



KubeCon



CloudNativeCon

Europe 2018

Scaling Distributed Deep Learning with Service Discovery

Yong Tang

Principal Software Engineer, Infoblox Inc.

GitHub: [yongtang](#)



Agenda



KubeCon



CloudNativeCon

Europe 2018

- Myself:
 - GitHub: github.com/yongtang
 - Committer: CoreDNS, Docker (Moby), TensorFlow
 - Working on machine learning projects in Infoblox
- My talk
 - DNS security with machine learning in Infoblox
 - Helping distributed deep learning with service discovery (CoreDNS)

Machine Learning @ Infoblox

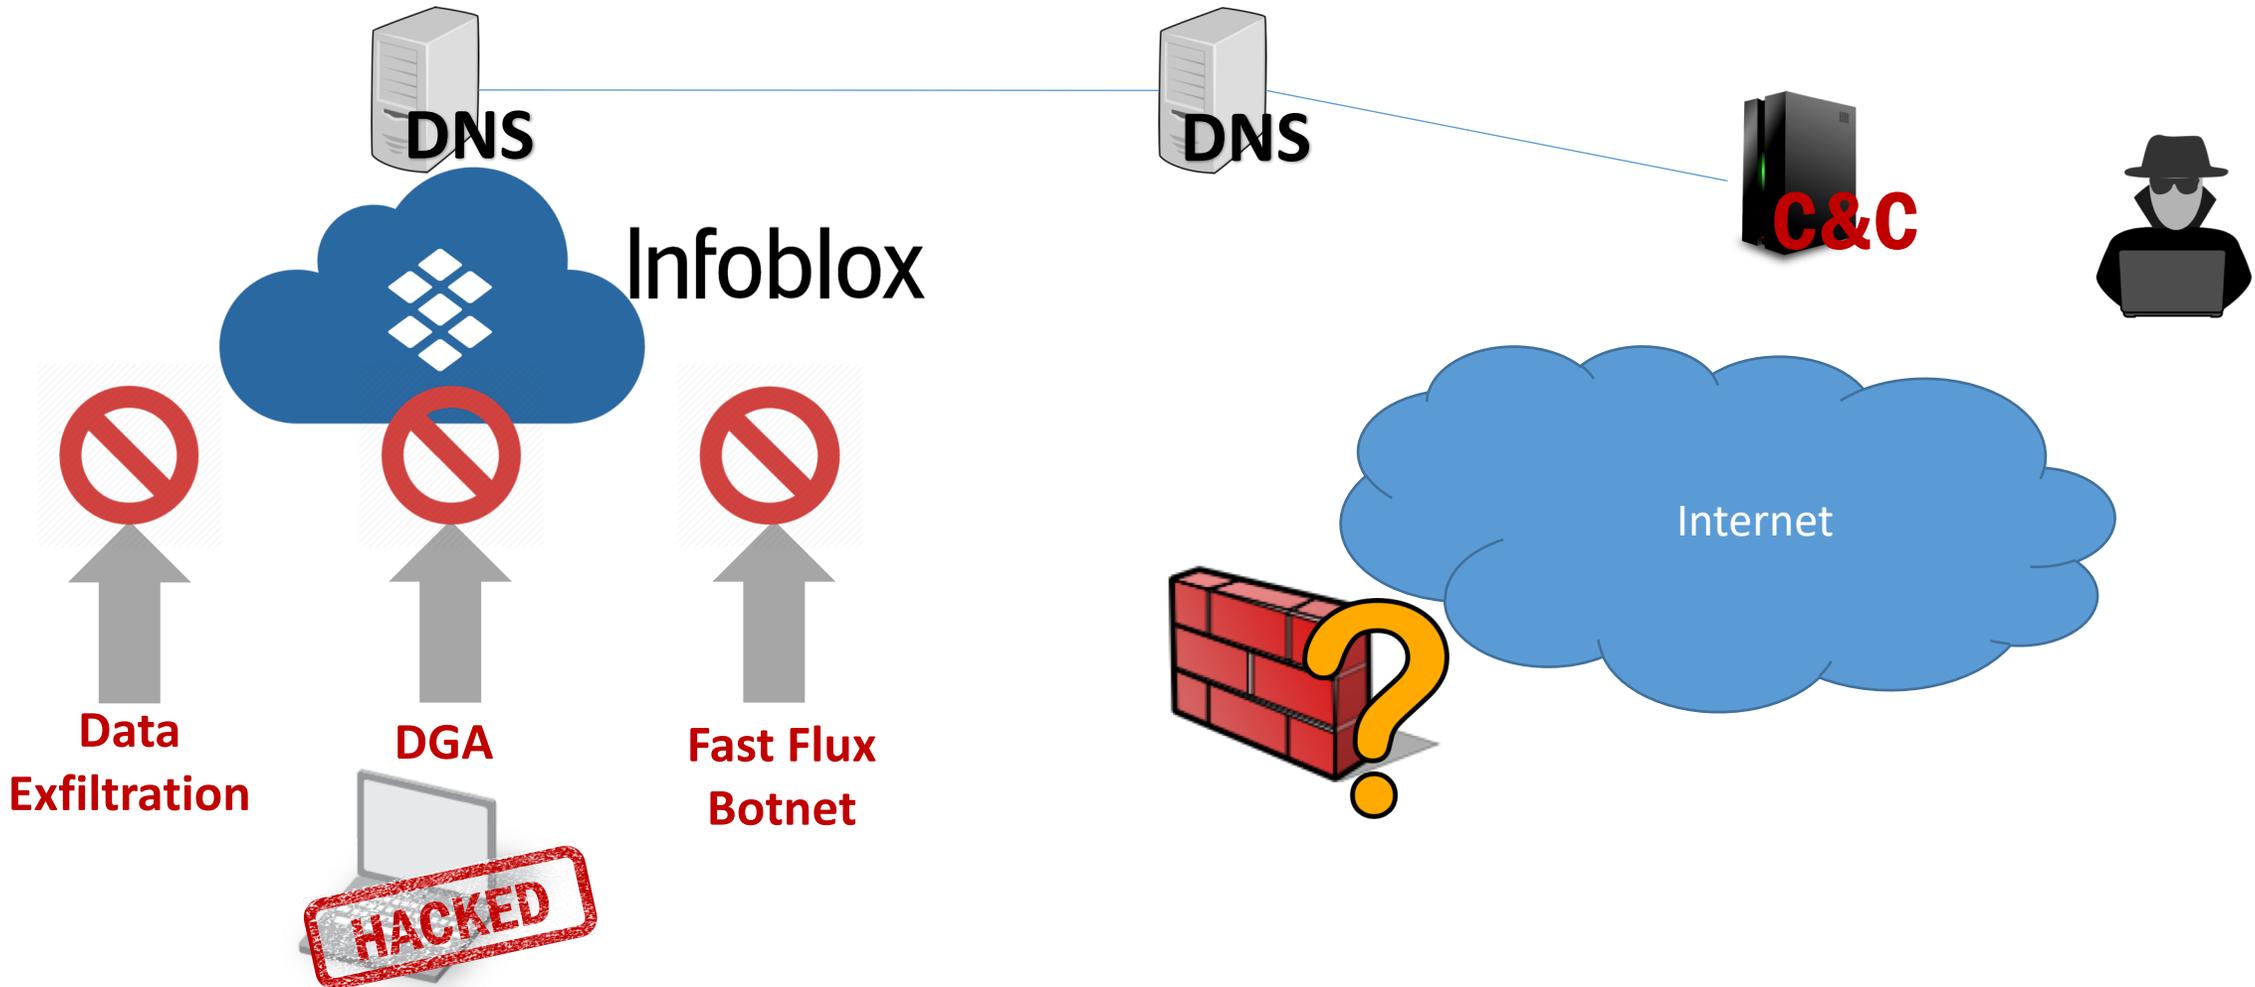


KubeCon



CloudNativeCon

Europe 2018



Machine Learning Infrastructure



KubeCon



CloudNativeCon

Europe 2018

- Small ML team, no dedicated ML DevOps/Infra team
 - Services provided by cloud vendors (SageMaker, EMR)
- Deploy our own TensorFlow clusters
 - Latest version of  TensorFlow
 - User defined ops (e.g., KafkaReader, PR [#14098](#))
- Scope
 - Bring up/down a TensorFlow cluster with one or two commands
 - Use TensorFlow once cluster is up and running
 - No dynamic scaling up/down (unnecessary complexity)

CoreDNS and Plugins



KubeCon



CloudNativeCon

Europe 2018

- DNS server with a focus on service discovery
- A CNCF incubating project, Infoblox is a sponsor
- Integration with k8s, default DNS server for k8s (expected)
- **Plugin-based architecture, easily extended (customized)**



CoreDNS



**CLOUD NATIVE
COMPUTING FOUNDATION**



CoreDNS

Infoblox 

Service Discovery with CoreDNS



KubeCon



CloudNativeCon

Europe 2018

- Worker: GPU, Parameter Server: CPU
- Whole cluster information for every nodes specified in advance
- Instance Metadata (ami-launch-index) -> DNS record

Parameter Server
ami-launch-index: 0



ps0.tensorflow.local

Parameter Server
ami-launch-index: 1



ps1.tensorflow.local

Worker

ami-launch-index: 0



worker0.tensorflow.local

Worker

ami-launch-index: 1



worker1.tensorflow.local

Worker

ami-launch-index: 2



worker2.tensorflow.local



CoreDNS

Infoblox

Service Discovery with CoreDNS



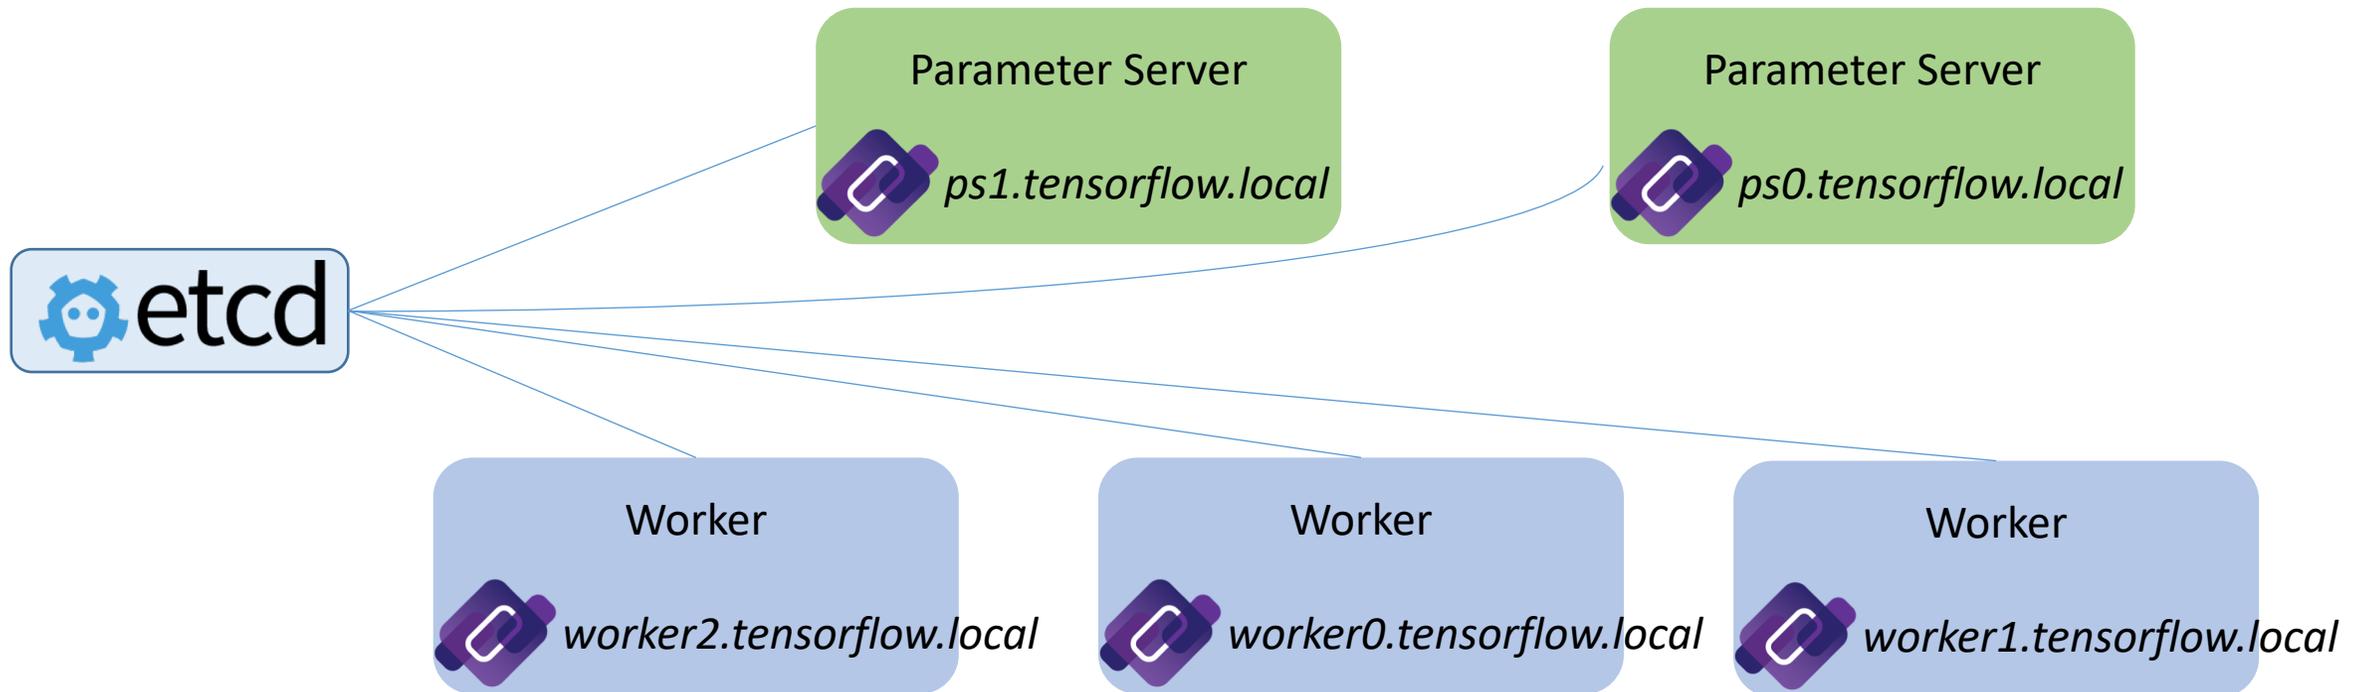
KubeCon



CloudNativeCon

Europe 2018

- Distributed K-V store (e.g., etcd/zookeeper)
- Exposed through CoreDNS



Service Discovery with CoreDNS



KubeCon



CloudNativeCon

Europe 2018

- Pack TensorFlow and CoreDNS into Cloud-Init script
- TensorFlow cluster could be deployed with one command
- 80% of a distributed system at a fraction of the cost



CoreDNS

Service Discovery with CoreDNS



KubeCon



CloudNativeCon

Europe 2018

THANK YOU