

# **Apache Airflow and Ray**

## **Orchestrating ML at Scale**

# Background

- Strategy Engineer at [astronomer.io](https://astronomer.io), Airflow PMC member, co-creator of K8sExecutor
- Previously: Building data science platforms at Bloomberg LP
- Obsessed with Data Science Tooling, and building distributed systems



# The Airflow Data Science Story

# The Airflow Data Science Story

- Airflow is the tool to take you from experiment to production model
- Monitoring and scheduling ensure your models update in time for SLAs
- Connection handling for easily switching between dev -> prod data sources
- Fault tolerant scheduler that can retry jobs in case of failure

# The Airflow Data Science Story



# The (Traditional) Airflow Data Science Story

**Experiment**



**Parameterize**



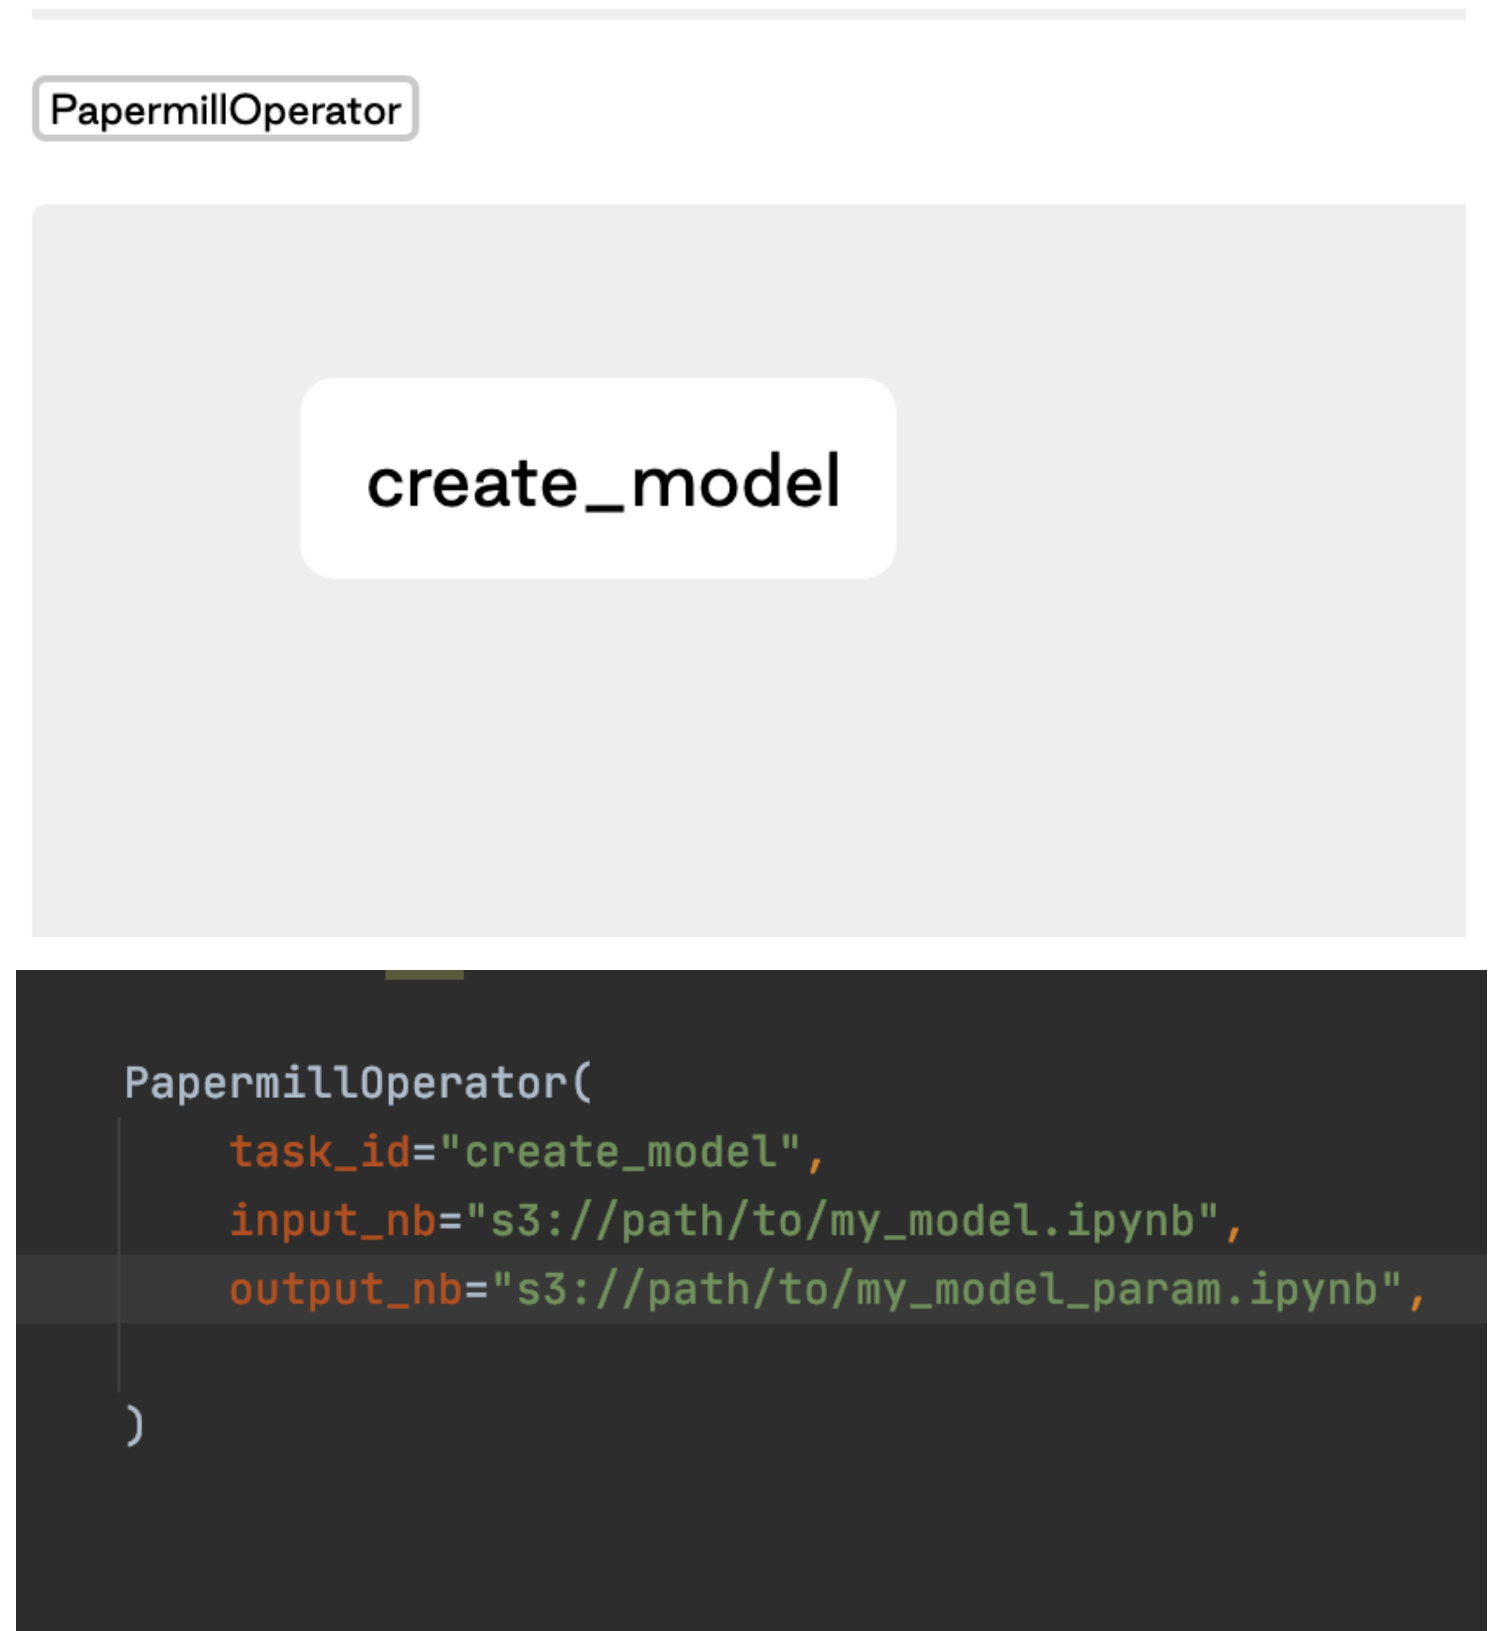
**Productionize**



# Papermill

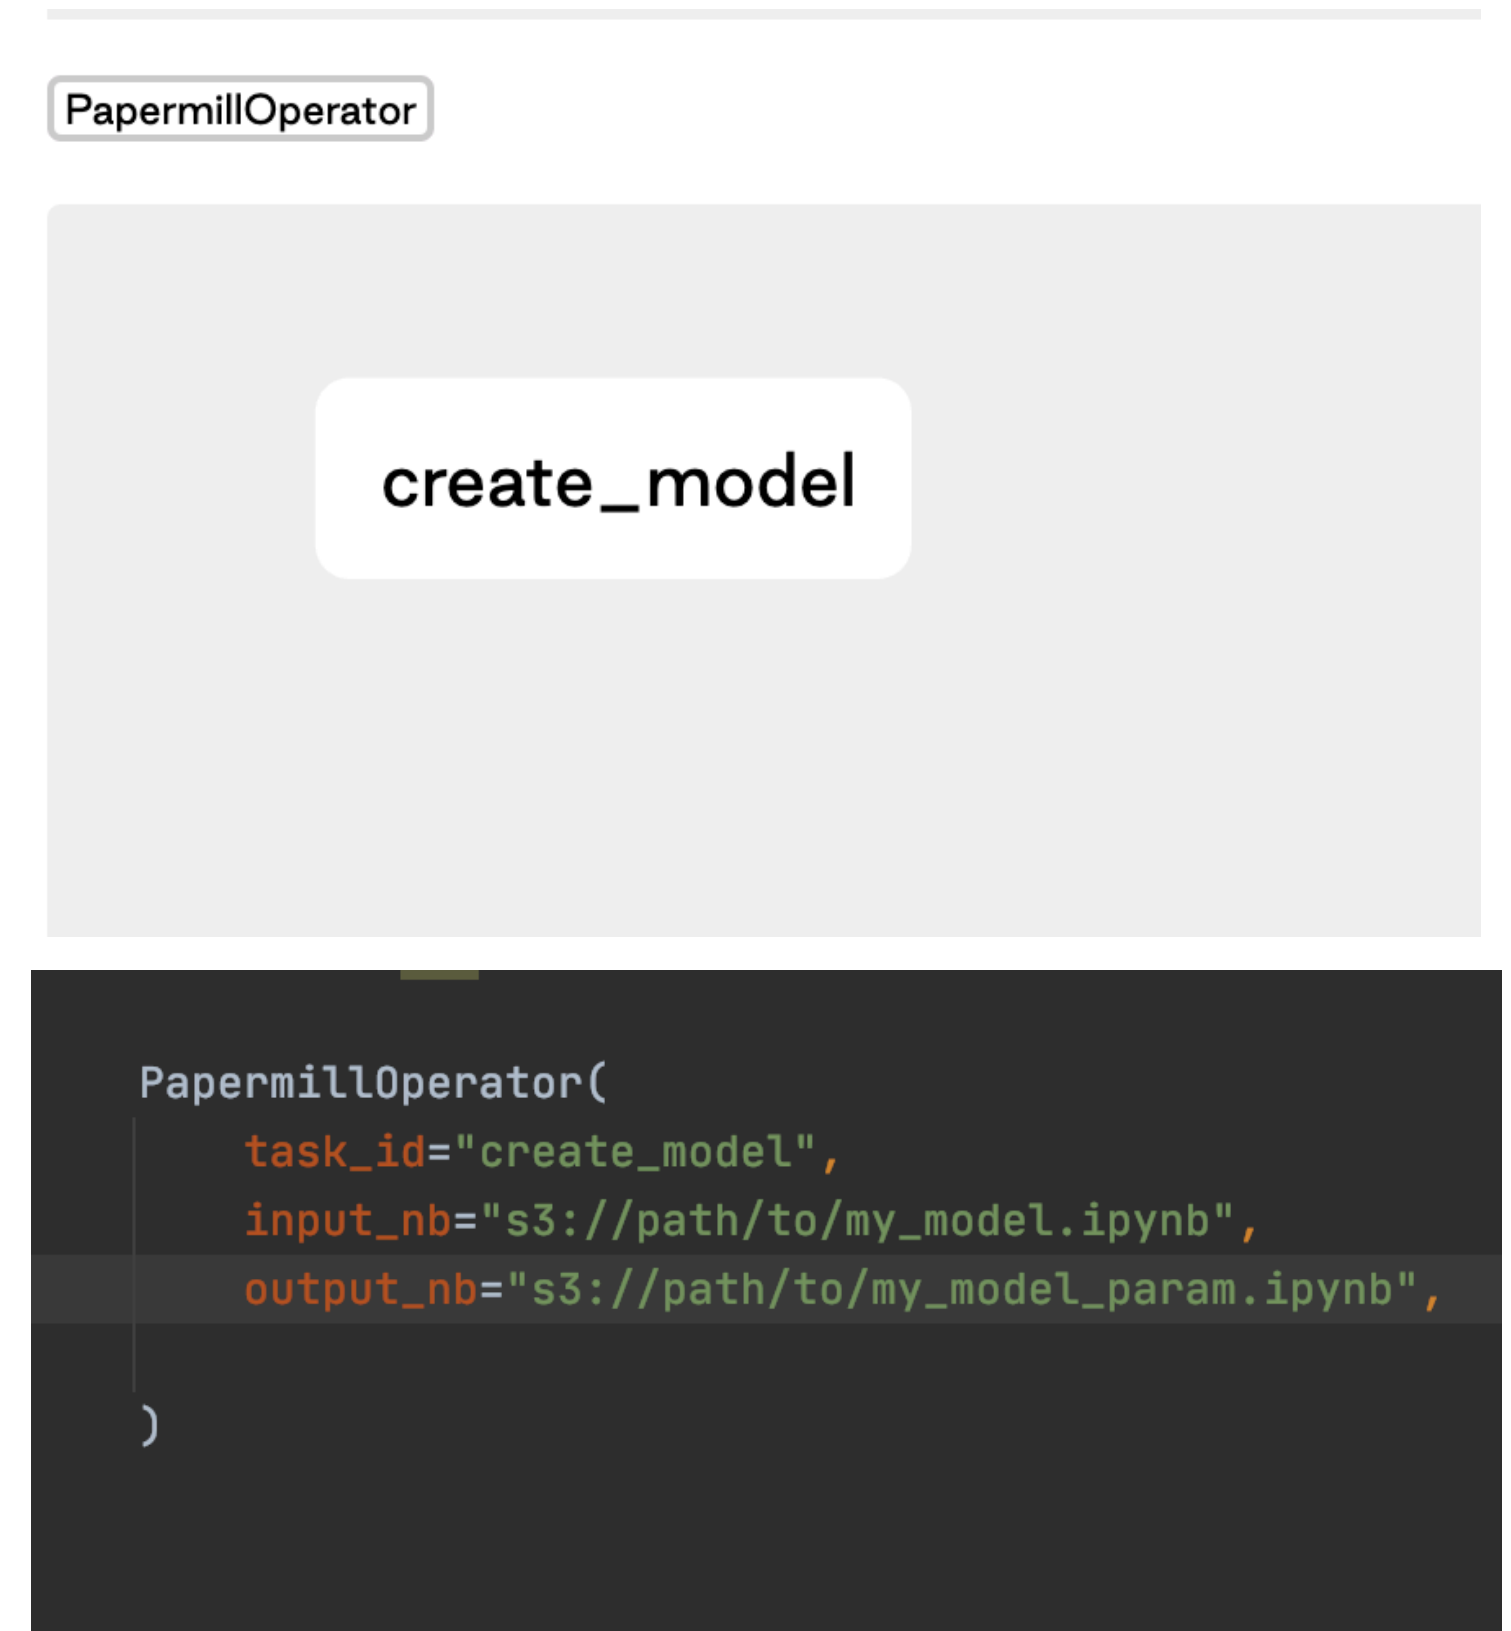


- Parameterize Notebooks through cell tagging
- Stores intermediate notebooks
- Execute using Python API or CLI
- Stores notebooks to S3/GCS



# Issues with this approach

- Entire notebook executes as a single task
- Low visibility, no fault tolerance
- Code is in multiple locations
- Experimentation becomes difficult
- Repeatability becomes messy



# The Next Gen Airflow Data Science Story?

**Experiment**



**Parameterize**



**Productionize**



# The Ideal Story

- Minimal conversion from Jupyter notebook -> Airflow DAG
- Moving large datasets between different tasks should be trivial.
- Should be able to request dedicated resources for the compute job
  - GPU, RAM, CPU, etc.
- Register & deploy, and replicate the resulting models.
- Maintain orchestration and monitoring at scale.

# Enter Ray

“[A] distributed execution framework that makes it easy to scale your applications and to leverage state of the art machine learning libraries.



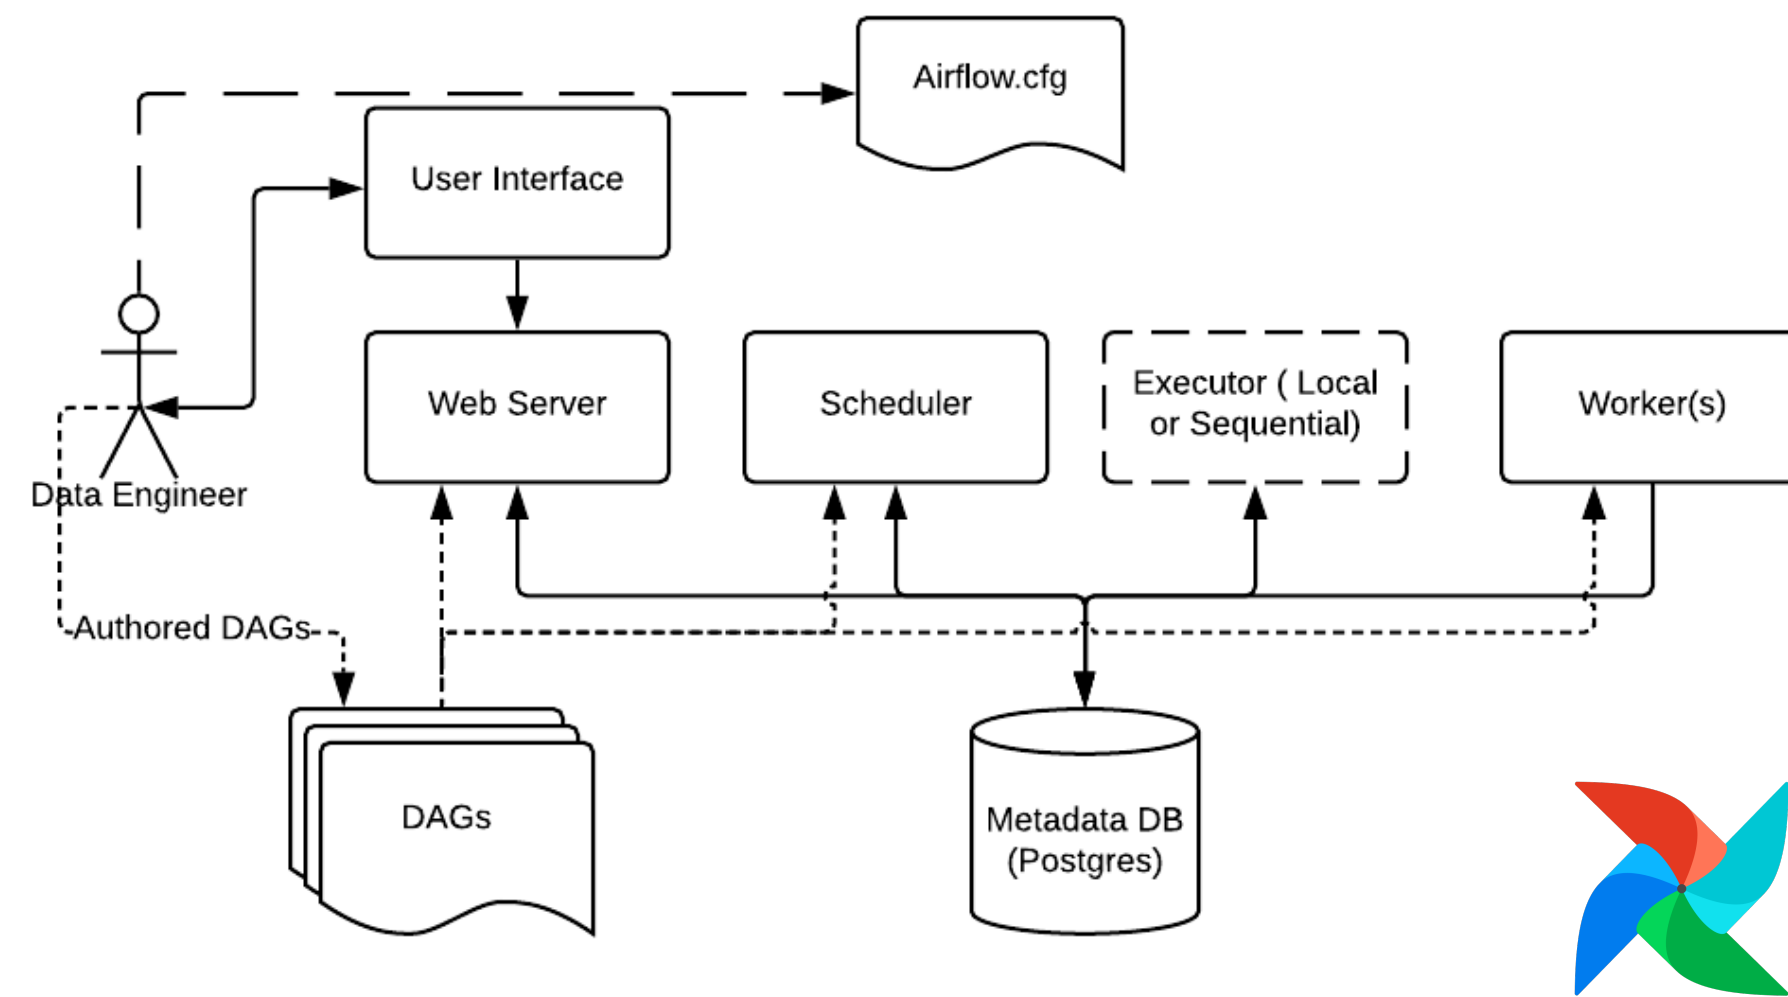
# Enter Ray



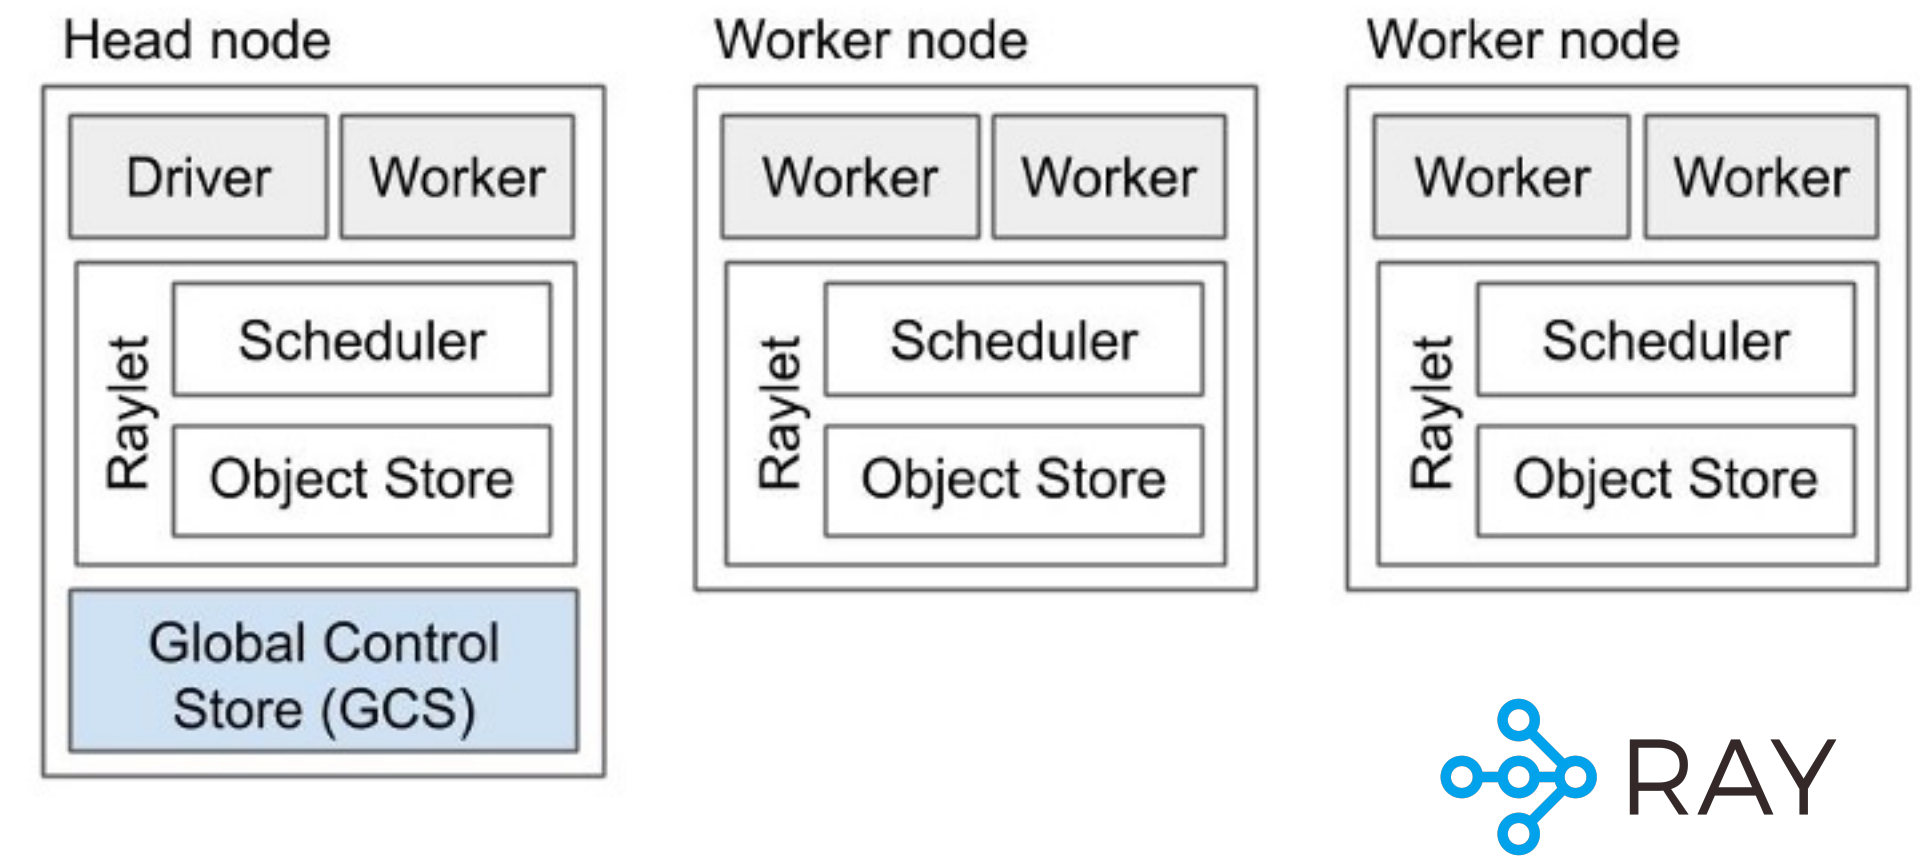
- Run the same code on your local machine, on an EC2 VM, a hardware machine, etc. with no code change!
- Native Integrations with many ML projects
- Simple setup and native pythonic library
- Options for distributed computation (dask, spark, modin)
- Ray Serve for model serving



# The Ideal Story



+



# The Taskflow API

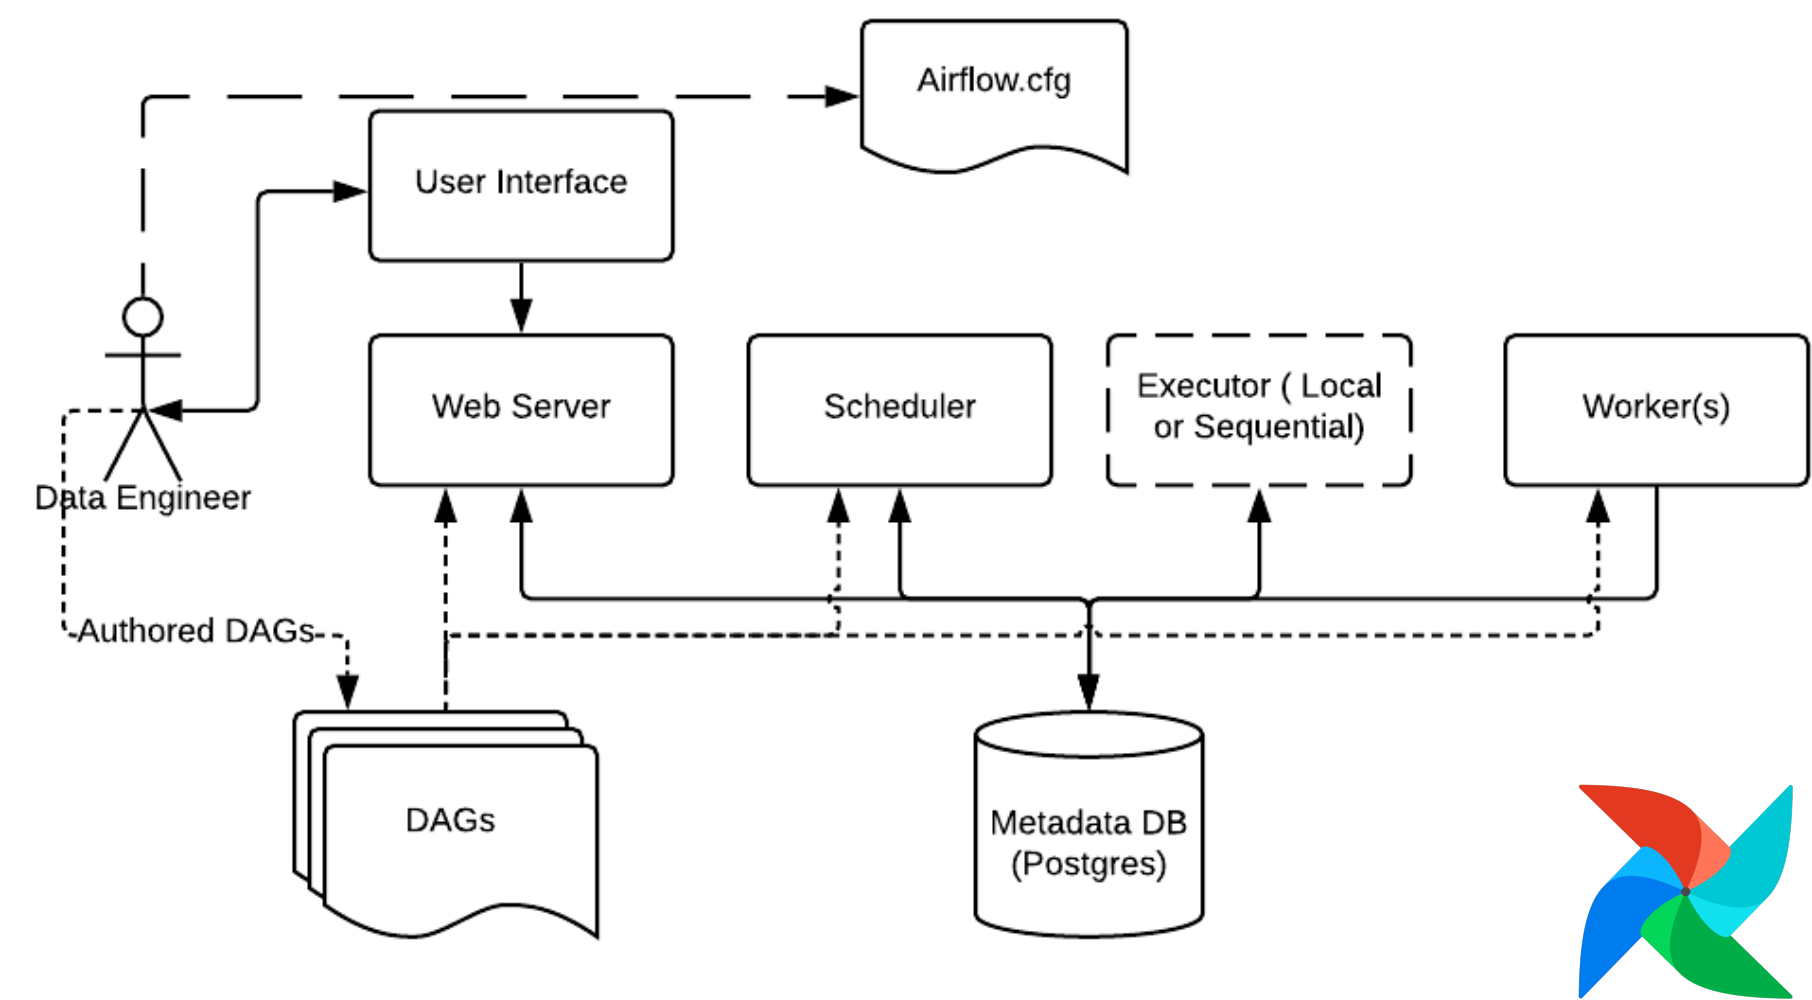
- Introduced Airflow 2.0
- Convert a python function to an Airflow task using just a single decorator!
- Pass data between tasks using functional composition
- But what if we could add the power of Ray?

```
@task
def get_initial_number():
    return 1
```

```
@task
def add_one(value):
    return value + 1
```

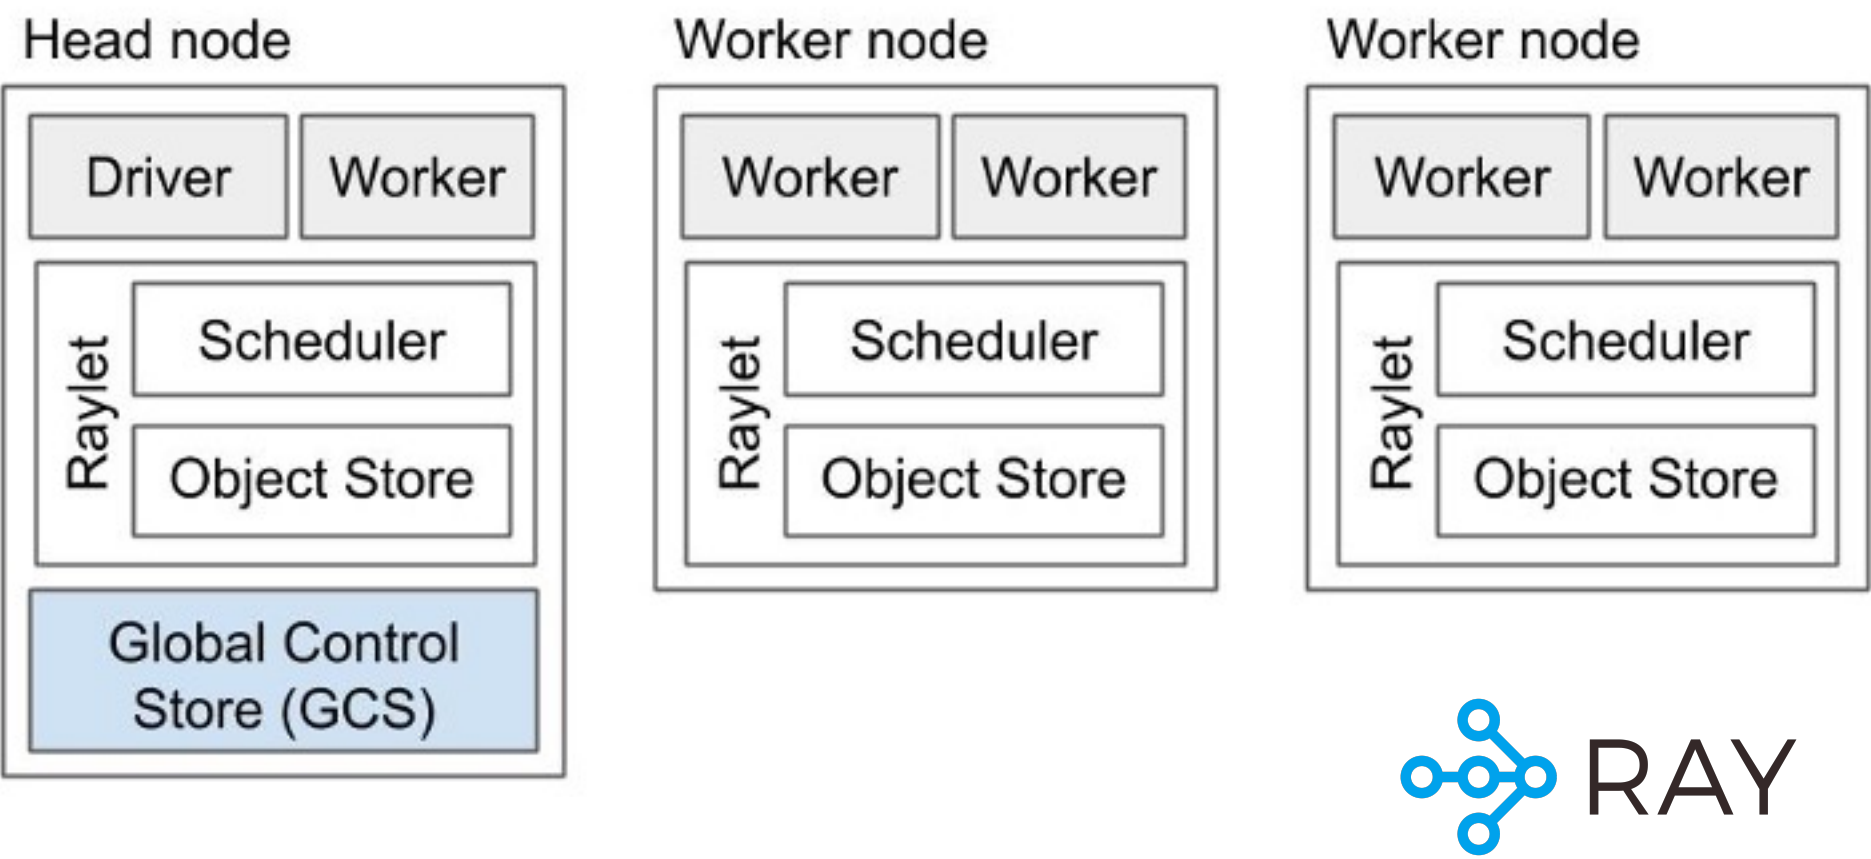
```
@dag(dag_kwargs=dag_kwargs)
def dag():
    value = get_initial_number()
    for i in range(5):
        value = add_one(value)
```

# The Ideal Story



@task

+



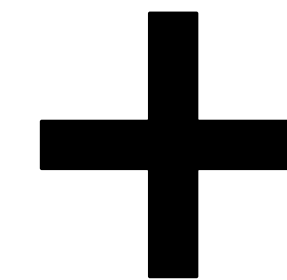
@ray.remote(num\_cpu=2)

# Introducing The Ray Decorator!

```
@ray_task
def get_initial_number():
    return 1

@ray_task(num_cpu=2)
def add_one(value):
    return value + 1

@dag(dag_kwargs=dag_kwargs)
def dag():
    value = get_initial_number()
    for i in range(5):
        value = add_one(value)
```



# Introducing The Ray Decorator!

- Automatically run your airflow tasks in your ray cluster with one line of code!
- Ability to dynamically size tasks and access large ray instances
- Intermediate values automatically stored in the plasma store for ease and data locality!

```
@ray_task
def get_initial_number():
    return 1
```

```
@ray_task(num_cpu=2)
def add_one(value):
    return value + 1
```

```
@dag(dag_kwargs=dag_kwargs)
def dag():
    value = get_initial_number()
    for i in range(5):
        value = add_one(value)
```

**The top-tier ML tooling of Ray with  
the Stability and Ecosystem of Airflow**

# From Notebook to Production

# Develop

```
In [ ]: @ray.remote
def load_dataframe() -> "ray.ObjectRef":
    """
    build dataframe from breast cancer dataset
    """

    import modin.pandas as mpd
    url = "https://archive.ics.uci.edu/ml/machine-learning-databases/" \
        "00280/HIGGS.csv.gz"
    colnames = ["label"] + ["feature-%02d" % i for i in range(1, 29)]
    data = mpd.read_csv(url, compression='gzip', names=colnames)
    print("loaded higgs")
    return data
```

```
In [ ]: @ray.remote
def split_train_test(data):
    print("Splitting Data to Train and Test Sets")
    df_train = data[(data['feature-01'] < 0.4)]
    colnames = ["label"] + ["feature-%02d" % i for i in range(1, 29)]
    train_set = xgbr.RayDMatrix(df_train, label="label", columns=colnames)
    df_validation = data[(data['feature-01'] >= 0.4) & (data['feature-01'] < 0.8)]
    test_set = xgbr.RayDMatrix(df_validation, label="label")
    print("finished data matrix")
    return train_set, test_set
```

```
In [ ]: def train_model(
    config,
    checkpoint_dir=None,
    data_dir=None,
    data=()
):
    logfile = open("/tmp/ray/session_latest/custom.log", "w")
    def write(msg):
        logfile.write(f"{msg}\n")
        logfile.flush()

    dtrain, dvalidation = data
    evallist = [(dvalidation, 'eval')]
    # evals_result = {}
    config = {
        "tree_method": "hist",
        "eval_metric": ["logloss", "error"],
    }
    print("Start training")
    bst = xgbr.train(
        params=config,
        dtrain=dtrain,
        ray_params=RAY_PARAMS,
        num_boost_round=100,
        evals=evallist,
        callbacks=[TuneReportCheckpointCallback(filename=f"model.xgb")])
```



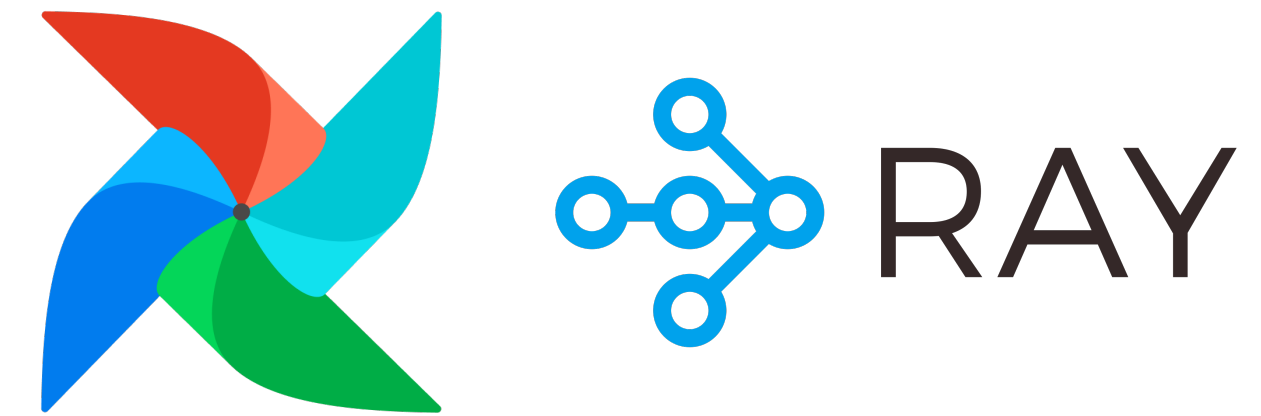
```
@ray.remote
def do_some_stuff():
    ...
```

```
do_some_stuff.remote()
```

```
@ray.remote
def do_some_stuff():
    ...
    return data
```

```
do_some_stuff.remote()
```

# Experiment



```
1  @ray_task(**task_args)
2  def train_model(
3      data
4  ):
5      train_df, validation_df = data
6      evallist = [(validation_df, 'eval')]
7      evals_result = {}
8      config = {
9          "tree_method": "hist",
10         "eval_metric": ["logloss", "error"],
11     }
12     bst = xgb.train(
13         params=config,
14         dtrain=train_df,
15         evals_result=evals_result,
16         ray_params=xgb.RayParams(max_actor_restarts=1, num_actors=8, cpus_per_actor=2),
17         num_boost_round=100,
18         evals=evallist)
19     return bst
```

```
@ray_task
def train_model():
    ...

@dag(...)
def my_dag():
    train_model()
```

# Parameterize



```
1 @ray_task(**task_args)
2 def train_model(
3     data
4 ):
5     train_df, validation_df = data
6     evallist = [(validation_df, 'eval')]
7     evals_result = {}
8     config = {
9         "tree_method": "hist",
10        "eval_metric": ["logloss", "error"],
11    }
12    bst = xgb.train(
13        params=config,
14        dtrain=train_df,
15        evals_result=evals_result,
16        ray_params=xgb.RayParams(max_actor_restarts=1, num_actors=8, cpus_per_actor=2),
17        num_boost_round=100,
18        evals=evallist)
19    return bst
```

```
data_path = \
    "{{ conf.data_path }}"
```

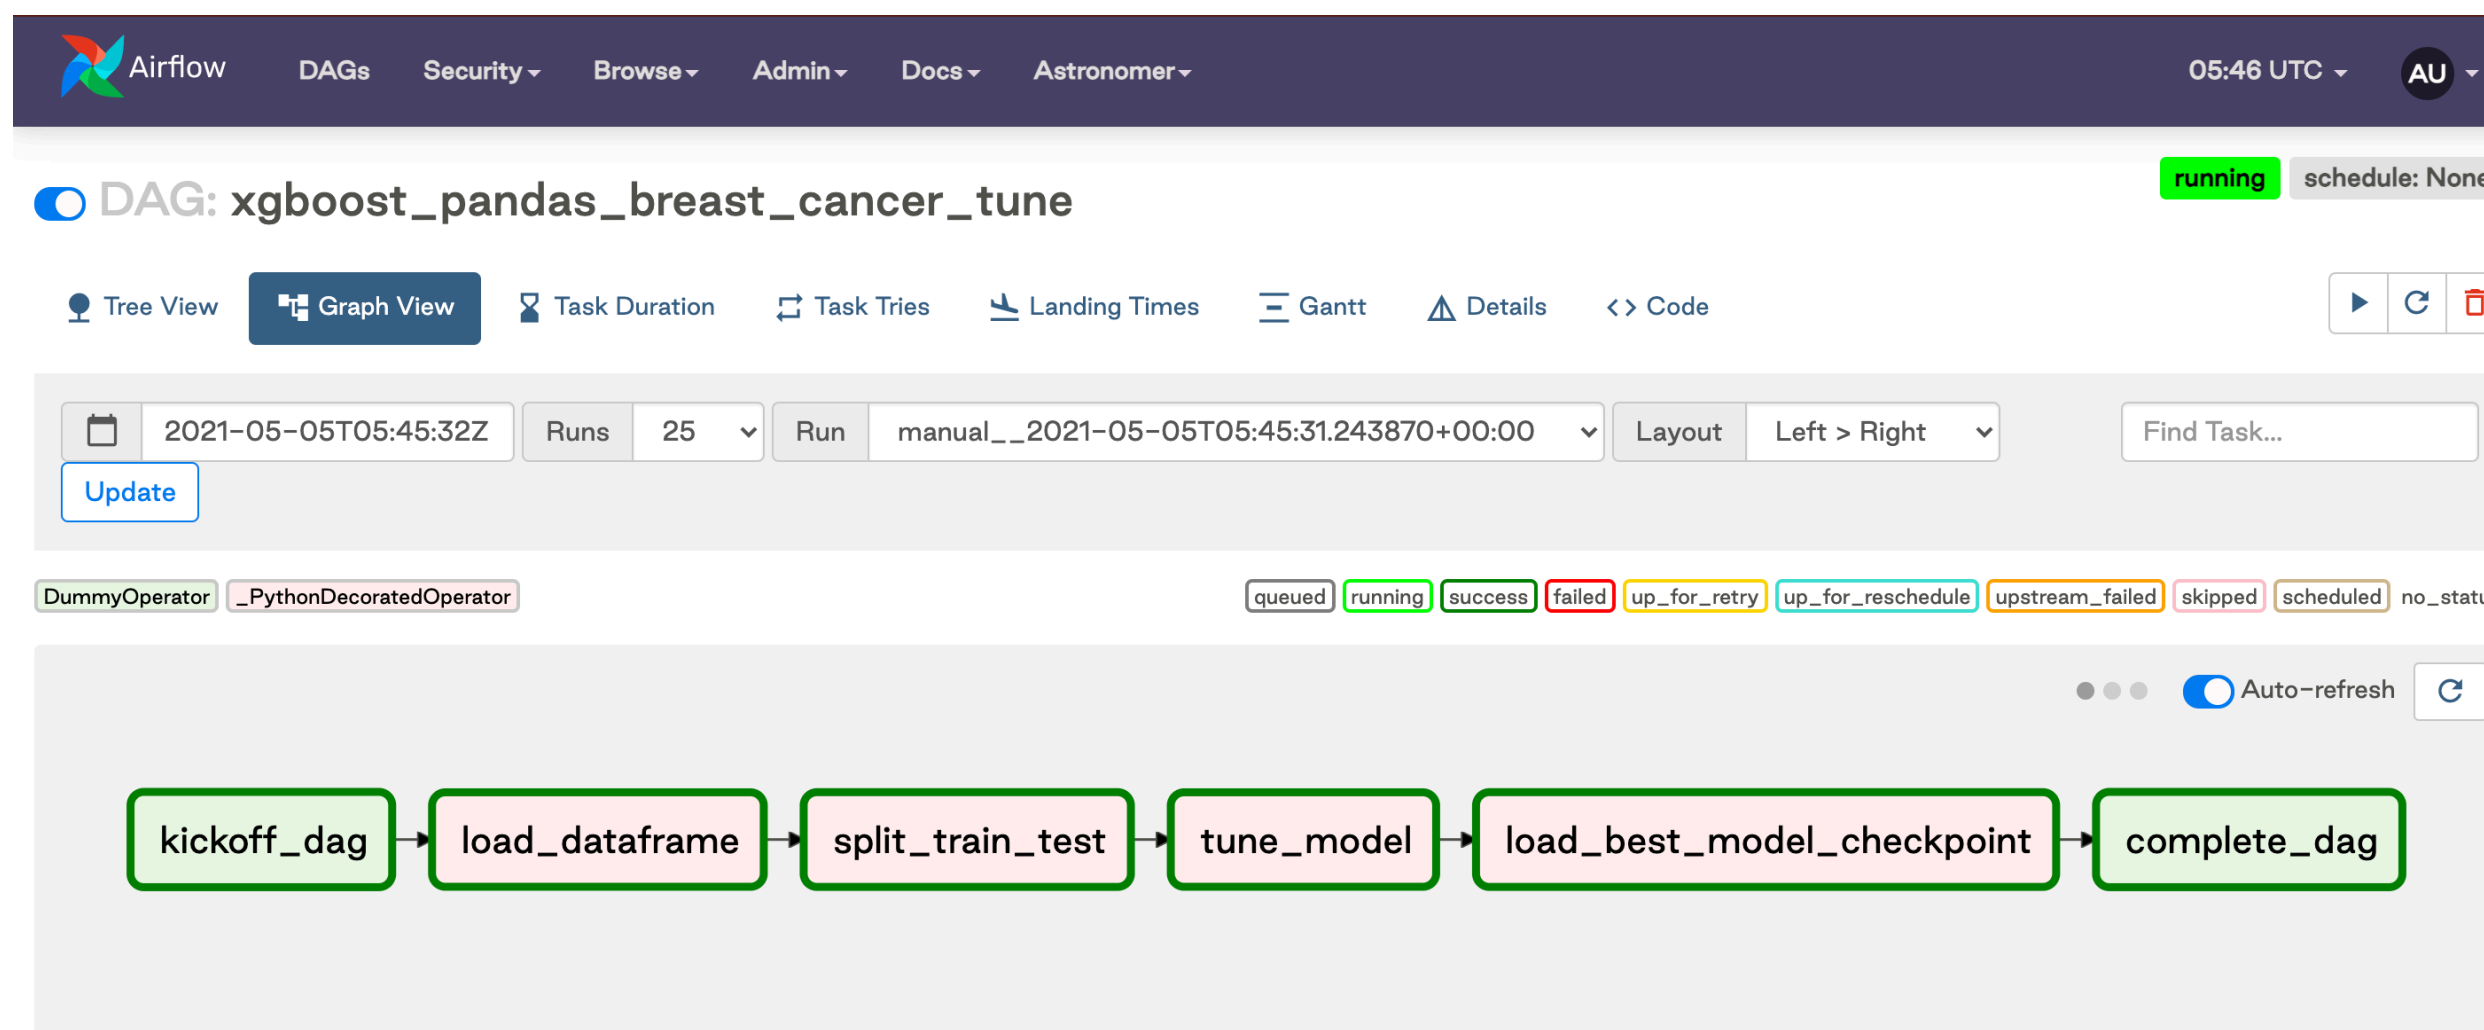
```
@ray_task
def train_model(path: str):
    ...
```

```
@dag(...)
def my_dag():
    train_model(data_path)
```

# Productionize

## Deploy your DAG!

```
1 @dag(default_args=default_args, schedule_interval=None, start_date=days_ago(2), tags=['finished-modin-ex:
2 def task_flow_xgboost_modin():
3     build_raw_df = load_dataframe()
4     data = create_data(build_raw_df)
5     trained_model = train_model(data)
6
7 task_flow_xgboost_modin = task_flow_xgboost_modin()
```



Now it's easy to:

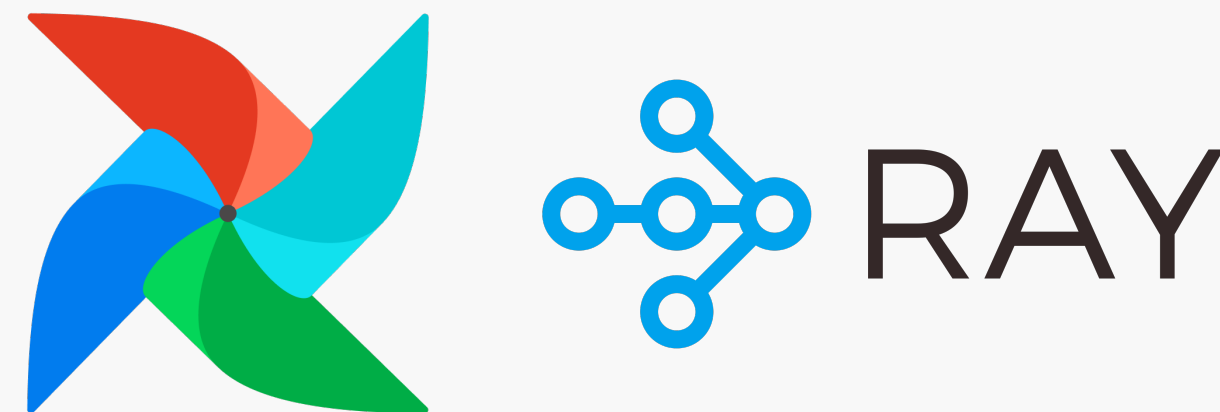
- Add more tasks & parallelize
- Tune the model(s)
- Schedule fresh updates
- Monitor for failures
- (Re)Deploy the best model(s)
- Connect to the ecosystem

# The Next Gen Airflow Data Science Story

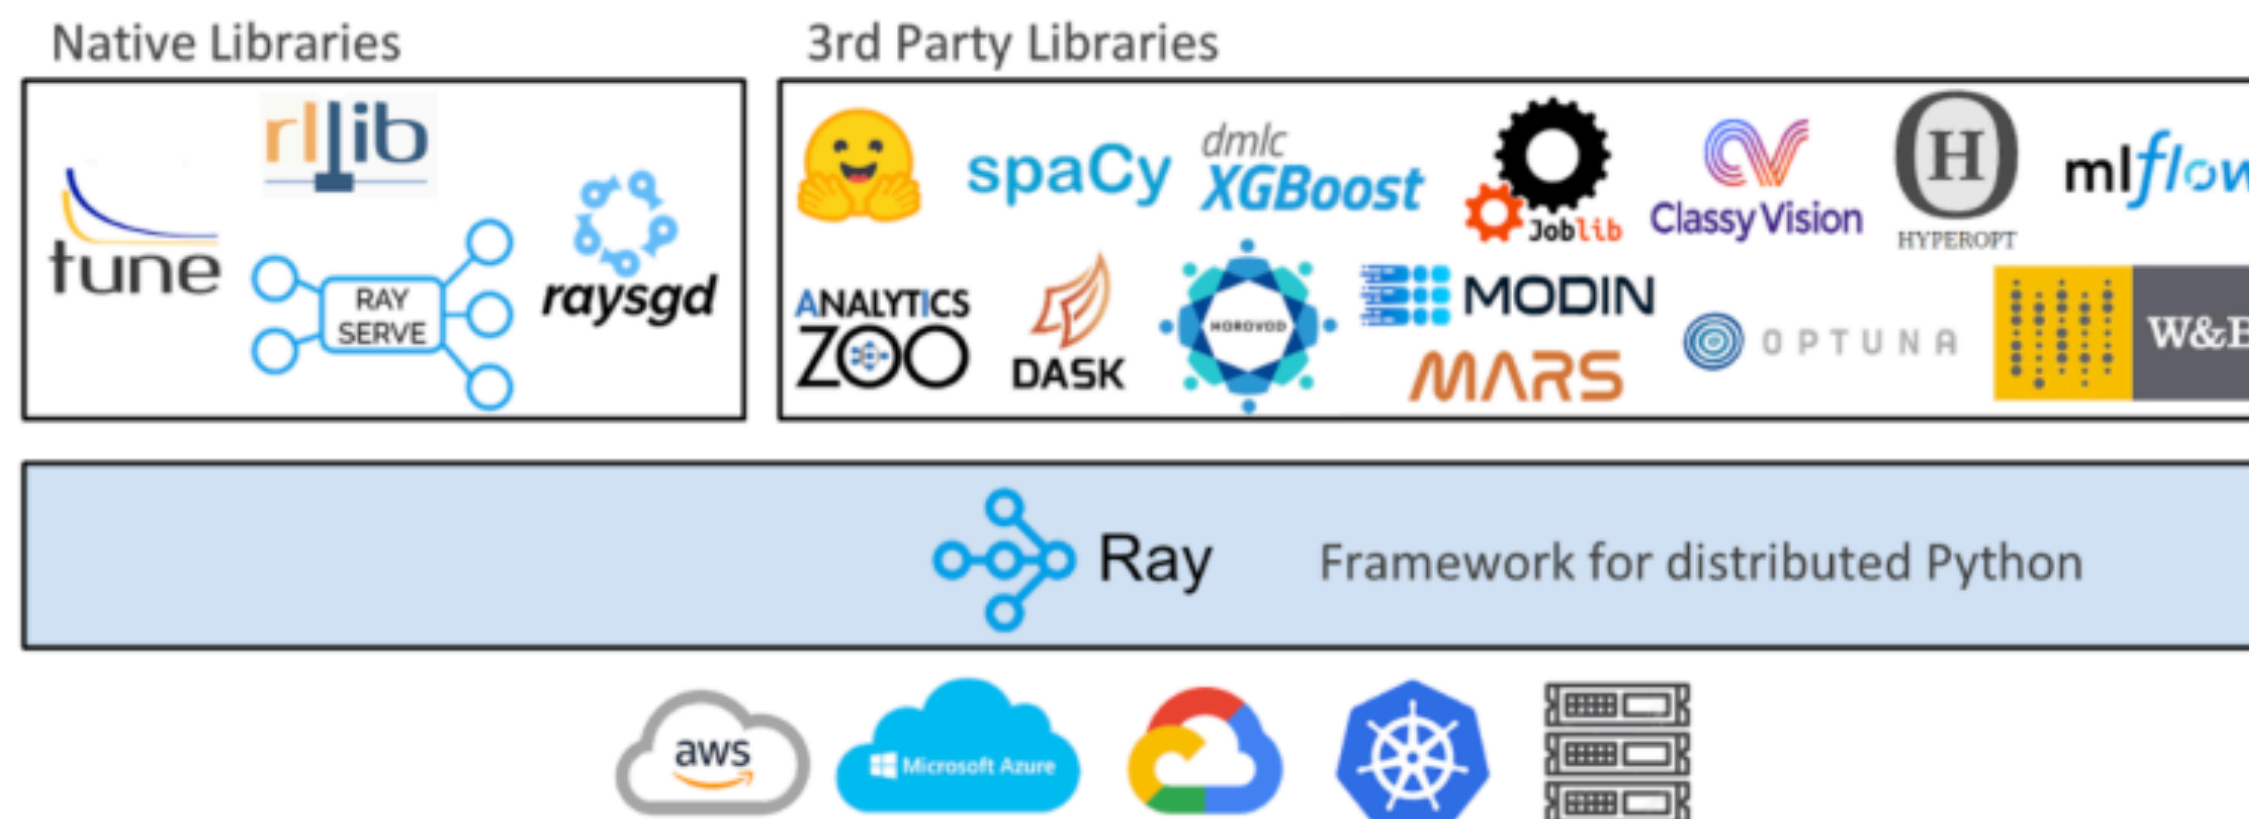
**Develop**



**Experiment &  
Parameterize**

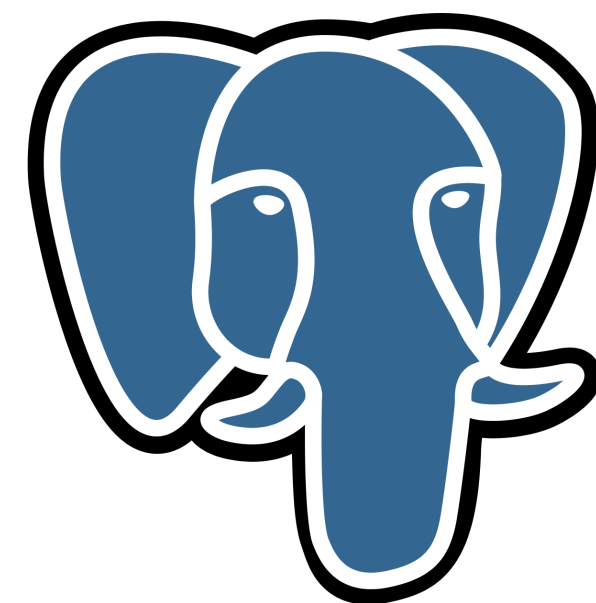


**Productionize**

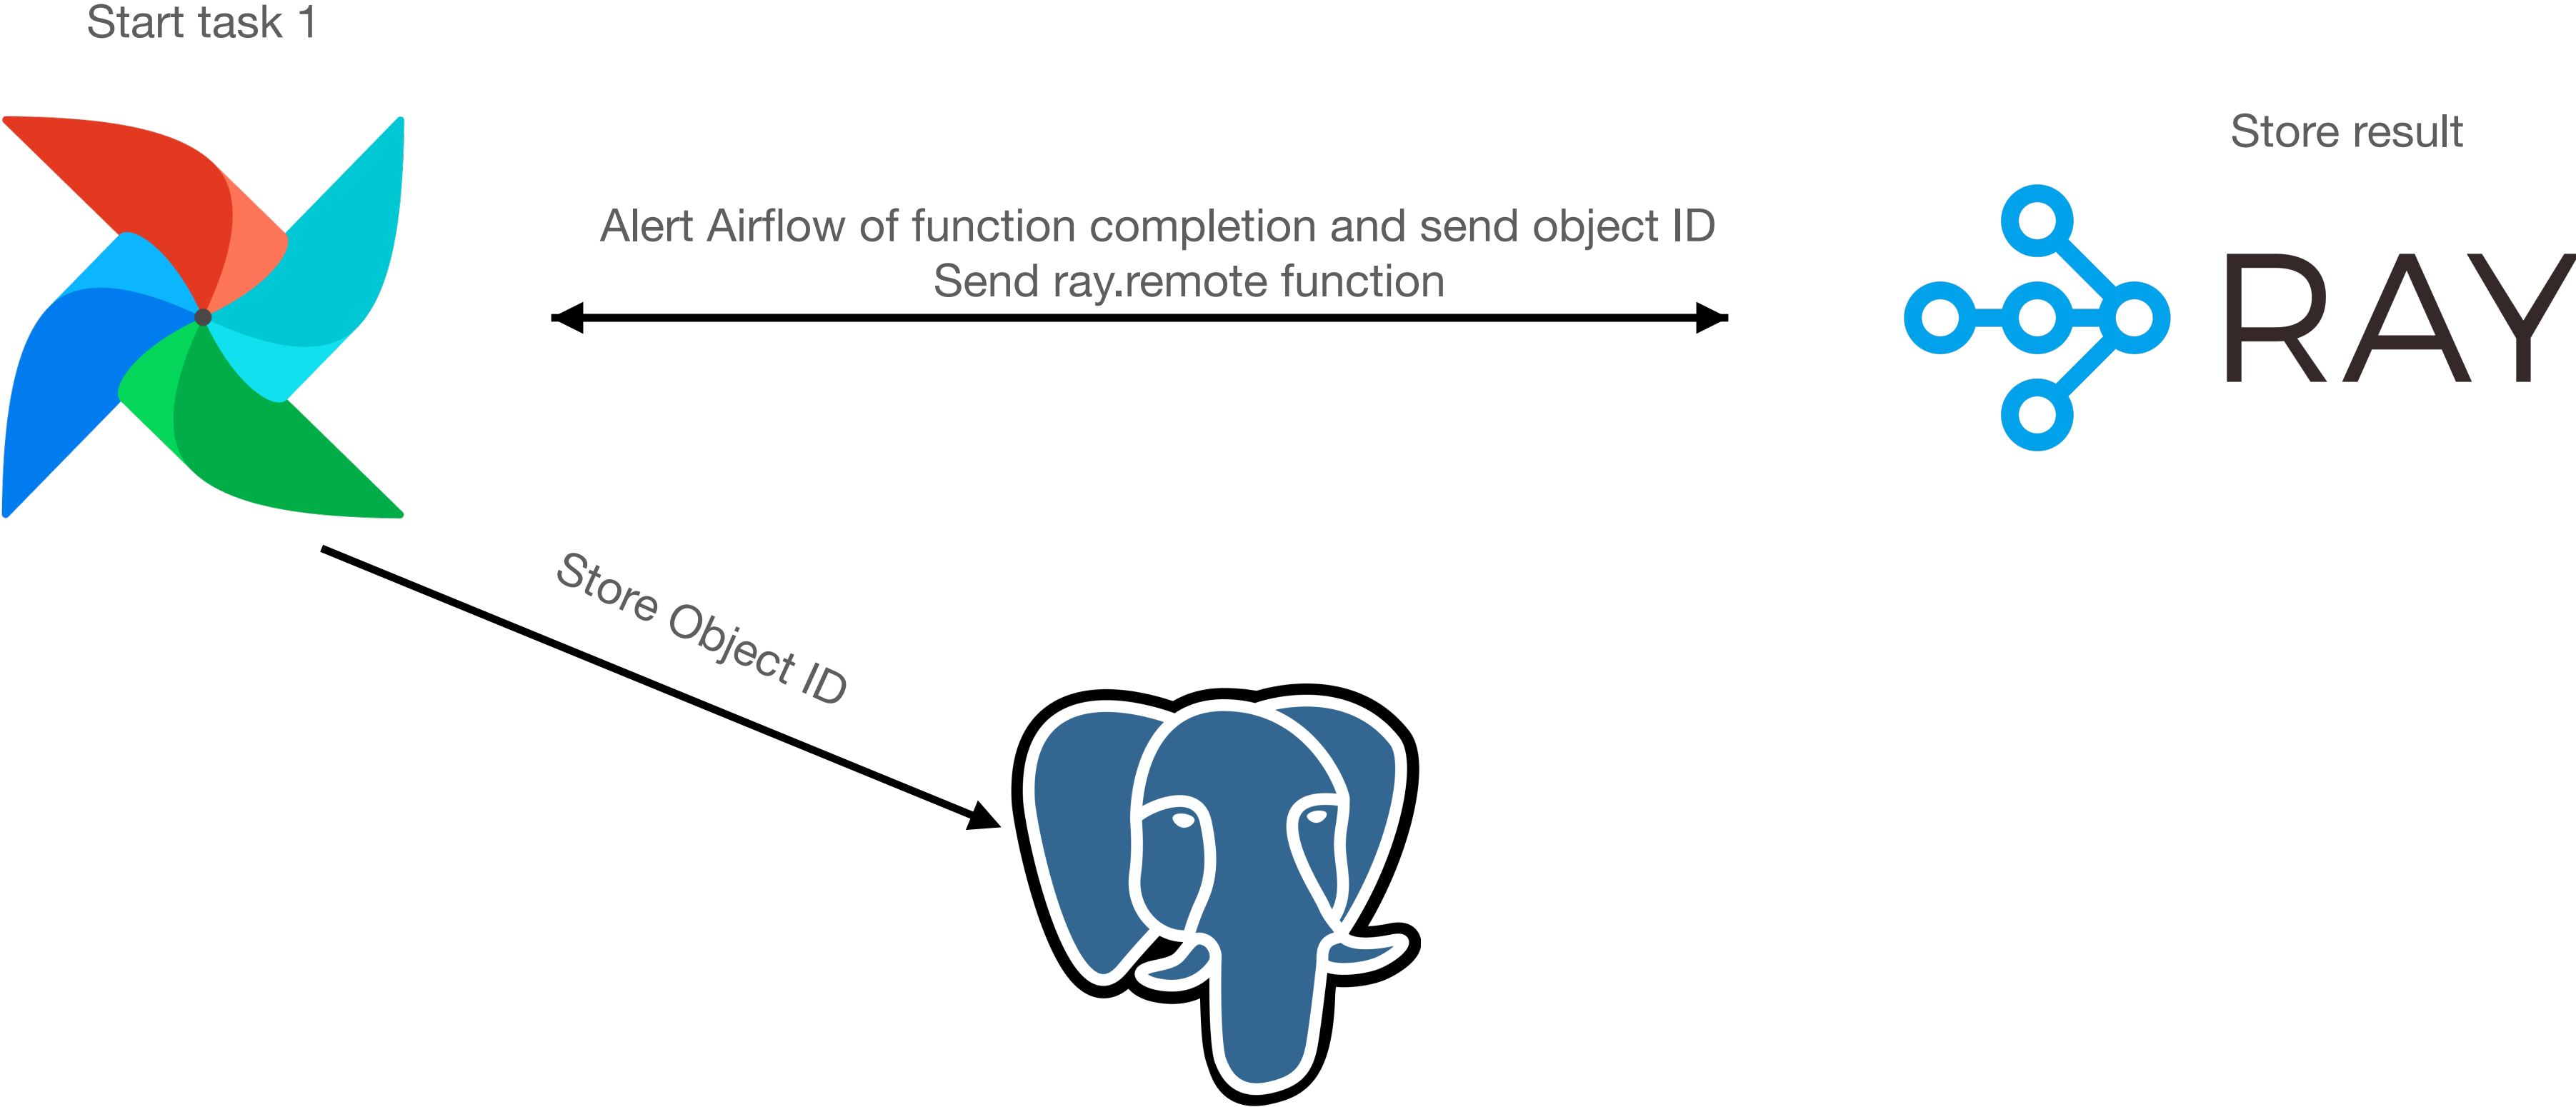


# How It Works

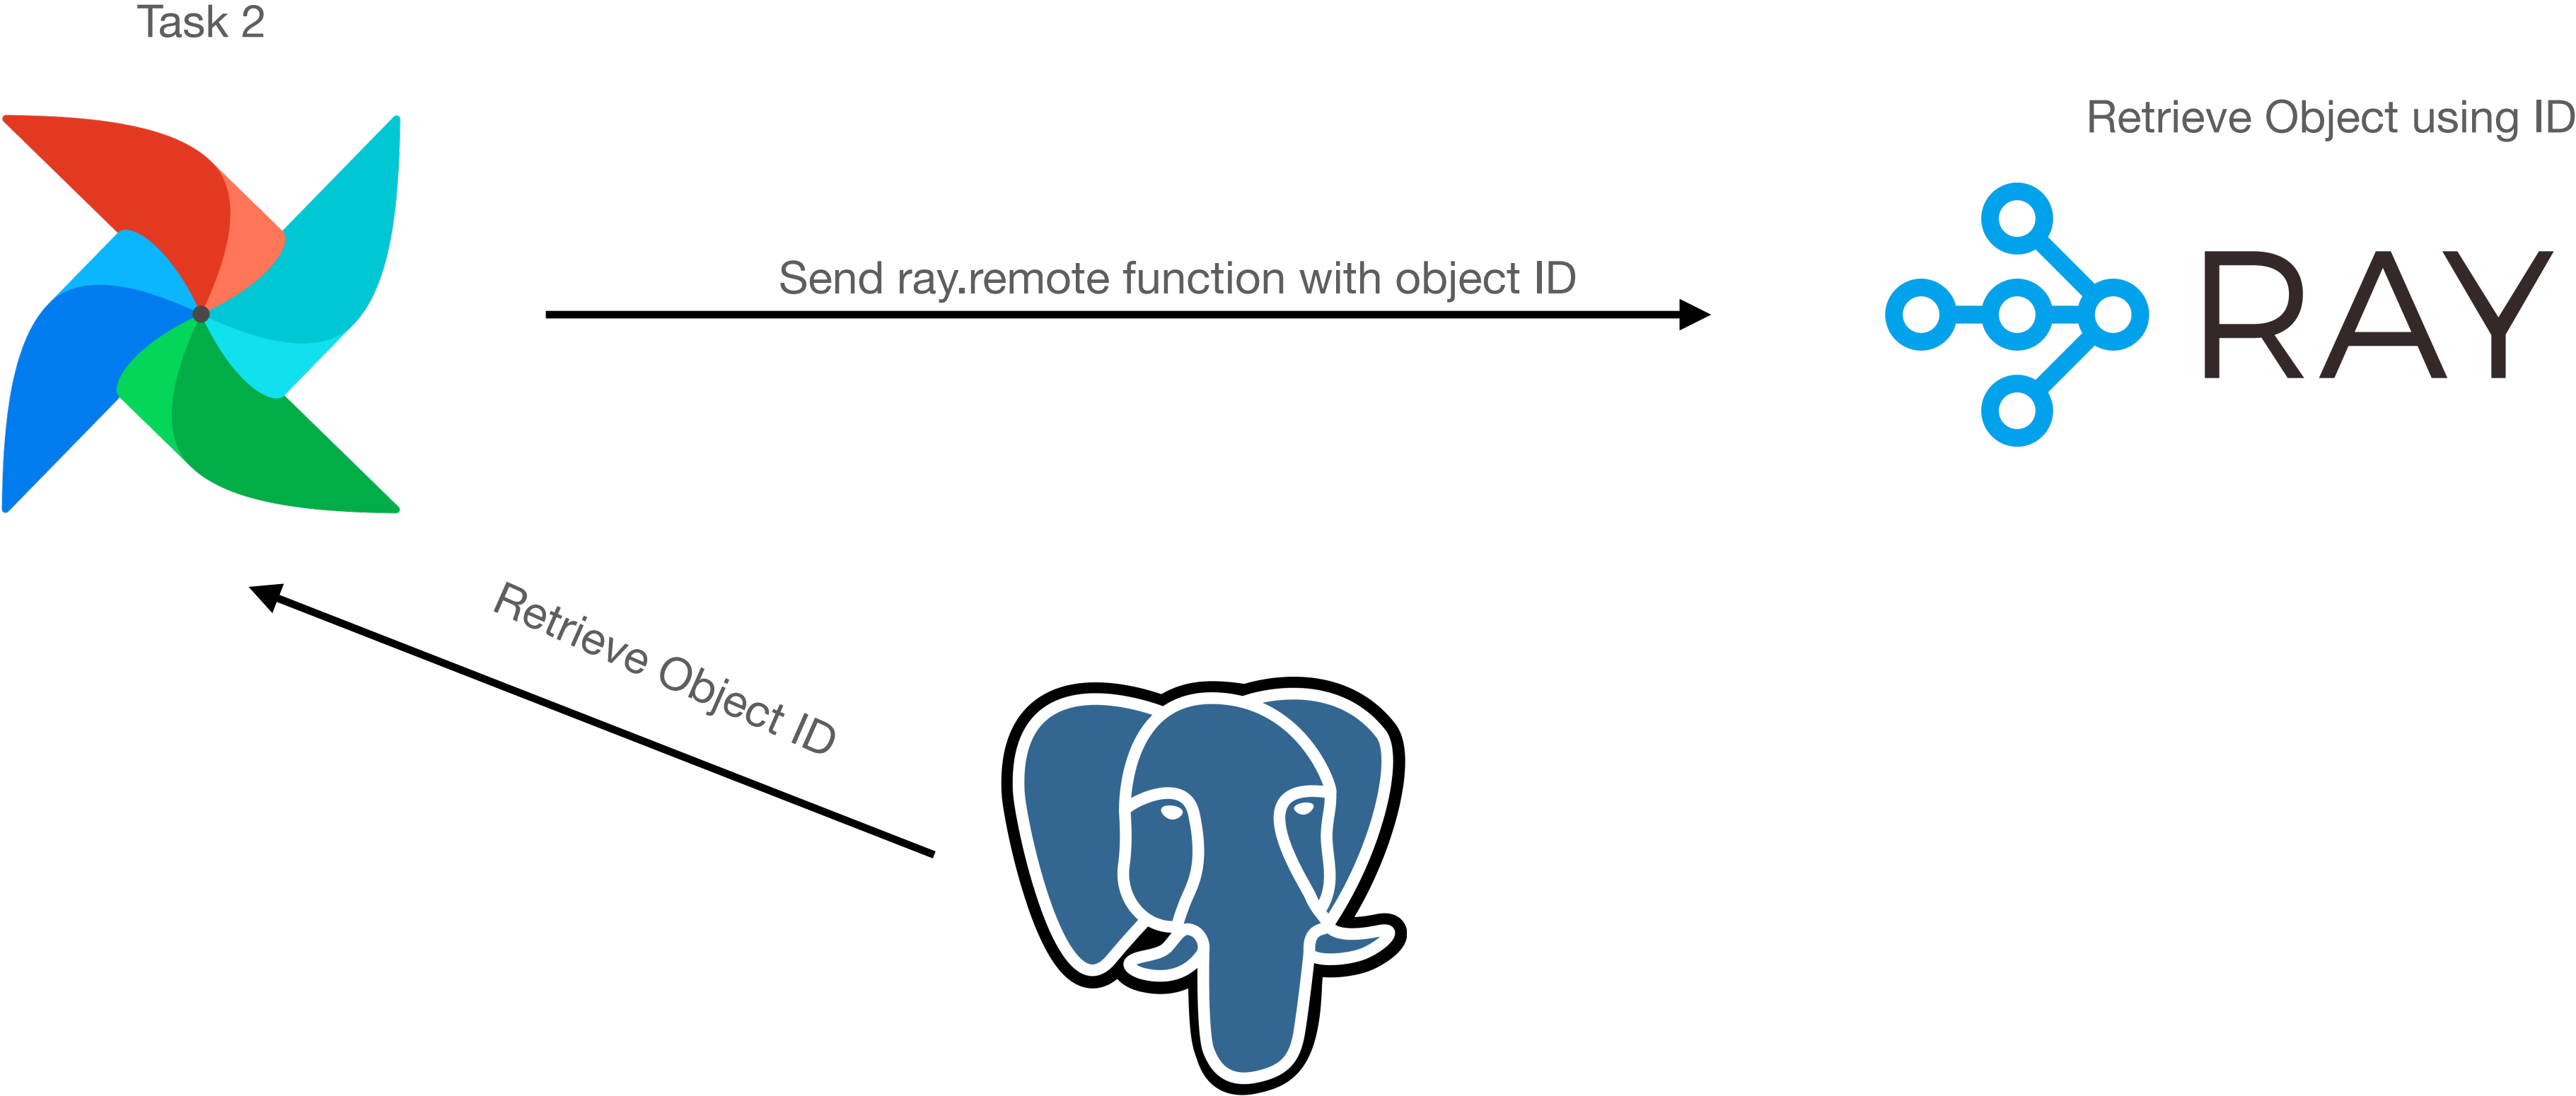
# How It Works



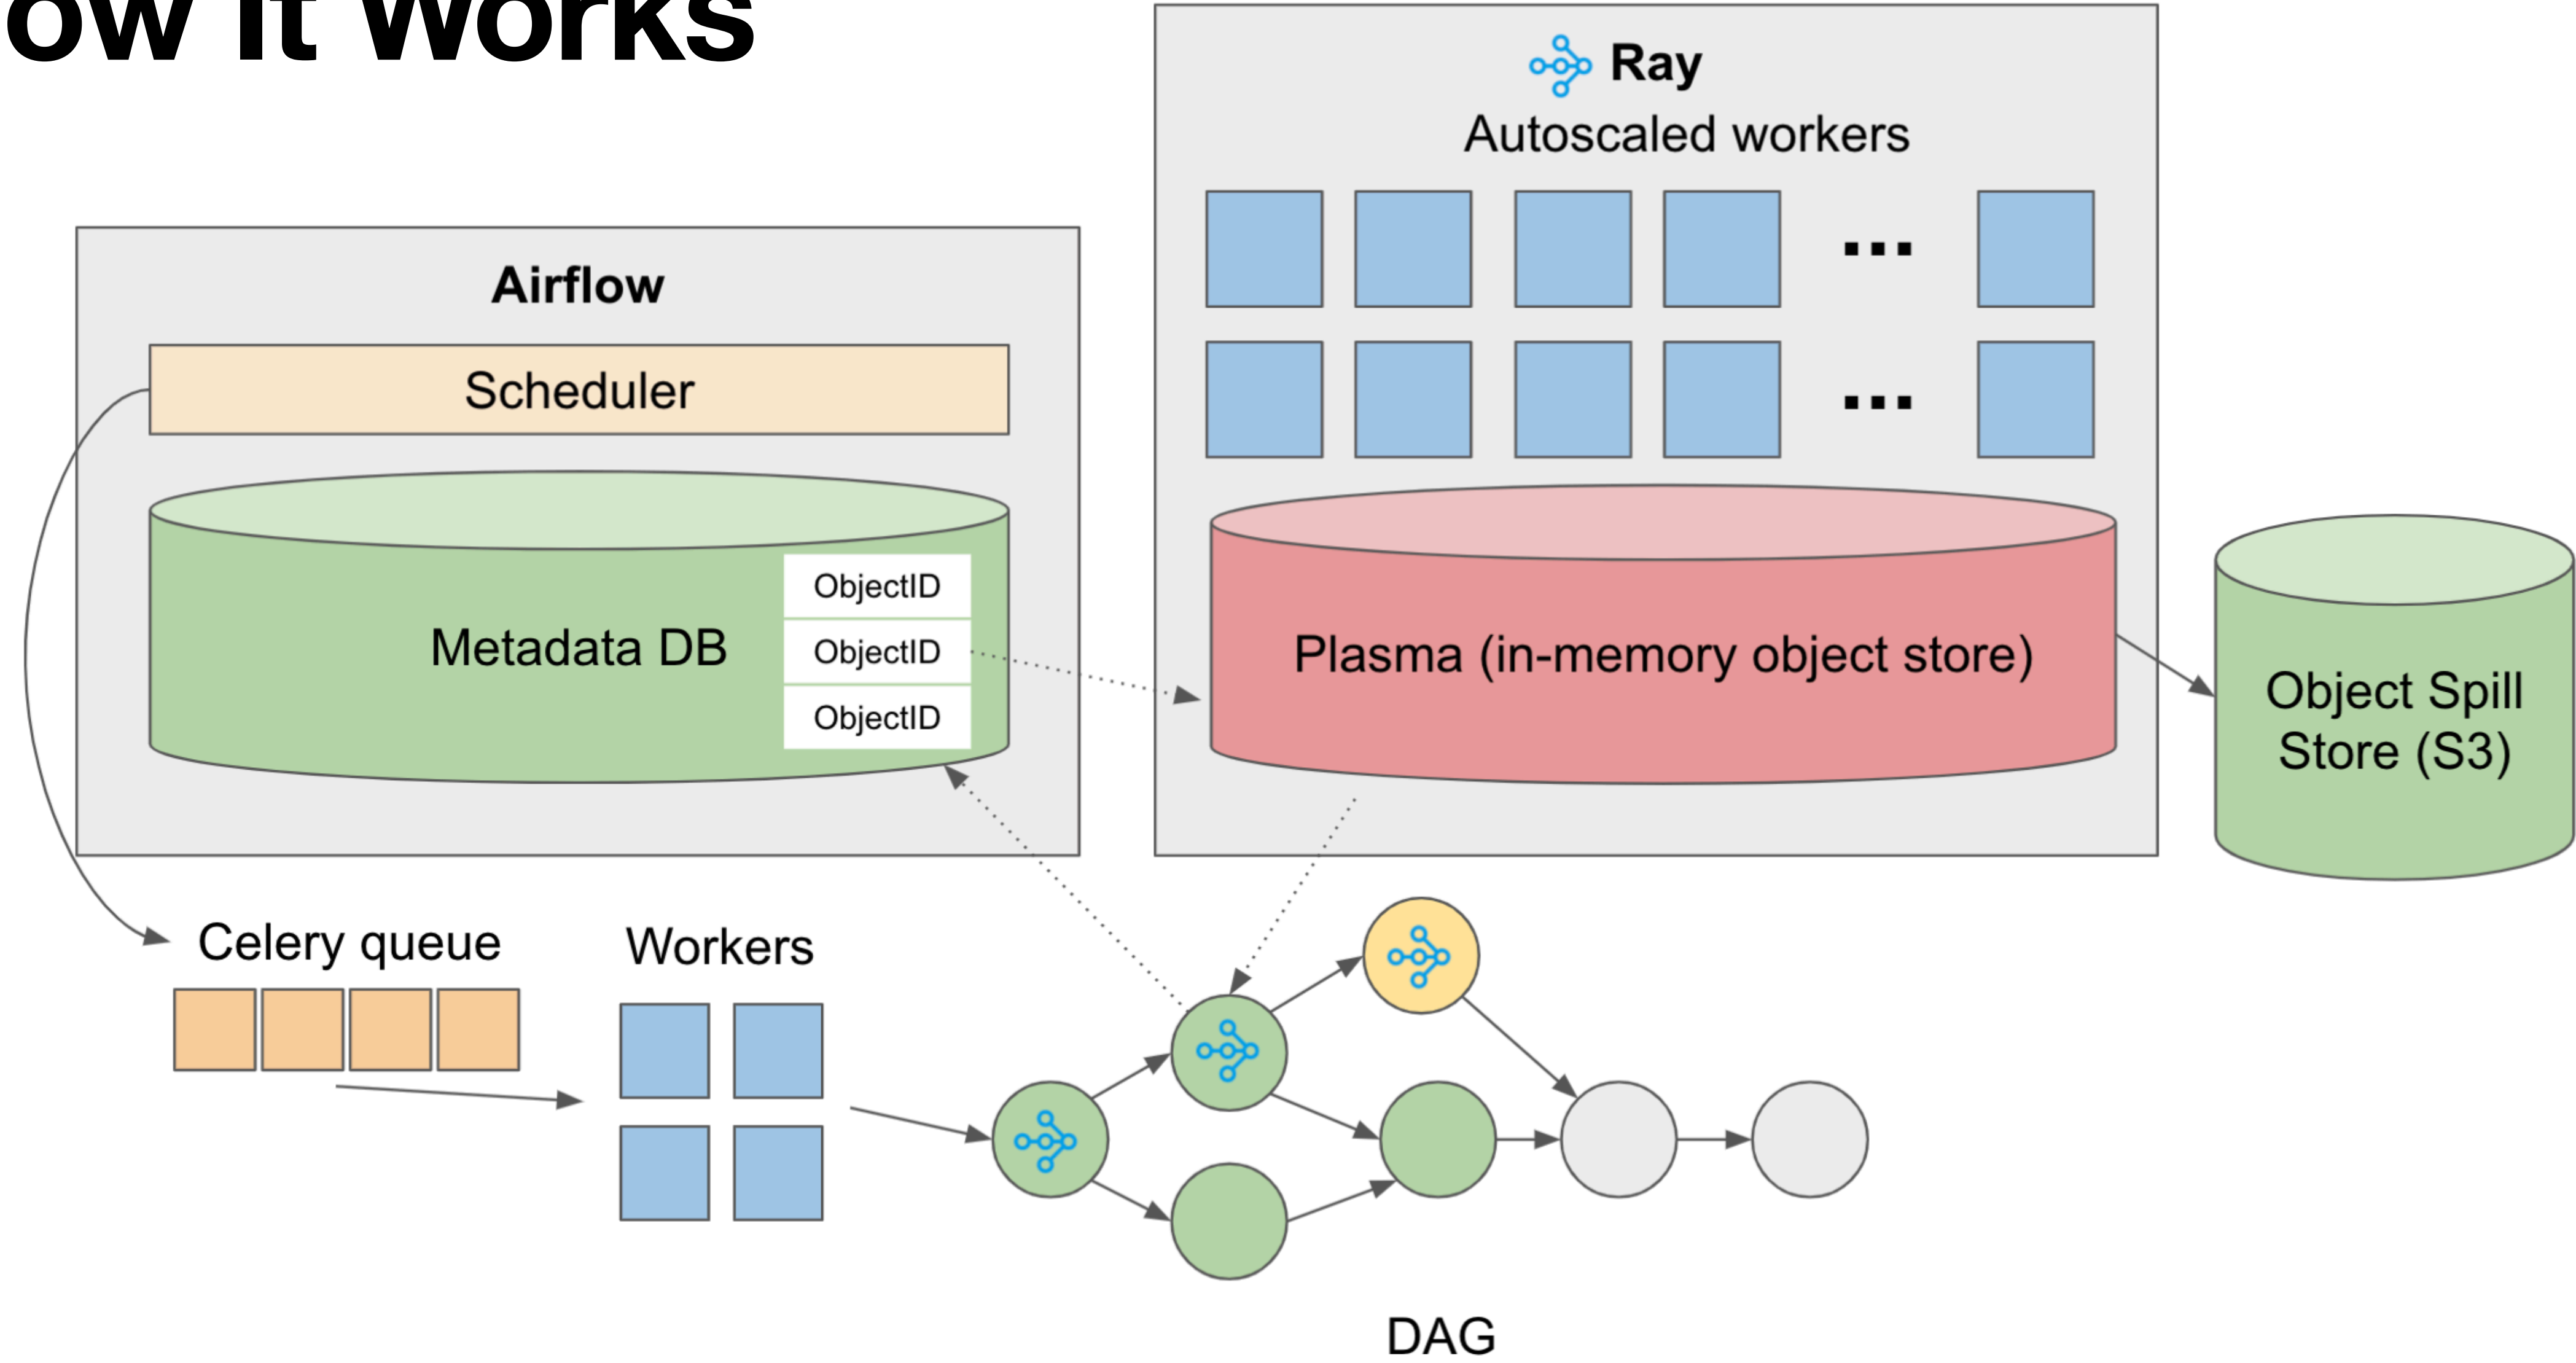
# How It Works



# How It Works



# How It Works



# Next Steps

# Checkpointing

- Store intermediate data in external data stores
- Re-run failed tasks
- Plug Tune checkpoints to model registries and experiment tracking libs
- Tweak Experiments so even if your ray cluster crashes, you will be able to restart DAG from checkpoint

```
@ray_task(checkpoint=True)
def really_long_model():
    ...
```

```
@ray_serve_task
def serve_model():
    ...
```

```
@dag(dag_kwarg=dag_kwarg)
def dag():
    model = really_long_model()
    serve_model(model)
```

# Ray serve decorator

- Deploy models to your ray cluster via airflow DAGs for instance prediction endpoints
- Composed Models = Multiple models based on business logic
- Parallelize multi-model training with Airflow

```
@ray_task
def create_model_1():
```

```
    ...
```

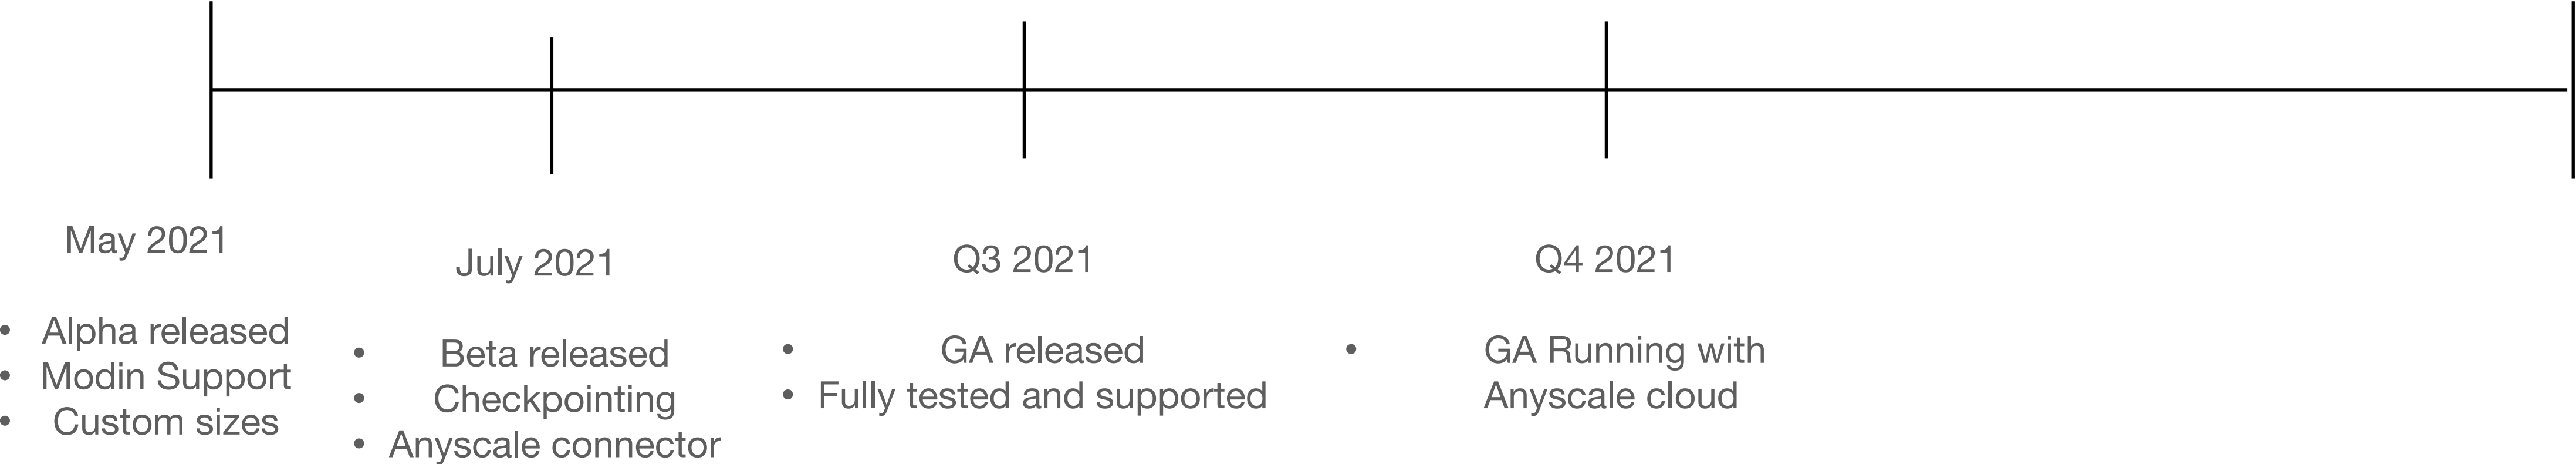
```
@ray_task
def create_model_2():
```

```
    ...
```

```
@ray_serve_task
def serve_model():
    ComposedModel.deploy()
```

```
@dag(dag_kwargs=dag_kwargs)
def dag():
    model = create_model()
    serve_model(model)
```

# Road Map

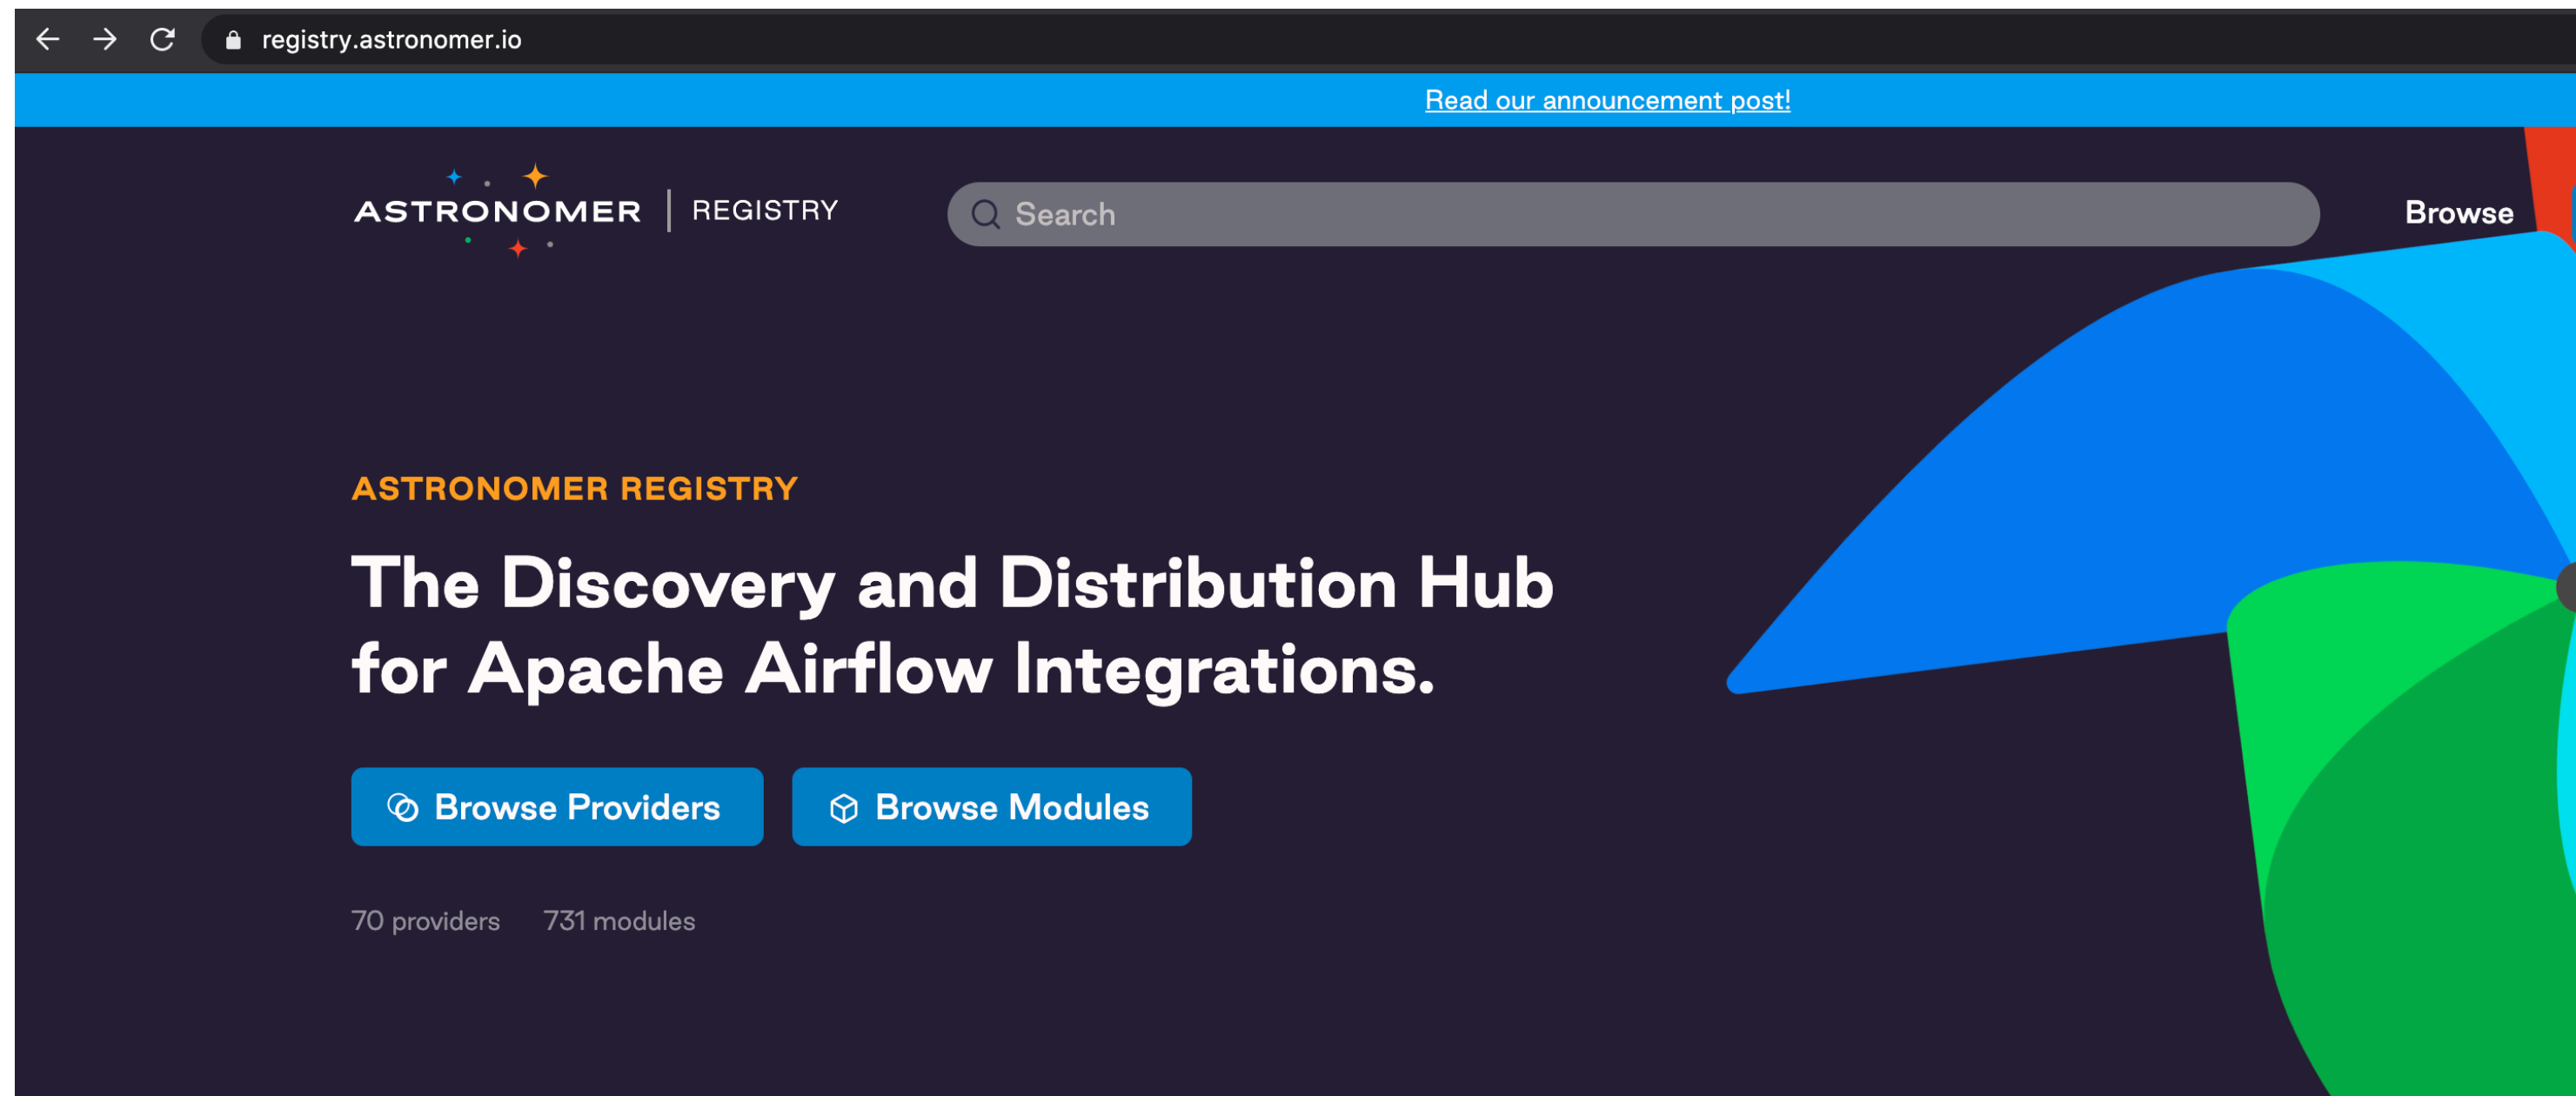


# How to Get the Ray Provider

# How to Get the Ray Provider

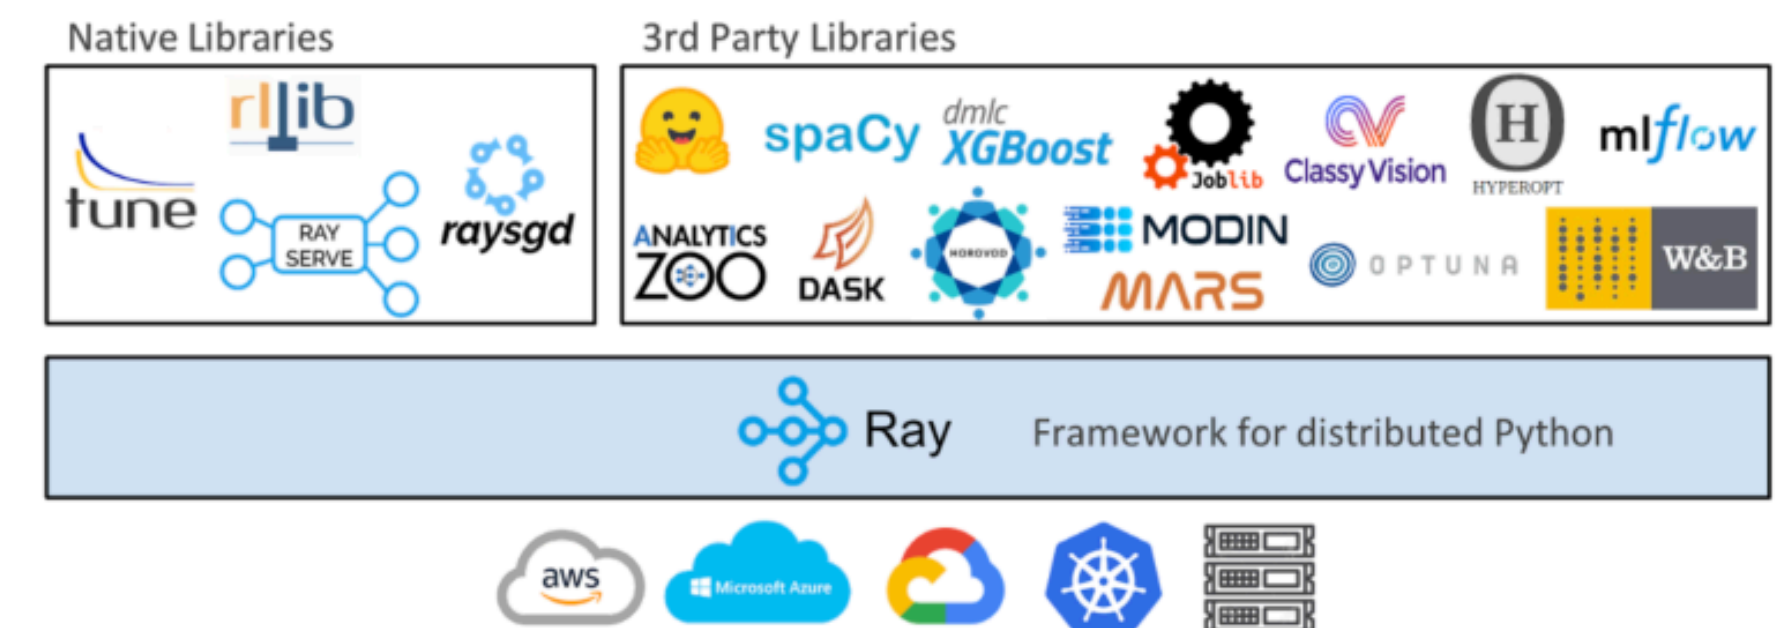


Head to <https://registry.astronomer.io/>



## Providers

Python packages containing all relevant Airflow modules for a third-party service.



# How to Get the Ray Provider



```
pip install airflow-provider-ray
```

# Thank You



@danimberman  
@ApacheAirflow

[astronomer.io](https://astronomer.io)  
@astronomerio

Special thanks to:

- Richard Liaw
- Will Drevo
- Charles Greer
- Pete DeJoy
- Rob Deeb
- Plinio Guzman

