

Machine Learning the Kubernetes Way

**KubeCon + CloudNativeCon North America 2018
December 13, 2018**

**Ania Musial, Senior Software Engineer
Keith Laban, Senior Infrastructure Engineer**

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Engineering

Bloomberg

About Us

- Ania Musial
 - Senior Software Engineer
 - Machine Learning Platform, AI Group
- Keith Laban
 - Senior Infrastructure Engineer
 - Data & Analytics Infrastructure
- <https://www.techatbloomberg.com/post-topic/data-science/>

How to Build a Data Science Platform

How to draw an owl

1.



1. ~~Start with~~
Kubernetes

2.



Draw the rest of the ~~platform~~

<https://knowyourmeme.com/memes/how-to-draw-an-owl>

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering



KubeCon + CloudNativeCon North America 2018

Machine Learning at Bloomberg

<https://www.techatbloomberg.com/post-topic/data-science>

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

The Bloomberg Terminal (1982)



<https://www.bloomberg.com/professional/blog/look-back-bloomberg-keyboard/>

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

The Bloomberg Terminal (2018)



The ***Bloomberg Terminal*** delivers a diverse array of information on a single platform to facilitate financial decision-making.

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Example Function: News Trends



TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Natural Language Processing & Events



TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Over \$136 Billion Wiped Out in Minutes

AP The Associated Press  @AP




Breaking: Two Explosions in the White House and Barack Obama is injured

 Reply  Retweet  Favorite  Buffer

3,242 RETWEETS **153** FAVORITES

12:07 PM - 23 Apr 13



TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

The Possibilities are Infinite!

- Text Processing
- Anomaly Detection
- Image Processing
- Quant/Financial Risk Modeling
- Market Data Forecasting
- Audio Transcriptions
- Search Result Relevance & Ranking
- And more...

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering



Embarking on a New Problem: What?

- What **data** is relevant to my business problem?
- What does the data **look** like?
- What **features** should I create from my data?
- What metric should I use to **measure** my model performance?
- What **algorithms** or **parameters** should I test out?
- How is my model doing on live data in **production**?

What ML Practitioners Say They Want

Compute Infrastructure

ML Workflow & Ease of Use

Data Liquidity

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering



Who is Involved in Building Such a System?

- ML Practitioners
 - Dev Ops
 - Infrastructure Engineers
 - Business
 - Security Officers
-
- And remember our (paying!) Bloomberg Terminal clients
 - Smarter Functionality, Privacy, Correctness, Uptime/Stability



KubeCon + CloudNativeCon North America 2018

Life Before Platform Era

Failures Prior Work

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Rules-based Systems → ML: Status Report

Data

- Application-specific databases
- Small-ish

Compute

- Age of monolithic servers
- Resource squatting
- `import multiprocessing`

ML Workflow

- Interactive
- Custom tooling around SDLC
- Lack of infrastructure composability
- Lack of cross-team model sharing

Support

- Self-regulated, but contained
- Left up to full-stack application teams or machine owners

Big Data Era: Status Report

Data

- Data lake in HDFS/Hadoop
- Highly secured
- But what about joining on data not in the lake?

Compute

- Larger, more distributed
- Limited by data locale

ML Workflow

- Spark ML only gets you so far
- Subject to cluster-wide dependencies
- Auth protocols :(

Support

- Self-regulated, subject to workflows
- Distributed systems or Spark experts, but not ML folks

GPU Era: Status Report

Data

- *Okay now how do I get this data lake stuff from one network over to this box over here...*

Compute

- Access to blazing fast hardware
- *Stop running so many notebooks, I need those devices*

ML Workflow

- Interactive
- Neural network packages
- But what about ETL?

Support

- Scheduling is self-regulated, and on more expensive hardware
- Drivers; hardware/-software coupling

This is the “We Chose Kubernetes” Slide



kubernetes



**CLOUD NATIVE
COMPUTING FOUNDATION**

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering



KubeCon + CloudNativeCon North America 2018

Platform 101

Custom Resources

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Kubernetes Meets ML: TensorFlowJob

```
apiVersion: ds.bloomberg.com/v1
kind: TensorFlowJob
metadata:
  name: tf-test
  namespace: keithlaban
spec:
  framework: tensorflow-1.7-python-3
  pipPackages: [ai.bloomberg.com.myteam.gpu_tftraining, numpy]
  module: gpu_tftraining
  identities:
    - hadoop:
      id: keithlaban
  size: GpuLarge
  args:
    --data-dir hdfs://CLUSTER/projects/news/news_index
    --output-dir hdfs://CLUSTER/users/klaban1/abcde/1
```

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Kubernetes Meets ML: TensorFlowJob

```
apiVersion: ds.bloomberg.com/v1
kind: TensorFlowJob
metadata:
  name: tf-test
  namespace: keithlaban
spec:
  framework: tensorflow-1.7-python-3
  pipPackages: [ai.bloomberg.com.myteam.gpu_tftraining, numpy]
  module: gpu_tftraining
  identities:
    - hadoop:
      id: keithlaban
  size: GpuLarge
  args:
    --data-dir hdfs://CLUSTER/projects/news/news_index
    --output-dir hdfs://CLUSTER/users/klaban1/abcde/1
```

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Kubernetes Meets ML: TensorFlowJob

```
apiVersion: ds.bloomberg.com/v1
kind: TensorFlowJob
metadata:
  name: tf-test
  namespace: keithlaban
spec:
  framework: tensorflow-1.7-python-3
  pipPackages: [ai.bloomberg.com.myteam.gpu_tftraining, numpy]
  module: gpu_tftraining
  identities:
    - hadoop:
        id: keithlaban
  size: GpuLarge
  args:
    --data-dir hdfs://CLUSTER/projects/news/news_index
    --output-dir hdfs://CLUSTER/users/klaban1/abcde/1
```

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Kubernetes Meets ML: TensorFlowJob

```
apiVersion: ds.bloomberg.com/v1
kind: TensorFlowJob
metadata:
  name: tf-test
  namespace: keithlaban
spec:
  framework: tensorflow-1.7-python-3
  pipPackages: [ai.bloomberg.com.myteam.gpu_tftraining, numpy]
  module: gpu_tftraining
  identities:
    - hadoop:
      id: keithlaban
  size: GpuLarge
  args:
    --data-dir hdfs://CLUSTER/projects/news/news_index
    --output-dir hdfs://CLUSTER/users/klaban1/abcde/1
```

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Kubernetes Meets ML: TensorFlowJob

```
apiVersion: ds.bloomberg.com/v1
kind: TensorFlowJob
metadata:
  name: tf-test
  namespace: keithlaban
spec:
  framework: tensorflow-1.7-python-3
  pipPackages: [ai.bloomberg.com.myteam.gpu_tftraining, numpy]
  module: gpu_tftraining
  identities:
    - hadoop:
      id: keithlaban
  size: GpuLarge
  args:
    --data-dir hdfs://CLUSTER/projects/news/news_index
    --output-dir hdfs://CLUSTER/users/klaban1/abcde/1
```

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Kubernetes Meets ML: TensorFlowJob

```
apiVersion: ds.bloomberg.com/v1
kind: TensorFlowJob
metadata:
  name: tf-test
  namespace: keithlaban
spec:
  framework: tensorflow-1.7-python-3
  pipPackages: [ai.bloomberg.com.myteam.gpu_tftraining, numpy]
  module: gpu_tftraining
  identities:
    - hadoop:
      id: keithlaban
  size: GpuLarge
args:
  --data-dir hdfs://CLUSTER/projects/news/news_index
  --output-dir hdfs://CLUSTER/users/klaban1/abcde/1
```

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

A Platform Needs More (CRDs)

- **Single node**
 - *PythonJob*
 - *JVMJob*
- **Distributed ETL**
 - *SparkJob*
- **GPU**
 - *TensorFlowJob*
- **Hyperparameter Tuning**
 - *HypertuneJob*
- **Foundational Components**
 - Data access, privacy, other multitenancy woes...

Platform 101: Status Report

Data

- HDFS
- S3
- Databases
- Permission Controlled Access to Data Warehouses

Compute

- “Cloud”
- GPU
- CPU

ML Workflow

- CRDs are ML
- ETL -> GPU
- Declarative

Support

- Kubernetes monitoring
- Prometheus
- Cleaner split for DevOps



KubeCon + CloudNativeCon North America 2018

Nebulous and Misunderstood

Jupyter, Interactive Debugging, Tooling

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

The *Cloud* Is Hard

- No native UNIX CLIs
- Paradigm shift for code execution and debugging
- Container-unaware applications

The *Cloud* Doesn't Have to be Hard

- *JupyterJob*
- `kubectl exec`
- UI
- Education

Living in the Nebula: Status Report

Data

- HDFS
- S3
- Databases
- Permission Controlled
Access to Data
Warehouses

Compute

- “Cloud”
- GPU
- CPU
- **Interactive**

ML Workflow

- CRDs are ML
- ETL -> GPU
- Declarative
- **Jupyter**

Support

- Kubernetes monitoring
- Prometheus
- Cleaner split for DevOps
- **Worry about Squatters**



KubeCon + CloudNativeCon North America 2018

Mo' People Mo' Tenancy

Namespaces, Resource Quotas, Preemption, Deadlines

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Sharing is Caring

- *Workspace*
 - Namespace
 - RBAC
 - Service Account
 - Resource Quotas
 - Network Policy
- *Jobs*
 - Priority Classes
 - Deadlines
- Tiered SLAs
 - User vs. Shared

A Shared Status Report

Data

- HDFS
- S3
- Databases
- Permission Controlled Access to Data Warehouses

Compute

- “Cloud”
- GPU
- CPU
- Local
- **Preemptable**

ML Workflow

- CRDs are ML
- ETL -> GPU
- Declarative
- Jupyter

Support

- Kubernetes monitoring
- Prometheus
- Cleaner split for DevOps
- ~~Worry about Squatters~~
- **System handles capacity**



KubeCon + CloudNativeCon North America 2018

Model Development Life Cycle Robustness

History, Labels, Metrics, Experiment Management

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Getting More Sophisticated

- Job History
- Experiment Trial Aggregation
- Model Metrics

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering



MDLC: Status Report

Data

- HDFS
- S3
- Databases
- Permission Controlled Access to Data
- Warehouses

Compute

- “Cloud”
- GPU
- CPU
- Local
- Preemptable

ML Workflow

- CRDs are ML
- ETL -> GPU
- Declarative
- Jupyter
- Projects
- Metrics
- Retraining

Support

- Kubernetes monitoring
- Prometheus
- Cleaner split for DevOps
- System handles capacity



KubeCon + CloudNativeCon North America 2018

To Inference and Beyond

KNative, A Human-in-the-Loop, Spark Native Scheduling

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

What We're Working on Now (Models)

- Inference as a Service
 - Lambdas via KNative
- Connectivity to other components
 - Annotation Platform (Human in the Loop)

What We're Working on Now (ETL)

- Spark Native Kubernetes Scheduling (Spark 2.3)
 - <https://github.com/apache/spark/pull/19468>
- Spark Workload Autoscaling
 - <https://issues.apache.org/jira/browse/SPARK-25299>



KubeCon + CloudNativeCon North America 2018

Concluding Remarks

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Thanks!

- We hope you saw our teammates' talks (if not, watch for the videos):
 - *Airflow on Kubernetes, Dynamic Workflows Simplified* (Daniel Imberman)
 - *Visualise Cluster Connectivity with Goldpinger: Smash it with PowerfulSeal* (Mikolaj Pawlikowski)
- Be our teammate
 - <https://www.bloomberg.com/careers/technology/engineering/data-science-teams/>



KubeCon + CloudNativeCon North America 2018

Questions?

TechAtBloomberg.com

© 2018 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering