


Using AWS Deep Learning Containers on Amazon Elastic Kubernetes Service



Naina Prasad & Arjuna Keshavan
AWS AI/ML

27th September 2021

After watching this webinar, you will be able to...

- Use AWS Deep Learning Containers (DLCs) to manage your machine learning (ML) workflows on Amazon Elastic Kubernetes Service (EKS)
- Use training DLCs to run training workflows on EKS
- Use inference DLCs to host models on EKS

Agenda

- Introduce Machine Learning on Kubernetes
- Dive into AWS Deep Learning Containers (DLCs)
 - Train a model on EKS with training DLCs
 - Host a model on EKS with inference DLCs
- Conclusion and Additional resources

Be familiar with...

- Docker containers
- Kubernetes (K8s)
- Training and Inference ML workflows
- ML frameworks (TensorFlow, PyTorch, MXNet)

Different ways to run ML workflows on AWS

- Use a managed ML service (Amazon SageMaker)
- Do it yourself (i.e. Amazon EKS, ECS, EC2)


Advantages to Machine Learning via K8s



Customizable



Portable

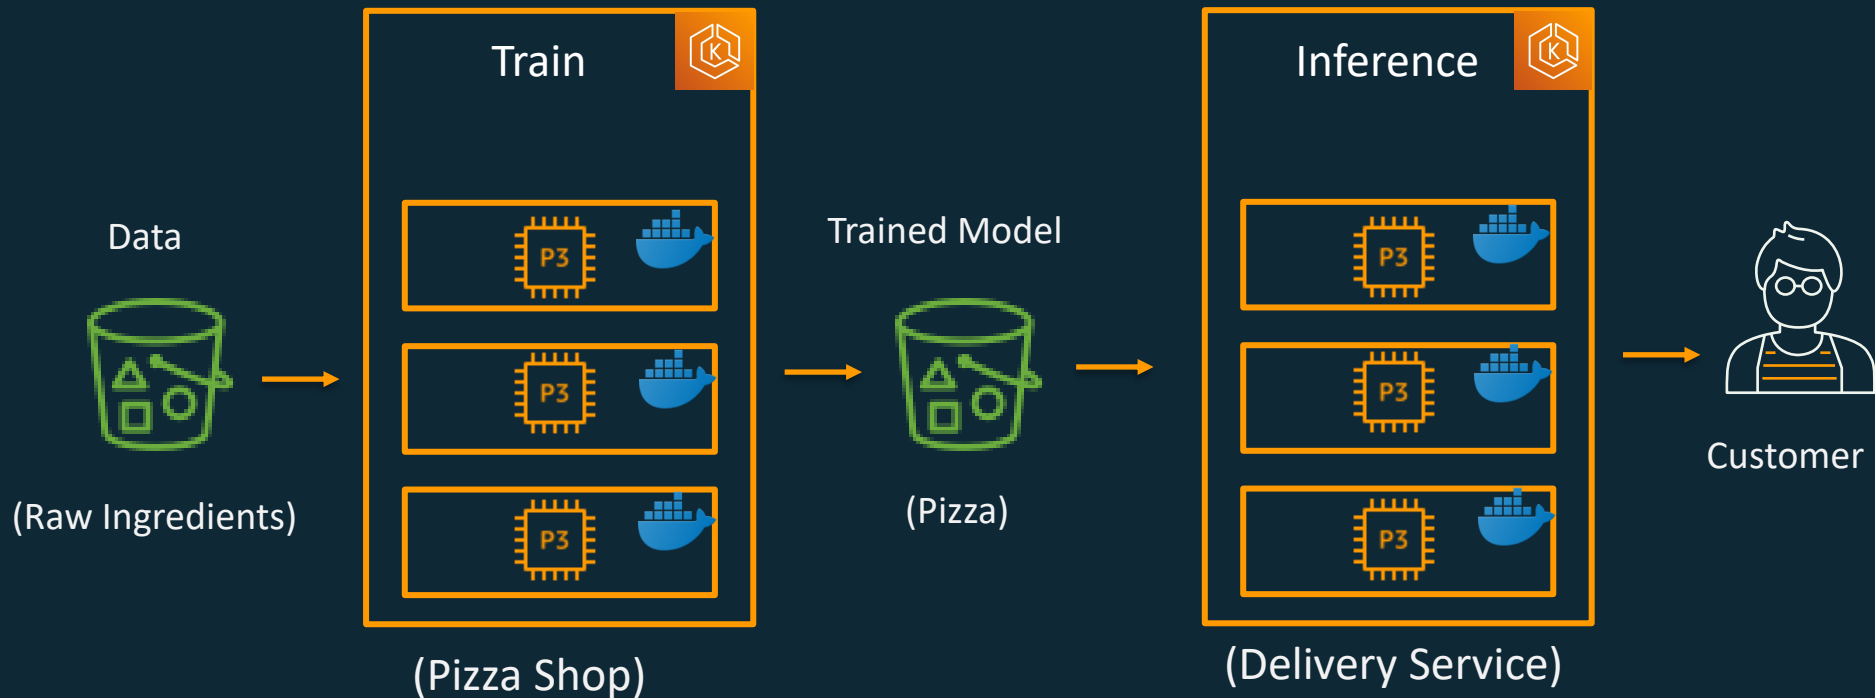


Scalable

The Pizza Delivery Workflow



ML on K8s Workflow



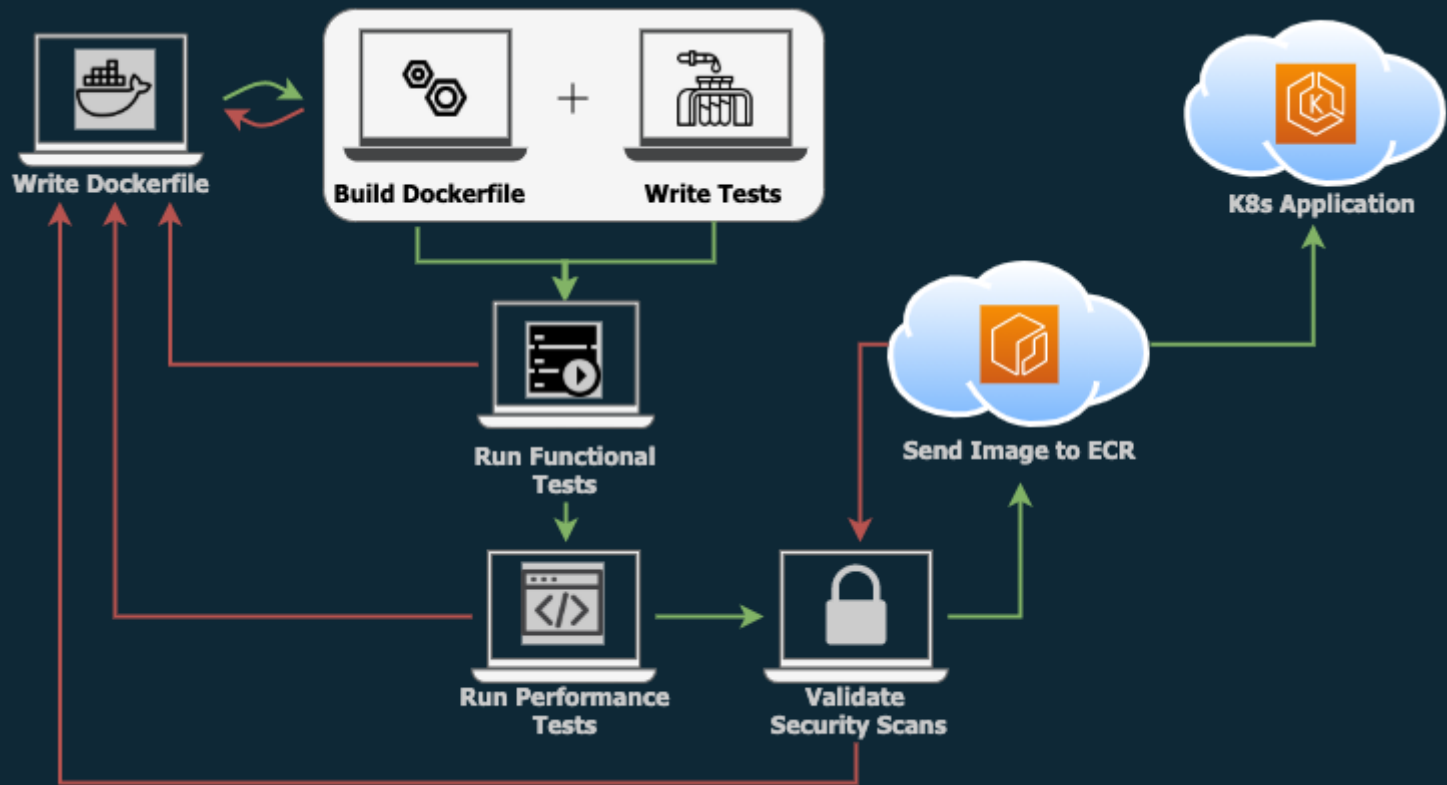
Pizza Kitchen

- Needs to be clean
- Appliances need to perform well
- Appliances need to be updated regularly

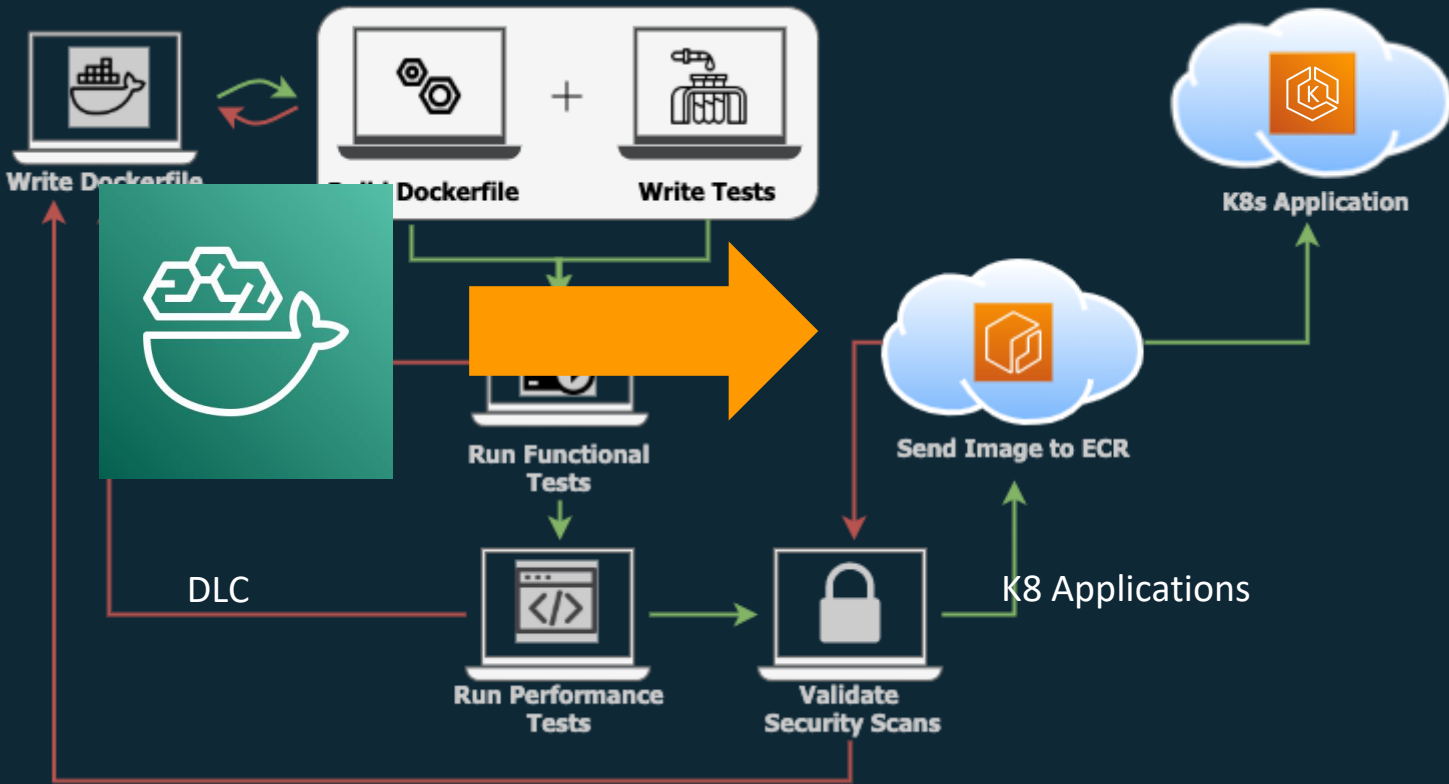
ML Container

- Needs to be secure
- Must optimize for performance and scale
- Rebuild and re-optimize for new framework versions

Managing your own container



AWS Deep Learning Containers (DLCs)



AWS Deep Learning Containers (DLCs)

Performance-tested
framework binaries

Open source!



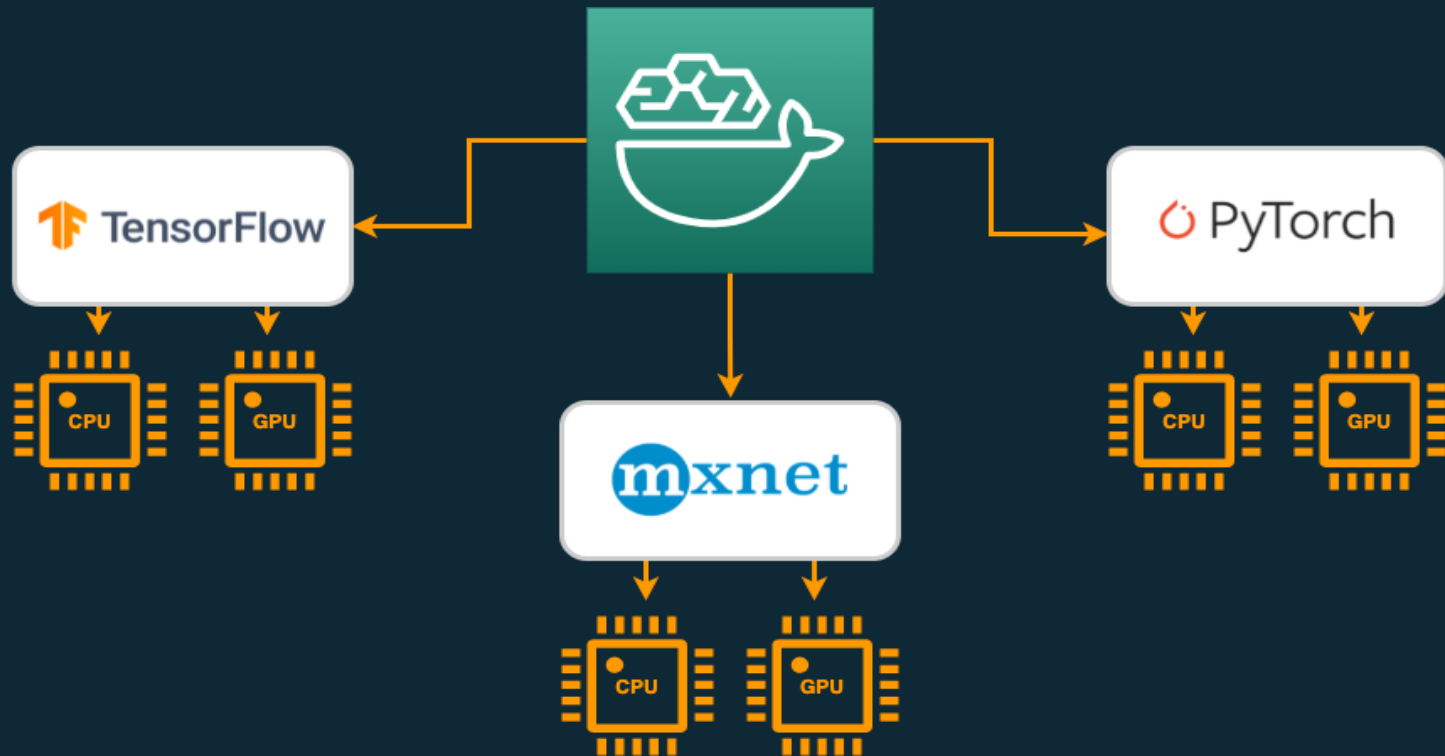
DLC

Useful ML utility packages

Regular security patches

Recent Framework Versions

DLC Training Images



TensorFlow GPU training images

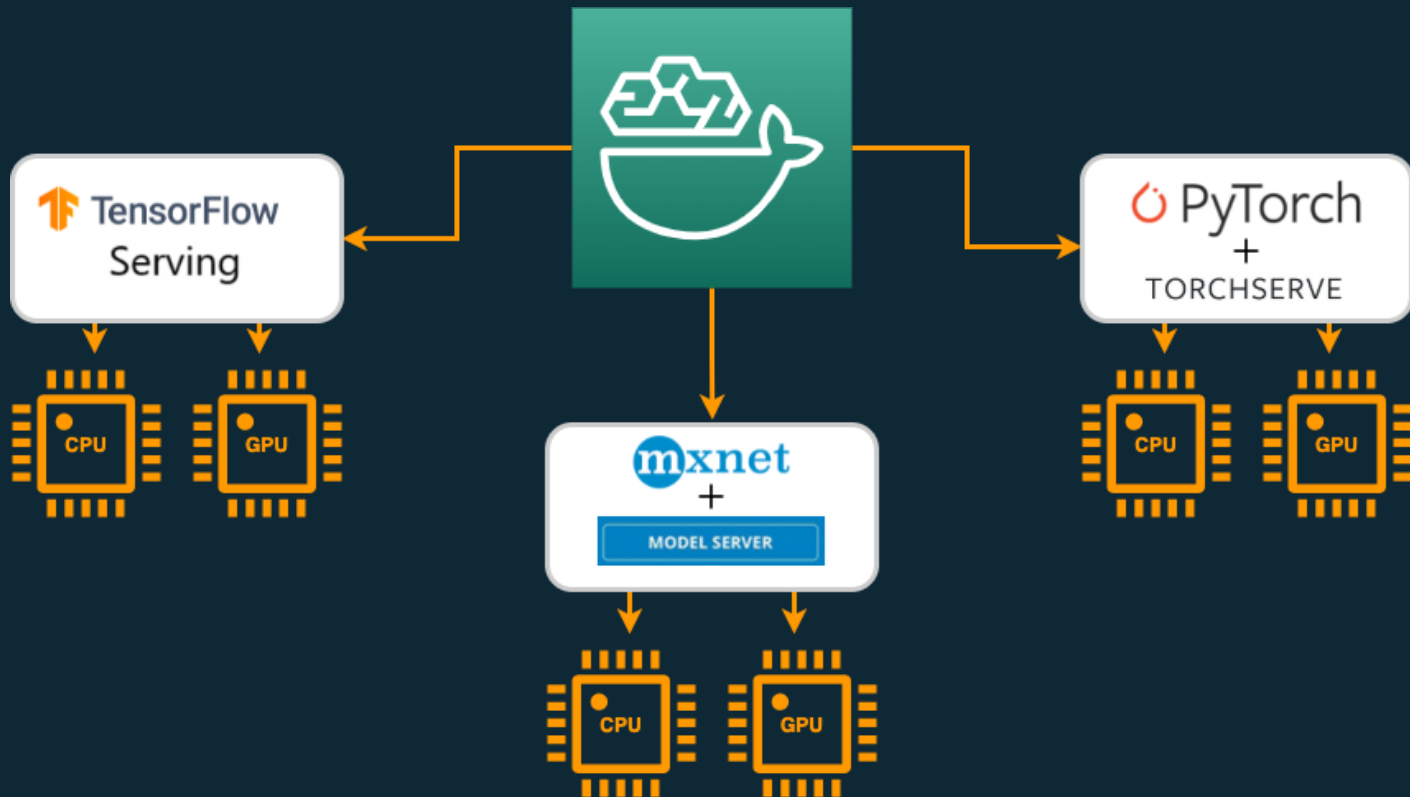
Python Packages

- tensorflow-gpu
- awscli
- horovod
- h5py
- numpy
- opencv-python
- pandas
- scikit-learn
- scipy

Apt | Deb Packages

- cuda dependencies and packages
- low level helper libraries

DLC Inference Images



TensorFlow GPU inference images

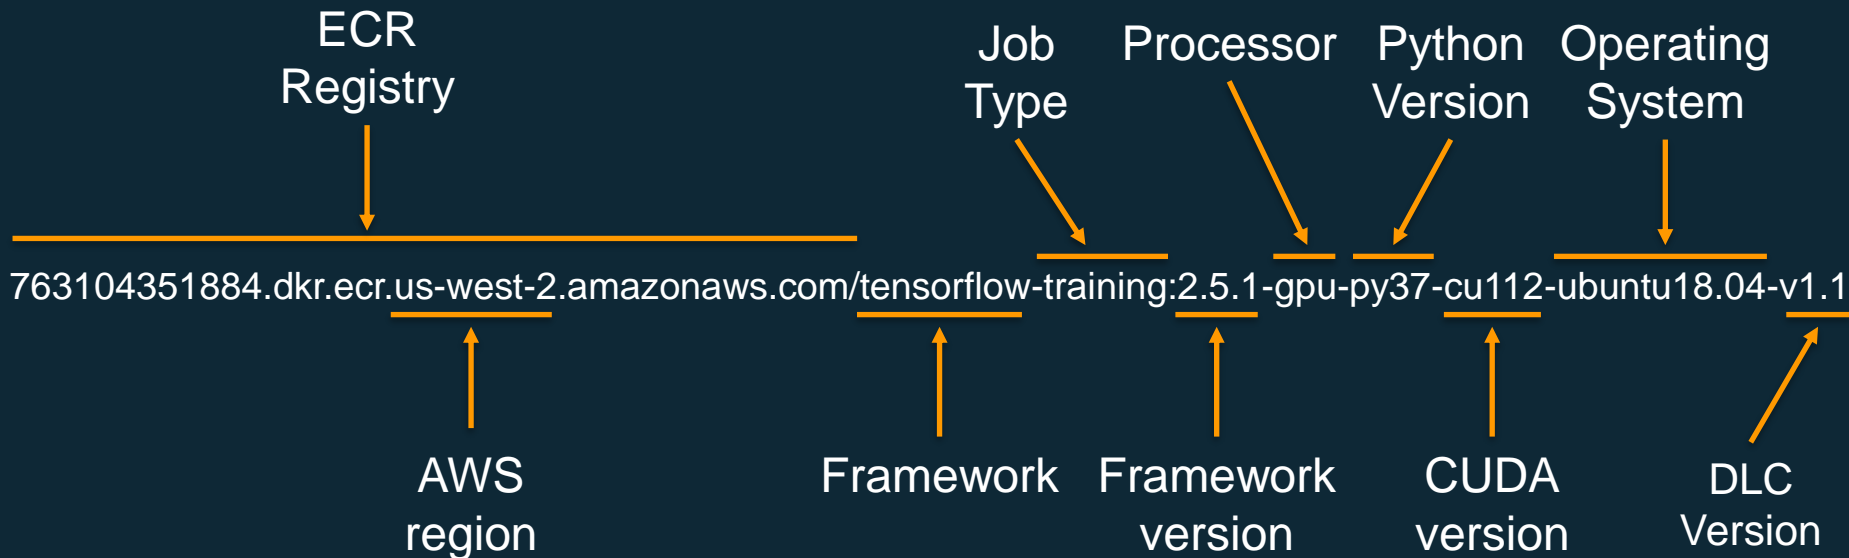
Python Packages

- awscli
- tensorflow-serving-api-gpu
- requests

Apt | Deb Packages

- cuda dependencies and packages
- low level helper libraries

DLC Tags

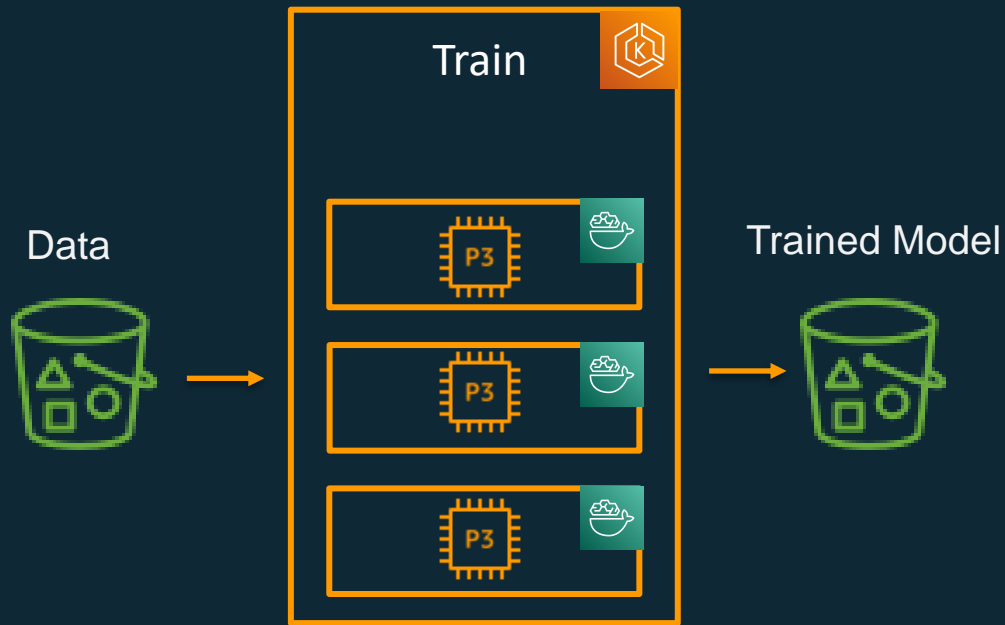


Building on top of DLCs

```
1 FROM 763104351884.dkr.ecr.us-west-2.amazonaws.com/tensorflow-training:2.5-gpu-py37-cu112-ubuntu18.04-v1
2
3 RUN git clone https://github.com/aws/deep-learning-containers.git
4
5 ...
6 █
```

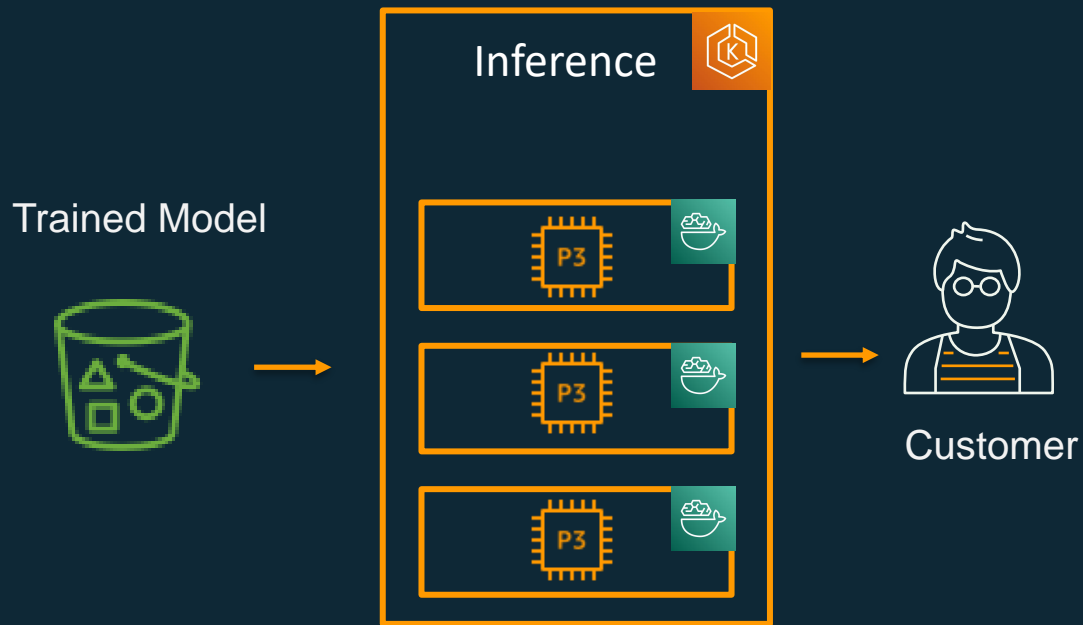
Training on EKS with DLC

Let's see it in action!



Inference on EKS with DLC

Let's see it in action!



Conclusion

- Simplify and structure ML workflows using DLCs with EKS
- Easier to customize, port, and scale your ML workflows
- DLC are pre-tested with EKS so you can get right to meaningful work
- Open source & free to use

Additional Resources

- [AWS Documentation Pages](#)
 - [EKS Features](#)
 - [Deep Learning Containers Developer Guide](#)
 - [Deep Learning Containers Release Notes](#)
- [DLC GitHub Repository](#)
- [ML on EKS Tech Talk](#)

Thank you!