



Kubeflow End-to-End: GitHub Issue Summarization

Google Cloud

Amy Unruh
Michelle Casbon
Kubecon North America
Seattle
December 11, 2018

TAs


Dan Sanche



Sarah Maddox




Gonzalo Gasca Meza



Steve Greenberg










Kubeflow is a curated set of compatible tools and artifacts that lays a foundation for running production ML apps

Enables consistency across deployments by providing Kubernetes object templates that bring together disparate components



Contributors

- Kubeflow is open
 - Open community
 - Open design
 - Open source
 - Open to ideas
- Get involved
 - github.com/kubeflow
 - kubeflow.slack.com 
 - @kubeflow 
 - kubeflow-discuss@googlegroups.com
 - Community call Tuesdays alternating 8:30am and 5:30pm Pacific
 - **Kubeflow Contributor Summit**
 - Q1 2019



<https://github.com/kubeflow/kubeflow>

Agenda g.co/code1labs/kubecon18

1

Set up the environment

2

Create a Kubeflow cluster

3

Run a pipeline from the Kubeflow Pipelines dashboard

4

Run a pipeline from a Jupyter notebook

5

Clean up

Agenda g.co/code1labs/kubecon18

Zones: us-west1-b, us-east1-c

1

Set up the environment

2

Create a Kubeflow cluster

3

Run a pipeline from the Kubeflow Pipelines dashboard

4

Run a pipeline from a Jupyter notebook

5

Clean up

PROBLEM

Moving from local to production

GitHub IS

Kubeflow

GCP



SOLUTION

Portability

Package infrastructure
components together





PROBLEM

Complexity

GitHub IS


Kubeflow

GCP

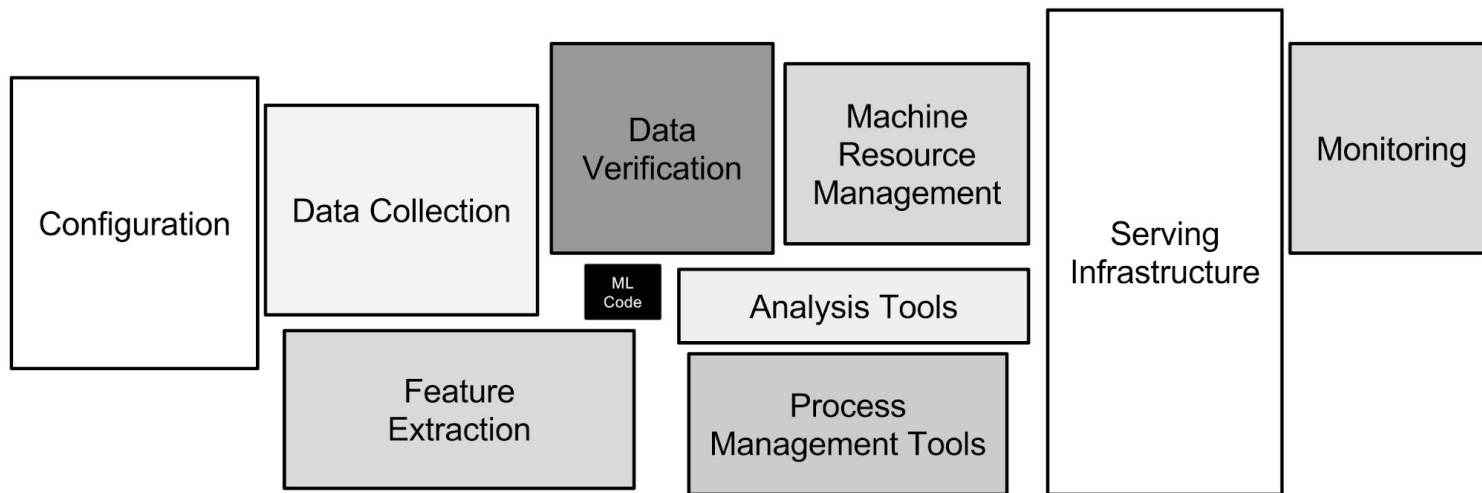


Perception

ML
Code



Reality



Data

Data Ingestion

Data Exploration

Data
Transformation

Data Validation

Data Analysis

Training Data
Segmentation**Featurization**Feature
Extraction**Training**

Model Building

Model Validation

Model Versioning

Model Auditing

Distributed
TrainingContinuous
Training**Application**Serving
Infrastructure

Business Logic

UI

Load Balancing

Platform

Configuration

Process
ManagementResource
Management

Monitoring

Logging

Continuous
DeliveryAuthentication/
Authorization

PROBLEM

Complexity



SOLUTION

Composability

Logical groupings

Reusable components

GitHub IS

Kubeflow

GCP

PROBLEM

Maintainability

- Error resolution, recovery, & prevention
- Speed of iteration
- Versioning

SOLUTION

Composability

Shorten the development
lifecycle

Automation

GitHub IS

Kubeflow

GCP

PROBLEM

Capacity Planning

- Usage patterns
- Demand spikes
- Efficient resource usage

SOLUTION

Scalability

Kubernetes

Autoprovisioning

GitHub IS

Kubeflow

GCP



Make it easy for everyone to develop,
deploy, and manage portable,
scalable ML everywhere

Kubeflow

Portability

Entire stack

Scalability

Native to k8s

Reduce variability
between services
& environments

Composability


Single, unified tool
for common
processes

Full product lifecycle

Support specialized hardware, like GPUs & TPUs

Reduce costs

Improve model
performance



Kubeflow

Who

Data scientists

ML researchers

Software engineers

Product managers

What

Portable ML products on k8s

v0.3.4 release

Why

Because building a platform is too big of a problem to tackle alone

<https://github.com/kubeflow/kubeflow>

Kubeflow

Kubernetes-native platform for ML

Run wherever k8s runs

Use k8s to manage ML tasks

CRDs for distributed training

Adopt k8s patterns

Microservices

Manage infra declaratively

Package infrastructure components together

Ksonnet

Move between local -> dev -> test -> prod -> onprem



Support multiple ML frameworks

TensorFlow

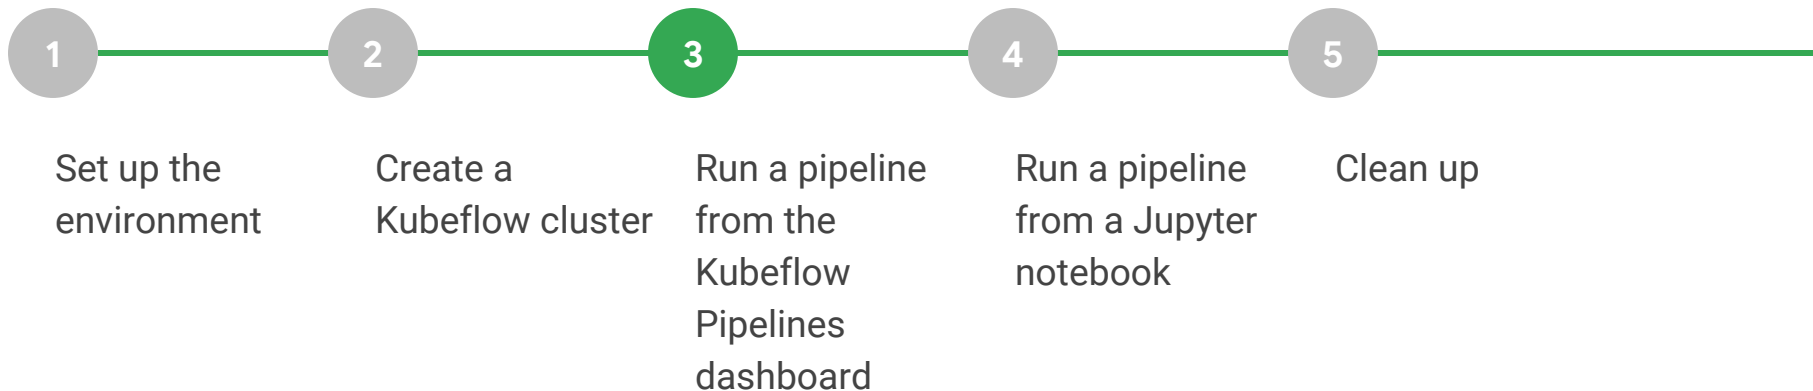
Pytorch

Scikit

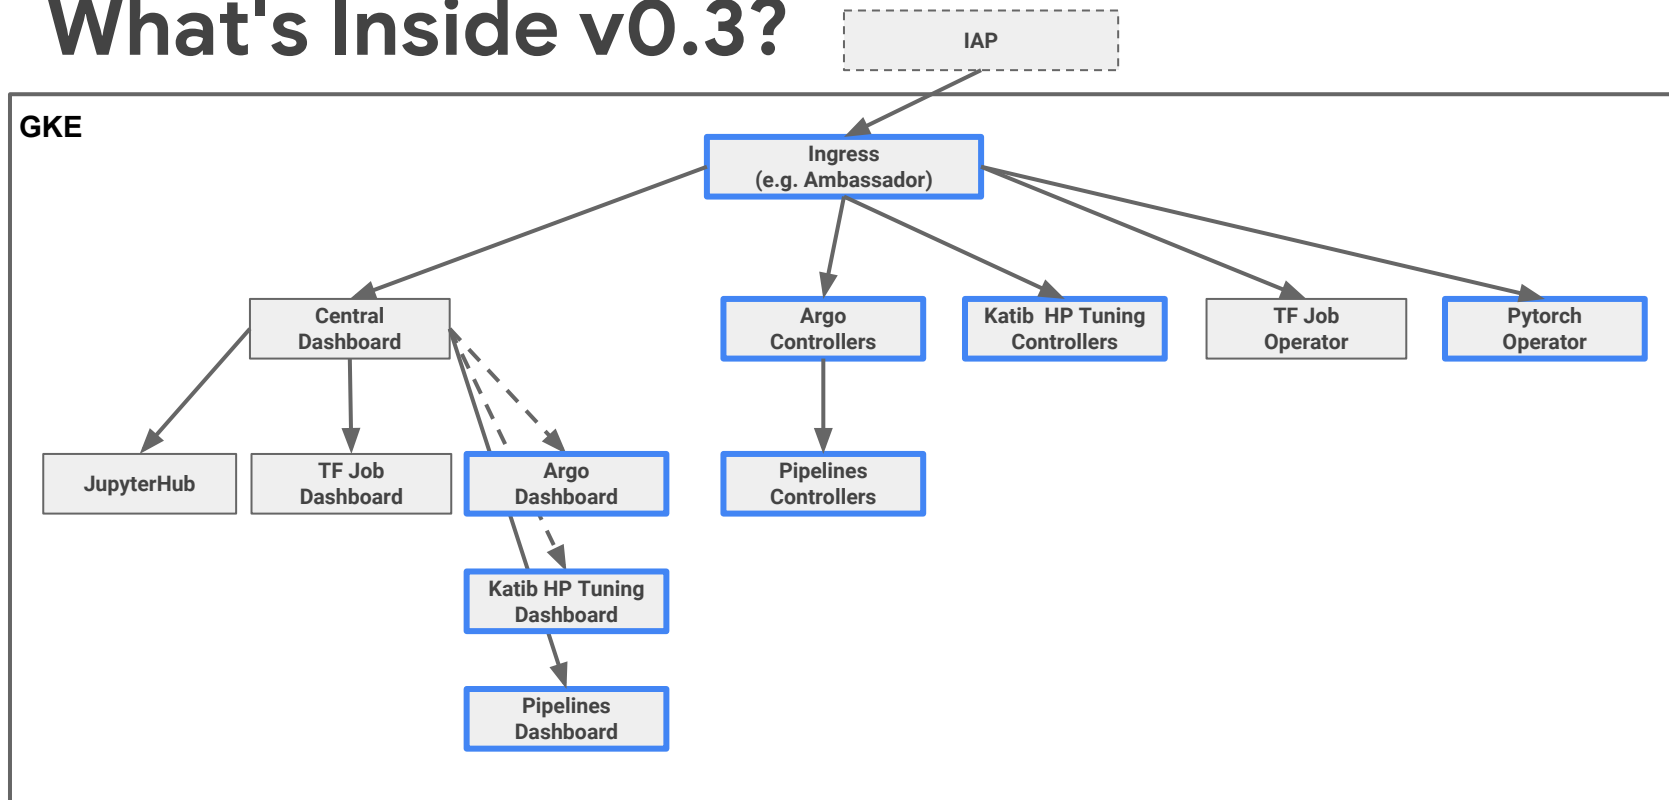
Xgboost

Et al.

Agenda g.co/code1labs/kubecon18



What's Inside v0.3?




What's new in v0.3?

- Deploy
 - Click-to-deploy
 - CLI (`kfctl`)
- Develop
 - Argo
 - Pytorch operator
 - Hyperparameter tuning StudyJob CRD
 - Kubeflow Pipelines

Pipelines

- End-to-end ML workflows
- Orchestration
- Service integration
- Components & sharing
- Job tracking, experimentation, monitoring
- Notebook integration



Agenda g.co/code1labs/kubecon18

1

Set up the environment

2

Create a Kubeflow cluster

3

Run a pipeline from the Kubeflow Pipelines dashboard

4

Run a pipeline from a Jupyter notebook

5

Clean up

Agenda g.co/code1abs/kubecon18

1

Set up the environment

2

Create a Kubeflow cluster

3

Run a pipeline from the Kubeflow Pipelines dashboard

4

Run a pipeline from a Jupyter notebook

5



Clean up

Roadmap

- v0.4 RC out Dec. 17
- Better Jupyter Notebook integration
 - Spawner with PVC support
 - Fairing alpha
- TFJob & Pytorch operators in beta
- Application CRD
- Katib TFJob support
- Pipeline experiment comparison
- v1.0 Enterprise readiness
 - Model management
 - Hardened APIs
 - Clean deployments, upgrades
- **You tell us!** (Or better yet, help!)



Contributors

- Kubeflow is open
 - Open community
 - Open design
 - Open source
 - Open to ideas
- Get involved
 - github.com/kubeflow
 - kubeflow.slack.com 
 - @kubeflow 
 - kubeflow-discuss@googlegroups.com
 - Community call Tuesdays alternating 8:30am and 5:30pm Pacific
 - **Kubeflow Contributor Summit**
 - Q1 2019



<https://github.com/kubeflow/kubeflow>



Questions?