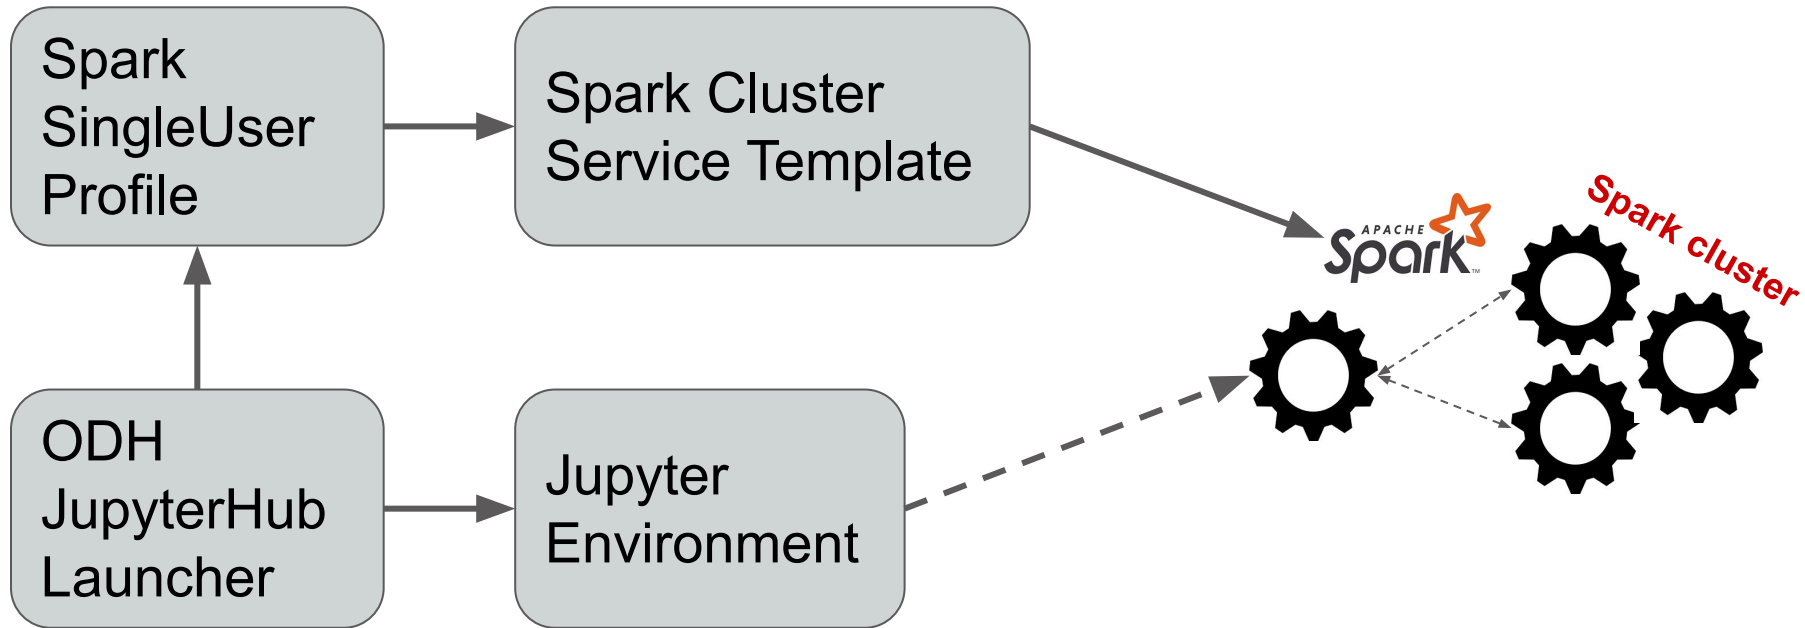

Ray on ODH

Erik Erlandson, Red Hat, Inc.

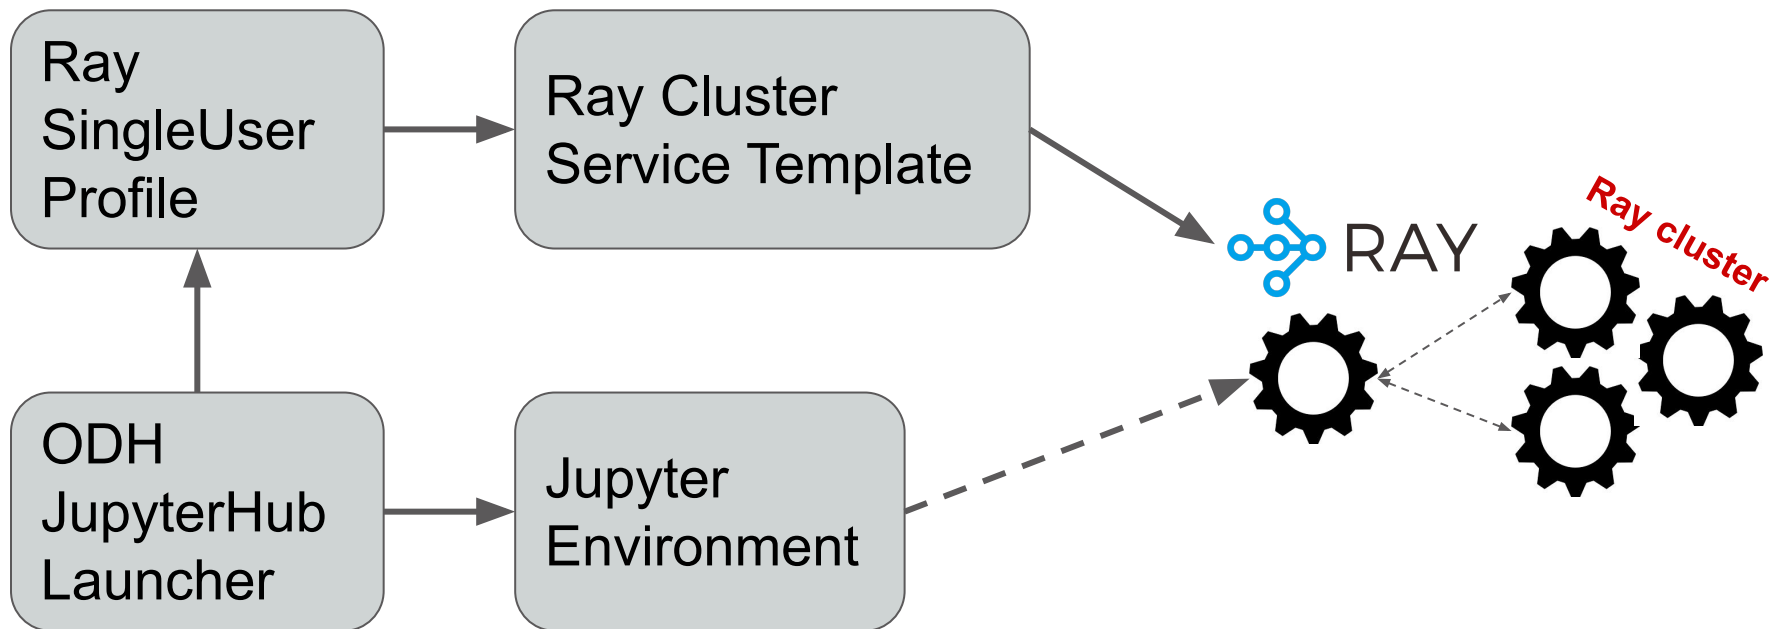
eje@redhat.com

@ManyAngled

Model: Spark on ODH



Ray on ODH?



Ecosystem Niche



Low Level

High Level

Ray's Compute Model

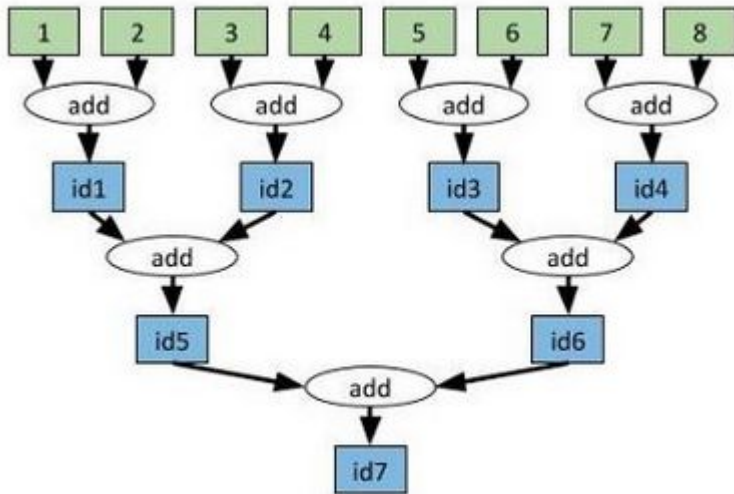
Tasks

```
@ray.remote  
def f(x):  
    # body
```

Actors

```
@ray.remote  
class foo(object):  
    # body
```

Dependency DAG



```
id1 = add.remote(1, 2)
id2 = add.remote(3, 4)
id3 = add.remote(5, 6)
id4 = add.remote(7, 8)
id5 = add.remote(id1, id2)
id6 = add.remote(id3, id4)
id7 = add.remote(id5, id6)
result = ray.get(id7)
```

<https://towardsdatascience.com/modern-parallel-and-distributed-python-a-quick-tutorial-on-ray-99f8d70369b8>

Ray's Data Model

- Plasma Object Store
- Typeless and Schemaless
- “Local First”
 - Pull remote data when needed
 - Read/Write is all local to worker node

Ray's Scheduling Model

- “Local First”
 - Prefers to run tasks on local scheduler
 - Submits to global scheduler if necessary

Native Ray Libraries

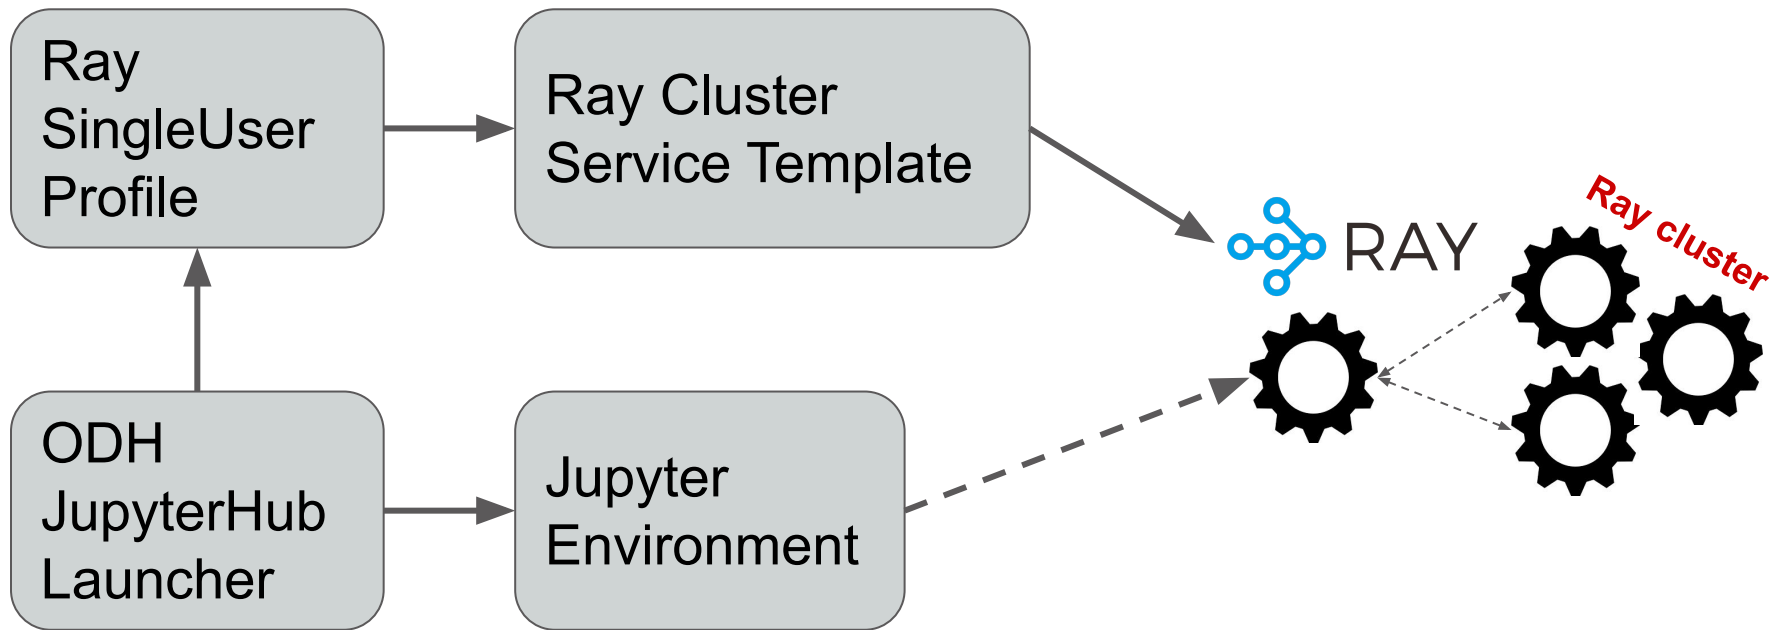
- **Tune**: Scalable Hyperparameter Tuning
- **RLlib**: Scalable Reinforcement Learning
- **RaySGD**: Distributed Training Wrappers
- **Ray Serve**: Scalable and Programmable Serving

Ray Community Integrations

- XGBoost
- Dask
- Horovod
- sklearn
- Spacy
- huggingface

<https://docs.ray.io/en/master/ray-libraries.html>

Demo: Ray on ODH!



Future

- Migrate to Operator
- Autoscaling from Operator
- Community Operator in Catalog?
- UBI based Ray images

Call To Action

- Install Demo On Your ODH
- Play with Ray on Jupyter
- Report Back!

<https://github.com/erikerlandson/ray-odh-demo>