Automating Netflix ML Pipelines With Meson

QCon SF 2017 | Eugen Cepoi, Davis Shepherd

NETFLIX


NETFLIX
Browse
Kids
DVD
Search
OneRingTo...
NETFLIX ORIGINAL
STRANGER THINGS
98% Match 2016 TV-14 1 Season
Season 2 Coming October 27
When a young boy vanishes, a small town uncovers a mystery involving secret experiments, terrifying supernatural forces and one strange little girl.
Starring: Winona Ryder, David Harbour, Finn Wolfhard, Millie Bobby Brown, Gaten Matarazzo
Creator: The Duffer Brothers
Genres: Halloween Favorites, Scary Halloween Favorites, TV Shows
MY LIST
OVERVIEW
EPISODES
TRAILERS & MORE
MORE LIKE THIS
DETAILS
Popular on Netflix
MINDHUNTER
NARCOS
BIG MOUTH
Breaking Bad
THE BABYSITTER

Goal

Create a **personalized experience** to help members find content to **watch** and **enjoy**

Recommendation Context

Member streaming data

Recommendation Context

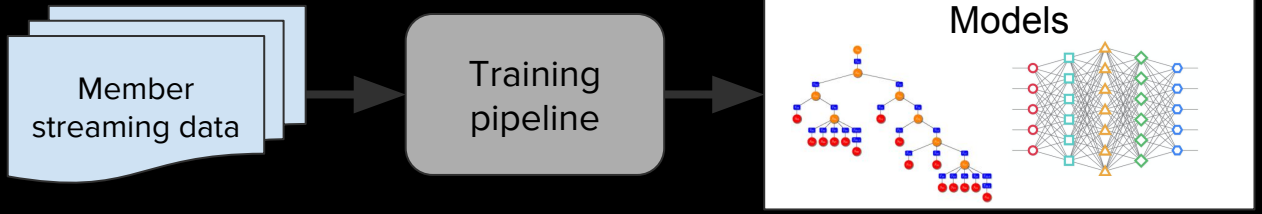

Recommendation Context

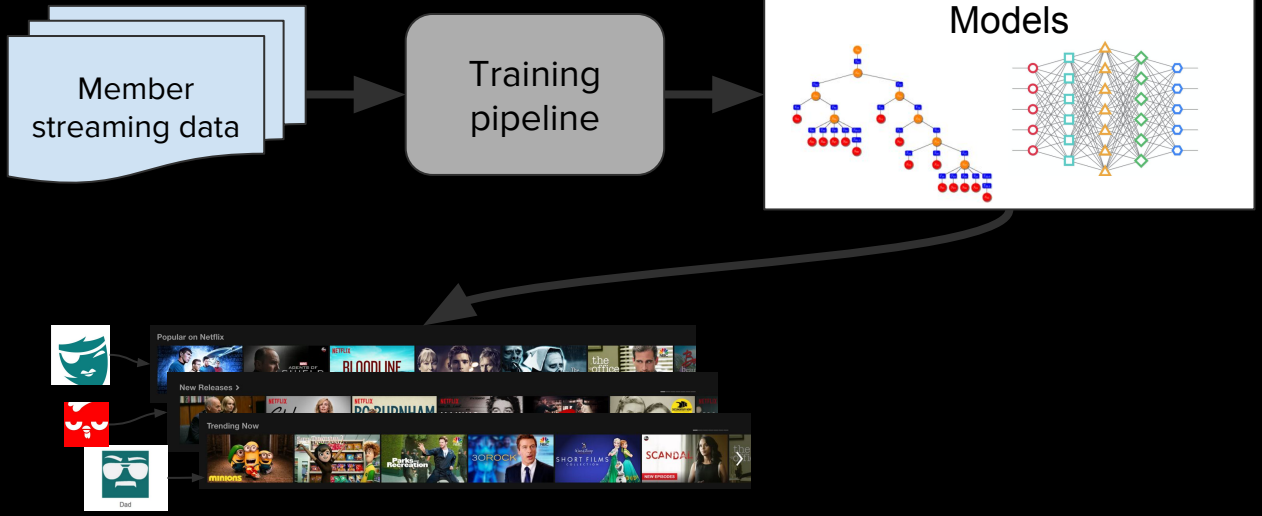

Recommendation Context

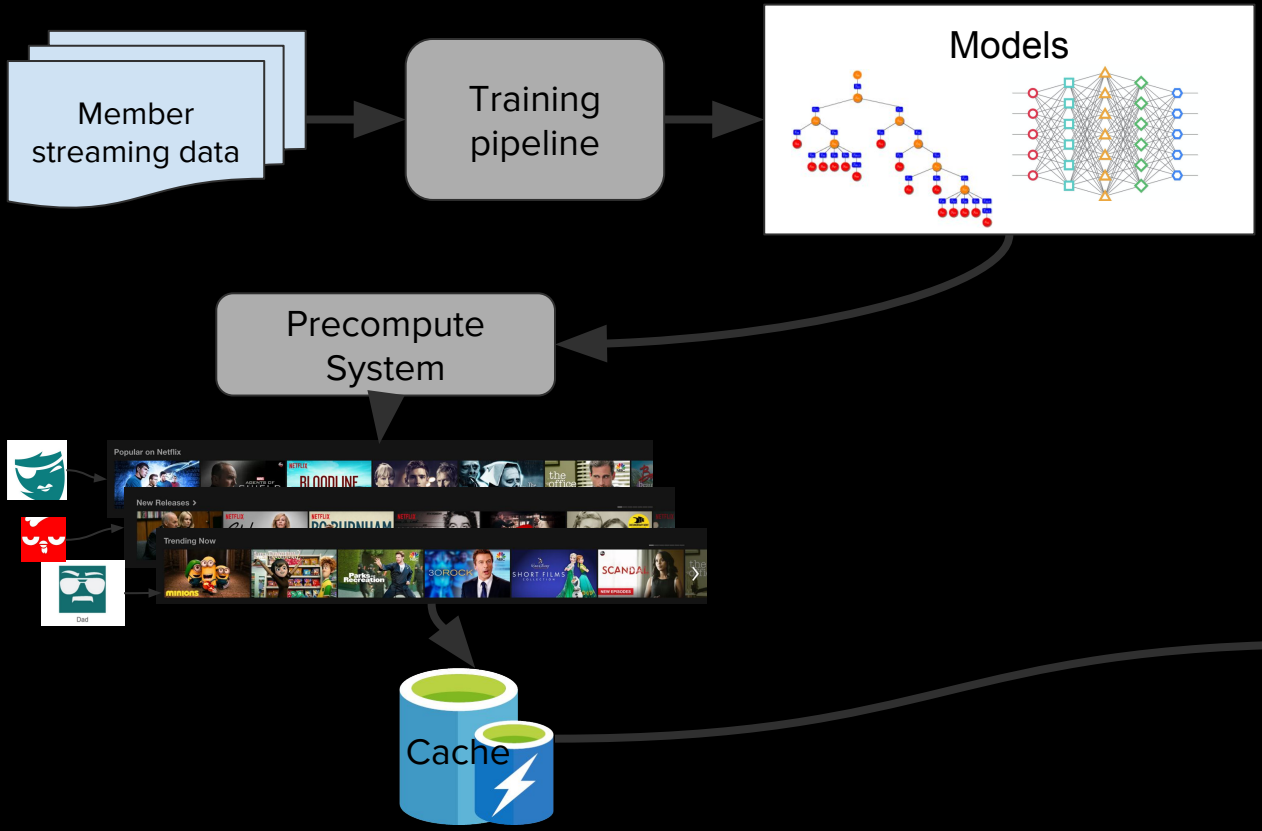

Recommendation Context

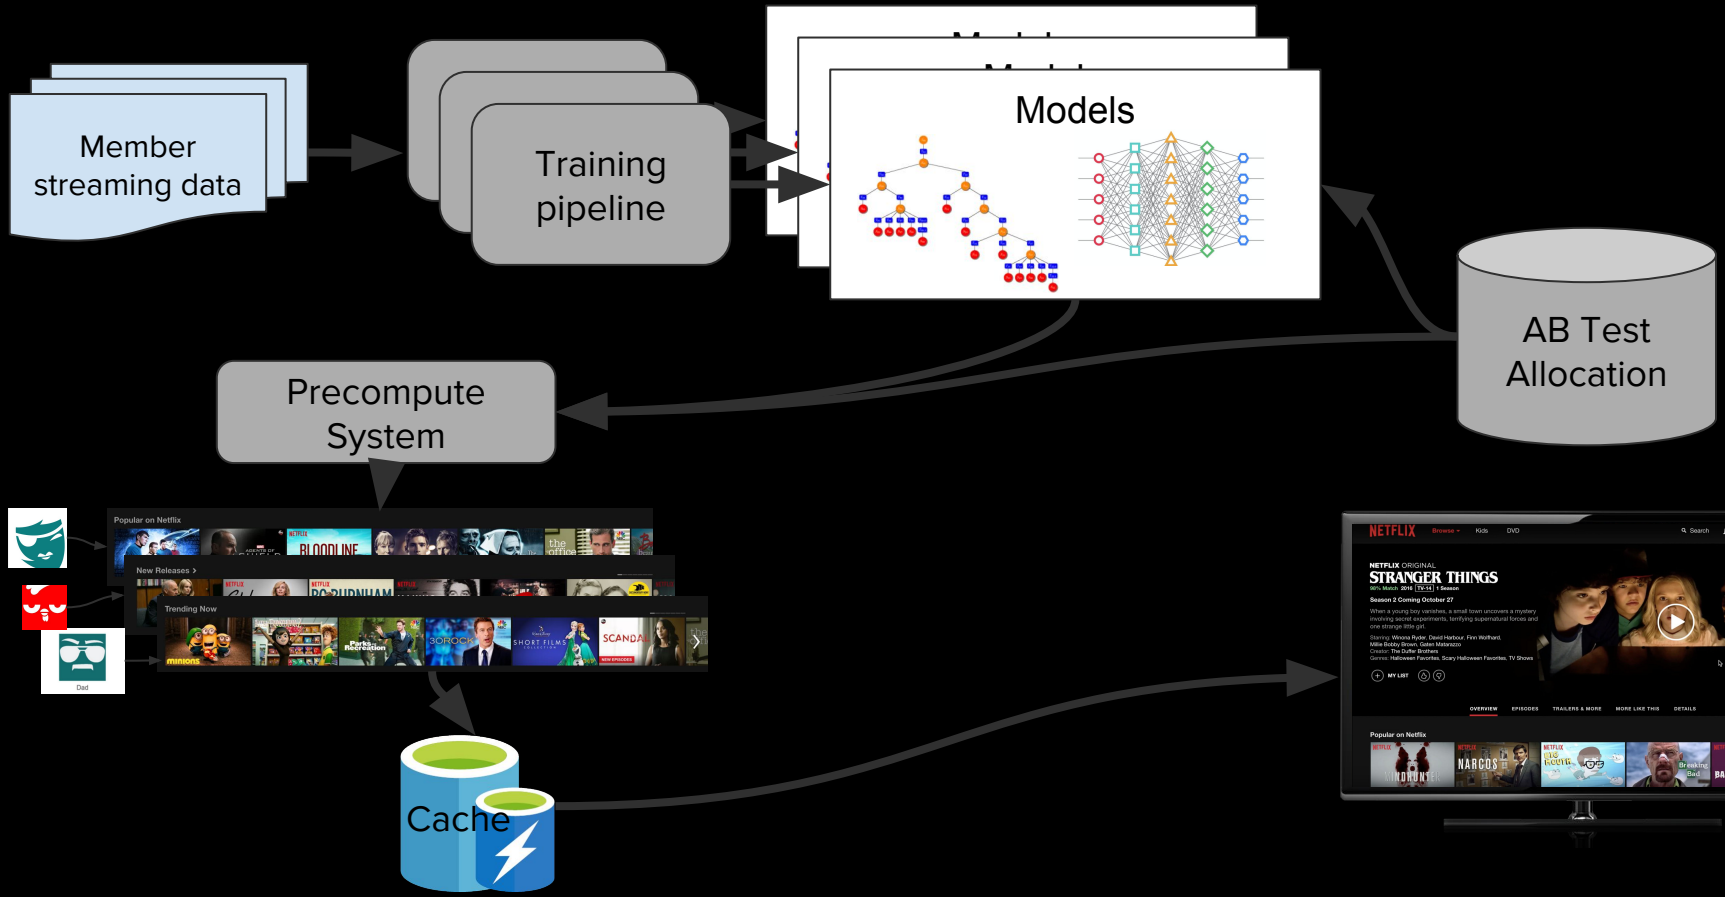

Innovation is driven by experimentation

Training Pipelines

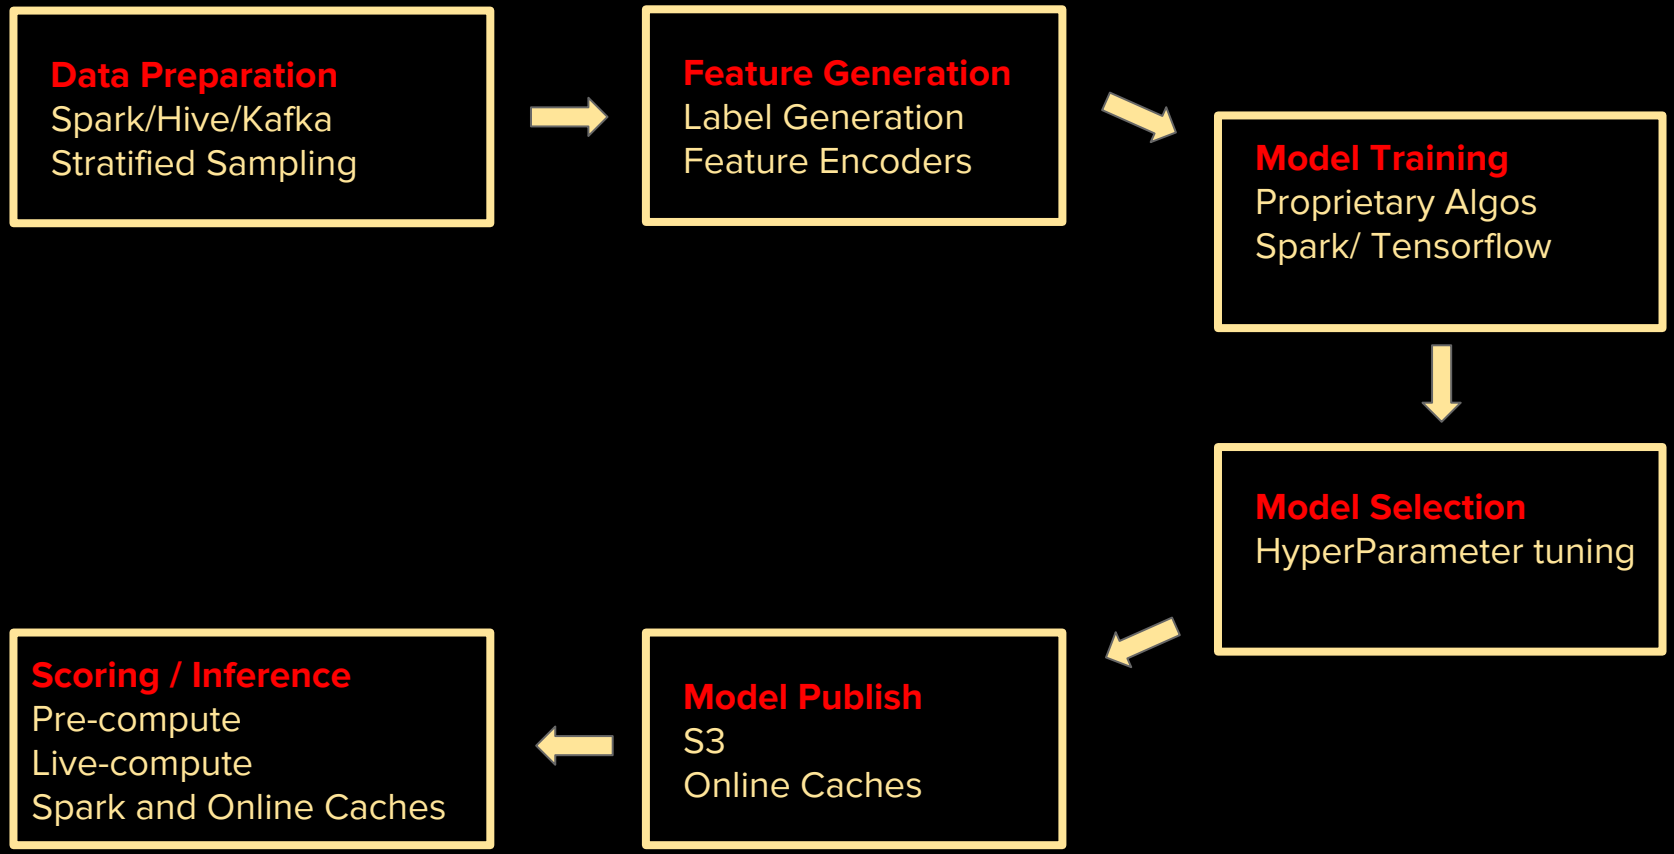

Before Meson

- A collection of operators
- Little to no orchestration
- Often limited to single machine

Desired Properties

- Support Heterogeneous systems
- Highly flexible generic orchestration
- Handle failures
- Provide Reproducibility
- Support Multi-tenancy
- Support External Triggers

Why didn’t an out of the box solution work?

- Spark and Scala support was paramount
- Options available didn’t have the flexibility and scalability that we needed

Meson overview

Meson Overview

General purpose workflow orchestration engine

Delegates execution to Mesos

Initially built for Machine Learning pipelines for personalization

Supports complex workflow patterns (branching, loops, foreach)

Concepts

Workflow Directed Graph of steps, global parameters, triggers...

Step Describes a job and its configuration

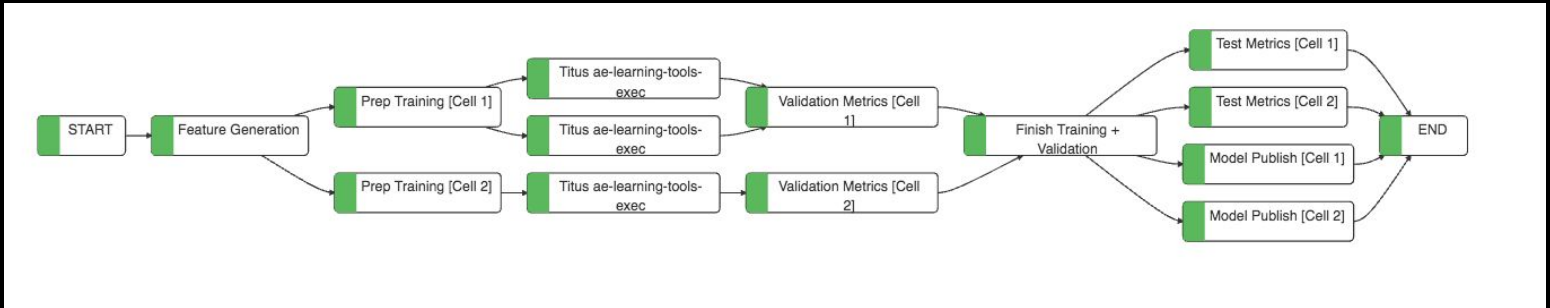

Defining workflows

Scala DSL

Python DSL

UI

REST API

```
val sparkStep = Spark(
  jobClass = "netflix.MySparkJob",
  sparkArgs = Seq("--driver-memory" -> "8g")
)

val cpStep = DistCp(from = "...", to = "...")

Workflow(
  id = "my-cool-workflow",
  triggers = Seq(Trigger.CRON("0 0 0 1/1 * ? *")),
  notification = Notification(
    whenStart = true, whenFinish = true, whoCustom = Option("me@gmail.com")
  ),
  parameters = Seq(Parameter.STRING(name = "country", value = "US")),
  definition = sparkStep.sequence(cpStep).end()
)
```

Parameters

Used to configure steps, job arguments, and step transitions

MVEL expression to derive parameter values at runtime

Predefined macros

```
Workflow(
  id = "my-cool-workflow",
  parameters = Seq(Parameter.LONG(name = "dateint", expression = Option("tsToDateInt(RUN_TS)"))),
  definition = Step(
    name = "jobX",
    parameters = Seq(
      Parameter.LONG(name = "day", expression = Option("dayFromDateInt(dateint)"))
    )).end()
)
```

Workflow patterns

Branches (OR, XOR, AND)

Loops with XOR

Foreach

Using parameters and MVEL

```
Step("a").or(
  Condition("x<5").branch(Step("b")),
  Condition("x<10").branch(Step("c")),
  Condition("x>10").branch(Step("d"))
).sequence(Step("e"))
```

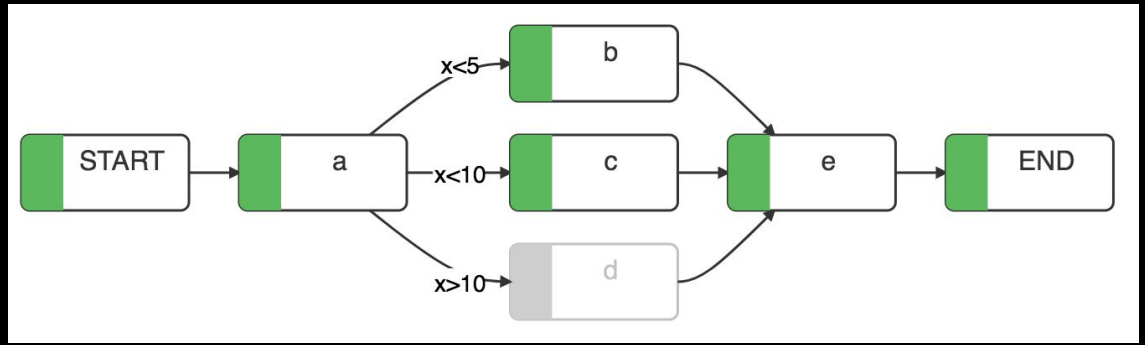

Data artifacts

Data artifact defined by a name and a set of partitions (parameters)

Cross workflow dependencies

External triggers

```
Step(
  name = "someJob",
  inputs = Seq(
    DataArtifact(
      name = "views",
      params = Seq(Parameter.STRING(name = "country", value = "US"))
    )
  ),
  outputs = Seq(
    DataArtifact(
      name = "aggViews",
      params = Seq(
        Parameter.STRING(name = "ts", value = "$RUN_TS"),
        Parameter.STRING(name = "country", value = "US")
      )
    )
  )
)
```

Workflow versions

Workflows have immutable versions

Enables:

- Better collaboration
- Rollbacks
- Reproducibility

Deploying workflows with the Gradle plugin

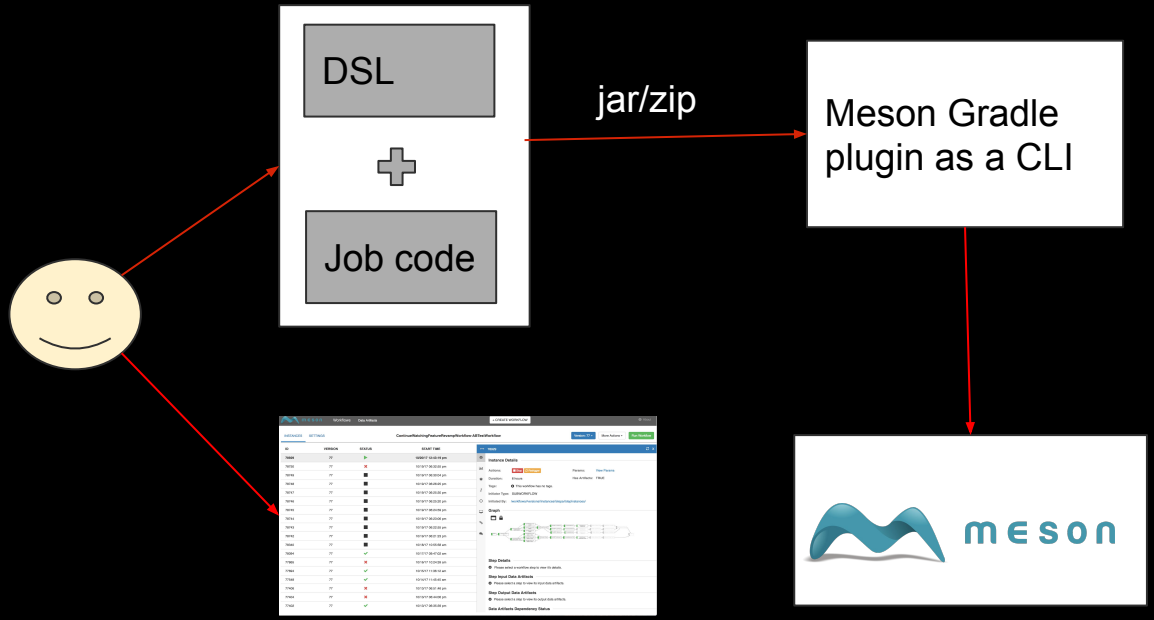

Link and version workflow with job code

Add metadata (SCM)

Validate

Deploy

Run

Automated releases with canary workflows

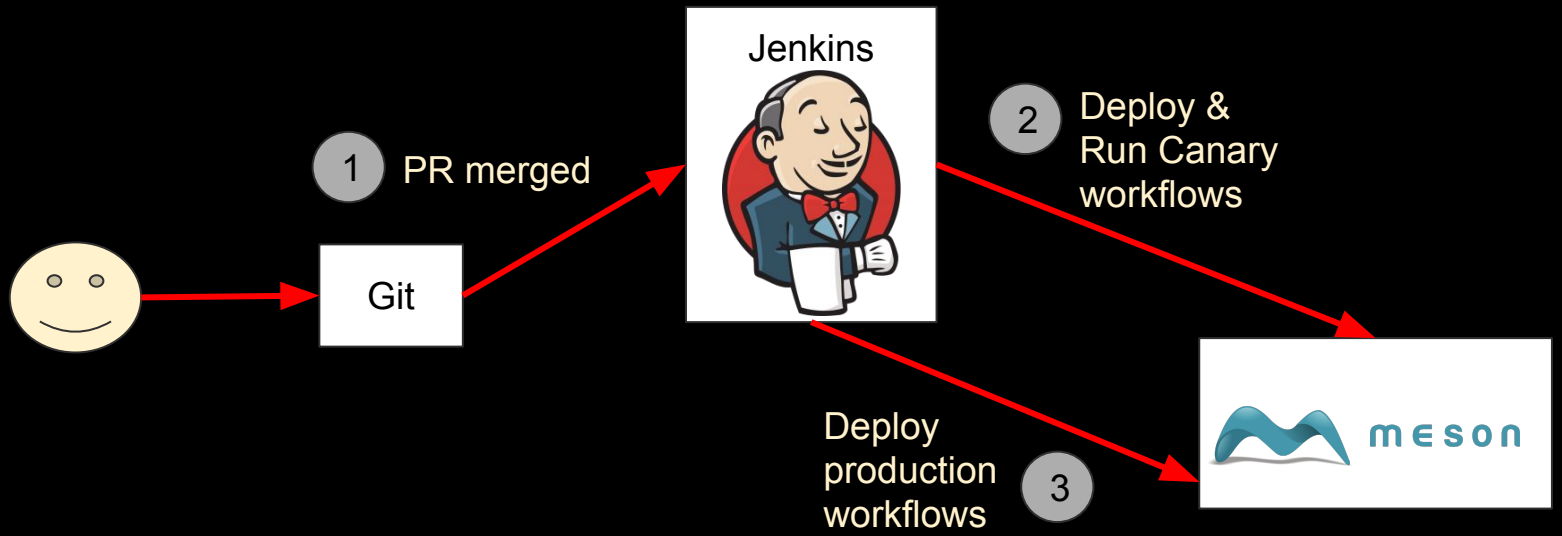

Monitor & Debug

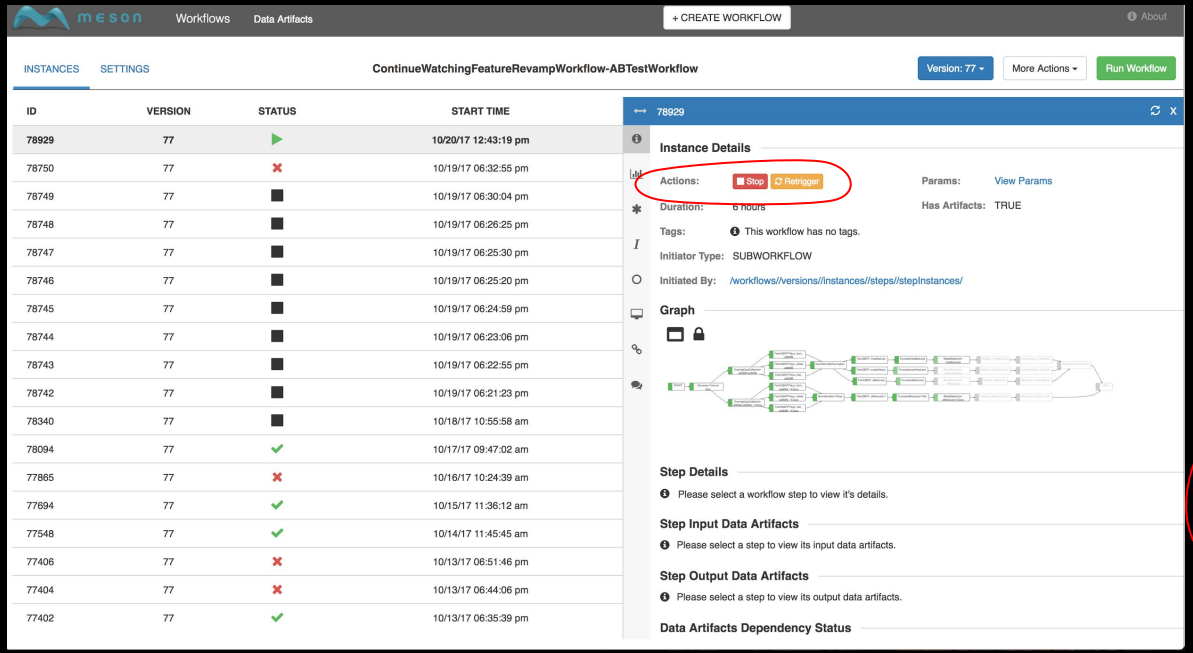

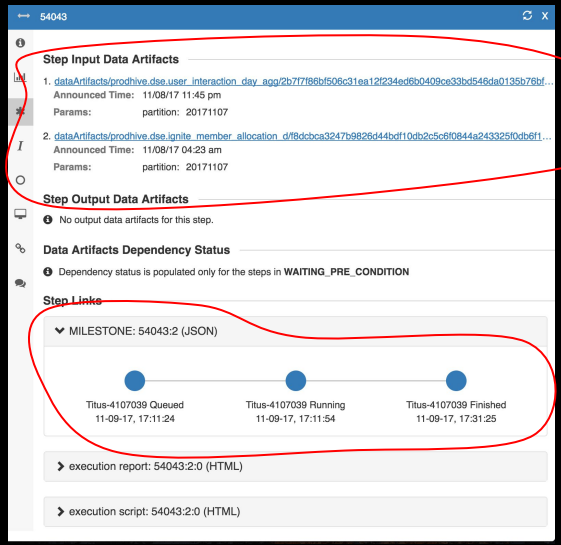

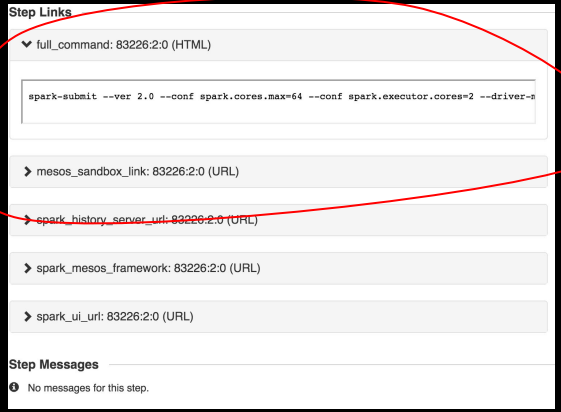

Architecture

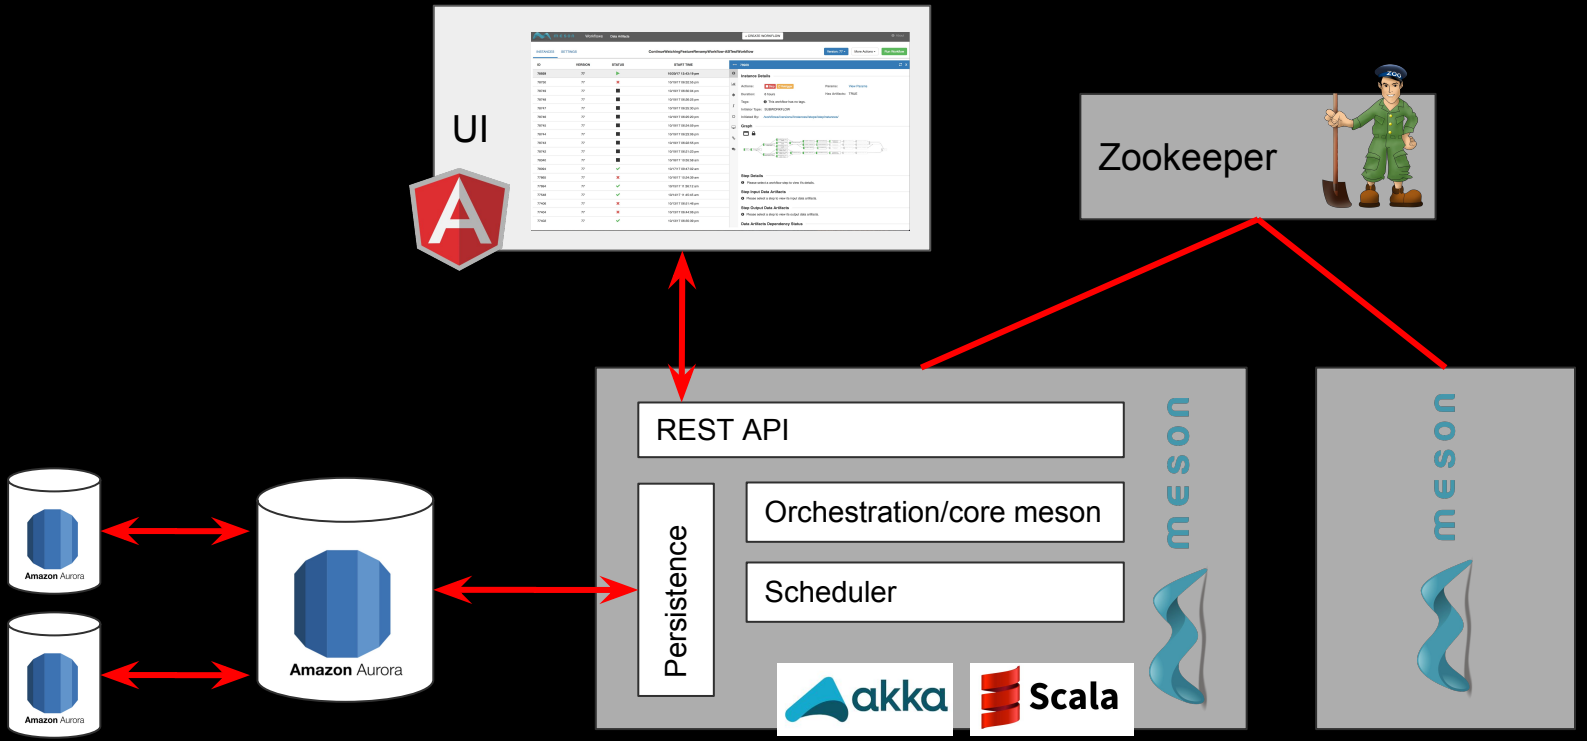

Scheduling

Meson as a Mesos Framework

Mesos offers resources and runs the steps

Fenzo (Netflix OSS) makes scheduling decisions

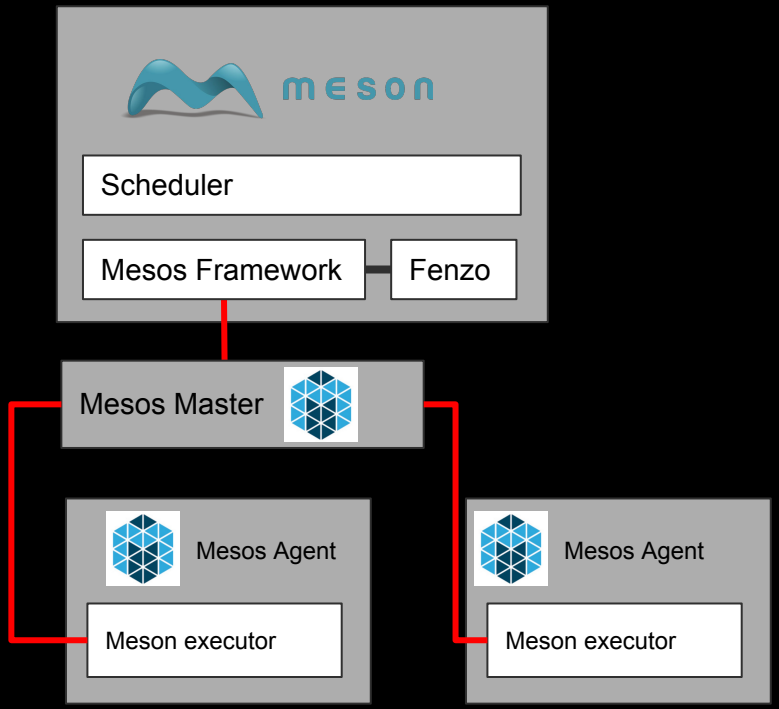

Execution

Custom executor code for different runtime systems (spark, bash...)

Publish runtime debug infos (logs, url to monitoring UIs...)

Meson executors survive Meson scheduler failure

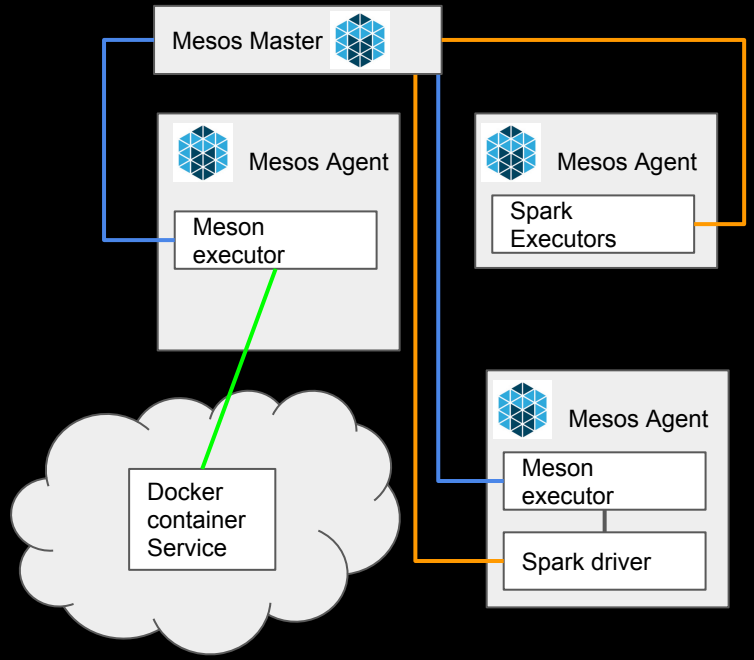

Looking Forward

Closing the loop

Interact with Meson from the running job

Output parameters to leverage loops and foreach

Expose debugging information through Artifacts

- Progress Milestones, Links, Counters, Images, etc.

A day in the life of a workflow...

Backfills, work prioritization and parallelism

Avoid re-doing work after fixing a bug and re-deploying a workflow

Explicit (data) lineage

Looking back

Adoption

2+ years in production

10+ managed and self-service deployed clusters

1000+ daily Production and A/B Test ML pipelines

2000+ EC2 instances in Spark/Mesos compute pool

20000+ of steps run per day

One Abstraction doesn't fit all

Evidenced by the many names:

- Workflow
- ProcessFlow
- Pipeline
- DAG
- DataFlow

Over specialization will inevitably weaken other use cases.

One Abstraction doesn't fit all

With the REST API Meson provides "workflows as a service".
Enables many domain specialized abstractions:

- A/B test orchestration
- ML orchestration
- ETL pipelines
- Notebook Automation
- And more..

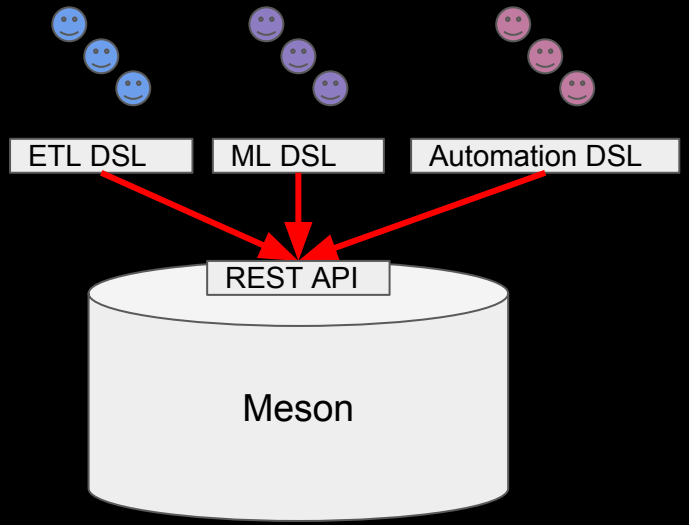

Questions?