Introduction to using and scaling dagster

Aleksandar Milicevic
Georg Heiler

Agenda

1 **Welcome and introduction to**
- What is a data platform
- What is the role of dagster

2 Crash course on Dagster concepts
- Asset based
- Metadata-created pipelines
- Resources and IO managers

3 Hands-on lab (local or GitHub Codespace)
- Spin up the environment
- Tour Dagster UI and run first asset
- Understand, run and extend the above Dagster concepts

4. Dagster @Magenta
- Open source implementation with local-data-stack

QA & Wrap-up

About us

Data expert in academia and industry Magenta Telecom

- meetup organizer and conference speaker

- data architecture, multimodal and complex data challenges

geoheil

geoheil

@geoheil.com

Data engineer at Magenta Telecom

- database internals, data platform engineering

milicevica23

milicevica23

@milicevica23.bsky.social

Data platform is there to help understand business process improve service

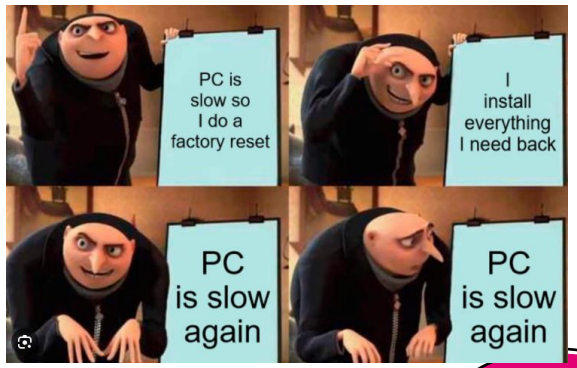

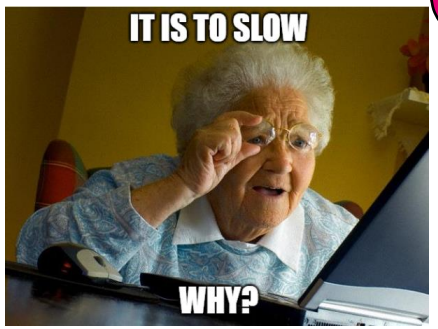

- customer facing performance is important – reporting is second priority (in a short run)
- different technology for different use case – oltp vs olap databases
- company data is an asset as car, phone or internet cable

Data platform components

What do I need for a bare minimum E2E reporting?

How to make it last?
Modern tooling, governance, development…

Data platform?

How to make it appealing?
AI…

…

Data catalog

Collibra

data visualisation/reporting

tableau

matplotlib

AI tooling

data transformation framework

SQLMesh

dbt

innovator

Modeling tool

data transformation engine

Apache Spark

Google BigQuery

Storage

Orchestration

Apache Airflow

Infrastructure

aws

…

Understanding tool silos

What should i do to get E2E reporting use case done?

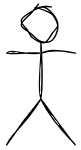

Developer

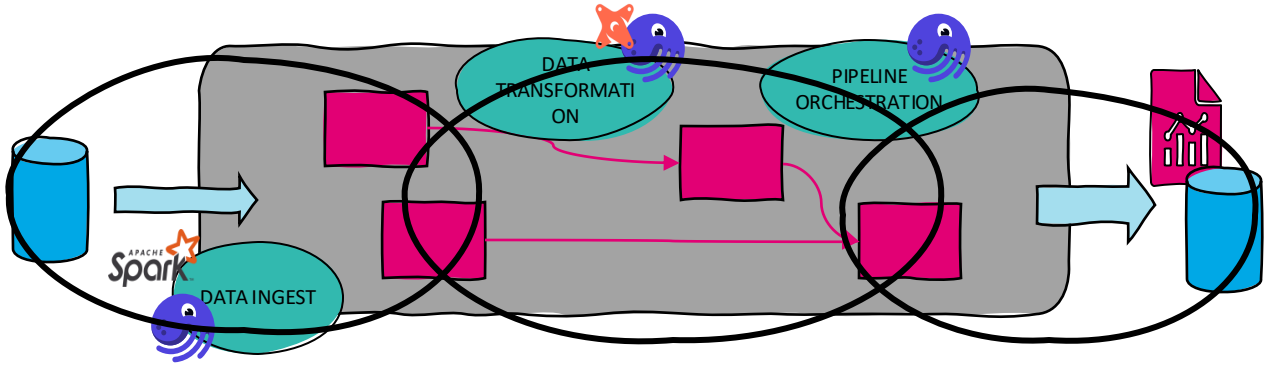

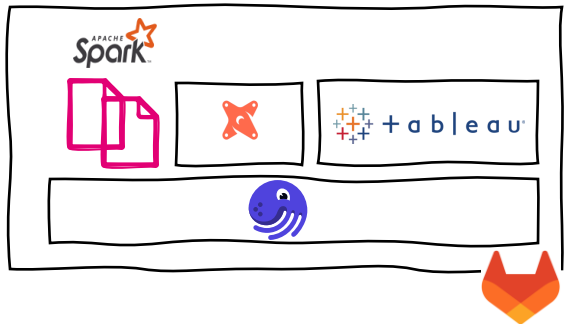

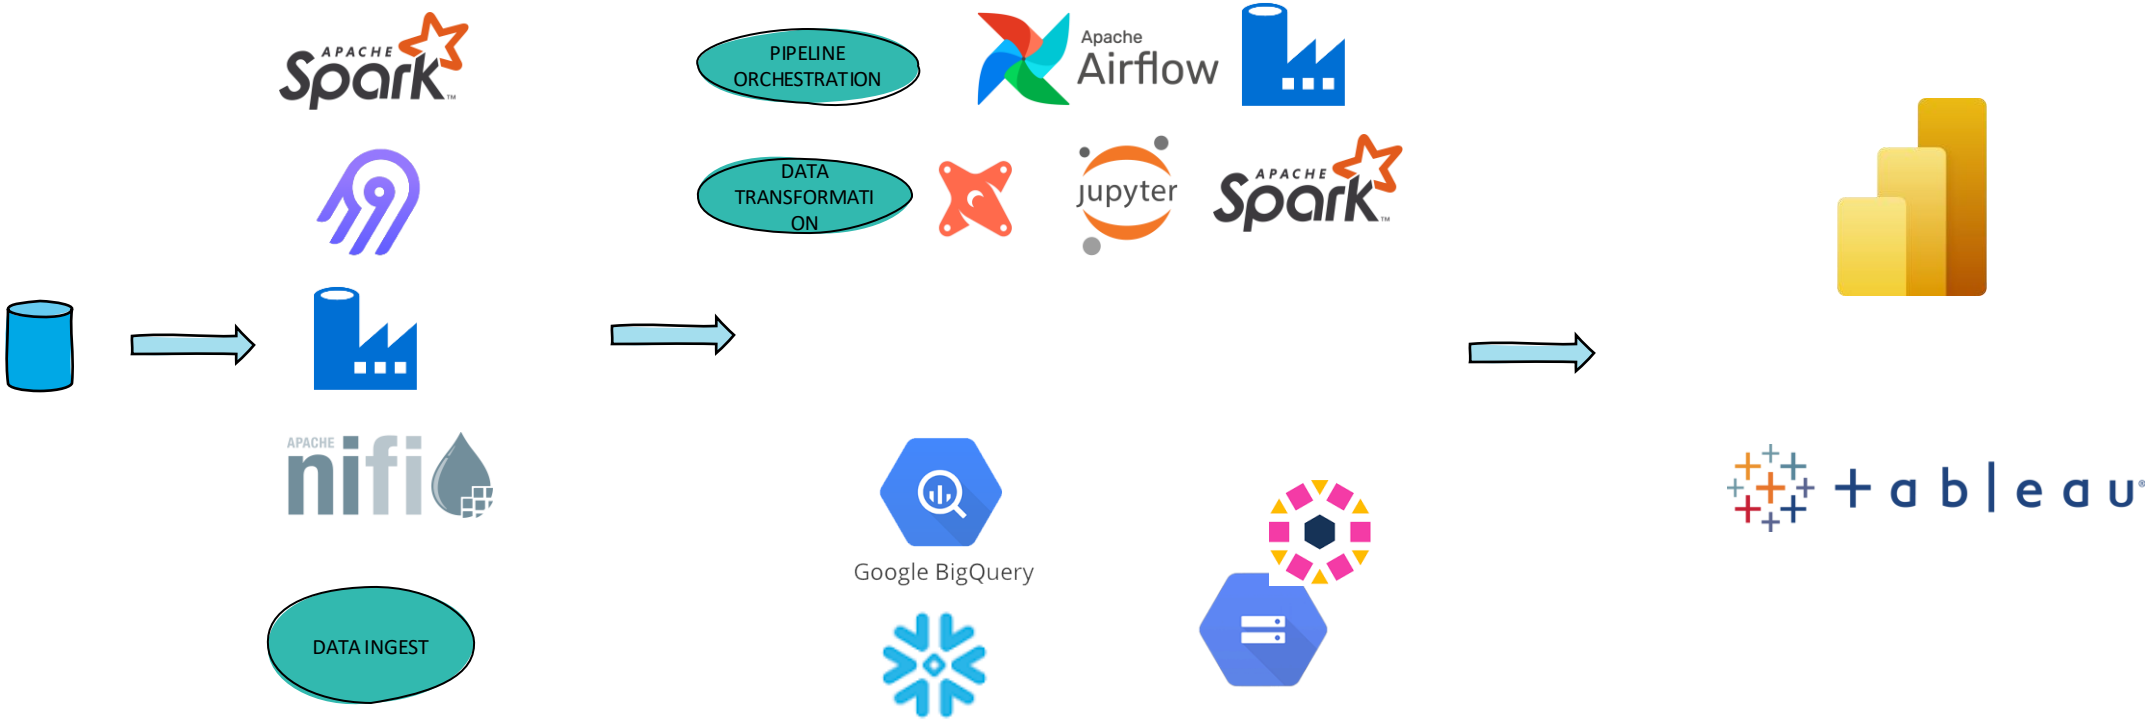

Dagster as the core of the platform

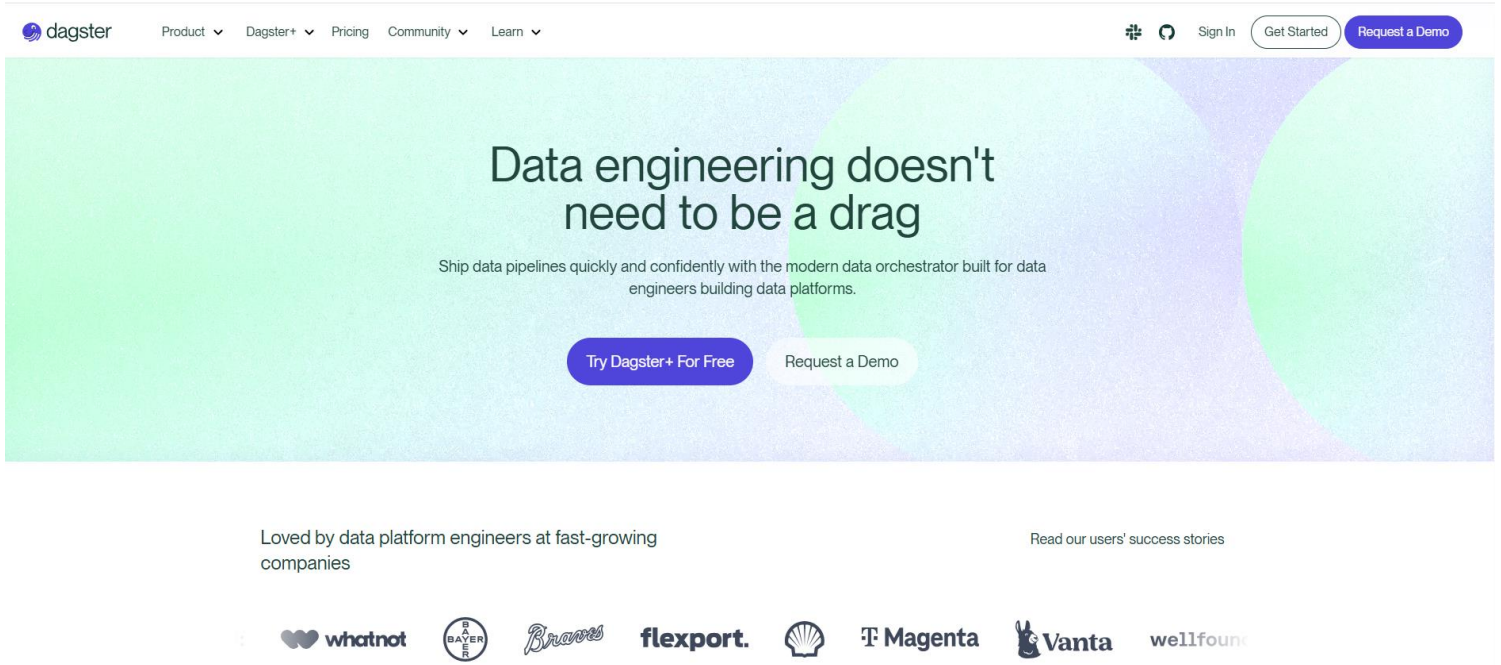

- at Magenta we decided to build around the orchestrator
  - hybrid deployment – controlplane SaaS – runtime in our k8s
  - software engineering best practices for project development and deployment
  - asset-based mindset for data flows (graph like a calculator for data dependencies)
- new concepts in orchestration...

New enabled concepts

docs.dagster.io/getting-started/concepts

- Asset based graph
- Metadata driven pipeline creation
- Reusable objects (resource, io manager)

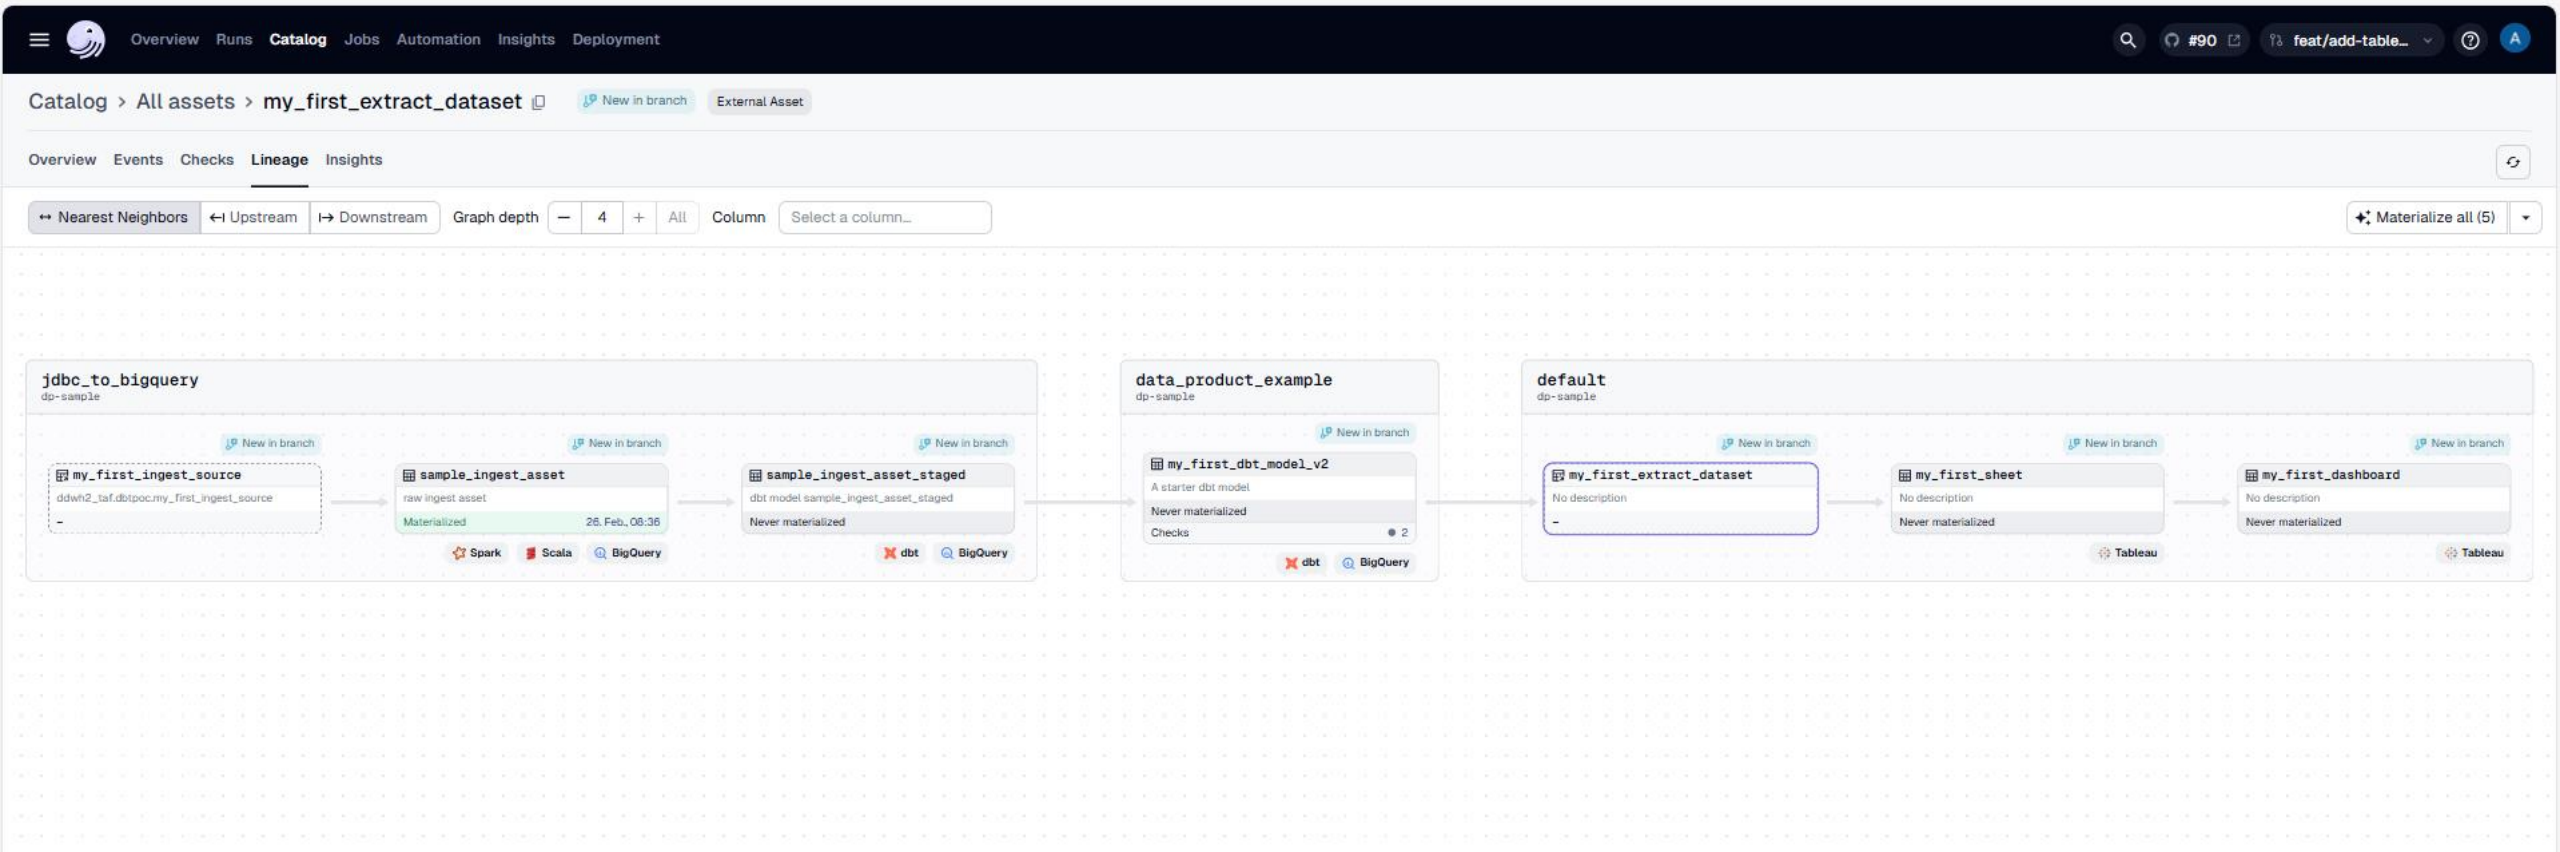

Asset and Task based orchestration: Chocolate cookie example

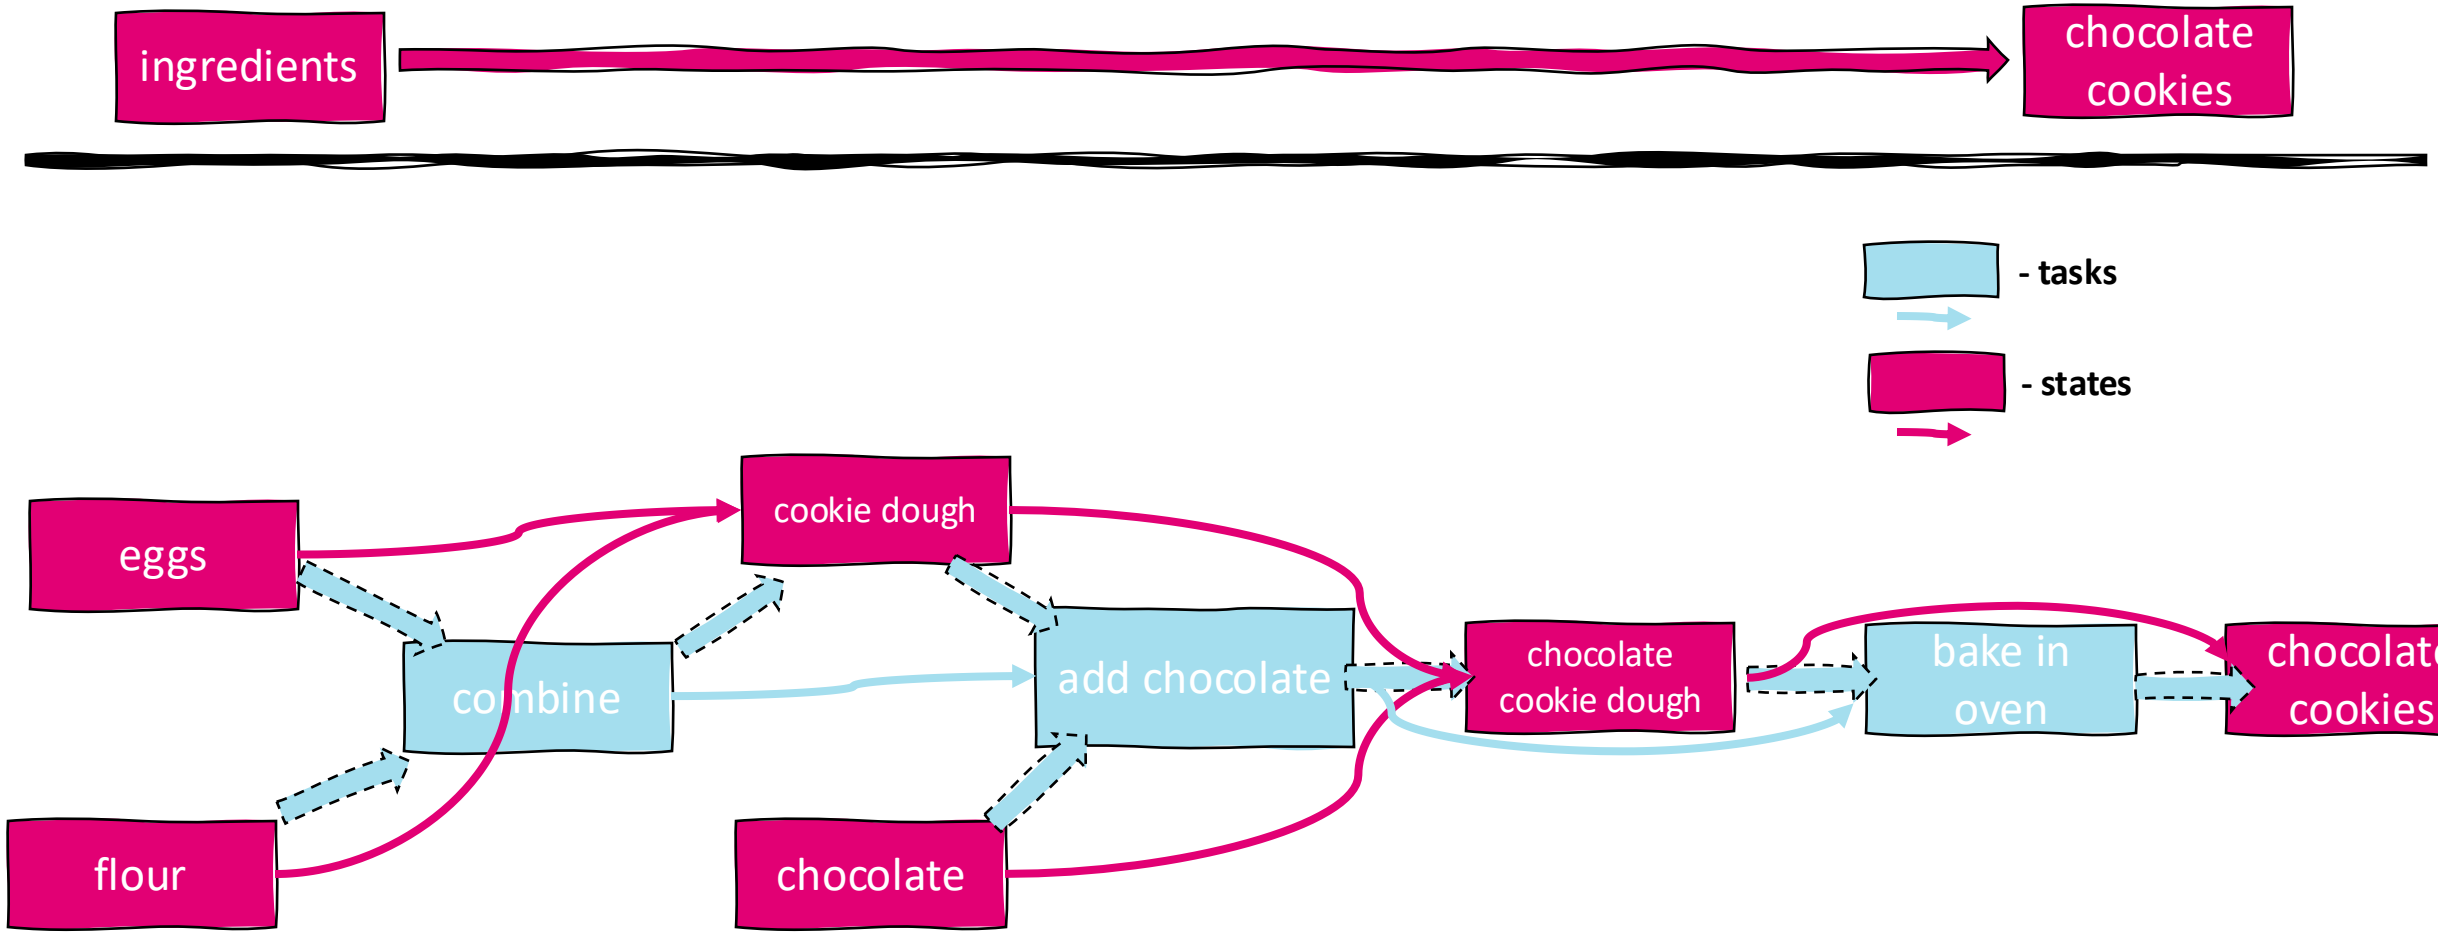

Asset based orchestration

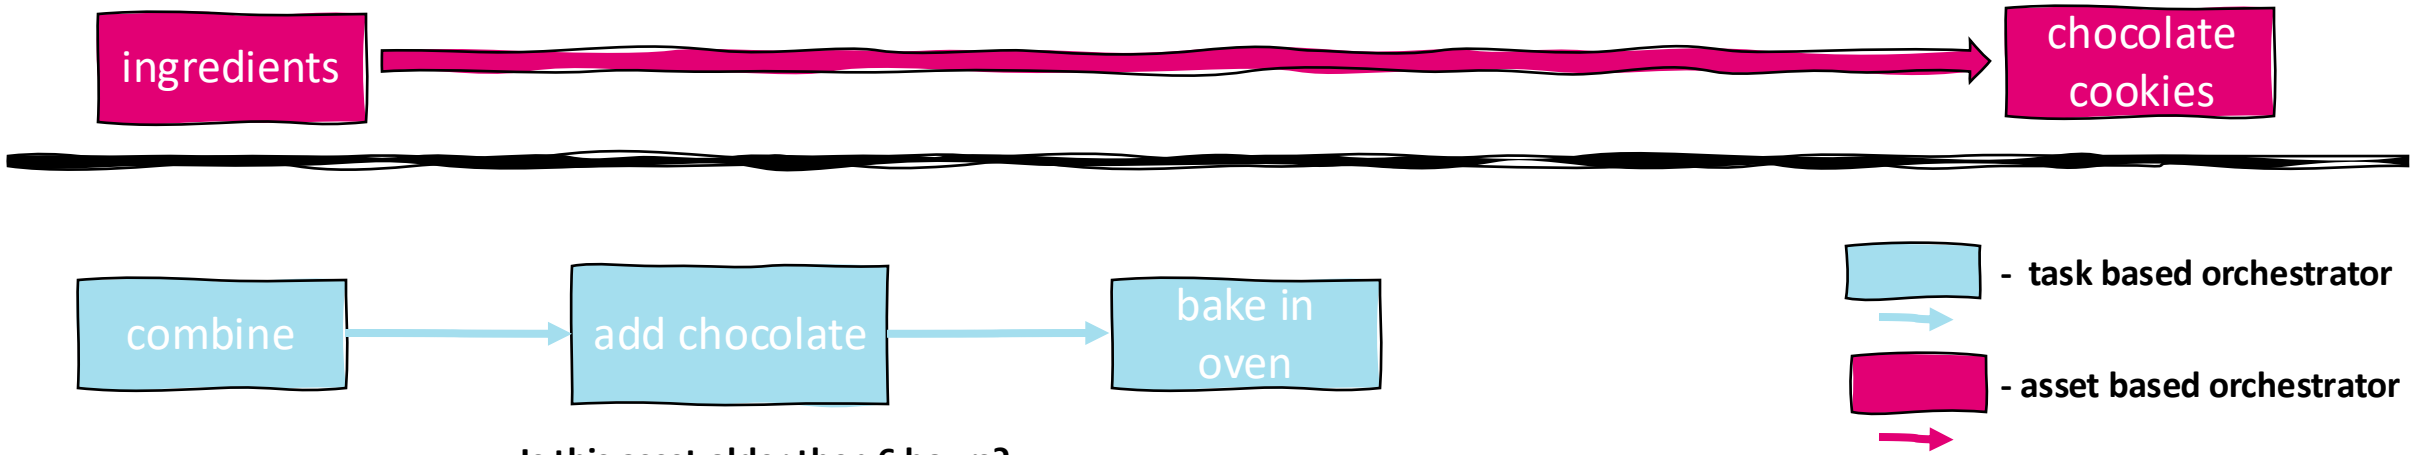

Advantages of asset-based orchestration:

- Asset testing

- Asset freshness

- Asset dependecy graph with granular declarative scheduling approach

Machine-readable metadata pipeline generation

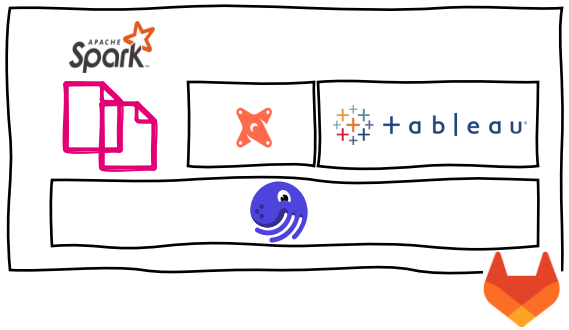


```python
metadata = collect_metadata()

list_of_assets = []
for each_metadata_node in metadata:
    #.. use_metadata
    @asset
    def asset():
        #..use_metadata

    list_of_assets.append(asset)

Definitions(assets = list_of_assets)
```

```python
tableau_server = TableauServer(creds)
#send rest api call to tableau server and get information
tableau_metadata = tableau_server.get_metadata()
list_of_assets = []
for each_tableau_object in tableau_metadata:
    tableau_object_deps = each_tableau_object.deps
    tableau_object_name = each_tableau_object.deps
    if each_tableau_object.type == extract_datasource:
        @asset(
            name = tableau_object_name,
            deps = tableau_object_deps
        )
        def refresh_extract(tableau_server):
            #send api call to refresh object
            tableau_server.refresh_extract(tableau_object_name)
        list_of_assets.append(run_dbt_asset)
    else:
        @asset(name = tableau_object_deps)
        def asset():
            pass
        list_of_assets.append(asset)

Definitions(assets = list_of_assets)
```

Reusable components

- Reusable objects
  - Resources → Encapsulate complex logic to interact with external systems
  - IO manager → Make complex IO interactions substitutable & testable
- Benefits
  - Dependency injection
  - Day 1 productivity: Scale the data pipeline down to a single laptop
  - Increase self-service: Business/DS focus not required to handle complex IO

```python
@asset(
    io_manager_key="bigquery_io_manager",
)
def awesome_ml_model(context, reference_addresses: pd.DataFrame, bigquery: BigQueryResource) -> pd.DataFrame:
    # simple normal python code here
    # IO is abstracted
    context.log.info(f"from source: \n{reference_addresses.head()}")
    # auth & complexity (imagine web API) is abstracted
    with bigquery.get_client() as client:
        job = client.query("select * from example.upstream")
        query_result = job.result().to_dataframe()
        context.log.info(f"direct query: \n{query_result.head()}")
    return pd.DataFrame({"foo": [1,2,3]})
```

Lab - Demo

- basic asset and UI overview
- resources
- metadata pipeline creation
- components

dagster - pip install dagster

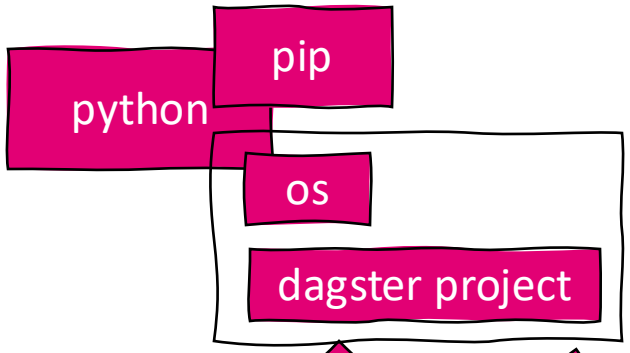

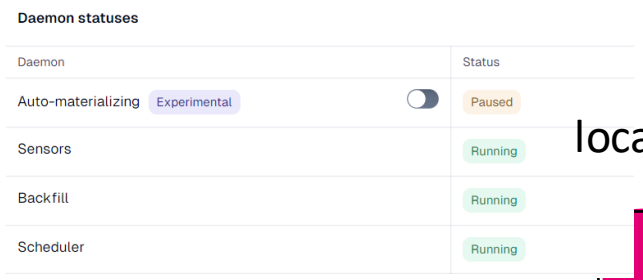

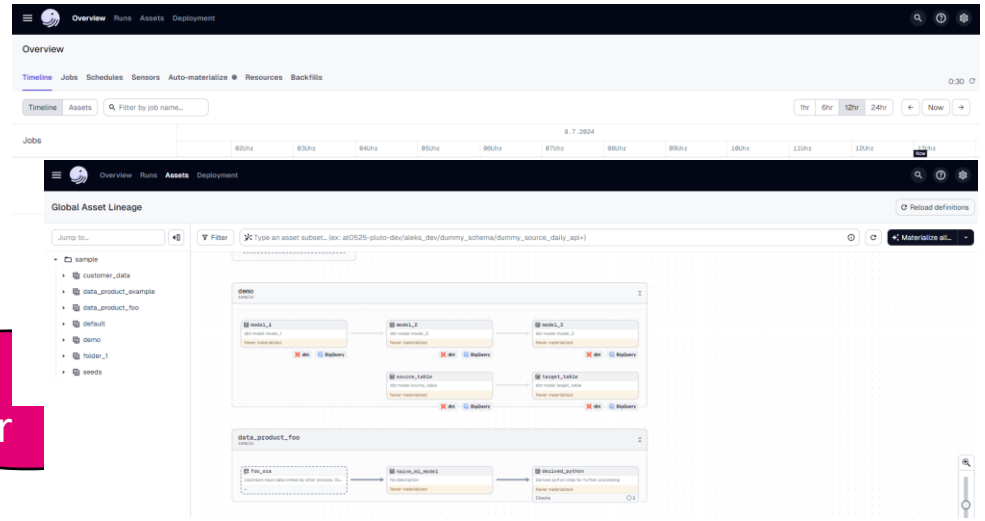

localhost:3000

pip install –e .

dagster dev

dagster webserver

Dagster demons

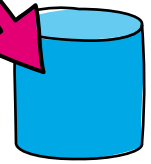

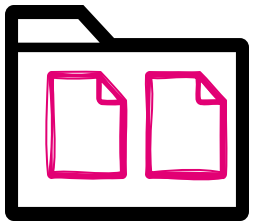

Dagster project

```
load_from:
  python_module: dagster_project
```

wokspace.yaml

code location

dagster project

- storage for dagster runs and metadata
  - sqlite/postgres

- grpc server serving the assets to the dagster deamon
  - needs definition object from the module

```
CMD ["dagster", "api", "grpc", "-h", "0.0.0.0", "-p", "4000", "-f", "sample/__init__.py"]
```

Takeaways

- Integrated asset-based graph is key (from ingest, transformation, reporting, tests – to AI)
  - Event driven connection
  - Better collaboration (scaling)
- Software engineering principles enable business self service
  - Blueprint
  - Automate all the things: CI/CD (stateful & stateless)
  - DRY: build tested foundation – dependency injection
  - Make business departments part of the key processes and pipelines
- Executable specification (metadata, contracts)
  - Interface Mangement
  - Preserve semantics
  - Preserve compliance (security classification, PII, retention)

Q&A

Dagster @Magenta
Local Modern Data Stack

Open-source contributions: L-MDS + Tableau integration + VertexAI integration

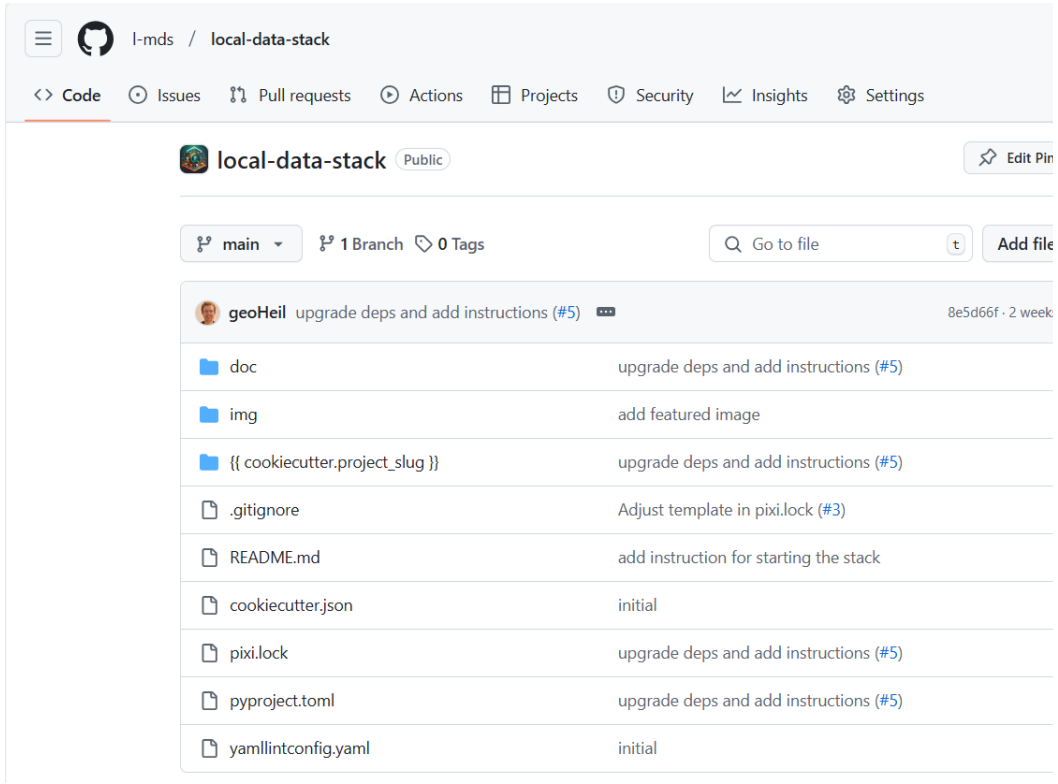

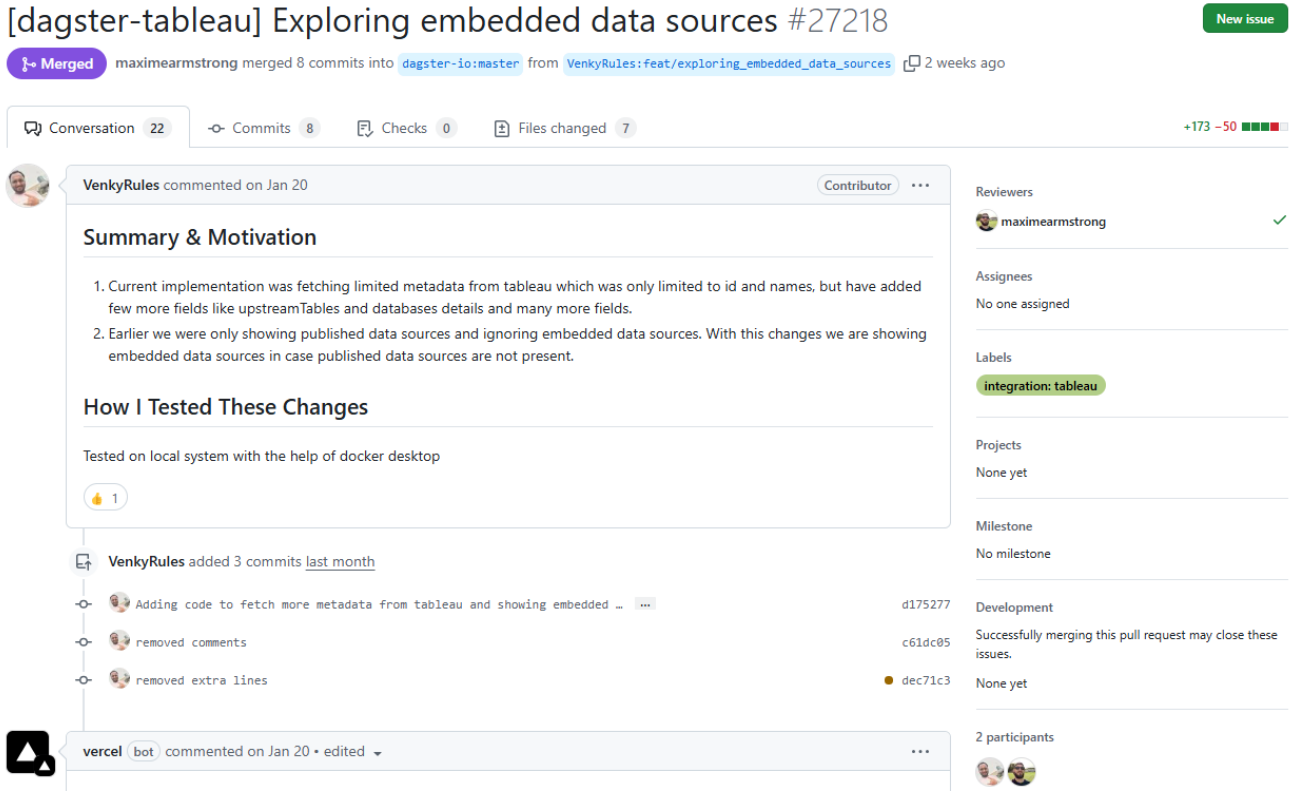

Building a data platform is team work

Further material

- public live stream about the basics: georgheiler.com/event/magenta-pixi-25
- n-depth Magenta implementation georgheiler.com/event/magenta-data-architecture-25
- https://courses.dagster.io/