

Research Review 2017

Measuring Performance of Big Learning Workloads

Dr. Scott McMillan, Senior Research Scientist

Copyright 2017 Carnegie Mellon University. All Rights Reserved.

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

The view, opinions, and/or findings contained in this material are those of the author(s) and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

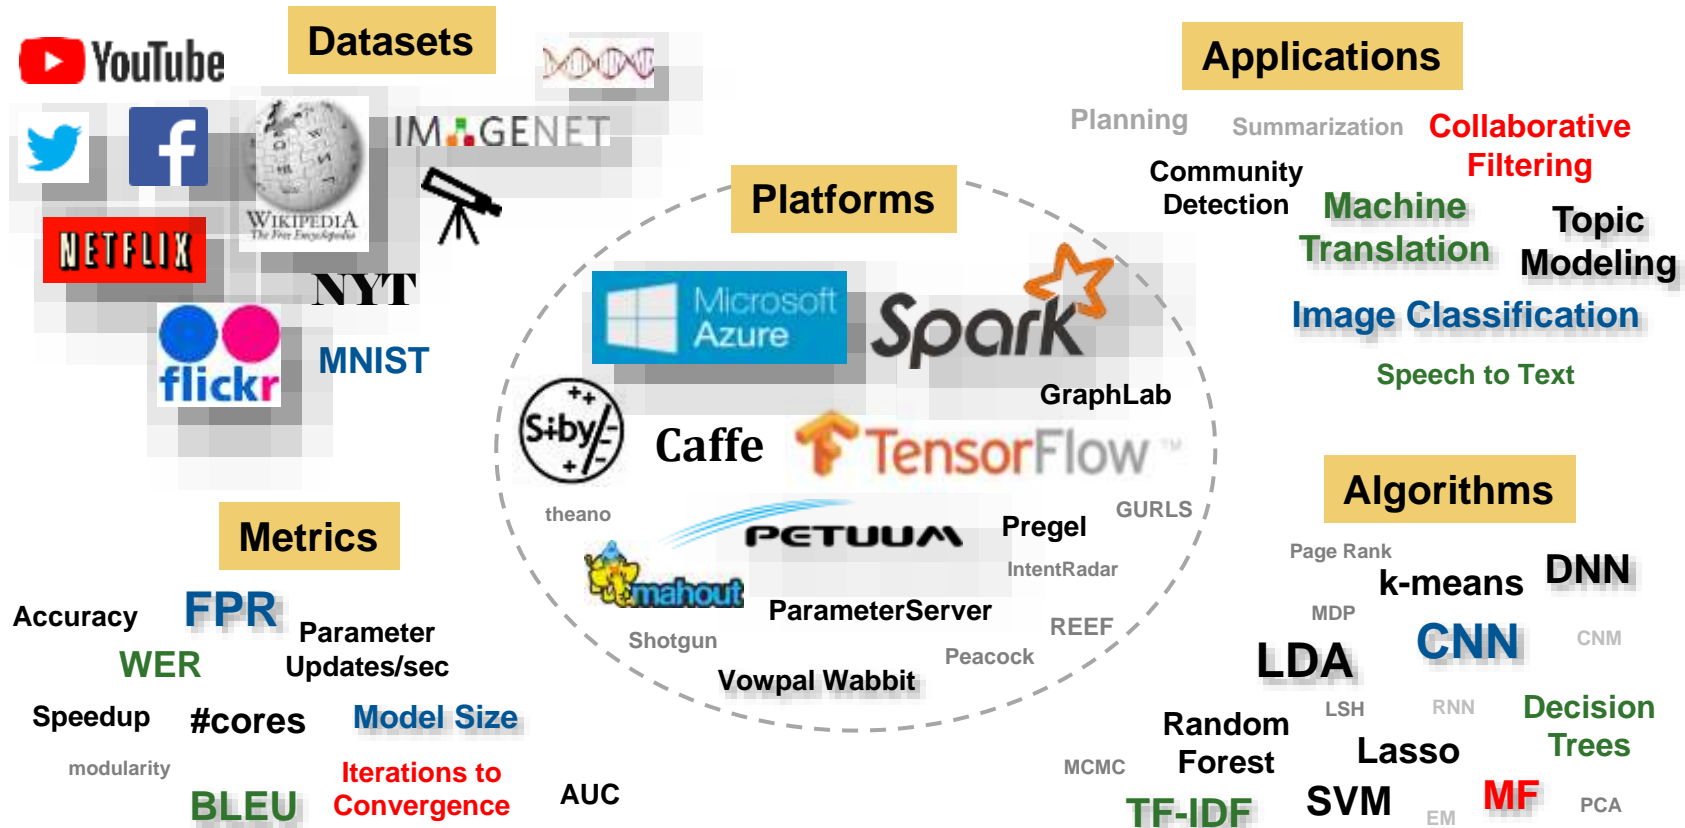
[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.

This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

Carnegie Mellon® is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

DM17-0783

Project Introduction: Big Learning Landscape



Big Learning: Large-Scale Machine Learning on Big Data

Problem:

- Over 1,000 papers are published each year in Machine Learning.
- Most are empirical studies and few (if any) provide enough detail to reproduce the results.
- Complexity of systems begets complexity in metrics – often partially reported.
- Slows adoption by DoD of new advances in Machine Learning.

Solution:

- Facilitate consistent research comparisons and advancements of Big Learning systems by providing sound **reproducible** ways to measure and report performance.

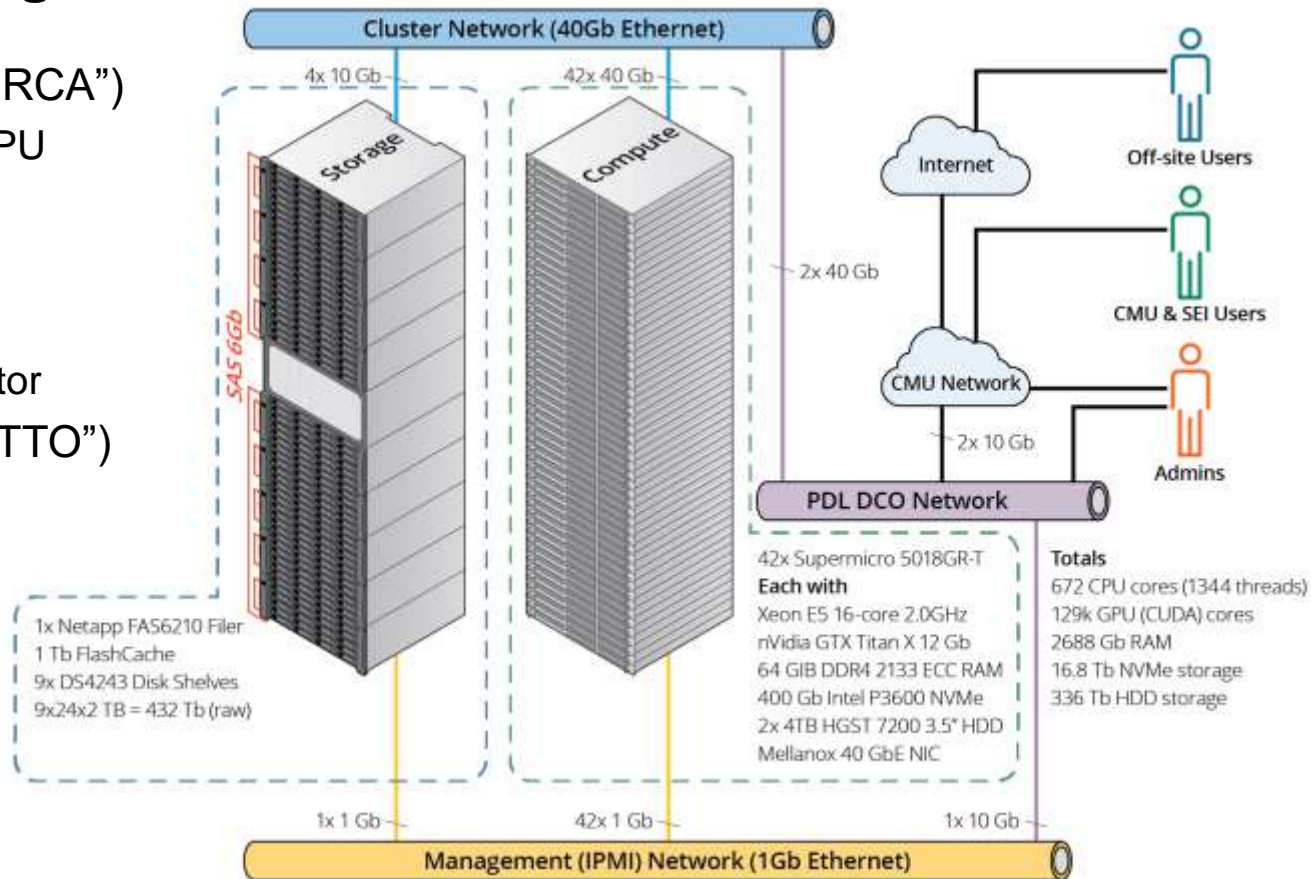
Approach:

- Develop technology platform for consistent (and complete) evaluation.
 - Performance of the computing system
 - Performance of the ML application
- Evaluate relevant Big Learning platforms using this benchmark: Spark+MLlib, Tensorflow, Petuum
- Collaborate with CMU's Big Learning Research Group

¹FACT SHEET: National Strategic Computing Initiative, 29 July 2015.

ORCA: Big Learning Cluster

- 42 compute nodes (“ORCA”)
 - 16 core (32 thread) CPU
 - 64GB RAM
 - 400GB NVMe
 - 8TB HDD
 - Titan X GPU accelerator
- Persistent storage (“OTTO”)
 - 400+ TB storage
- 40Ge networking



Performance Measurement Workbench (PMW) Architecture

Persistent Services

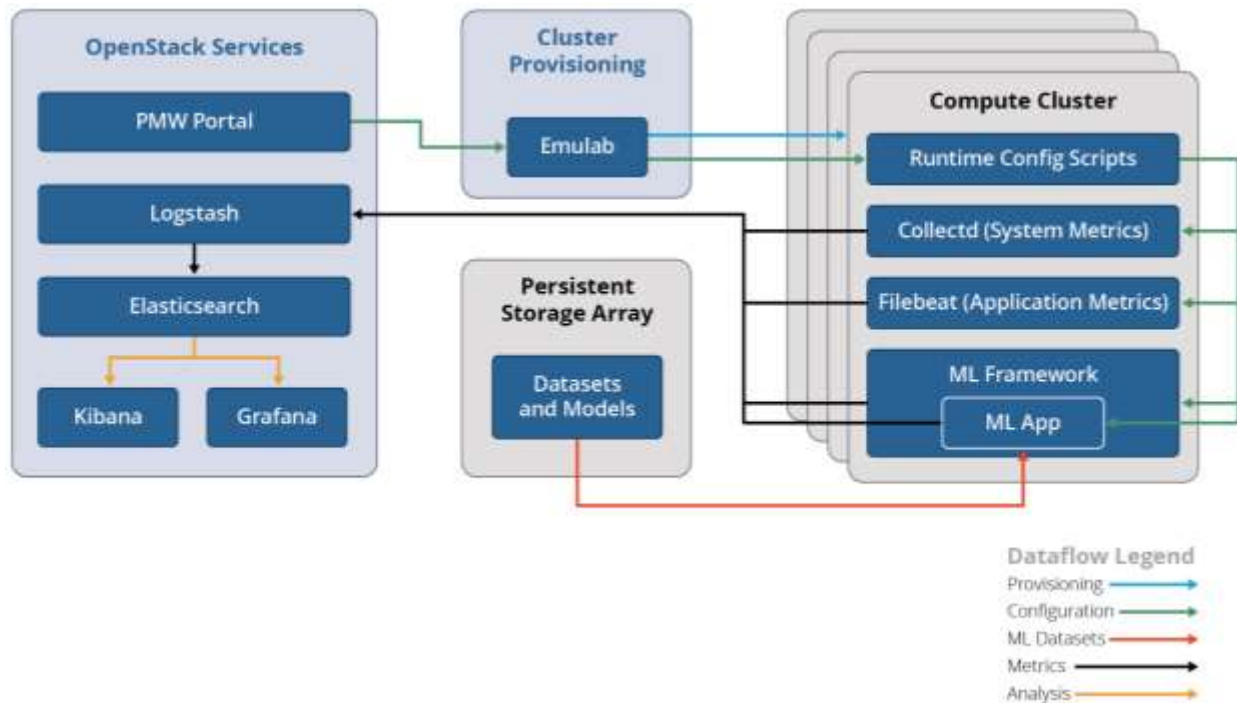
- OpenStack
- Web portal
 - simplifies provisioning
 - coordinates tools
- Data collection/analysis
 - “Elastic Stack”
 - Grafana

Provisioning

- Emulab
- “bare-metal” provisioning

Hardware Resources

- Compute cluster
- Data storage



Performance Measurement Workbench (PMW) Architecture

Persistent Services

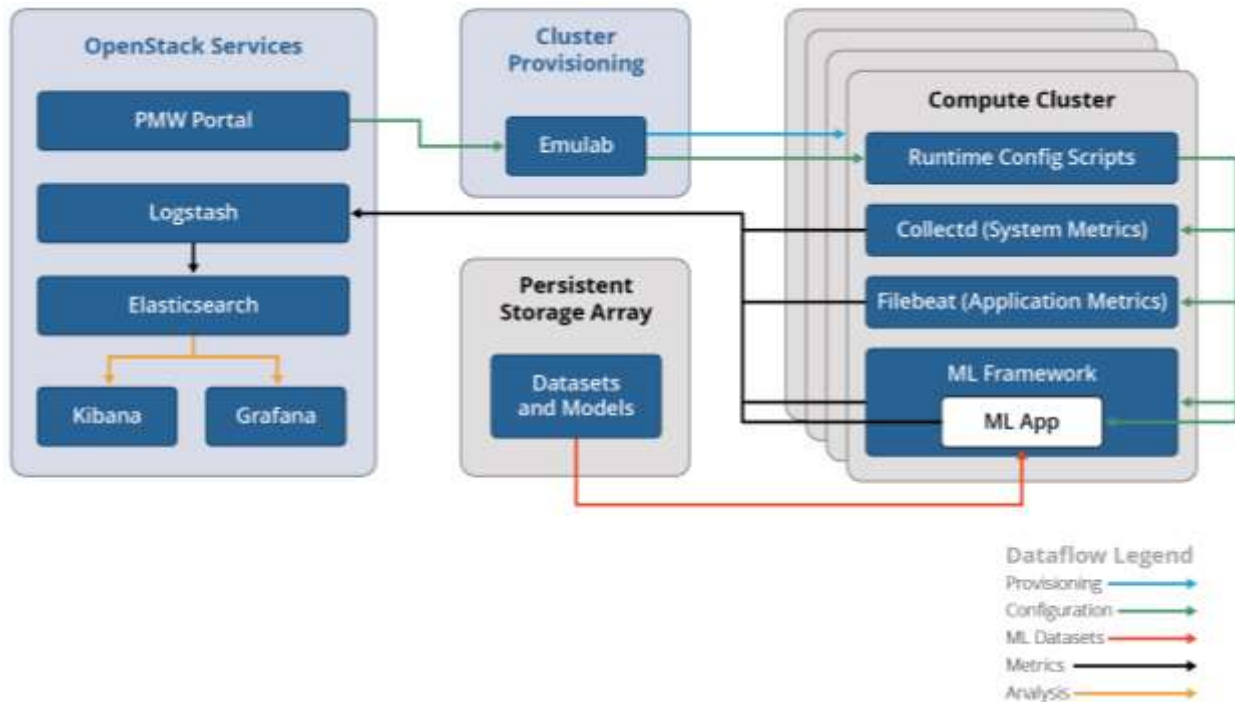
- OpenStack
- Web portal
 - simplifies provisioning
 - coordinates tools
- Data collection/analysis
 - “Elastic Stack”
 - Grafana

Provisioning

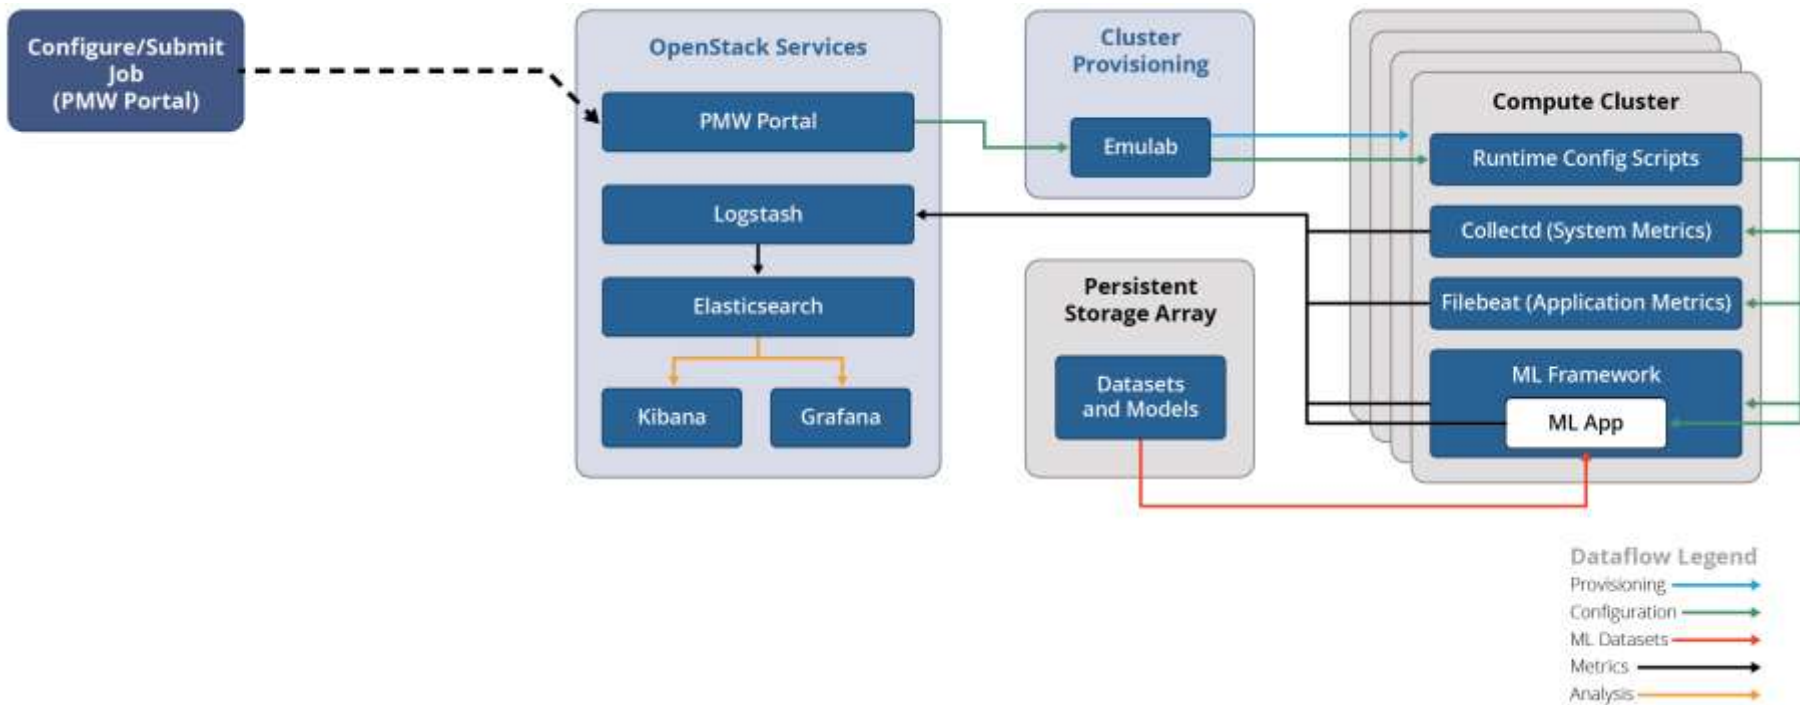
- Emulab
- “bare-metal” provisioning

Hardware Resources

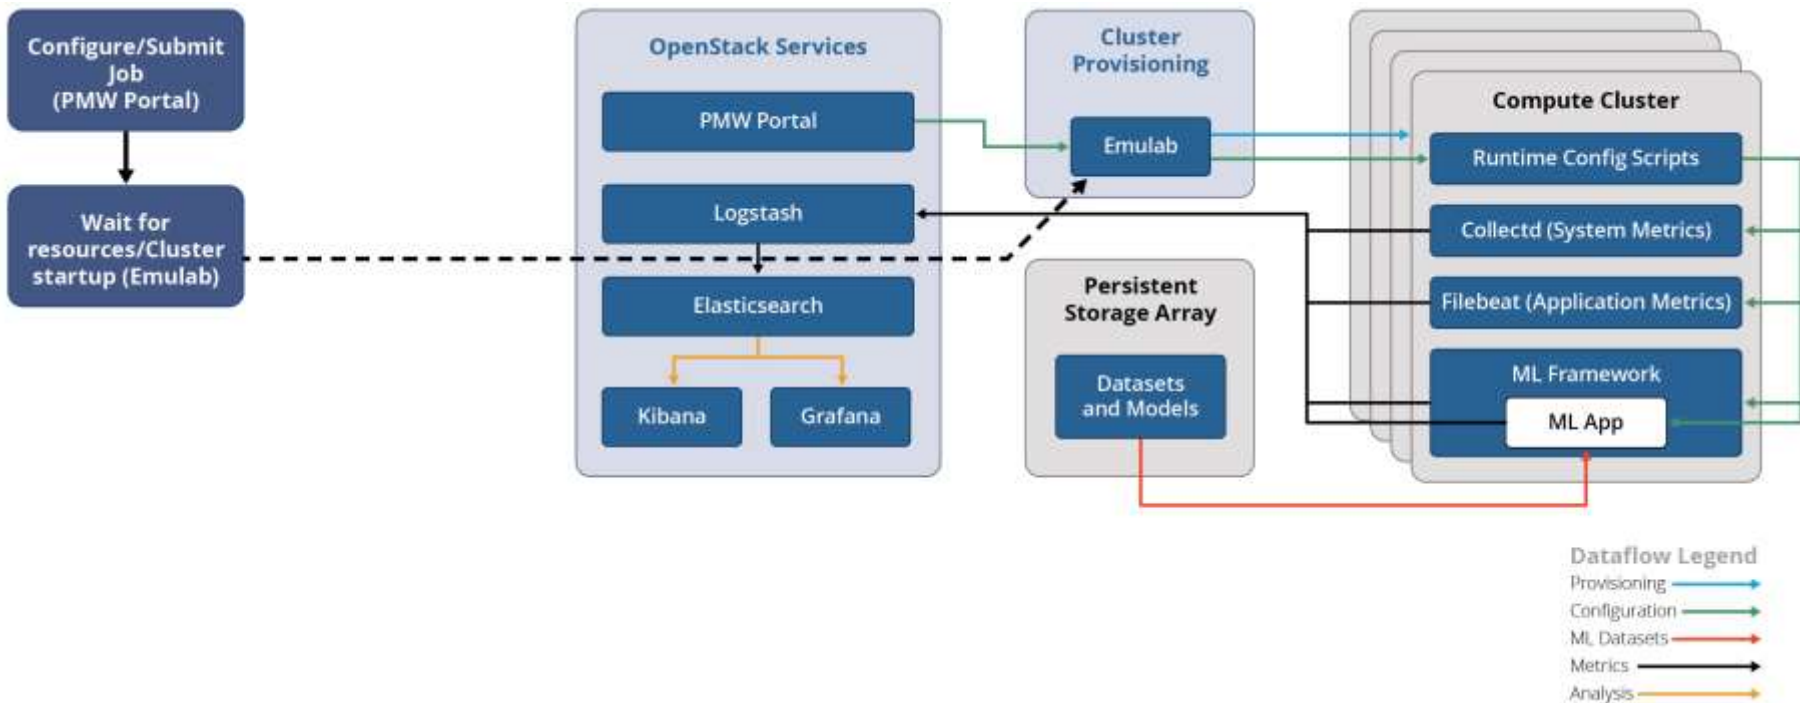
- Compute cluster
- Data storage



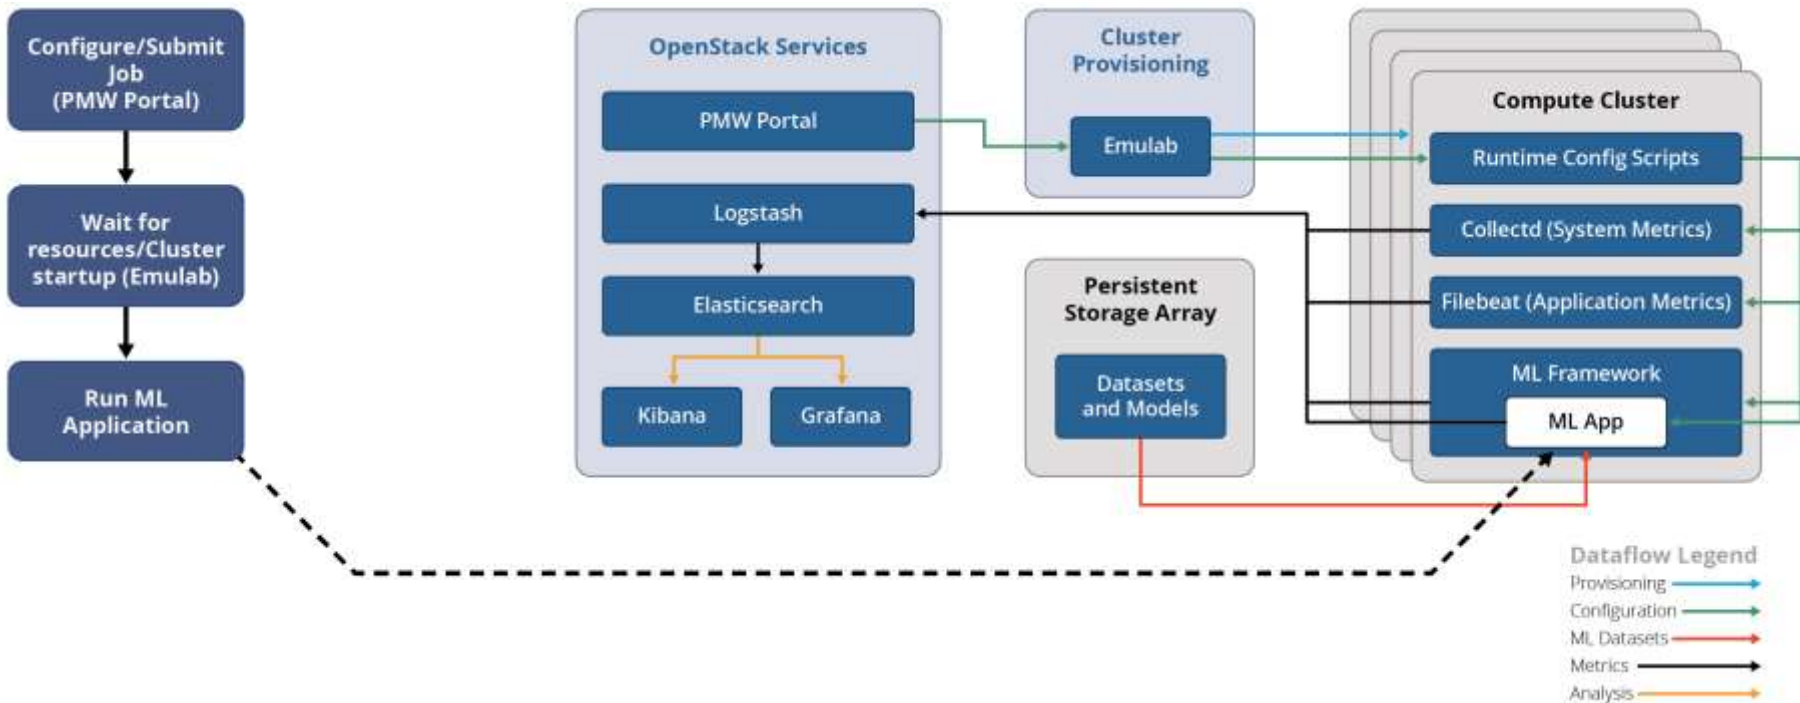
PMW Workflow



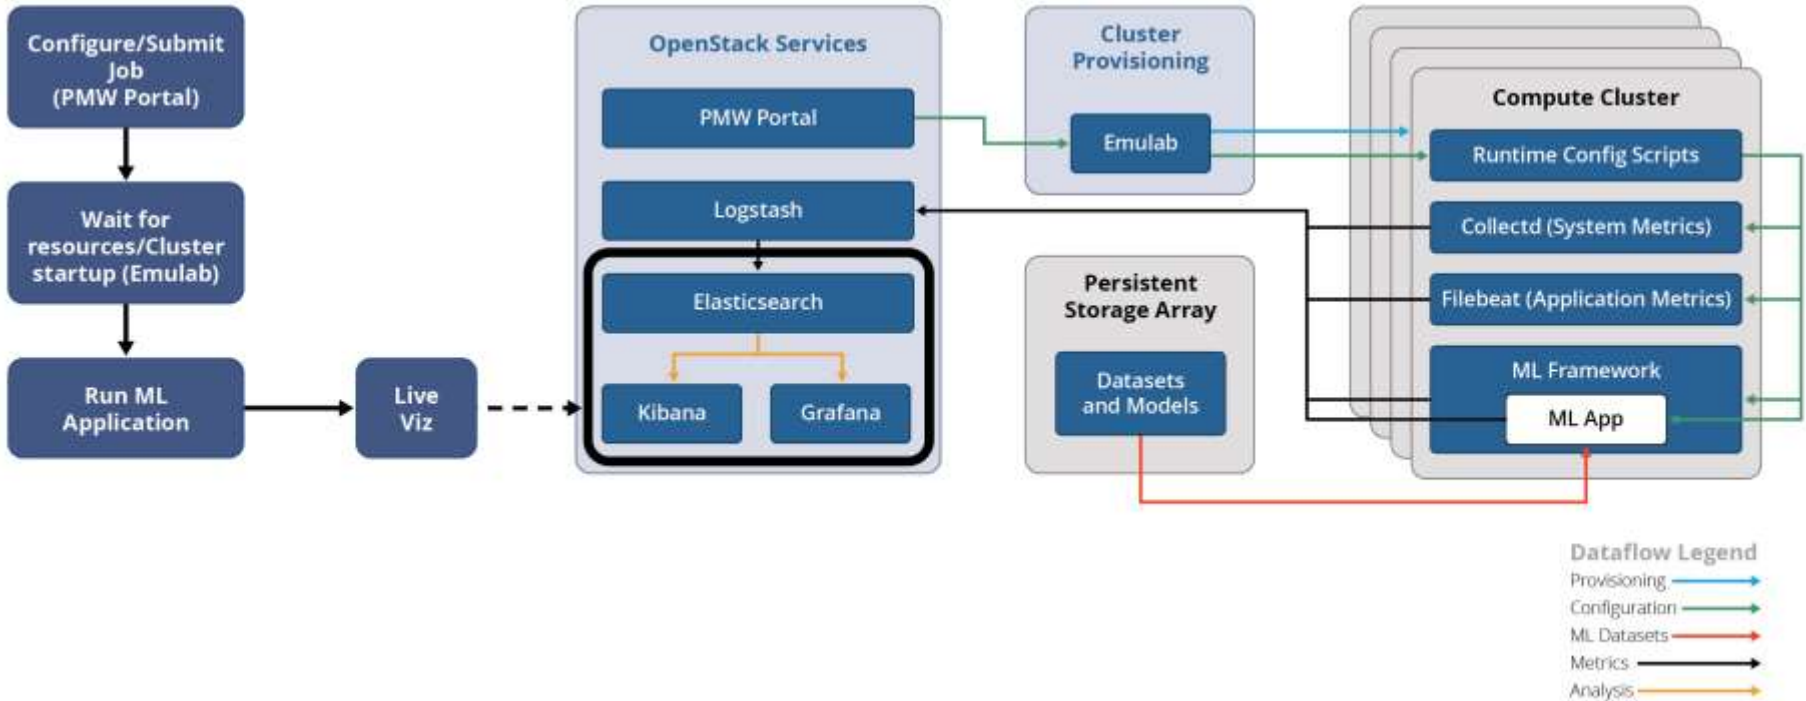
PMW Workflow



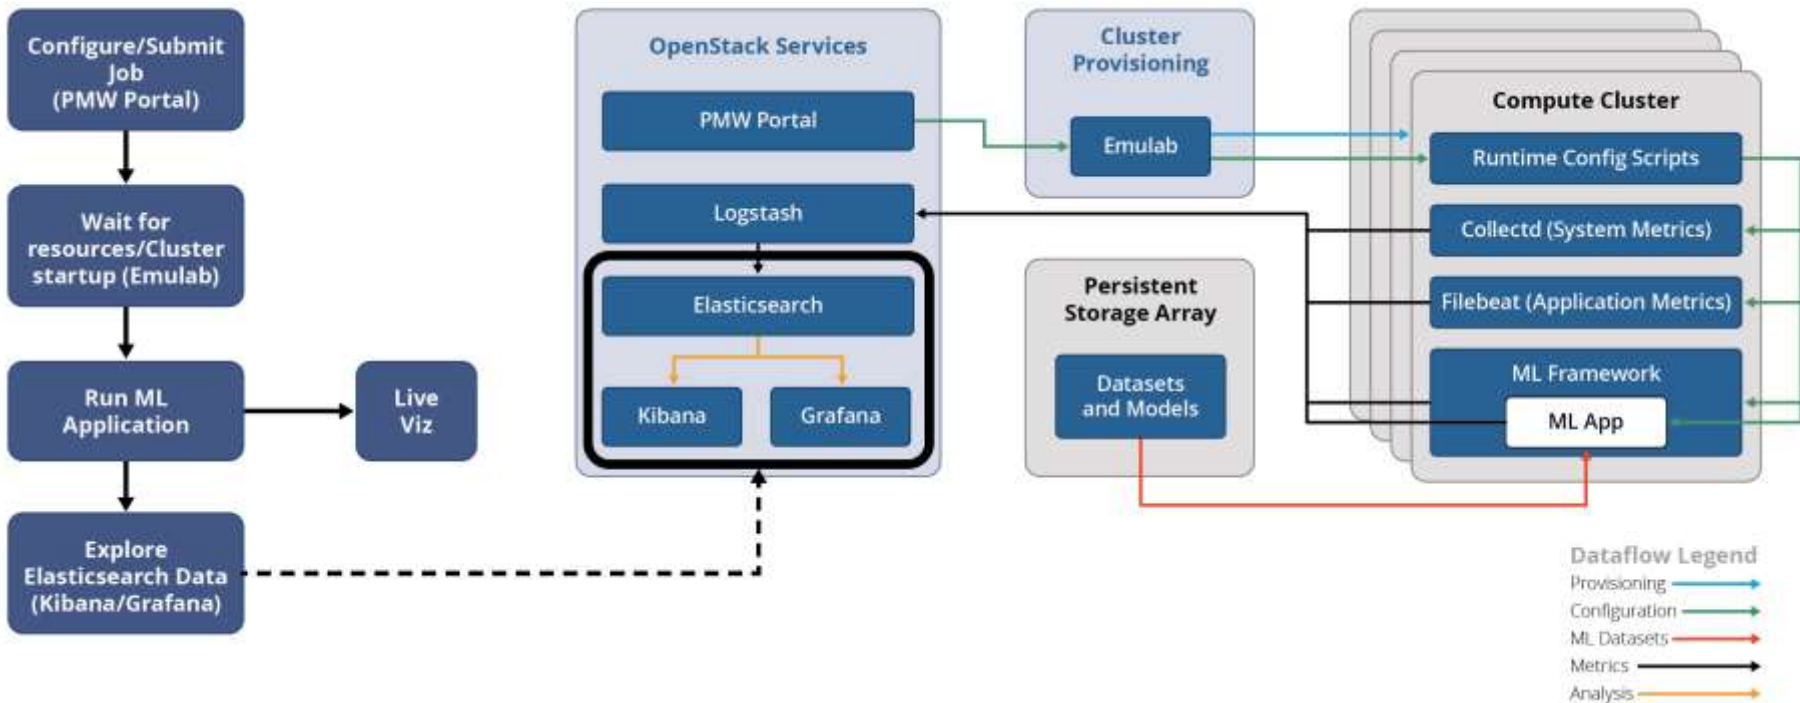
PMW Workflow










PMW Workflow



PMW Workflow



Dashboard (Live Display and Historic)




In-Depth Analysis

framework	algorithm	data_set	nodes_count *	load_training_phase_time	fit_phase_time	load_testing_phase_time	validate_phase_time
spark	linear_regression	5000_by_500000	3	68.713	386.479	18.952	321.56
spark	linear_regression	5000_by_500000	5	35.767	218.848	16.905	199.124
spark	linear_regression	5000_by_500000	9	26.86	161.769	14.948	164.846
spark	linear_regression	5000_by_500000	17	25.815	149.232	14.173	147.584

- Supports arbitrary queries
- Tabular data format
- Database format allows for queries across experiments
- Example: producing scaling plots

Strong Scaling, Learning (Fit) Phase



Approaching Complete Reproducibility

Configuration tracking

- Complete OS Image(s) used
 - Distribution version
 - Installed package versions
- ML Platform configuration/tuning parameters (Spark, Tensorflow, etc.)
- The ML application code command line parameters
- Dataset(s) used

- Caveats (future work)
 - Not tracking hardware firmware levels.
 - Application code not yet integrated with revision control system
 - Datasets assumed to be tracked independently