

# TensorFlowOnSpark

Scaled ML@Stanford  
March 25, 2017

# Machine Learning, AI & No Free Lunch

- Five key ingredients for ML towards AI
  1. Lots & lots of data
  2. Very flexible models
  3. Enough computing power
  4. Computationally efficient inference
  5. **Powerful priors that can defeat the curse of dimensionality**

\* Yoshua Bengio @ ICDM 2016

# What is TensorFlowOnSpark?



YAHOO!

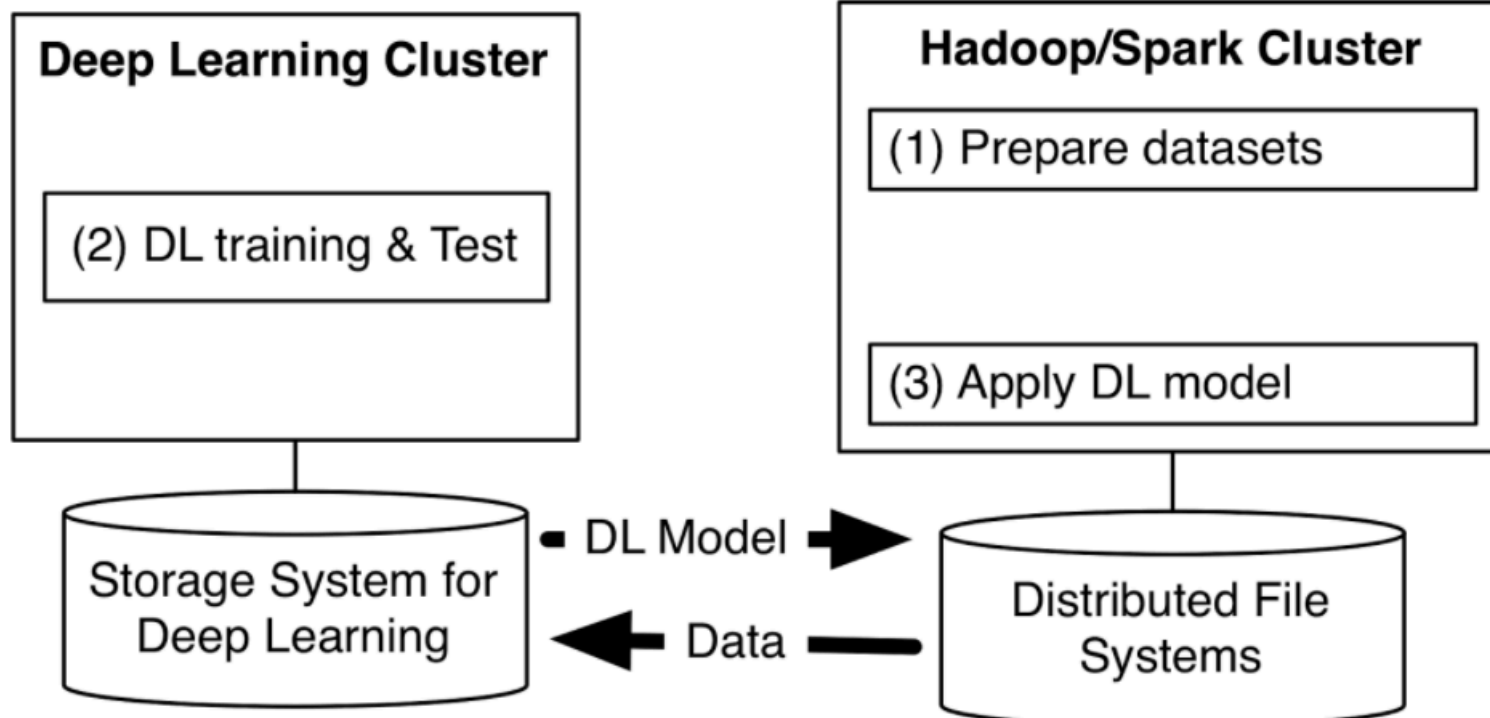
# What's TensorFlowOnSpark?

- Scale up TensorFlow apps with minimal changes
- Support all TensorFlow functionalities
  - Model/data parallelism, Synchron/Asynchron, TensorBoard
- Integrate with existing data & pipeline
  - ex. HDFS, SQL, MLlib
- Deployable on cloud or on-premise

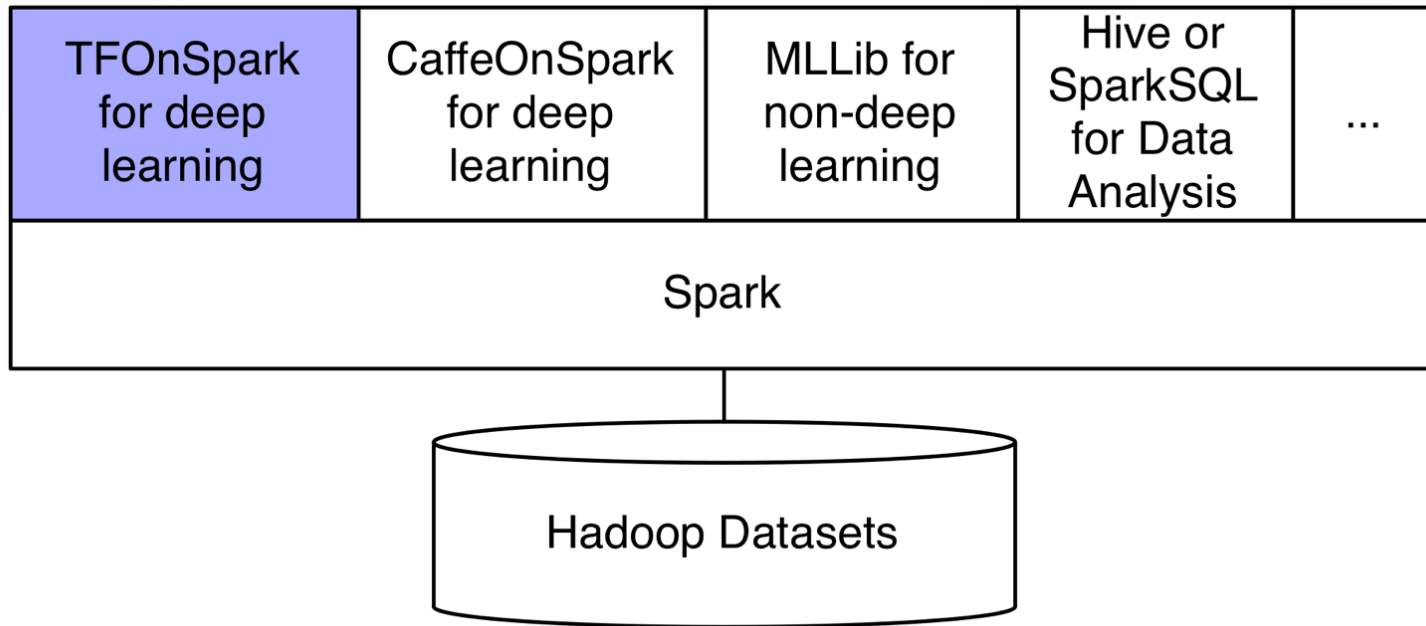
# Why TensorFlowOnSpark at Yahoo?

- Major player of open-source ecosystem
  - Birth place of Apache Hadoop
  - Adopter/contributor of Spark since 2013
- Large clusters in house
  - Tens of clusters
  - Thousands of nodes per cluster
- Massive amount of data
  - Petabytes of data

# Why TensorFlowOnSpark?



# TensorFlowOnSpark



# Open Source: github.com/yahoo/TensorFlowOnSpark

yahoo / TensorFlowOnSpark

Unwatch 128 Unstar 1,215 Fork 246

Code Issues 9 Pull requests 3 Projects 0 Wiki Pulse Graphs Settings

TensorFlowOnSpark brings TensorFlow programs onto Apache Spark clusters

Edit

tensorflow spark yahoo machine-learning cluster Manage topics

83 commits 4 branches 0 releases 6 contributors Apache-2.0

Branch: master New pull request Create new file Upload files Find file Clone or download

anfeng committed on GitHub Merge pull request #47 from yahoo/leewyang\_py3\_compat Latest commit 5863dfc 18 days ago

examples	python3 compatibility	18 days ago
scripts	remove local-setup-hadoop.sh	a month ago
src/com	python3 compatibility	18 days ago
tensorflow @ 22fd3aa	add tensorflow submodule	2 months ago
.gitmodules	add tensorflow submodule	2 months ago
LICENSE	initial check-in	2 months ago
README.md	README w/ a blog link	a month ago



# TensorFlowOnSpark

- Launches TF clusters using Spark executors
- Supports TF data ingestion modes
  - Spark – RDD.mapPartitions()
  - TensorFlow – directly access HDFS
- Supports TensorBoard during/after training
- Generally agnostic to Spark/TF versions

# TFoS Basics

1. **Launch** TensorFlow cluster
2. **Feed data** to TensorFlow app
3. **Shutdown** TensorFlow cluster

# TFoS Python API

```
cluster = TFCluster.run(sc, map_fn, args, num_executors,  
num_ps, tensorboard, input_mode)
```

```
cluster.train(dataRDD, num_epochs=0)
```

```
cluster.inference(dataRDD)
```

```
cluster.shutdown()
```

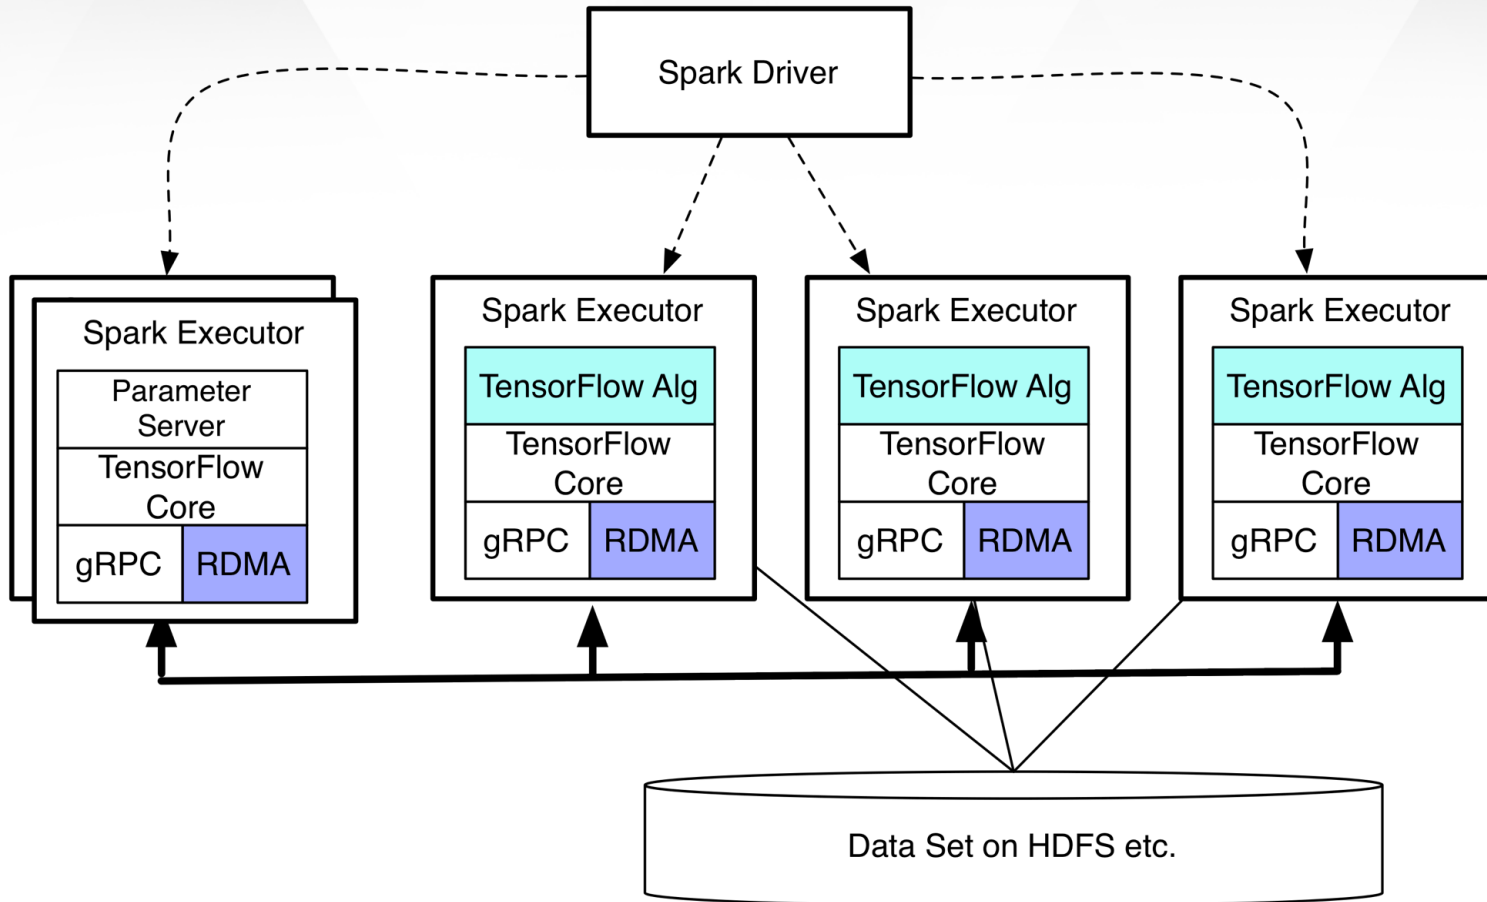
# TFoS: Minimum Code Changes

```
# diff -w eval_image_classifier.py
20a21,27
> from pyspark.context import SparkContext
> from pyspark.conf import SparkConf
> from com.yahoo.ml.tf import TFCluster, TFNode
> import sys
>
> def main_fun(argv, ctx):
27a35,36
>     sys.argv = argv
84,85d92
< def main(_):
88a96,97
>     cluster_spec, server = TFNode.start_cluster_server(ctx)
191c200,204
<     tf.app.run()
---
>     sc = SparkContext(conf=SparkConf().setAppName("eval_image_classifier"))
>     num_executors = int(sc._conf.get("spark.executor.instances"))
>     cluster = TFCluster.run(sc, main_fun, sys.argv, num_executors, 0, False, TFCluster.InputMode.TENSORFLOW)
>     cluster.shutdown()
```

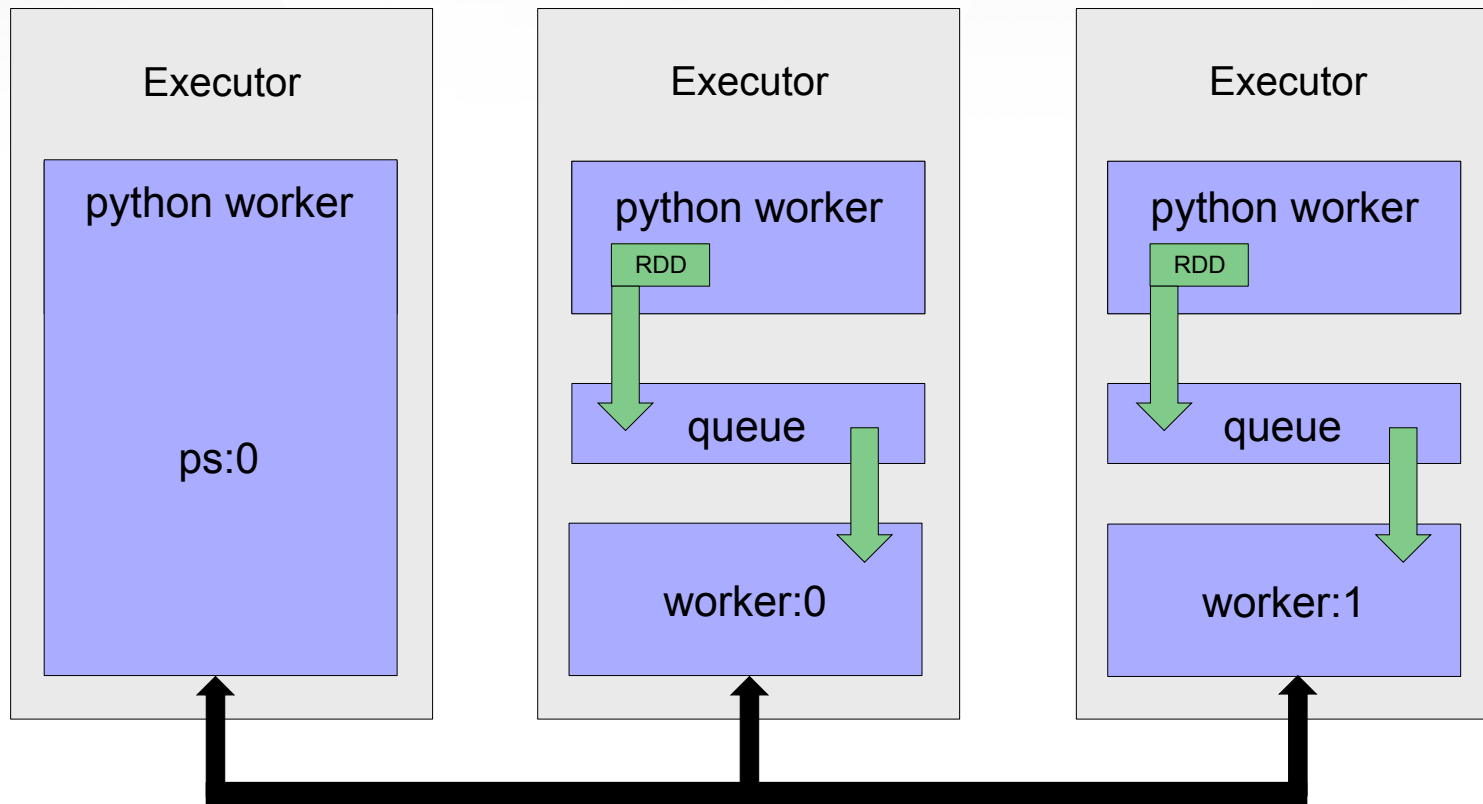
# TFoS Input Modes

- **InputMode.SPARK**
  - feed\_dict
  - Small-medium scale data
  - Fed via `RDD.mapPartitions()`
- **InputMode.TENSORFLOW**
  - Reader + QueueRunner
  - Large scale data
  - Reads directly from HDFS

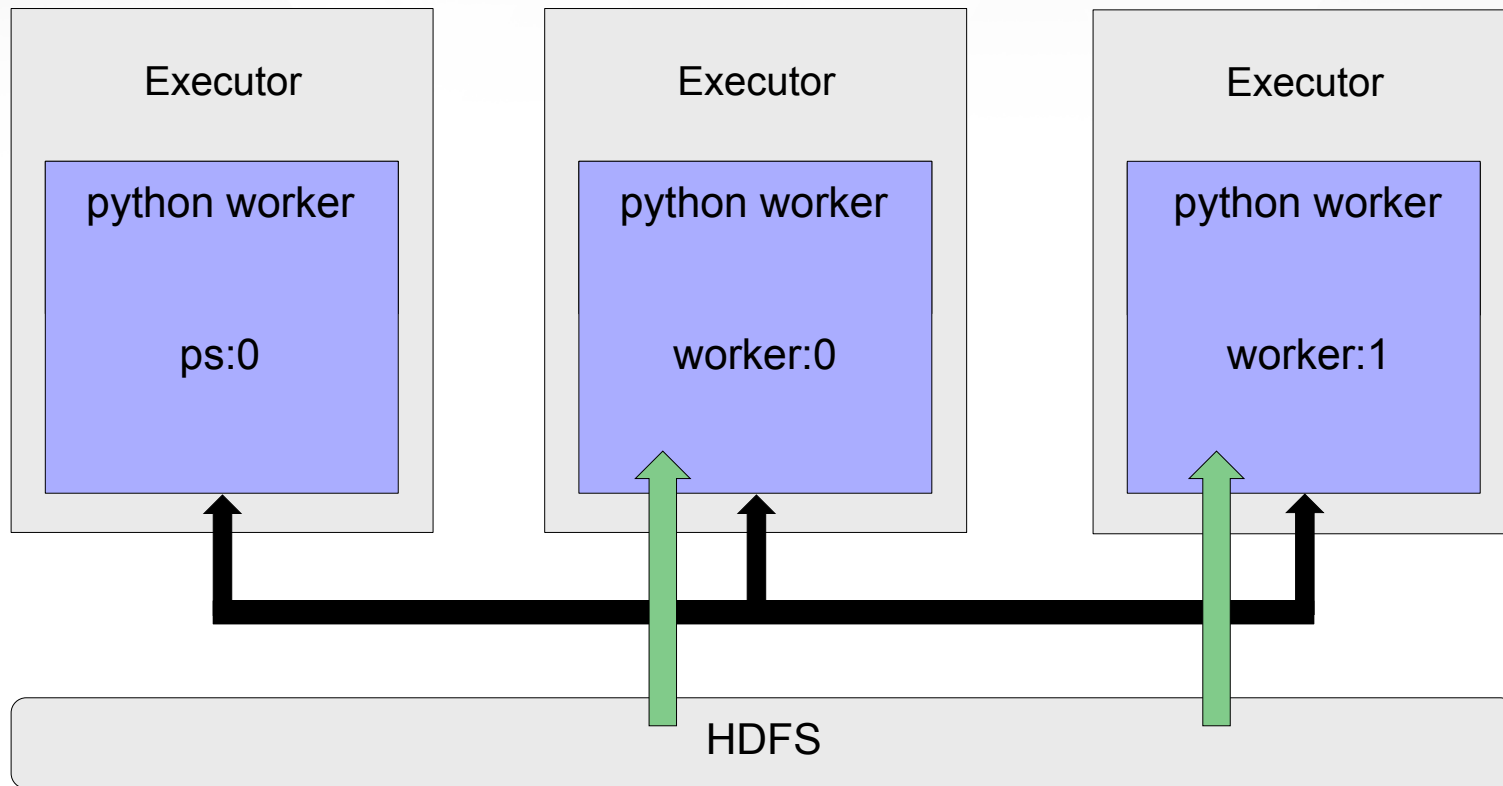
# TFoS Architecture



# TFoS: InputMode.SPARK



# TFoS: InputMode.TENSORFLOW





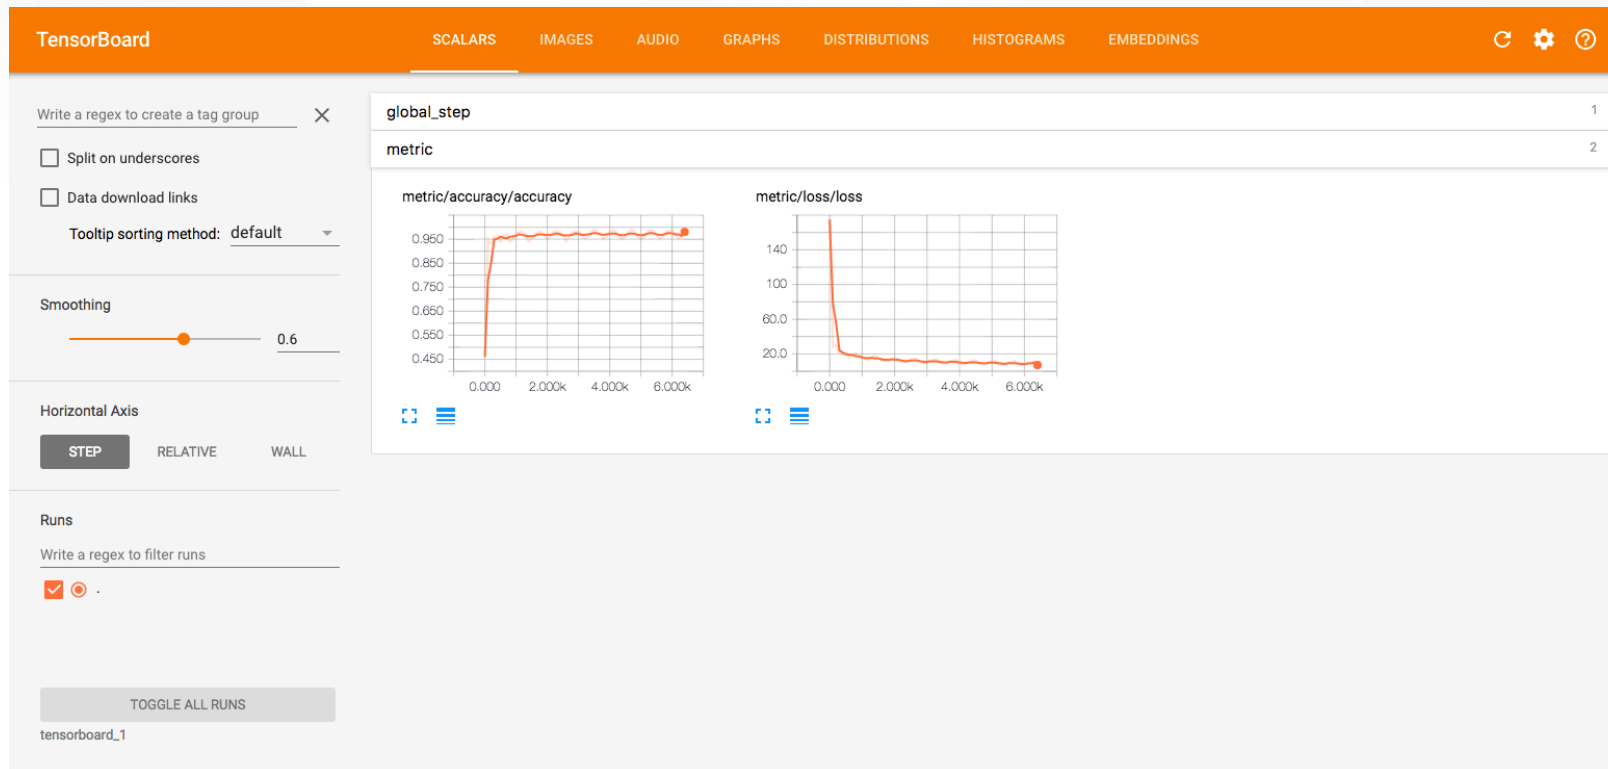
# TFoS: Failure Recovery

- TF Checkpoints written to HDFS
- **InputMode.SPARK**
  - TF worker runs in background
  - RDD data feeding tasks can be retried
  - However, TF worker failures will be “hidden” from Spark
- **InputMode.TENSORFLOW**
  - TF worker runs in foreground
  - TF worker failures will be retried as Spark task
  - TF worker restores from checkpoint

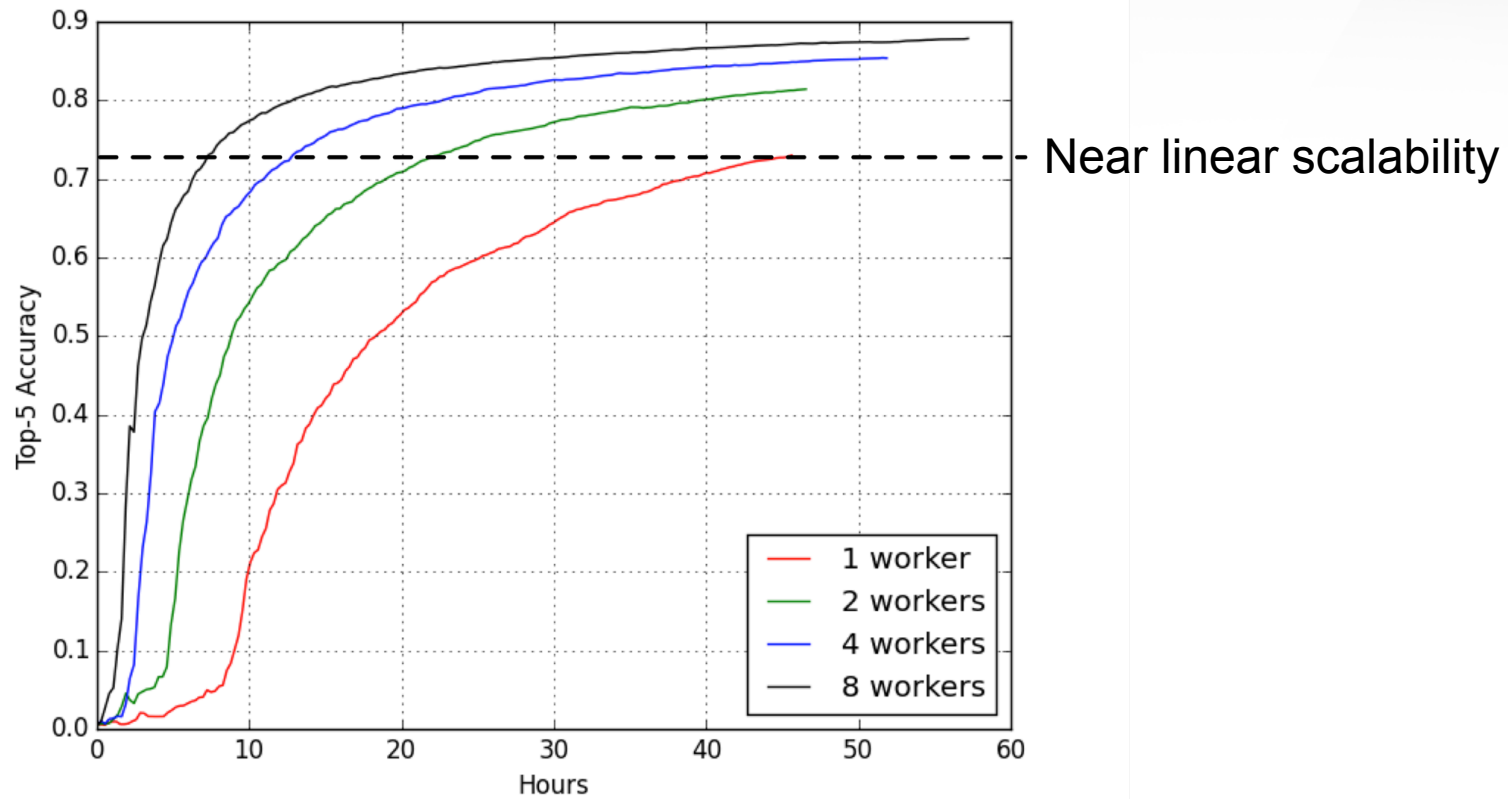
# TFoS: Failure Recovery

- **Executor failures are problematic**
  - TF cluster\_spec is statically-defined
  - YARN doesn't re-allocate on same node
  - Port may no longer be available
- **Need dynamic cluster membership**
  - Explore options w/ TensorFlow team

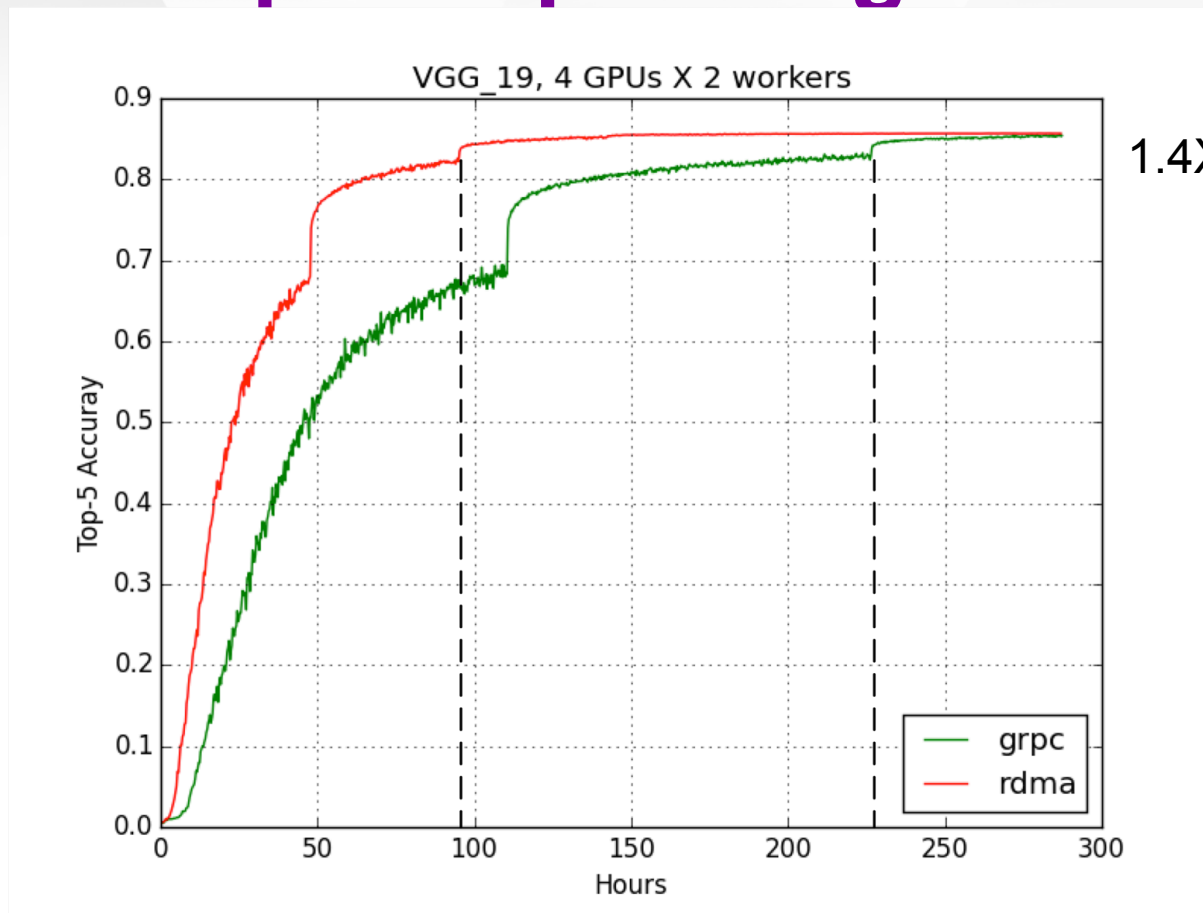
# TensorBoard on TFoS



# TFoS Scaling



# RDMA Speedup over gRPC



1.4X speedup

## Related Work

	SparkNet	TensorFrames	TFonS
Programming Language	Scala	Python, Scala	Python
Migration	Major	Medium	Minor
Parallelism	Data Parallelism	Data Parallelism	Data + Model Parallelism
Distributed Training	Synchronous	Synchronous	Synchronous + Asynchronous
TensorBoard	X	X	✓
Scalability	Driver bottleneck	Driver bottleneck	✓

# Summary

- TFoS brings deep learning to big-data clusters
  - TensorFlow: 0.12 -1.0
  - Spark: 1.6-2.x
  - Cluster manager: YARN, Standalone, Mesos
  - EC2 image provided
- RDMA enhancement for faster training
  - PR for [github/tensorflow](https://github.com/tensorflow/tensorflow) repo

**Questions?**

<https://github.com/yahoo/TensorFlowOnSpark>