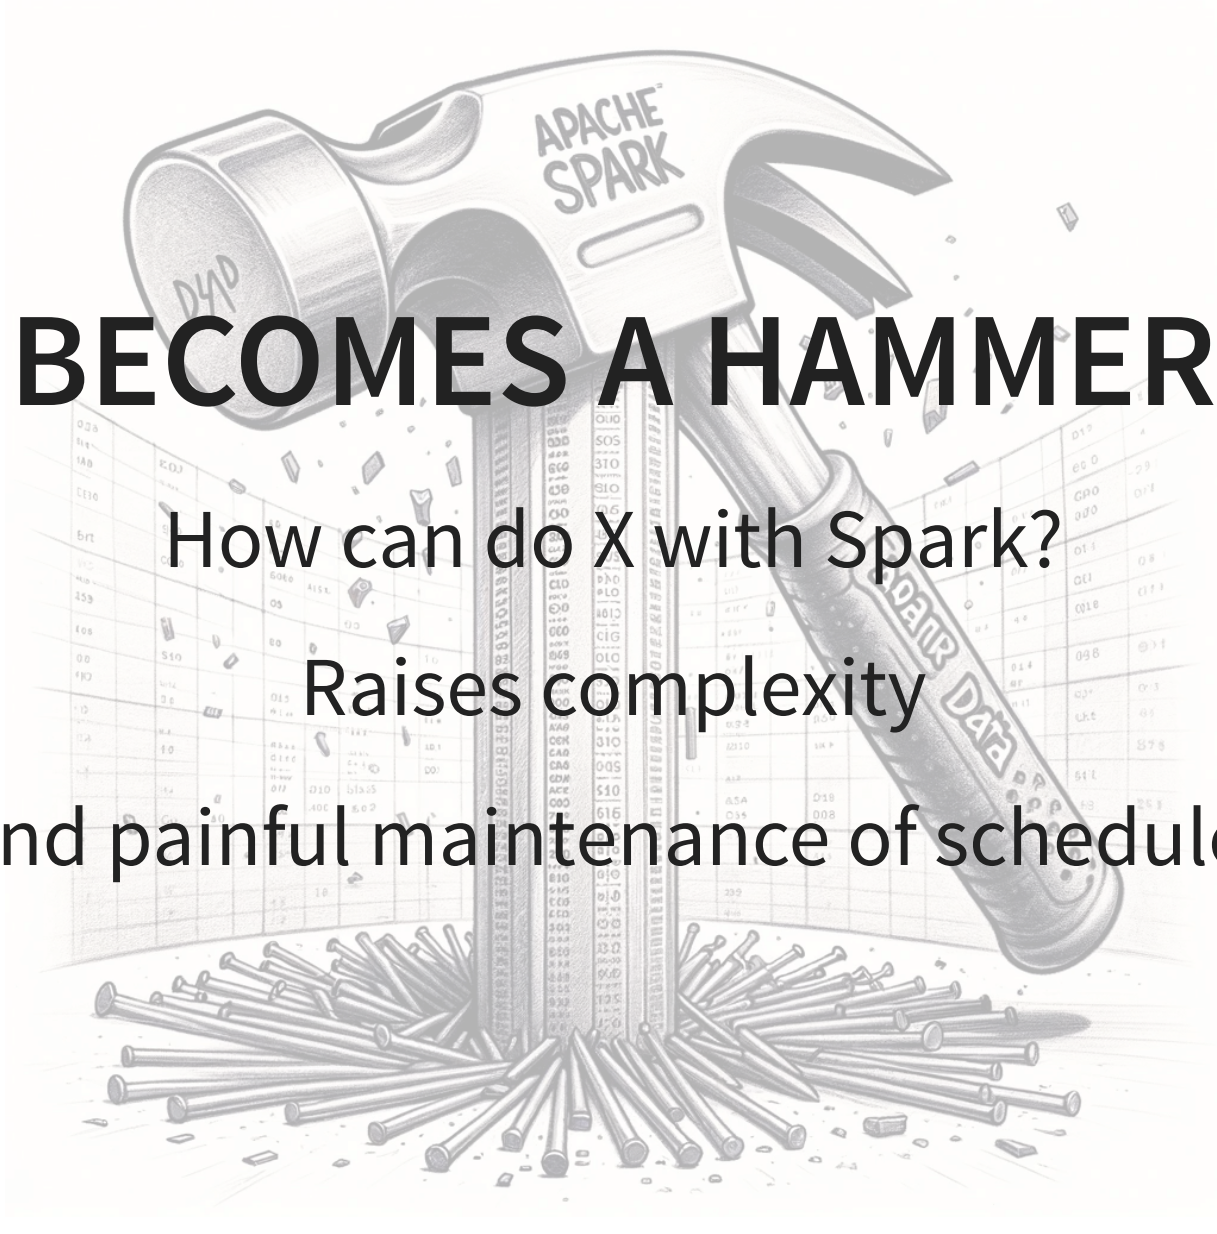

BECOMES A HAMMER

How can do X with Spark?

Raises complexity

Hard and painful maintenance of scheduled jobs

TAKEAWAY

TAKEAWAY

If it works out of the box, you probably don't need it.

TAKEAWAY

If it works out of the box, you probably don't need it.

Limit Spark's use to projects larger than a terabyte scale.

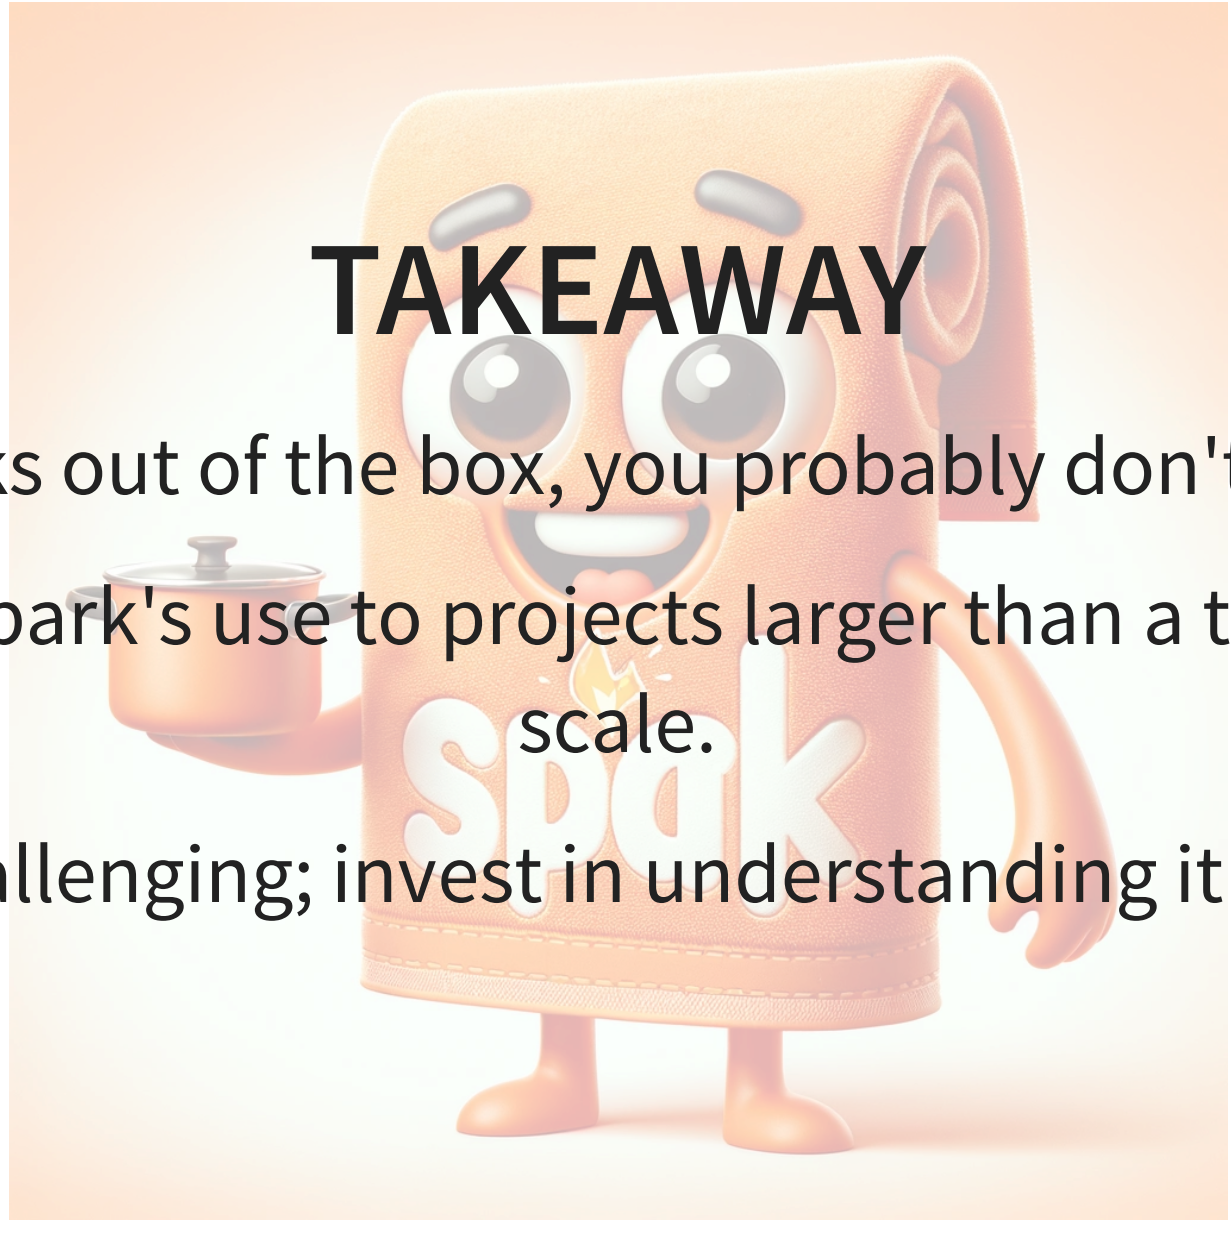

TAKEAWAY

If it works out of the box, you probably don't need it.

Limit Spark's use to projects larger than a terabyte scale.

It's challenging; invest in understanding it better.

TAKEAWAY

If it works out of the box, you probably don't need it.

Limit Spark's use to projects larger than a terabyte scale.

It's challenging; invest in understanding it better.

Remember to always bring a towel.

THANK YOU!

Questions?