

# Ray: The Next Generation Compute Runtime for ML Applications

---

October 24, 2022

Zhe Zhang, Head of Open Source Engineering,  
Anyscale

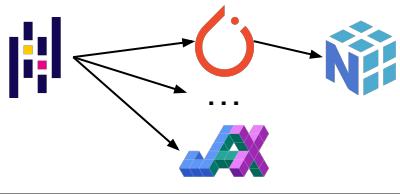


How do you manage your compute  
resources for Machine Learning?

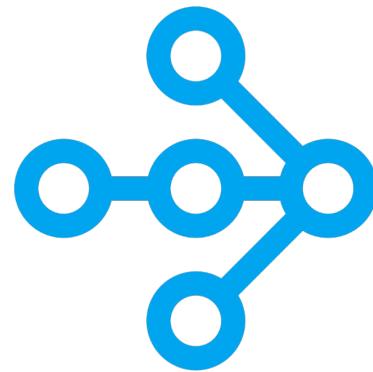
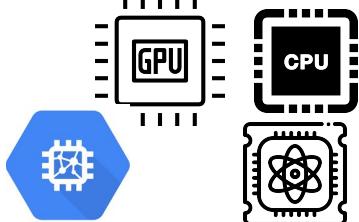
## Scalability



## Unification



## Heterogeneous Hardware



**Simple &  
Flexible API**



# Simple yet Flexible API for Distributed Computing

---



# Minimalist API



---

<code>ray.init()</code>	Initialize Ray context.
<code>@ray.remote</code>	Function or class decorator specifying that the function will be executed as a task or the class as an actor in a different process.
<code>.remote</code>	Postfix to every remote function, remote class declaration, or invocation of a remote class method. Remote operations are <i>asynchronous</i> .
<code>ray.put()</code>	Store object in object store, and return its ID. This ID can be used to pass object as an argument to any remote function or method call. This is a <i>synchronous</i> operation.
<code>ray.get()</code>	Return an object or list of objects from the object ID or list of object IDs. This is a <i>synchronous</i> (i.e., blocking) operation.
<code>ray.wait()</code>	From a list of object IDs returns (1) the list of IDs of the objects that are ready, and (2) the list of IDs of the objects that are not ready yet.

# Ray in a nutshell

Functions → Tasks : stateless computations (essentially RPC)

Classes → Actors : stateful computations

Distributed futures : enables parallelism

In-memory object store: enables passing args/results by reference

Extension of existing languages, rather than a new language

# Function

```
def read_array(file):
    # read ndarray "a"
    # from "file"
    return a

def add(a, b):
    return np.add(a, b)

a = read_array(file1)
b = read_array(file2)
sum = add(a, b)
```

# Class

```
class Counter(object):
    def __init__(self):
        self.value = 0
    def inc(self):
        self.value += 1
    return self.value

c = Counter()
c.inc()
c.inc()
```



# Function → Task

```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
a = read_array(file1)  
b = read_array(file2)  
sum = add(a, b)
```

# Class → Actor

```
@ray.remote  
class Counter(object):  
    def __init__(self):  
        self.value = 0  
    def inc(self):  
        self.value += 1  
        return self.value
```

```
c = Counter()  
c.inc()  
c.inc()
```

# Function → Task

```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
id1 = read_array.remote(file1)  
id2 = read_array.remote(file2)  
id = add.remote(id1, id2)  
sum = ray.get(id)
```

# Object → Actor

```
@ray.remote  
class Counter(object):  
    def __init__(self):  
        self.value = 0  
    def inc(self):  
        self.value += 1  
        return self.value
```

```
c = Counter.remote()  
id4 = c.inc.remote()  
id5 = c.inc.remote()
```

# Task API

id1: distributed future (object id)

```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

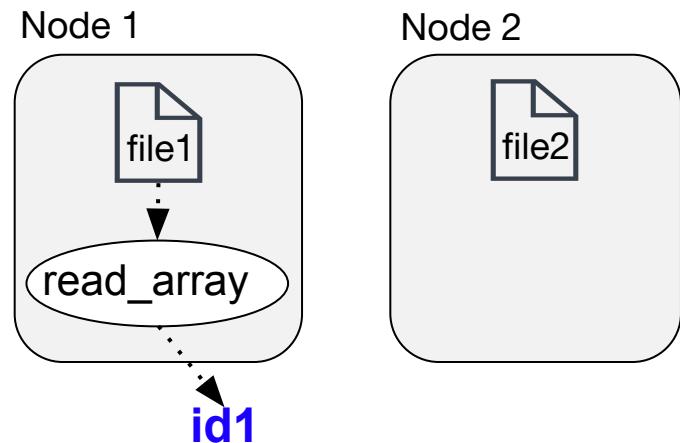
```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
id1 = read_array.remote(file1)
```

```
id2 = read_array.remote(file2)
```

```
id = add.remote(id1, id2)
```

```
sum = ray.get(id)
```



Return future id1; before  
read\_array() finishes

# Task API

```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

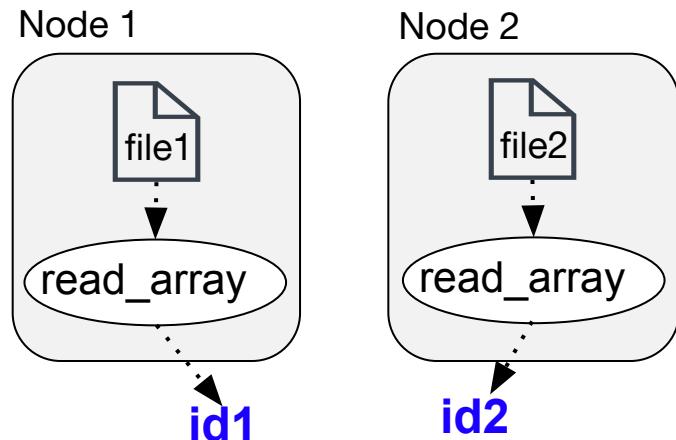
```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
id1 = read_array.remote(file1)
```

```
id2 = read_array.remote(file2)
```

```
id = add.remote(id1, id2)
```

```
sum = ray.get(id)
```



Dynamic task graph:  
build at runtime

# Task API

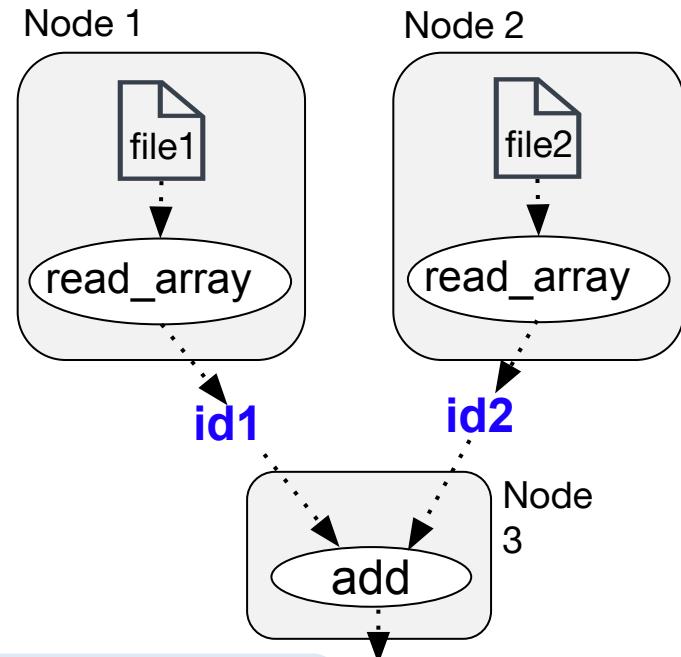
```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
id1 = read_array.remote(file1)  
id2 = read_array.remote(file2)  
id = add.remote(id1, id2)
```

```
sum = ray.get(id)
```

Every task submitted,  
but not finished yet



# Task API

```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

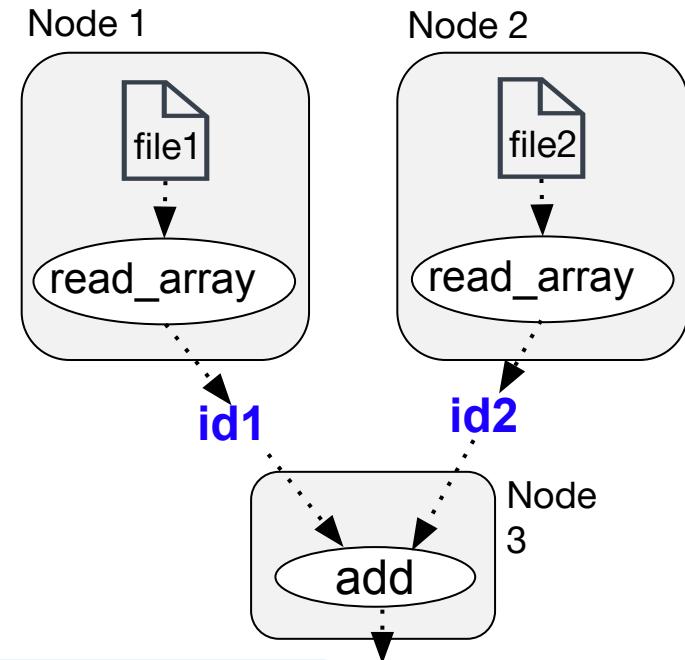
```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
id1 = read_array.remote(file1)  
id2 = read_array.remote(file2)  
id = add.remote(id1, id2)
```

```
sum = ray.get(id)
```

ray.get() block until  
result available

id



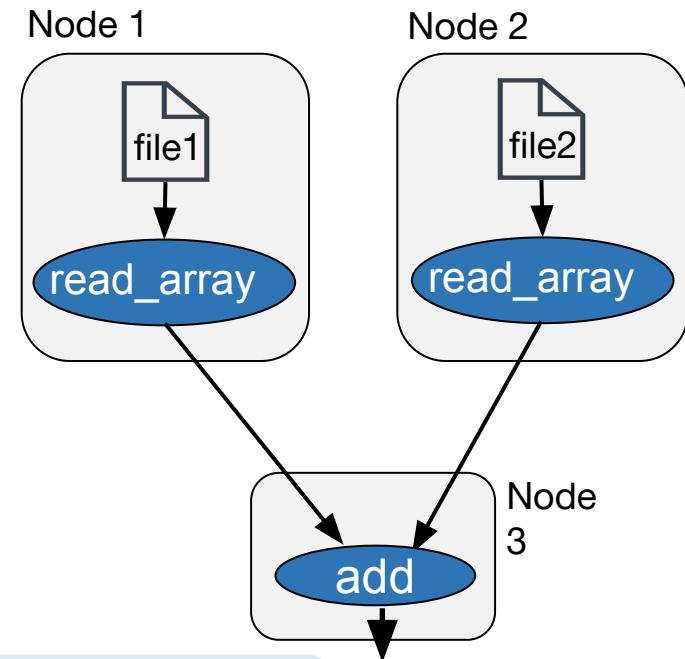
# Task API

```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
id1 = read_array.remote(file1)  
id2 = read_array.remote(file2)  
id = add.remote(id1, id2)
```

```
sum = ray.get(id)
```

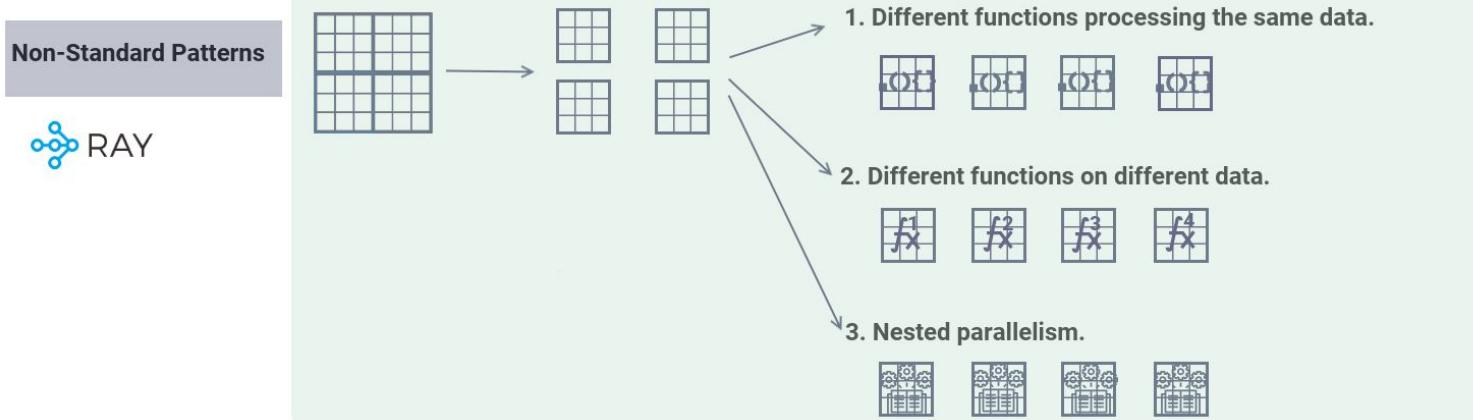
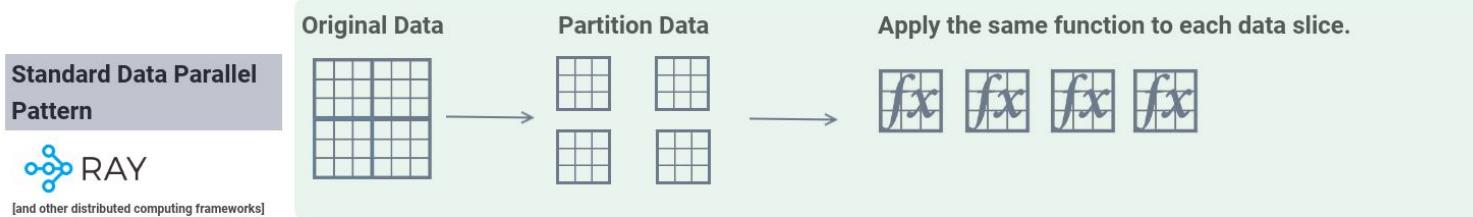


Task graph executed to  
compute sum



# Using Ray's API for Data Parallelism

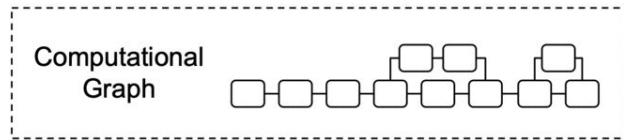
Ray can be used to parallelize complex computation patterns



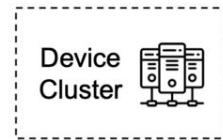
anyscale

# Advanced Usage of Ray's API

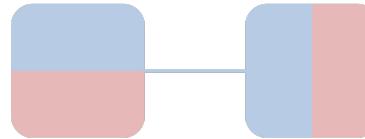
User Input



+



# Alpa



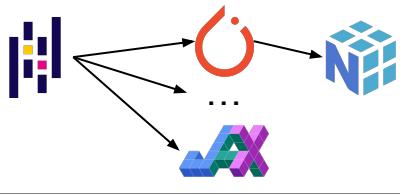
Source:

<https://ai.googleblog.com/2022/05/alpa-a-utomated-model-parallel-deep.html>

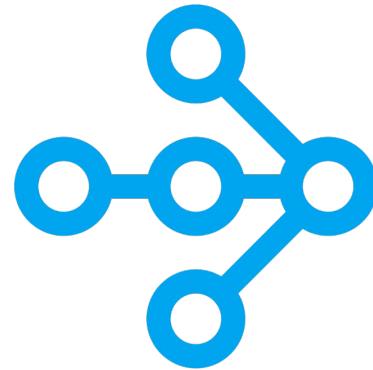
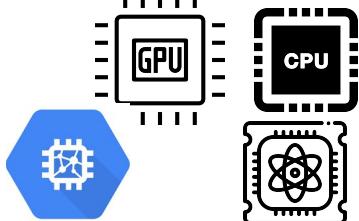
## Scalability



## Unification



## Heterogeneous Hardware



**Simple &  
Flexible API**



# RAY

# Scalability

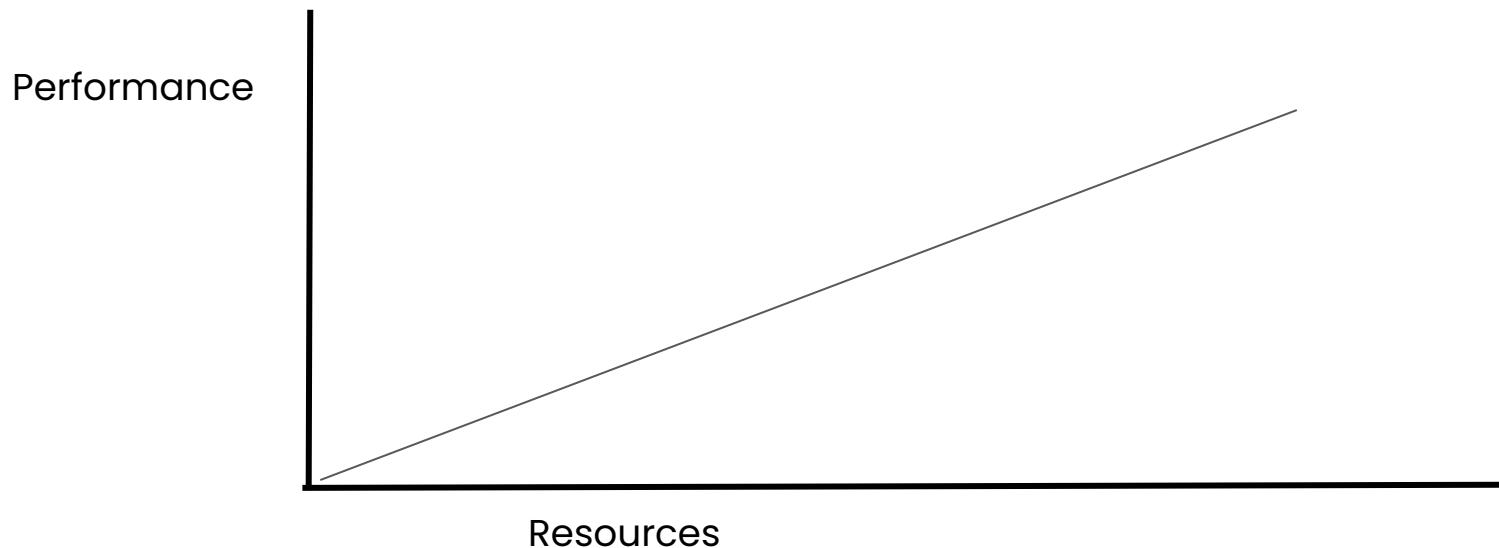
---



# Scalability

=

**Removing non-linear factors**



“We looked at a half-dozen distributed computing projects, and Ray was by far the winner. ... We are **using Ray to train our largest models**. It’s been very helpful to us to be able to scale up to unprecedented scale. ... **Ray owns a whole layer, and not in an opaque way.**”

– *Greg Brockman, President, OpenAI*

# **Removing Non-linear Factors**

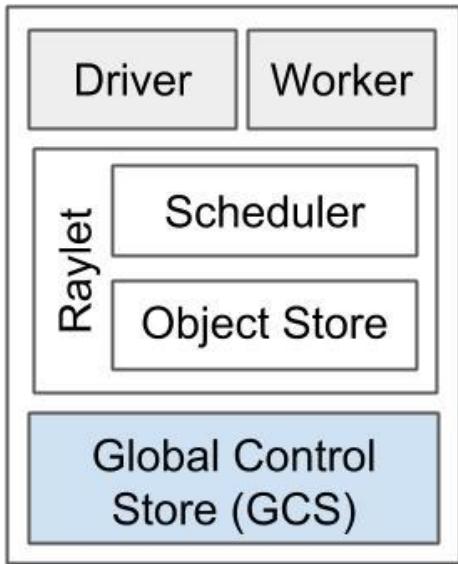
Flexible Fault Tolerance

Global Distributed Scheduling

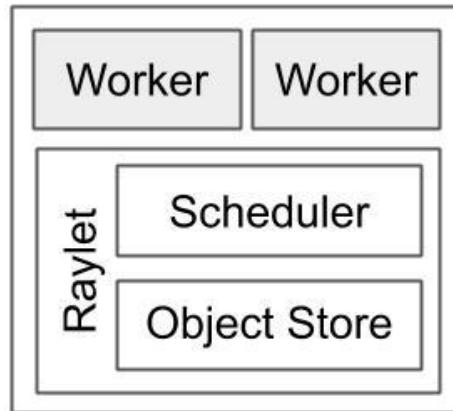


# Scalability – A Ray Cluster Looks Like...

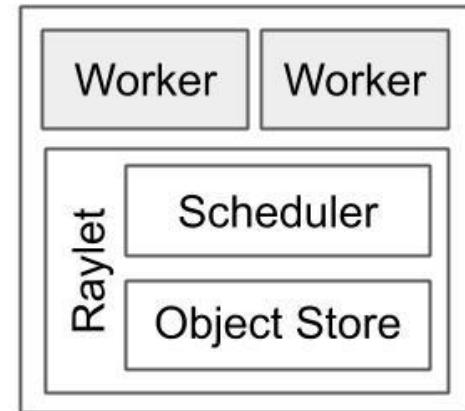
Head node



Worker node

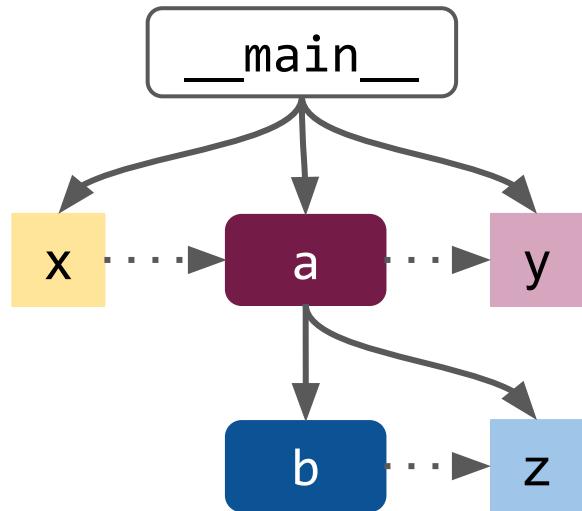


Worker node



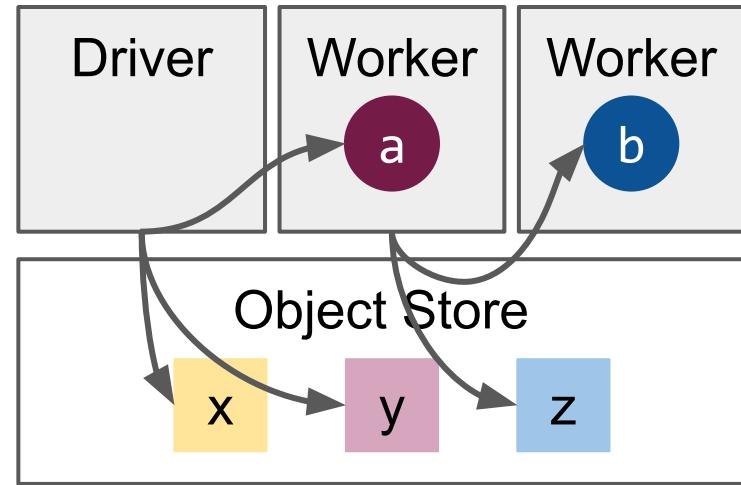
```
@ray.remote  
def b():  
    return  
  
@ray.remote  
def a(dep):  
    z = b.remote()  
  
x = ray.put(...)  
y = a.remote(x)
```

Program



Task graph

→ Ownership  
→ Dependency



Physical execution

# Ray's Fault Tolerance Model

**Tasks**: by default retry 3 times

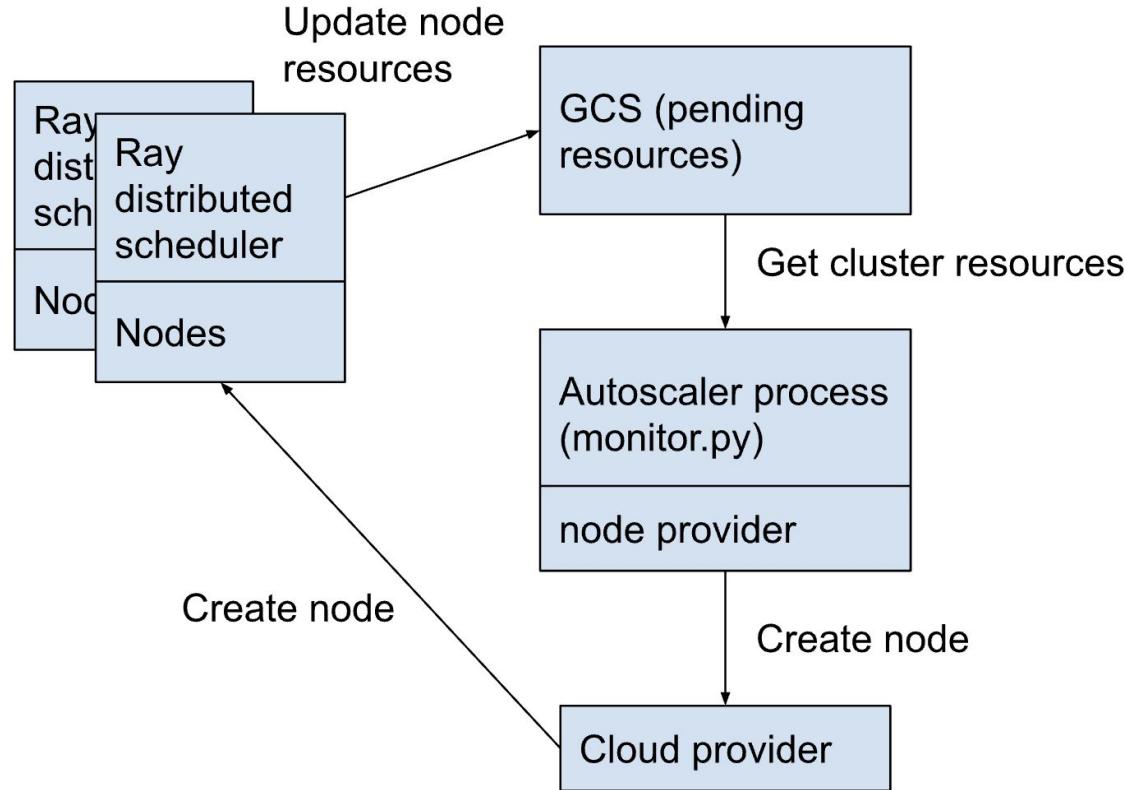
```
@ray.remote (max_retries=1)  
def read_array(file):
```

**Actors**: by default not restarted. But behavior can be customized.

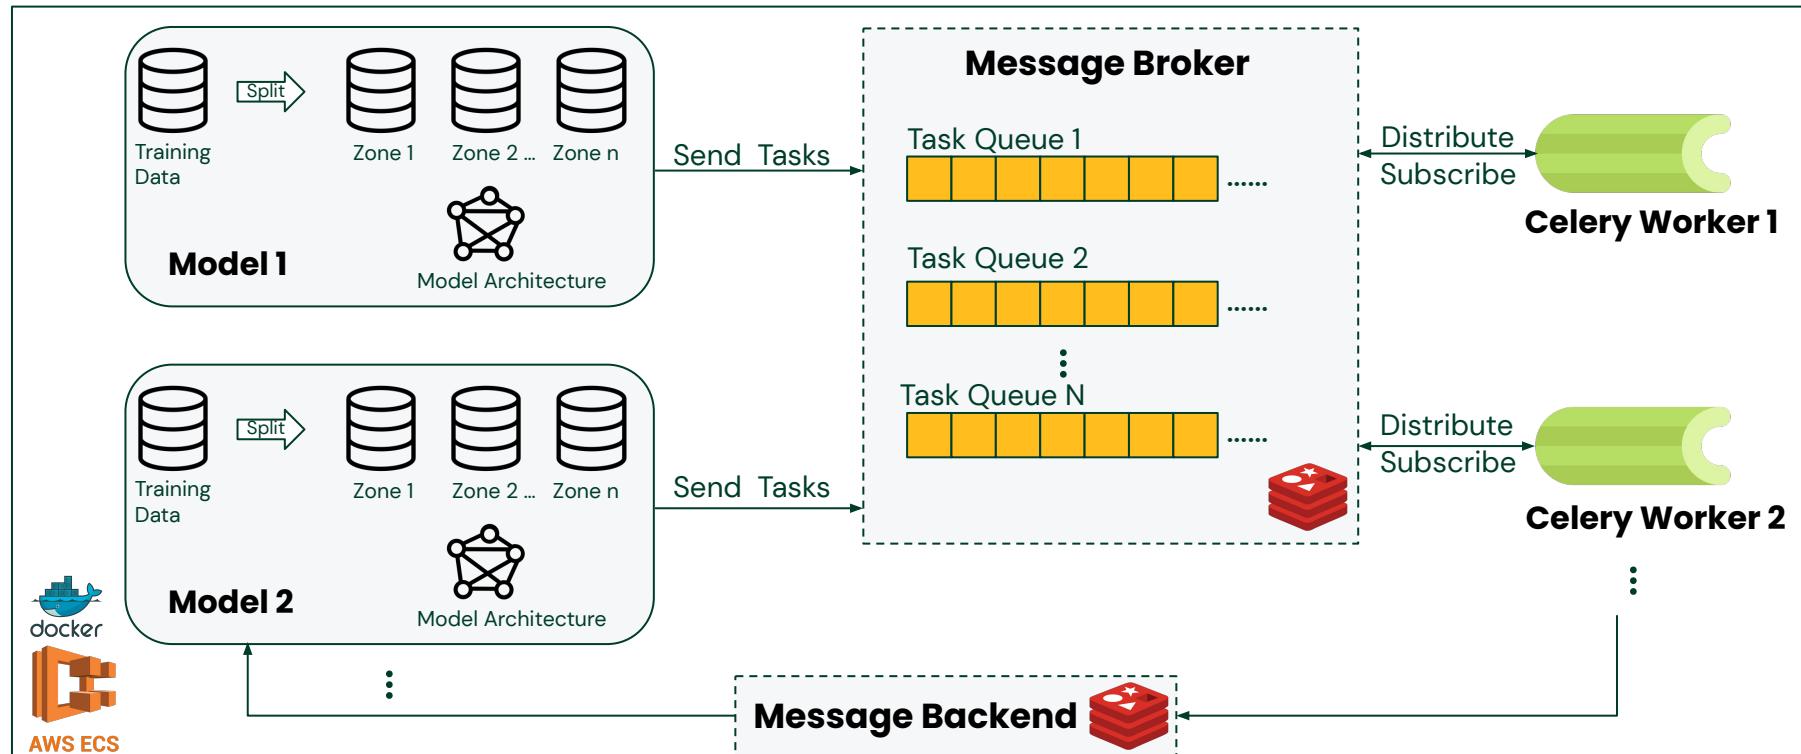
```
@ray.remote (max_restarts=5)  
class Counter(object):
```

**Objects**: lineage-based reconstruction

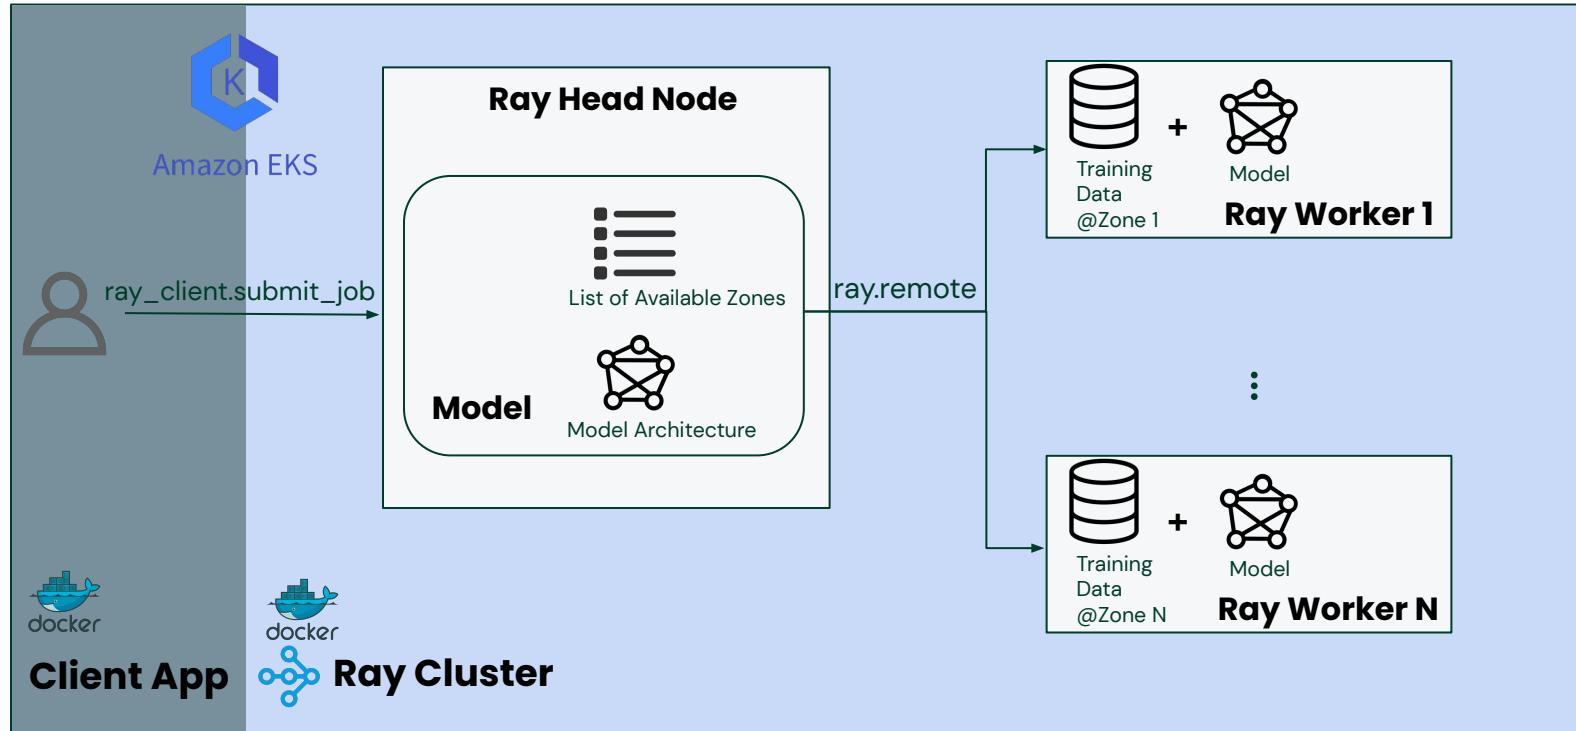
# Ray's Autoscaler



# Existing Solution: Celery



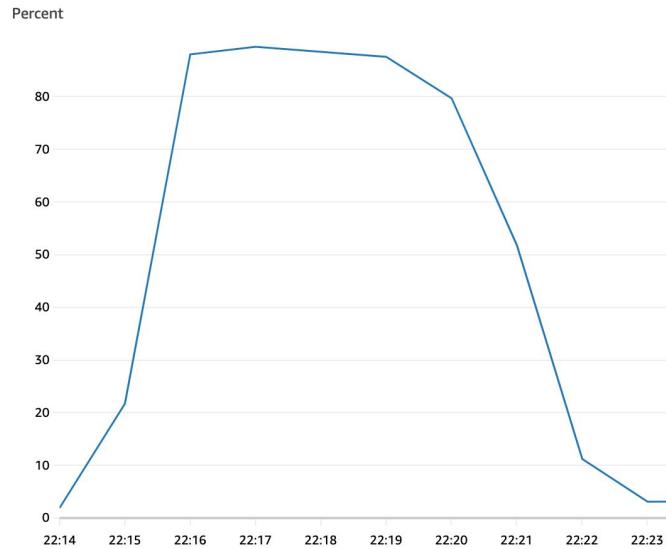
# Proposed Solution: Ray



# Results: Better Utilization



Before

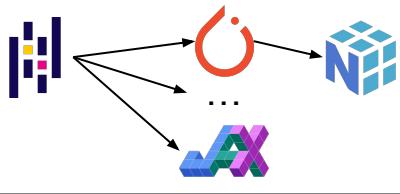


After

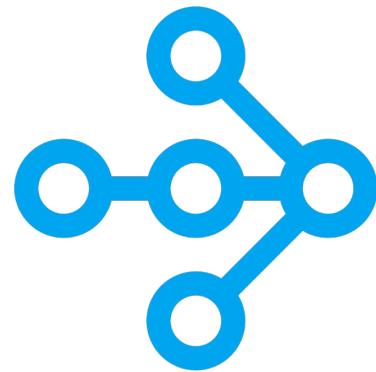
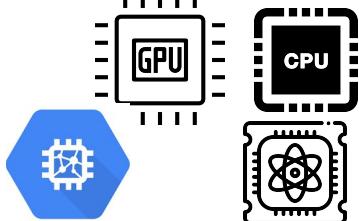
## Scalability



## Unification



## Heterogeneous Hardware



**Simple &  
Flexible API**



# RAY

# Unification – One Substrate for ML Compute

---



# Ray AI Runtime

Ray AI Runtime

Datasets

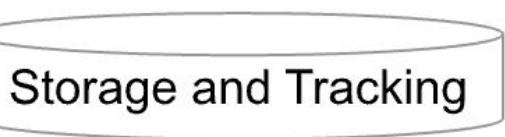
Training

Tuning

Scoring

Serving

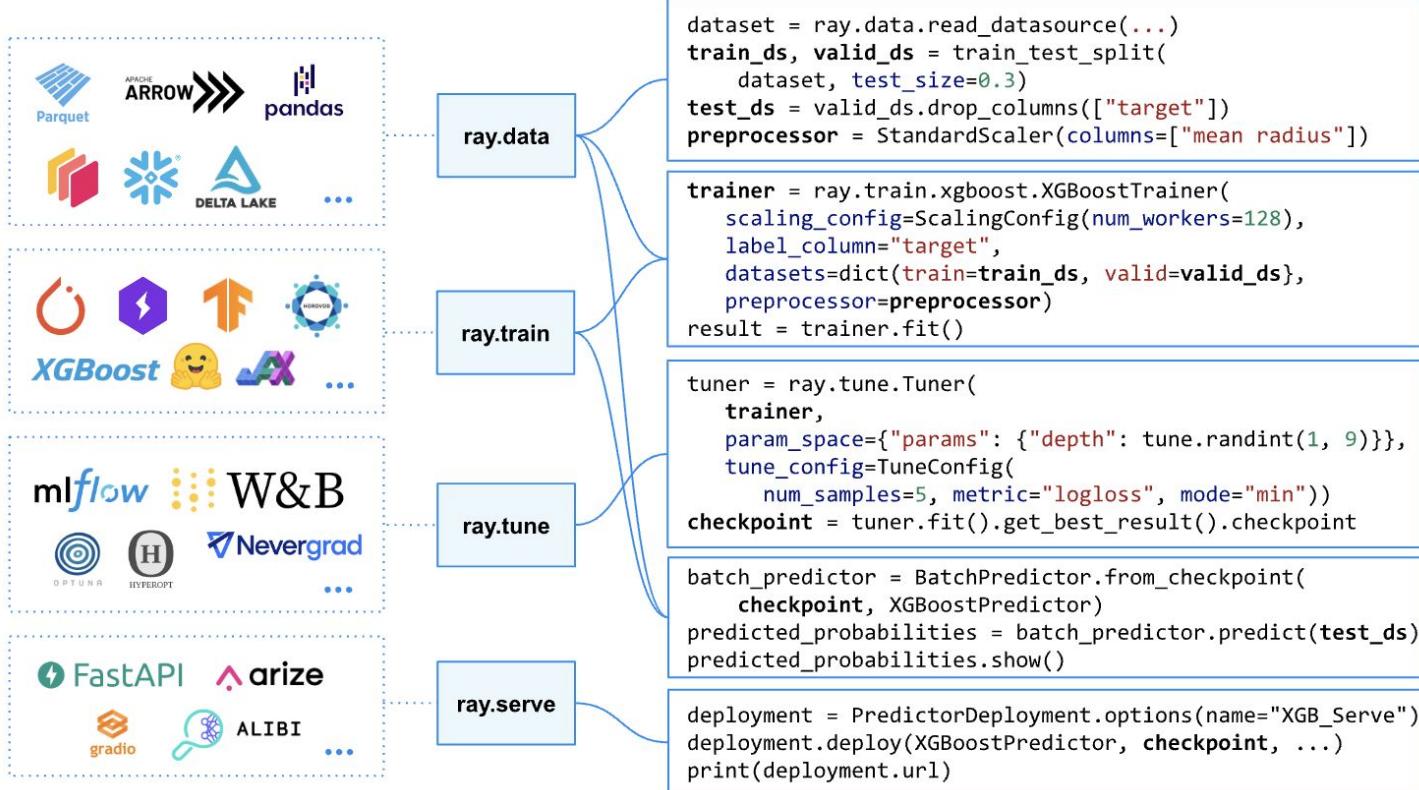
RL



Ray Core

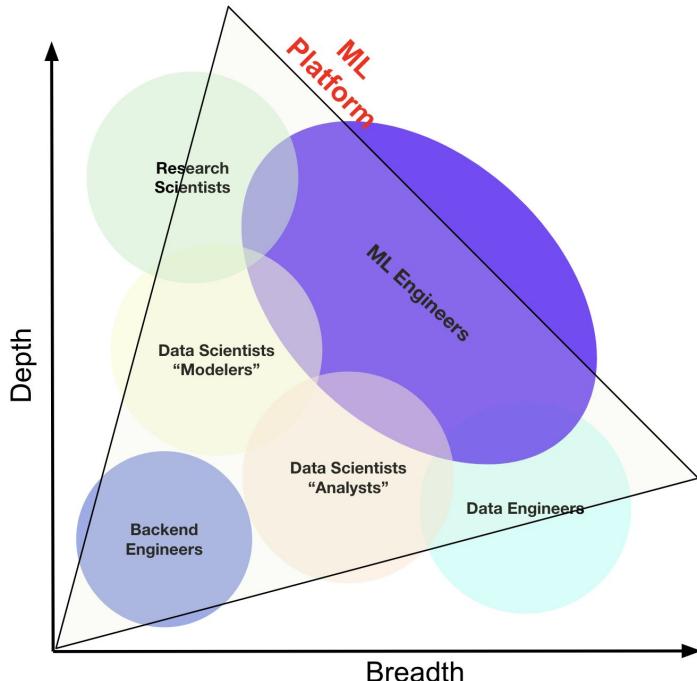


# Ray AI Runtime

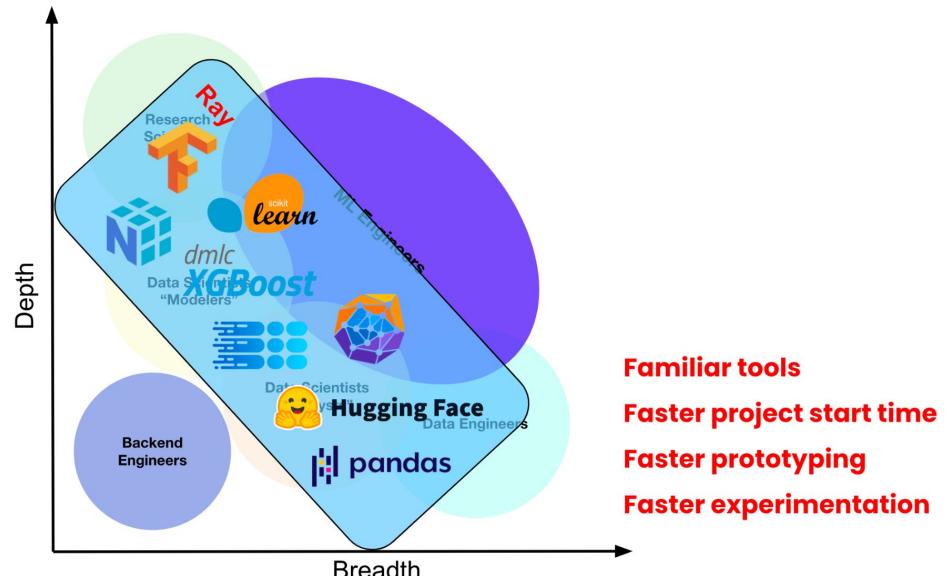


# Using Ray at Spotify

## Where we want to be



## What Ray's Ecosystem can help



## Sample of Companies Who Use Ray in their Machine Learning Platform

RIDECELL



Uber

cruise



Meta

Google

OpenAI

Microsoft



shopify

amazon

STITCH FIX

RIOT GAMES

NETFLIX

Spotify®

ByteDance

DOW®

intel®

IBM



Hugging Face

co:here

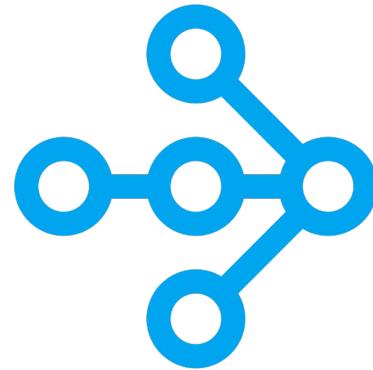
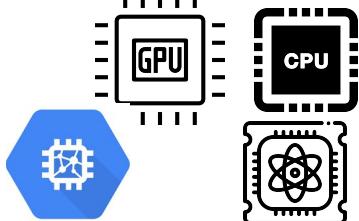
## Scalability



## Unification



## Heterogeneous Hardware



**Simple &  
Flexible API**



# Heterogeneous Hardware

---



# Function → Task

```
@ray.remote  
def read_array(file):  
    # read ndarray "a"  
    # from "file"  
    return a
```

```
@ray.remote  
def add(a, b):  
    return np.add(a, b)
```

```
id1 = read_array.remote(file1)  
id2 = read_array.remote(file2)  
id = add.remote(id1, id2)  
sum = ray.get(id)
```

# Class → Actor

```
@ray.remote(num_cpus=2, num_gpus=1)  
class Counter(Actor):
```

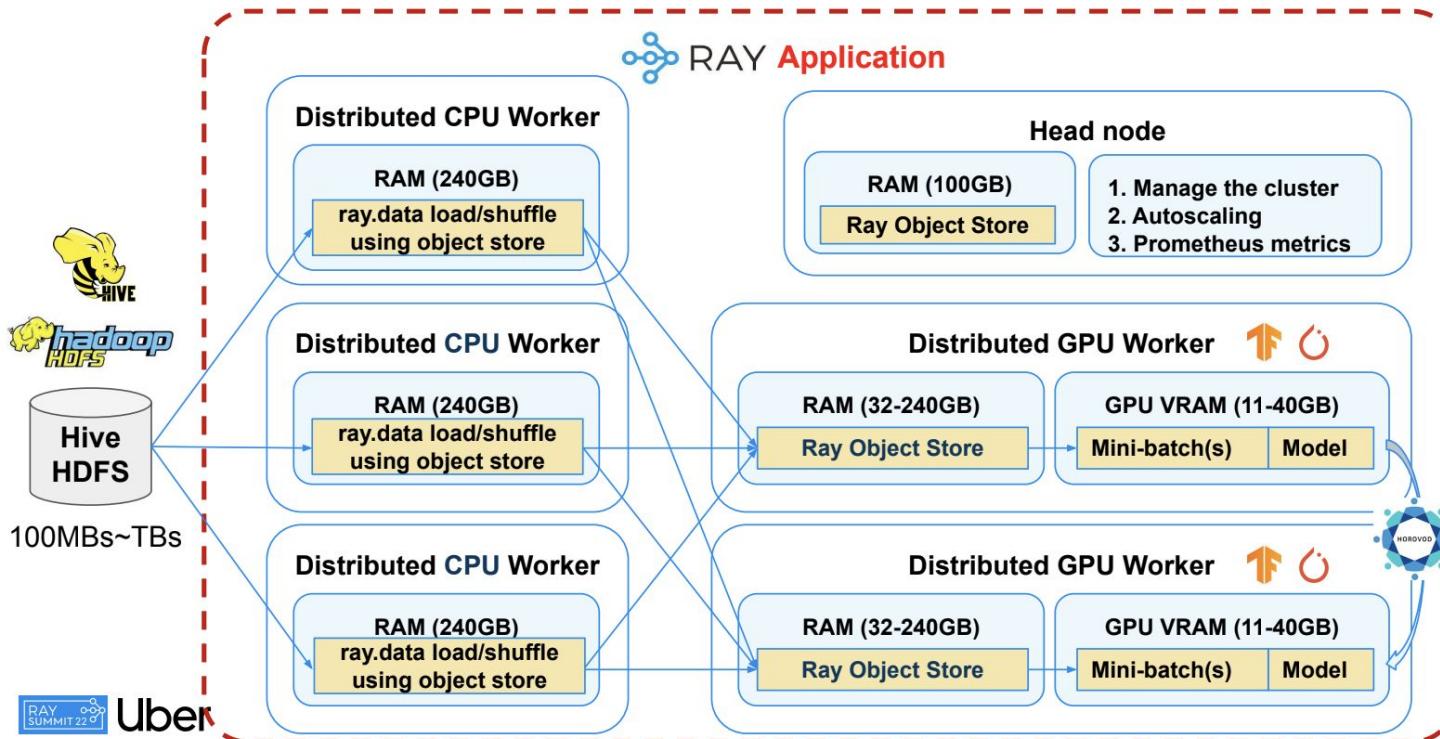
can specify  
resource demands;  
support heterogeneous hardware

```
self.value += 1  
return self.value
```

```
c = Counter.remote()  
id1 = c.inc.remote()  
id2 = c.inc.remote()  
val = ray.get(id2)
```

# Deep Learning Cluster at Uber

## Heterogeneous Ray Cluster (v2): Architecture



# IBM: Quantum + Classical Computing

Send program to execute

# Summary

Ray's **simple and flexible APIs** provides **easy access** to:

**Scalable** and **unified** distributed computing on **heterogeneous** hardware.

## Call to action

- Get started at [docs.ray.io](https://docs.ray.io)
- Dive deeper on <https://github.com/ray-project/ray/>
- See the Ray Summit Program
- Stay tuned on our meetup series