

Technical Advisory Council Meeting

April 23, 2020

 THE **LINUX** FOUNDATION

 **LF AI**

Antitrust Policy Notice

- › Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws.
- › Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at <http://www.linuxfoundation.org/antitrust-policy>. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrave of the firm of Gesmer Undergone LLP, which provides legal counsel to the Linux Foundation.

Recording of Calls

Reminder that TAC calls are recorded and available for viewing on the [TAC Wiki](#)

Reminder: LF AI Useful Links

Web site: lfai.foundation
Wiki: wiki.lfai.foundation
GitHub: github.com/lfai
Landscape: landscape.lfai.foundation or l.lfai.foundation
Mail Lists: <https://lists.lfai.foundation>

LF AI Logos: <https://github.com/lfai/artwork/tree/master/lfai>

LF AI Presentation Template:

https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk_-czASlz2GTBRZk2/view?usp=sharing

Events Page on LF AI Website: <https://lfai.foundation/events/>

Events Calendar on LF AI Wiki (subscribe available):

<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091544>

Event Wiki Pages: <https://wiki.lfai.foundation/display/DL/LF+AI+Foundation+Events>

Agenda

- › Roll Call
- › Approval of Minutes
- › TAC Chair Nominations Extension
- › TAC Sponsors Discussion
- › ODH Introduction
- › Open Data Initiative
- › LF AI Updates
- › Upcoming TAC Meetings
- › Open Discussion

TAC Voting Members

Member	TAC Contact	TAC Email
Amdocs	Ofer Hermoni*	oferher@gmail.com
AT&T	Jack Murray	jfm@research.att.com
Baidu	Daxiang Dong	dongdaxiang@baidu.com
Ericsson	Rani Yadav-Ranjan	rani.yadav-ranjan@ericsson.com
Huawei	Huang Zhipeng	huangzhipeng@huawei.com
Nokia	Pantelis Monogioudis	pantelis.monogioudis@nokia.com
Tech Mahindra	Nikunj Nirmal	nn006444@techmahindra.com
Tencent	Juniping Du	junipingdu@tencent.com
Zilliz	Jun Gu	jun.gu@zilliz.com
ZTE	Wei Meng	meng.wei2@zte.com.cn
Acumos AI Project	Anwar Aftab	anwar@research.att.com
Angel Project	Bruce Tao	brucetao@tencent.com
ONNX Project	Prasanth Pulavarthi	prasanth.pulavarthi@microsoft.com

Approval of Minutes

- › Draft minutes from the April 9th meeting of the TAC were previously distributed to the TAC members
- › **Proposed Resolution:**
 - › That the minutes of the April 9th meeting of the Technical Advisory Council of the LF AI Foundation are hereby approved

TAC Chairperson Nominations Extended

TAC Chairperson Nominations Extended





- › **Timeline:**
 - › ***Begin nomination period:*** April 9th, nominations include a short bio + statement
 - › ***End nomination period:*** April 23rd at 12 pm PT
 - › ***LF to send slate of respective nominees and voting instructions to TAC Voting Members*** - Hold election via [CIVS](#) if there are multiple nominees or via email, as applicable:
April 24th
 - › ***Voting period ends:*** May 1st at 12 pm PT
 - › ***Election winner announced:*** May 4th
- › **General TAC and Chairperson details can be viewed within the LF AI Charter under section 7 [here](#) - A summary of the role:**
 - › Represent the TAC as a voting member on the Governing Board (will attend monthly meeting)
 - › Lead TAC agenda and meetings with coordination among the TAC representatives and broader community
 - › Attend sync meetings with LF AI staff to discuss overall TAC activities and planning
 - › General representation of the TAC and the LF AI technical community





TAC Sponsors Discussion

Projects

Graduation

 Acumos LF Artificial Intelligence Foundation ★ 10	 Angel Angel-ML LF Artificial Intelligence Foundation ★ 5,743	 ONNX LF Artificial Intelligence Foundation ★ 8,078	
---	--	--	--

 ADLIK Adlik LF Artificial Intelligence Foundation ★ 88	 EDL Elastic Deep Learning LF Artificial Intelligence Foundation ★ 63	 ForestFlow ForestFlow LF Artificial Intelligence Foundation ★ 7	 HOROVOD Horovod LF Artificial Intelligence Foundation ★ 8,831
---	---	--	--

 MARQUEZ Marquez LF Artificial Intelligence Foundation ★ 239	 Milvus Milvus LF Artificial Intelligence Foundation ★ 2,902	 NNStreamer NNStreamer LF Artificial Intelligence Foundation ★ 185	 Pyro Pyro LF Artificial Intelligence Foundation ★ 6,087
--	--	--	--

 sparklyr sparklyr LF Artificial Intelligence Foundation ★ 724
--

Incubation

Companies hosting projects in LF AI

amazon

AT&T

Baidu 百度

DREAMWORKS

FACEBOOK

Microsoft

R Studio®

SAMSUNG

Tech Mahindra

Tencent 腾讯

Uber

wework

ZILLIZ
Reinvent Data Science

ZTE

TAC Sponsors Needed for Technical Projects

- › Per the [LF AI Projects and Lifecycle Document](#) all accepted projects receive an appointment of an existing TAC member that will act as a sponsor of the project and provide recommendations regarding governance best practices
 - › Activities include annual reviews, check ins for completion of [Technical Project Getting Started Checklist](#), and general mentorship
- › To date, [all 12 of the current projects](#) have defaulted to LF AI Executive Director, Ibrahim Haddad, as a sponsor and LF AI staffing for full support
- › We need TAC sponsor volunteers in order to continue support for these projects and for any new moving forward
- › **Please contact Ibrahim/Jacqueline to discuss supporting a project**

What does it involve?

- › Our goal as a Foundation is to support hosted projects, enable them to grow their base of users and contributors, and help them be a core project in the open source AI ecosystem
- › TAC Sponsor is a delegate from the TAC to a specific project who will guide the project through its journey in LF AI, be the mentor, help advise the project and be the advocate for it on the TAC, etc.
- › Collaborate with LF AI staff on big ticket items (need a website, need to host a virtual event, etc.)
- › Time allocation: 2-3 hours per month depending on the project needs

ODH Introduction

Juana Nakfour, jnakfour@redhat.com



OPEN DATA HUB

AI Platform powered by Open Source

Juana Nakfour
CTO AI Center of Excellence
Open Data Hub

AGENDA

How Open Data Hub Started

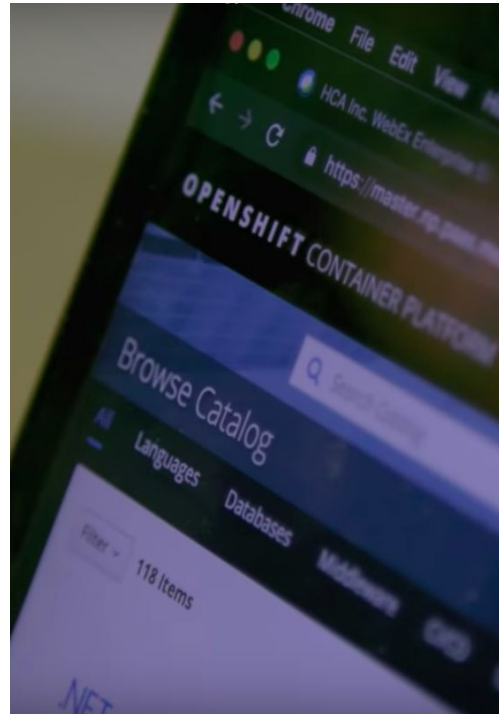
What is Open Data Hub

Where is Open Data Hub Used


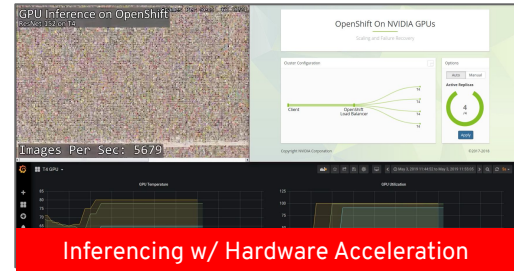
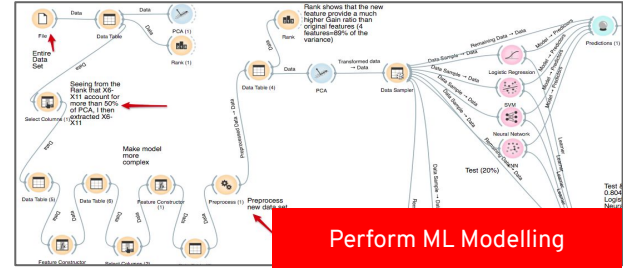
How Open Data Hub Started

BUILDING A PLATFORM FOR DATA SCIENCE

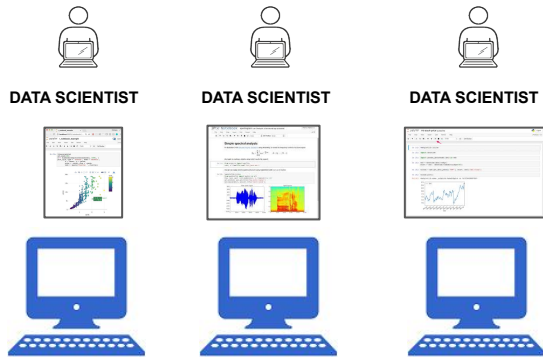
As a **Data Scientist**, I want a “self-service cloud like” experience for my Machine Learning projects, where I can access a rich set of modelling frameworks, data, and computational resources, **share** and **collaborate** with colleagues, and deliver my work into production with speed, agility and **repeatability** to drive business value!



Self Service Portal to Select ML Frameworks, and Data Sources



CHALLENGES WITH STATUS QUO

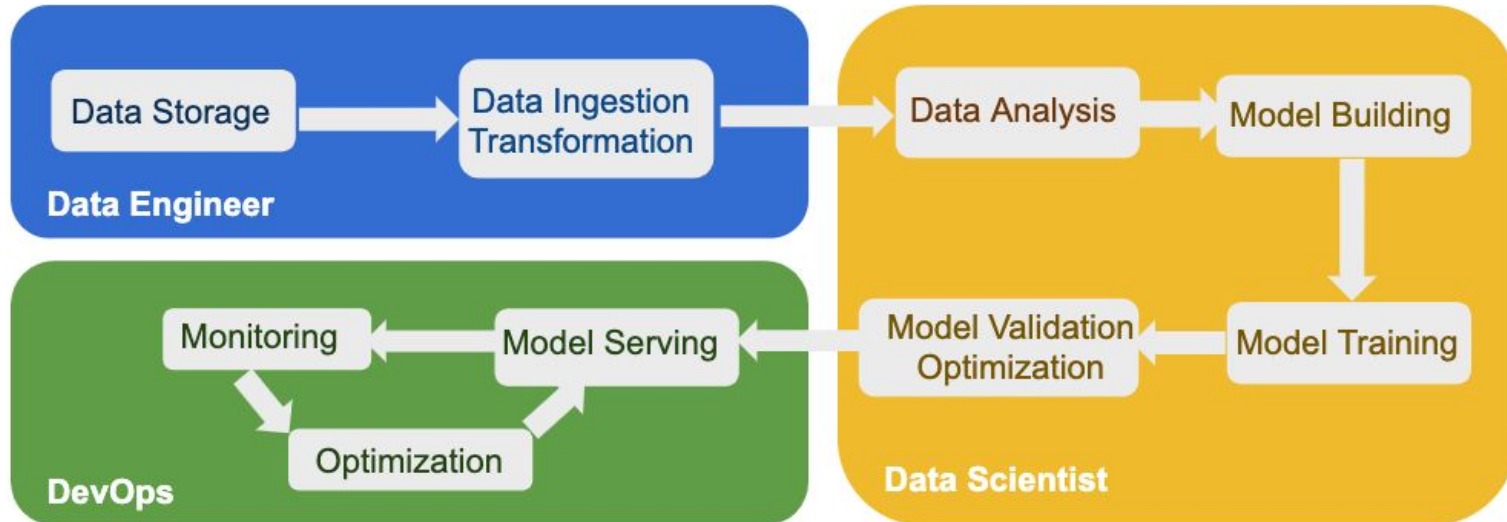


- **Team(s)** of Data Scientists and Developers
- **Sharing and collaboration**, if any, is **difficult**, manual, error prone and takes time
- Access to **limited non-shared resources** means modeling takes a long time or can't achieve desired accuracy
- Delivering models into **production** is a **challenge**

What is Open Data Hub

The Open Data Hub Community Project

- **Collaborate** on a Data & AI platform for the **Hybrid Cloud** - opendatahub.io/
- **Open Source** community **Meta-Project** to integrate open source projects into an **end-to-end AI/ML platform** on OpenShift
- Install today from the Openshift **OperatorHub Community Operators**

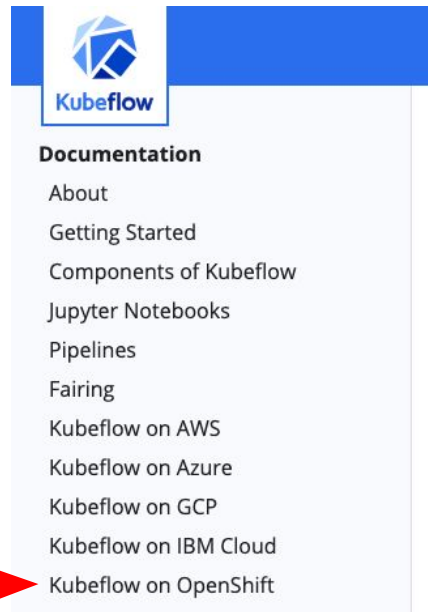


OPEN DATA HUB - Open Source Downstream/Upstream

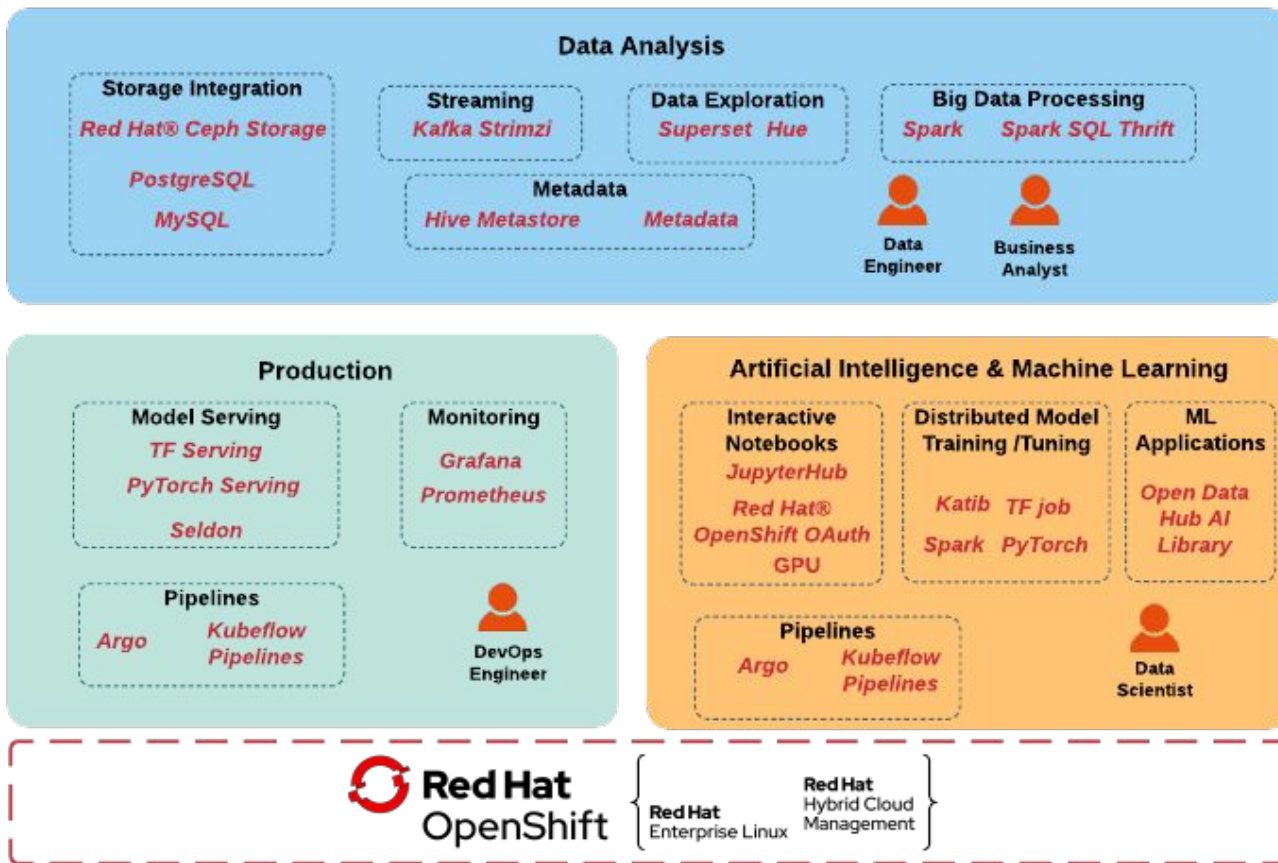


OPEN DATA HUB - Kubeflow


- Kubeflow - A new open source project dedicated to making deployments of machine learning (ML) workflows on Kubernetes **simple, portable and scalable** - kubeflow.org
- It is now integrated into Open Data Hub and runs on **OpenShift**.
- Kubeflow brings multiple new AI/ML capabilities and features
 - For **model training** Tensorflow and Pytorch
 - For **model serving** Seldon and KFServing
 - For **pipelines** Kubeflow Pipelines based on Argo



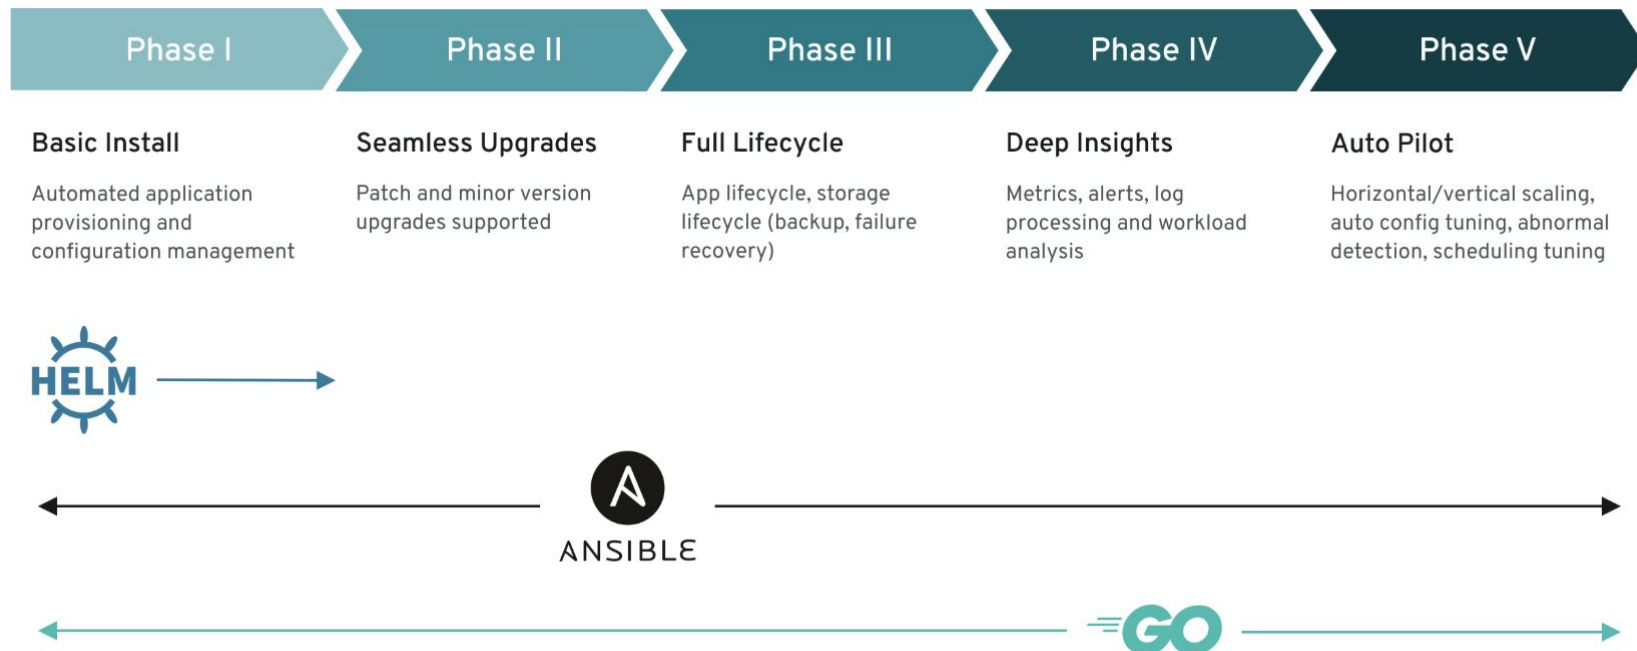
OPEN DATA HUB - 6.x



OPEN DATA HUB - Current Releases



OPEN DATA HUB - Operator



OPEN DATA HUB - Fraud Detection

- Provide **end-to-end AI/ML** platform
- Create an AI/ML **model** that can predict fraud transactions
- **Serve** model and collect model **metrics**
- Provide **monitoring** tools for model and services used by DevOps
- Provide **development** tools for Data Scientists
- Provide ETL tools used by **Data Engineers**
- Demo: https://youtu.be/lcQ2bhs_wkQ

Data Scientist



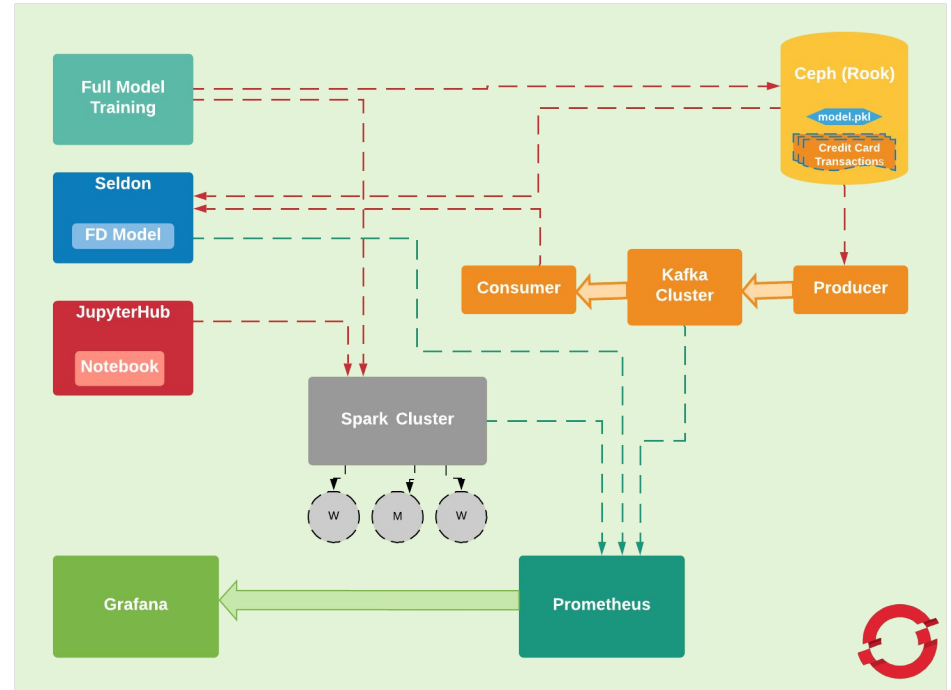
End User



DevOps



FRAUD DETECTION USING OPEN DATA HUB ON OPENSIFT



OPEN DATA HUB - Contact and Engagement

ODH site: opendatahub.io

ODH Gitlab: <https://gitlab.com/opendatahub>

ODH-Kubeflow Github: <https://github.com/opendatahub-io>

ODH Community Mailing lists: announcements@lists.opendatahub.io,
contributors@lists.opendatahub.io

ODH Community Meetings:

<https://gitlab.com/opendatahub/opendatahub-community>

AI/ML playlist on Openshift Commons YouTube channel



Where is Open Data Hub Used

OPEN DATA HUB - Red Hat Internal Customers

These are some significant users of the Internal Open Data Hub



Application Logs

Applications in the product release pipeline store their **runtime logs** in our system. These groups are also engaged for **anomaly detection**



Cluster Metrics

Operational metrics from OpenShift clusters. AIOps is engaged here.



Customer Support Data

Storage of customer data like **SOSReports, customer feedback**, etc.

OPEN DATA HUB - Adopters

ADVERTISEMENT

Governor Patrick Announces Funding to Launch Massachusetts Open Cloud Project

Mon, 04/28/2014 - 12:07pm
by Mass Open Cloud Project

Get the latest news in High Performance Computing, Informatics, Data Analysis So more - Sign up now!



Led by Boston University, the MOC is a collaborative effort among BU, Harvard, UMass Amherst, MIT, and Northeastern University, as well as the Massachusetts Green High-Performance Computing Center (MGHPCC) and Oak Ridge National Laboratory (ORNL).

It is supported by a broad alliance of industry partners, including Red Hat.



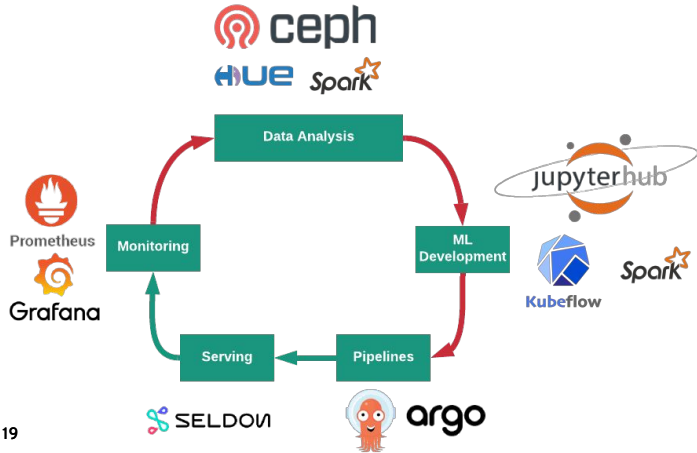
OPEN DATA HUB

AI Platform powered by Open Source

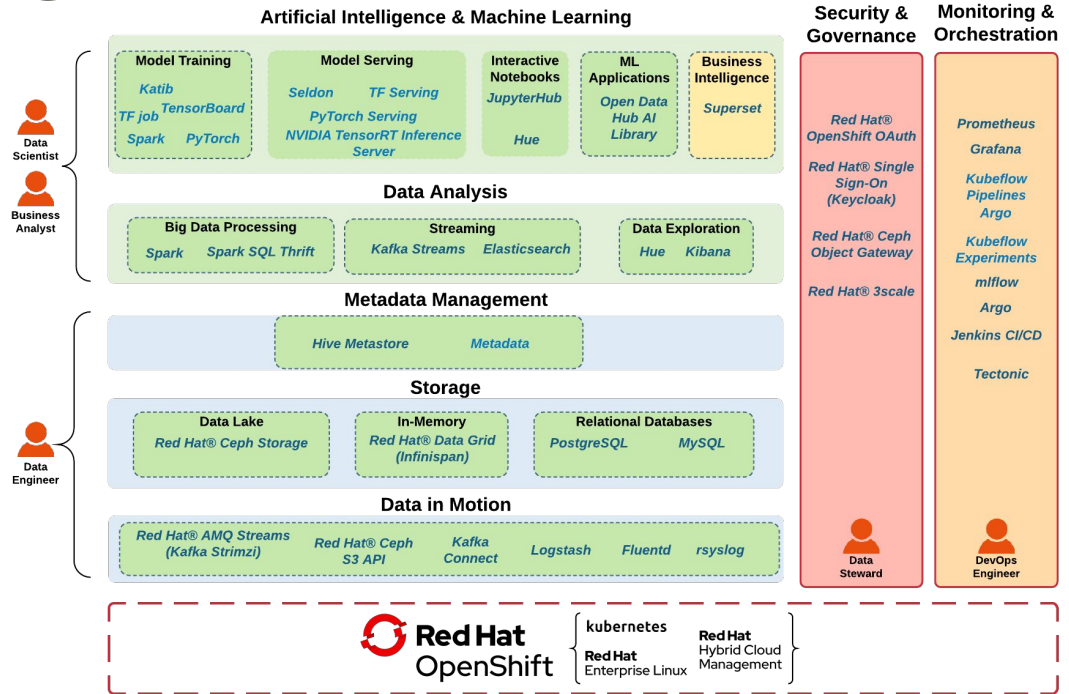
opendatahub.io

OPEN DATA HUB - Gaps and LF AI Collaboration

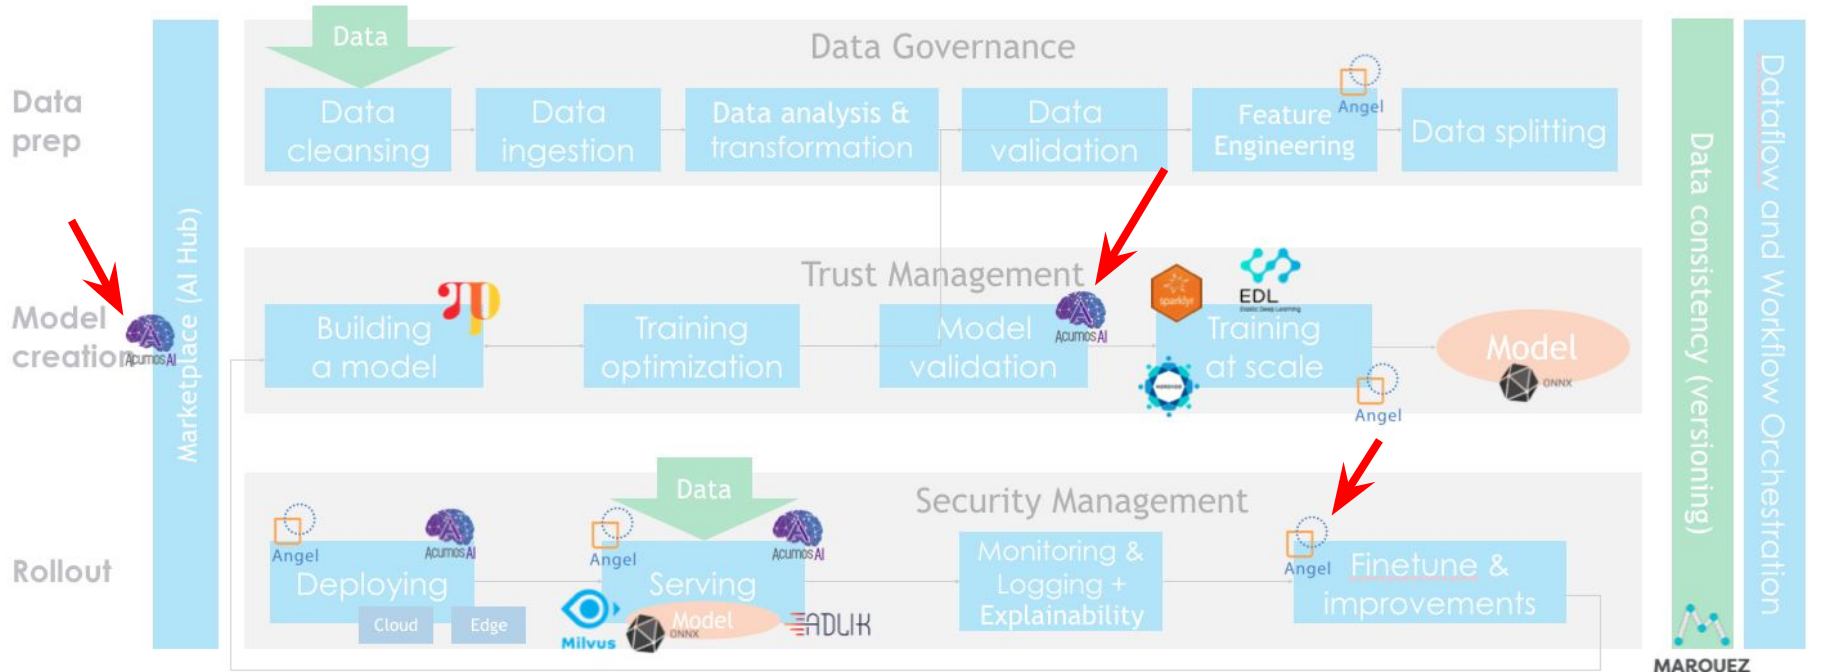
- Data and Security Governance
- Multi Hybrid Cloud Deployment
- Business Intelligence
- Seamless Development to Production Loop



Reference Architecture for AI on OpenShift



OPEN DATA HUB - Gaps and LF AI Collaboration



Open Data Initiative

Millo Avissar

David Wajnryt



Open Data Initiative

Privacy protected sharing of data
for AI in the context of covid-19

Millo Avissar

David Wajnryt

LFAI TAC Meeting 23/4/2020



Challenges



- ❑ Data stored in fragmented silos
- ❑ Public Open-data is sparse and anonymized
- ❑ Connecting and Sharing data requires a lot of effort
 - ❑ Privacy
 - ❑ Legal constraints
 - ❑ Recognition of business value of data
 - ❑ Unable to connect data at individual level



Covid-19, but not only...


- Data providers overwhelmed
- HR at home willing to help – enhanced but not unique to Covid-19
- More data allows for better models and study of rare cases



Goals

- Allow the creation of ML models based on the whole corpus of existing data instead of the data available at a specific organization or database in an environment where the data is held at different organizations that are not allowed to share it.
- Allow access to data for people who would like to train models and test them and create a platform that allows pre-trained models to run queries from around the world

...While protecting privacy




Use cases - Examples


- Data held at multiple organizations
 - Eg: various hospitals with a small number of cases each one – Covid, Parkinson's surgery
- Data from different countries
 - Eg: forecast outbreak based on data form all countries
- Link between data domains: medical, governmental, social, economic, etc
 - Eg: forecast children at risk using medical data related to outbreak cases in family members



Solution principles

- “Eyes-off training”
 - Assumptions
 - “Honest but curious” data providers
 - Data scientist can be malicious
 - Main Components
 - Private and secure cloud storage and computation
 - Federated Learning on secure cloud
 - Marketplace + federated modeling (Acumos?)
 - Privacy protection mechanism
- 

Proposed architecture






So far... Consultations round

- Domain experts and government agencies to validate needs
- Privacy and encryption experts to validate solution
- ML experts to understand needs and value proposition for data scientists
- Cloud and hardware providers to understand building blocks and assess feasibility – Both business and technical
- LFAI + Acumos



Next Stages

- A better understanding of what exists and what still needs to be developed
 - Roadmap
 - Project management and Backing
 - Business development
- 



Potential partners



- Cloud providers (AWS, Azure, Alibaba, IBM, etc)
- AI hardware providers (Intel, nVidia, etc)
- IT providers (SAP, IBM, etc)
- Data giants (AT&T, FB, Google, Apple, etc)
- Academia (John Hopkins , Stanford, etc)
- Open source evangelists

LF AI Updates

Project Updates

Releases Coming Up

- › ONNX 1.7
 - › <https://github.com/onnx/onnx/releases>
- › Sparklyr 1.2
 - › <https://github.com/sparklyr/sparklyr/wiki/Releases>
- › Acumos Demeter
 - › <https://wiki.acumos.org/display/REL>
- › Angel 3.1
 - › <https://github.com/Angel-ML/angel/blob/master/RELEASE.md>

Project releases will be announced via a blog post and promoted on LF AI and LF social media accounts

Note: If you are an LF AI hosted project and would like LF AI to promote your release, reach out to pr@lfai.foundation to coordinate in advance (min 2 wks) of your expected release date. Please reach out for questions and more details.

Outreach Committee

Recent Announcements

1. **NNStreamer Joins as New Incubation Project:**
<https://lfai.foundation/blog/2020/04/16/nstreamer-joins-ai-as-new-incubation-project/>
2. **Welcome New LF AI Members Q1 2020:**
<https://docs.google.com/document/d/1P540-yNQ7VFyO3QfwlZpr-gl2UmtxONUQB5Rzki7qY/edit>
3. **Enabling the Open Source AI Native Ecosystem with an AI Specific Computing Framework:**
<https://lfai.foundation/blog/2020/04/06/enabling-the-open-source-ai-native-ecosystem-with-an-ai-specific-computing-framework/>
4. **Milvus Joins as New Incubation Project:**
<https://lfai.foundation/blog/2020/04/02/milvus-joins-ai-as-new-incubation-project/>
5. **Happy 2nd Birthday LF AI:**
<https://lfai.foundation/blog/2020/03/18/happy-2nd-birthday-ai/>
6. **LF AI Partnering with ITU AI/ML in 5G Challenge:**
<https://lfai.foundation/blog/2020/03/25/ai-partnering-with-itu-ai-ml-in-5g-challenge/>
7. **LF AI Day ONNX Community Virtual Meetup:**
<https://lfai.foundation/blog/2020/03/16/ai-day-onnx-community-virtual-meetup-silicon-valley-2020/>
8. **Pyro 1.3 Now Available:**
<https://lfai.foundation/blog/2020/03/13/pyro-1-3-now-available/>

Upcoming Announcements

April:

Sparklyr 1.2 Release
Milvus - 0.8.0 Release
ForestFlow New Project
ONNX 1.7 Release
Marquez New Project
LF AI Day ONNX Virtual Meetup - Event Summary
Angel + Acumos Cross Project Collaboration

May:

Pyro M5 Starter Kit - Kaggle Competition
Angel 3.1 Release + Roadmap
Acumos Demeter Release
OSS NA - LF AI Booth, AI Track, Mini Summit Promotion

Request for PR/Comms Participation

- › Please follow LF AI via [Twitter](#) and/or [LinkedIn](#) and please help amplify LF AI news via your social networks - Please re-tweet and share
 - › Please keep an eye out for news updates via outreach-committee mail list to help spread the word about announcements
- › Open call to publish project updates and initiative/committee updates on [LF AI Blog](#)
 - › Please reach out to Jacqueline/ Ibrahim to follow up on getting featured with your efforts

LF AI Events

Event Leads Needed - Please Encourage Member Participation Across Your Organizations

Events Page on LF AI Website:

<https://lfai.foundation/events/>

Events Calendar on LF AI Wiki (subscribe available):

<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091544>

2020 Events Wiki:

<https://wiki.lfai.foundation/display/DL/2020+Events>

The [Outreach Committee](#) and LF AI Community need to identify Event Leads to drive these events as the Foundation is not resources/budgeted to manage them all. Please encourage participation within your organizations. LF AI provides [event planning templates/resources](#) and is available to help with guidance and questions along the way. Please email Jessica Kim (Outreach Committee Chair) at jessica.kim@huawei.com or contact LF AI staff at info@lfai.foundation to discuss participation.

Projects and Members- Please consider holding virtual events

Example: ONNX <https://lfai.foundation/event/lf-ai-day-onnx-community-virtual-meetup/>

LF AI Day ONNX Community Virtual Meetup

April 9 @ 9:00 am - 11:00 pm Free

IBM, ONNX, and the LF AI Foundation are pleased to sponsor the upcoming LF AI Day – ONNX Community Virtual Meetup – Silicon Valley 2020, to be held via Zoom on April 9th from 9am-12pm PT.

ONNX, an [LF AI Foundation Graduated Project](#), is an open format to represent deep learning models. With ONNX, AI developers can more easily move models between state-of-the-art tools and choose the combination that is best for them.

The virtual meetup will cover ONNX Community updates, partner/end-user stories, and SIG/WG updates. Check out the full agenda [here](#). If you are using ONNX in your services and applications, building software or hardware that supports ONNX, or contributing to ONNX, you should attend! This is a great opportunity to connect with and hear from people working with ONNX across many companies.

Registration is now open and the event is free to attend. Capacity will be 500 attendees. For up to date information on this virtual meetup, please visit the [event website](#).

+ GOOGLE CALENDAR

+ ICAL EXPORT

Details

Date:

April 9

Time:

9:00 am - 11:00 pm

Cost:

Free

Website:

<https://events.linuxfoundation.org/lf-ai-day-onnx-community-virtual-meetup/>

Venue

Zoom

Organizer

LF AI Foundation

Email:

info@lfai.foundation

Website:

<https://lfai.foundation>

Call to Participate in Ongoing Efforts

Trusted AI

- › **Leadership:**
Animesh Singh (IBM), Souad Ouali (Orange), and Jeff Cao (Tencent)
- › **Goal:** Create policies, guidelines, tooling and use cases by industry
- › **Github:**
<https://github.com/lfai/trusted-ai>
- › **Wiki:**
<https://wiki.lfai.foundation/display/DL/Trusted+AI+Committee>
- › **To participate:**
<https://lists.lfai.foundation/g/trustedai-committee/>
- › **Next call:** Bi-weekly on Thursdays at 7am PT, subscribe to group calendar on wiki
<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091895>

ML Workflow & Interop

- › **Leadership:**
Huang “Howard” Zhipeng (Huawei)
- › **Goal:**
Define an ML Workflow and promote cross project integration
- › **Wiki:**
<https://wiki.lfai.foundation/display/DL/ML+Workflow+Committee>
- › **To participate:**
<https://lists.lfai.foundation/g/mlworkflow-committee>
- › **Next call:** Every 4 weeks on Thursdays at 7:00 am PT, subscribe to group calendar on wiki
<https://wiki.lfai.foundation/pages/viewpage.action?pageId=18481242>

Upcoming TAC Meetings

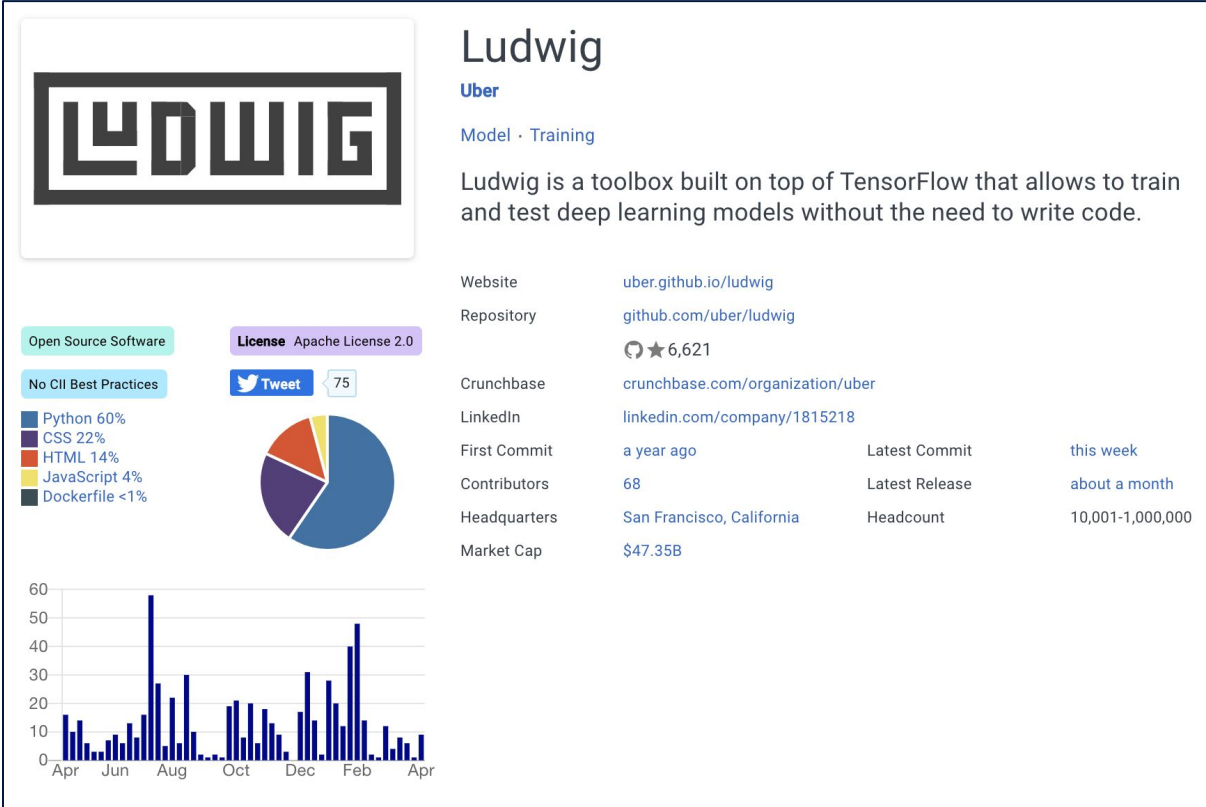
Upcoming TAC Meetings

May 7:

[Ludwig](#) Project Presentation + Proposal + TAC Vote

May 21:

To be announced on May 7th



Ludwig
Uber
Model · Training

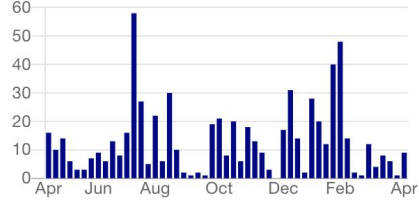
Ludwig is a toolbox built on top of TensorFlow that allows to train and test deep learning models without the need to write code.

Website: uber.github.io/ludwig
Repository: github.com/uber/ludwig
Stars: 6,621

Crunchbase: crunchbase.com/organization/uber
LinkedIn: linkedin.com/company/1815218
First Commit: a year ago
Latest Commit: this week
Contributors: 68
Latest Release: about a month
Headquarters: San Francisco, California
Headcount: 10,001-1,000,000
Market Cap: \$47.35B

Open Source Software
License: Apache License 2.0
No CII Best Practices
75 Tweets

Python 60%
CSS 22%
HTML 14%
JavaScript 4%
Dockerfile <1%



TAC Meeting Details

- › To subscribe to the TAC Group Calendar, visit the wiki: <https://wiki.lfai.foundation/x/XQB2>
- › Join from PC, Mac, Linux, iOS or Android: <https://zoom.us/j/430697670>
- › Or iPhone one-tap:
 - › US: +16465588656,,430697670# or +16699006833,,430697670#
- › Or Telephone:
 - › Dial(for higher quality, dial a number based on your current location):
 - › US: +1 646 558 8656 or +1 669 900 6833 or +1 855 880 1246 (Toll Free) or +1 877 369 0926 (Toll Free)
- › Meeting ID: 430 697 670
- › International numbers available: <https://zoom.us/u/achYtcw7uN>

Open Discussion

Legal Notices

- › The Linux Foundation, The Linux Foundation logos, and other marks that may be used herein are owned by The Linux Foundation or its affiliated entities, and are subject to The Linux Foundation's Trademark Usage Policy at <https://www.linuxfoundation.org/trademark-usage>, as may be modified from time to time.
- › Linux is a registered trademark of Linus Torvalds. Please see the Linux Mark Institute's trademark usage page at <https://lmi.linuxfoundation.org> for details regarding use of this trademark.
- › Some marks that may be used herein are owned by projects operating as separately incorporated entities managed by The Linux Foundation, and have their own trademarks, policies and usage guidelines.
- › TWITTER, TWEET, RETWEET and the Twitter logo are trademarks of Twitter, Inc. or its affiliates.
- › Facebook and the "f" logo are trademarks of Facebook or its affiliates.
- › LinkedIn, the LinkedIn logo, the IN logo and InMail are registered trademarks or trademarks of LinkedIn Corporation and its affiliates in the United States and/or other countries.
- › YouTube and the YouTube icon are trademarks of YouTube or its affiliates.
- › All other trademarks are the property of their respective owners. Use of such marks herein does not represent affiliation with or authorization, sponsorship or approval by such owners unless otherwise expressly specified.
- › The Linux Foundation is subject to other policies, including without limitation its Privacy Policy at <https://www.linuxfoundation.org/privacy> and its Antitrust Policy at <https://www.linuxfoundation.org/antitrust-policy>, each as may be modified from time to time. More information about The Linux Foundation's policies is available at <https://www.linuxfoundation.org>.
- › Please email legal@linuxfoundation.org with any questions about The Linux Foundation's policies or the notices set forth on this slide.