

# Meeting of the Technical Advisory Council (TAC)

July 15th, 2021

 **DLF** AI & DATA

# Antitrust Policy

- › 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:**

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

# Reminder: LF AI & Data Useful Links

- › Web site: [lfaidata.foundation](https://lfaidata.foundation)
- › Wiki: [wiki.lfaidata.foundation](https://wiki.lfaidata.foundation)
- › GitHub: [github.com/lfaidata](https://github.com/lfaidata)
- › Landscape: <https://landscape.lfaidata.foundation> or <https://l.lfaidata.foundation>
- › Mail Lists: <https://lists.lfaidata.foundation>
- › Slack: <https://slack.lfaidata.foundation>
- › Youtube: <https://www.youtube.com/channel/UCfasaeqXJBCAJMNO9HcHfbA>
- › LF AI Logos: <https://github.com/lfaidata/artwork/tree/master/lfaidata>
- › LF AI Presentation Template:  
[https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk\\_-czASlz2GTBRZk2/view?usp=sharing](https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk_-czASlz2GTBRZk2/view?usp=sharing)
- ›
- › Events Page on LF AI Website: <https://lfaidata.foundation/events/>
- › Events Calendar on LF AI Wiki (subscribe available):  
<https://wiki.lfaidata.foundation/pages/viewpage.action?pageId=12091544>
- › Event Wiki Pages: <https://wiki.lfaidata.foundation/display/DL/LF+AI+Data+Foundation+Events>

# Agenda

- › Roll Call (2 mins)
- › Approval of Minutes from June 17th (3 mins)
- › Invited Presentation - (35 minutes)
  - › Tony Project - Tony Project <https://github.com/linkedin/TonY> (Kequi Hu - [khu@linkedin.com](mailto:khu@linkedin.com));
- › Annual Project Review (15 minutes)
  - › EDL - <https://github.com/lfai/proposing-projects/blob/master/proposals/edl.adoc> ([zhouti@baidu.com](mailto:zhouti@baidu.com))
- › Annual Project Review
- › LF AI General Updates
- › Open Discussion

# TAC Voting Members

\* = still need backup specified on [wiki](#)

Board Member	Contact Person	Email
AT&T	Anwar Atfab*	<a href="mailto:anwar@research.att.com">anwar@research.att.com</a>
Baidu	Ti Zhou	<a href="mailto:zhouti@baidu.com">zhouti@baidu.com</a>
Ericsson	Rani Yadav-Ranjan*	<a href="mailto:rani.yadav-ranjan@ericsson.com">rani.yadav-ranjan@ericsson.com</a>
Huawei	Huang Zhipeng	<a href="mailto:huangzhipeng@huawei.com">huangzhipeng@huawei.com</a>
IBM	Susan Malaika	<a href="mailto:malaika@us.ibm.com">malaika@us.ibm.com</a>
Nokia	Jonne Soininen	<a href="mailto:jonne.soininen@nokia.com">jonne.soininen@nokia.com</a>
OPPO	Jimin Jia*	<a href="mailto:jjajimin@oppo.com">jjajimin@oppo.com</a>
SAS	Nancy Rausch	<a href="mailto:nancy.rausch@sas.com">nancy.rausch@sas.com</a>
Tech Mahindra	Amit Kumar	<a href="mailto:Kumar_Amit@techmahindra.com">Kumar_Amit@techmahindra.com</a>
Tencent	Bruce Tao	<a href="mailto:brucetao@tencent.com">brucetao@tencent.com</a>
Zilliz	Jun Gu	<a href="mailto:jun.gu@zilliz.com">jun.gu@zilliz.com</a>
ZTE	Wei Meng	<a href="mailto:meng.wei2@zte.com.cn">meng.wei2@zte.com.cn</a>
Graduate Project	Contact Person	Email
Acumos	Nat Subramanian	<a href="mailto:natarajan.subramanian@techmahindra.com">natarajan.subramanian@techmahindra.com</a>
Angel	Bruce Tao	<a href="mailto:brucetao@tencent.com">brucetao@tencent.com</a>
Egeria	Mandy Chessell	<a href="mailto:mandy_chessell@uk.ibm.com">mandy_chessell@uk.ibm.com</a>
Horovod	Travis Addair*	<a href="mailto:taddair@uber.com">taddair@uber.com</a>
Milvus	Xiaofan Luan	<a href="mailto:xiaofan.luan@zilliz.com">xiaofan.luan@zilliz.com</a>
ONNX	Jim Spohrer (Chair of TAC)	<a href="mailto:spohrer@us.ibm.com">spohrer@us.ibm.com</a>
Pyro	Fritz Obermeyer*	<a href="mailto:fritz.obermeyer@gmail.com">fritz.obermeyer@gmail.com</a>

# Approval of June 17th, 2021 Minutes

Draft minutes from the June 17<sup>th</sup> TAC call were previously distributed to the TAC members via the mailing list

## **Proposed Resolution:**

- › That the minutes of the June 17<sup>th</sup> meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

# Invited Presentation - Tony Project

LinkedIn - Tony Project

<https://github.com/linkedin/TonY>

(contact Keqiu Hu- [khu@linkedin.com](mailto:khu@linkedin.com));



# TonY @ LF AI

---

Keqiu Hu [khu@linkedin.com](mailto:khu@linkedin.com)

07/15/2021

# Why contribute TonY to Linux Foundation

## **Neutral holding ground**

- vendor-neutral, not for profit.

## **Open governance model**

- Transparent and open governance model
- Instill trust in contributors and adopters in the management of the project
- Neutral management of projects' assets by the foundation

## **Growing community**

- Increase visibility of project through LF ecosystem
- Increase contributors by converting new & existing users
- Opportunities to collaborate with other hosted projects



# Agenda

- 1 AI & Big Data @ LI
- 2 Training at Scale Challenges
- 3 How TonY Helps?
- 4 Next Steps

# LinkedIn's Vision

Create economic opportunity  
for every member of the global workforce



**750M**  
Members



**57M**  
Companies



**40M**  
Jobs

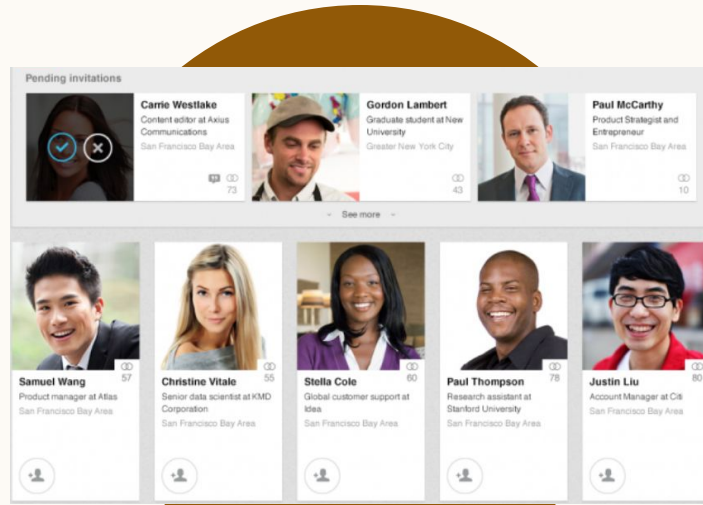


**38K**  
Skills

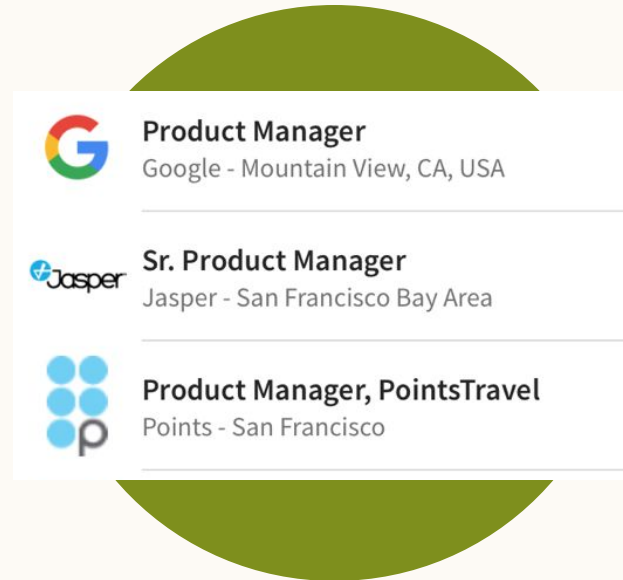


**120K**  
Schools

# AI @ LinkedIn



People You May Know



Job Recommendations



News Feed


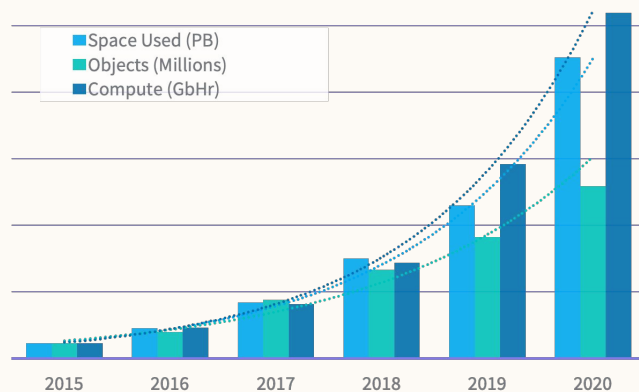
# Big Data Platform @ LinkedIn


## Our stack

- Hadoop as big data OS
- Spark + TonY/TF + Hive for compute engines
- Azkaban for workflow management

## Hadoop Cluster

- Single cluster 10K+ nodes, total 20K+ nodes.
- 2.2 PB compute memory
- 1 exabyte storage
- 5K+ users
- 3K concurrent apps.
- 2x YoY growth in capacity





Deep Learning

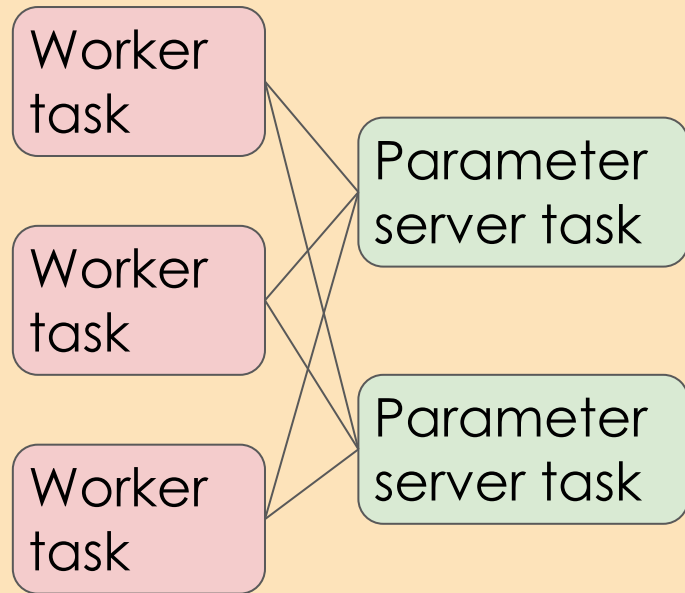
? Big Data Platform

TensorFlow / PyTorch

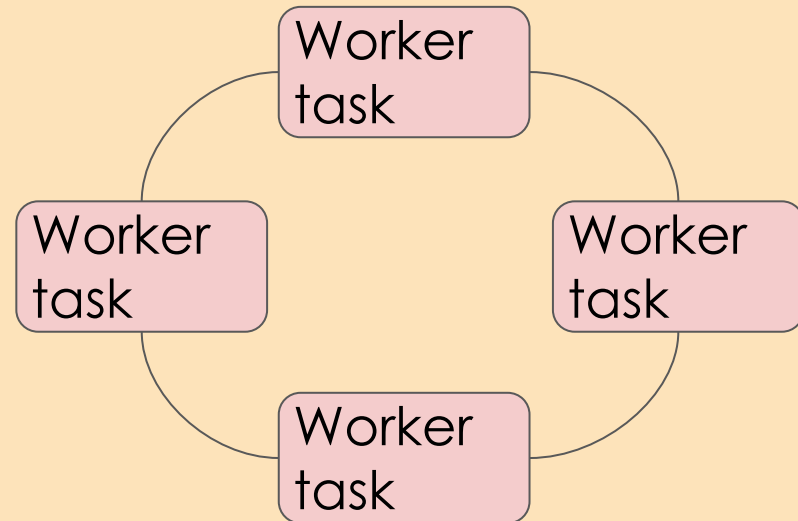
Hadoop / Kubernetes

# Distributed TensorFlow

## Worker/Parameter Server Model



## Ring All-Reduce Model







# Early days: how AI engineers did training

- Copy code and dependencies to each host
- Manually specify host and port of each process
- Customize arguments for each process

```
# On ps0.example.com:
$ python trainer.py \
  --ps_hosts=ps0.example.com:2222,ps1.example.com:2222 \
  --worker_hosts=worker0.example.com:2222,worker1.example.com:2222 \
  --job_name=ps --task_index=0
# On ps1.example.com:
$ python trainer.py \
  --ps_hosts=ps0.example.com:2222,ps1.example.com:2222 \
  --worker_hosts=worker0.example.com:2222,worker1.example.com:2222 \
  --job_name=ps --task_index=1
# On worker0.example.com:
$ python trainer.py \
  --ps_hosts=ps0.example.com:2222,ps1.example.com:2222 \
  --worker_hosts=worker0.example.com:2222,worker1.example.com:2222 \
  --job_name=worker --task_index=0
# On worker1.example.com:
$ python trainer.py \
  --ps_hosts=ps0.example.com:2222,ps1.example.com:2222 \
  --worker_hosts=worker0.example.com:2222,worker1.example.com:2222 \
  --job_name=worker --task_index=1
```

# Challenges of scaling up training


- Managing code and dependencies
- Orchestrating distributed training
- Resource contention (especially for GPUs)
- Managing an ML workflow (data preparation, training, deployment)
- Fault tolerance



Deep Learning

Big Data Platform

TensorFlow / PyTorch



Hadoop YARN

# Using TonY


## CLI

```
java -cp `hadoop
classpath`:tony-cli-0.3.7-all.jar \
com.linkedin.tony.cli.ClusterSubmitter \
--python_venv=venv.zip \
--src_dir=src \
--executes='venv/Python my_model.py' \
--conf_file=config.xml
```

## Conf File

```
<configuration>
  <property>
    <name>tony.worker.instances</name>
    <value>3</value>
  </property>
  <property>
    <name>tony.worker.gpus</name>
    <value>1</value>
  </property>
  <property>
    <name>tony.ps.instances</name>
    <value>1</value>
  </property>
</configuration>
```

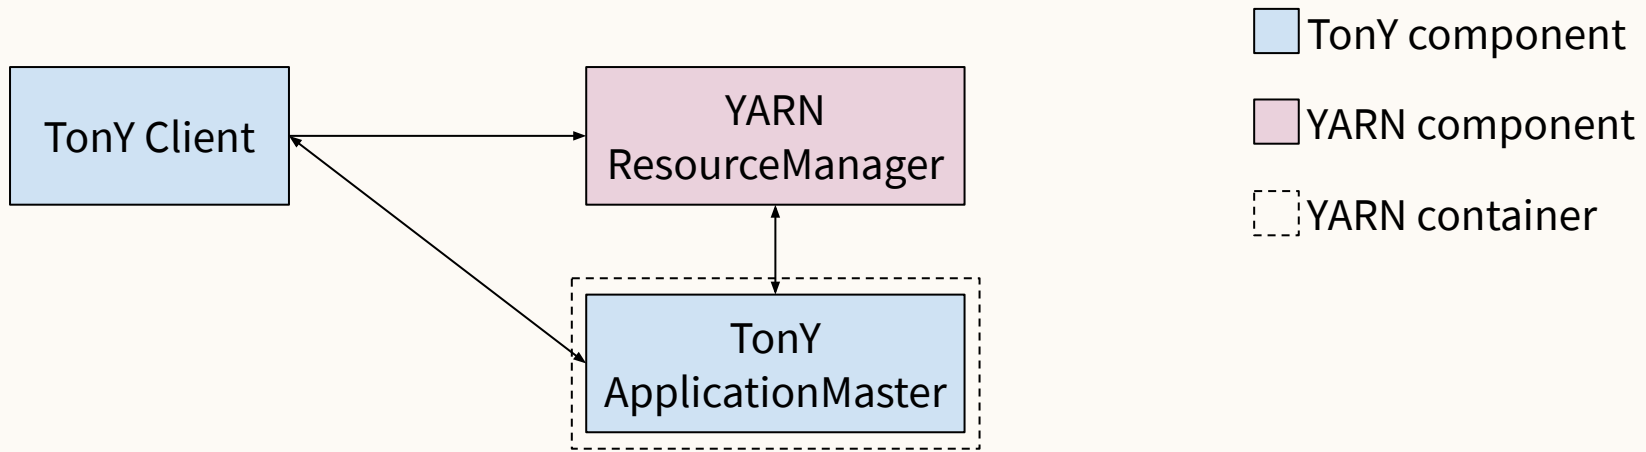
# TonY Under the Hood



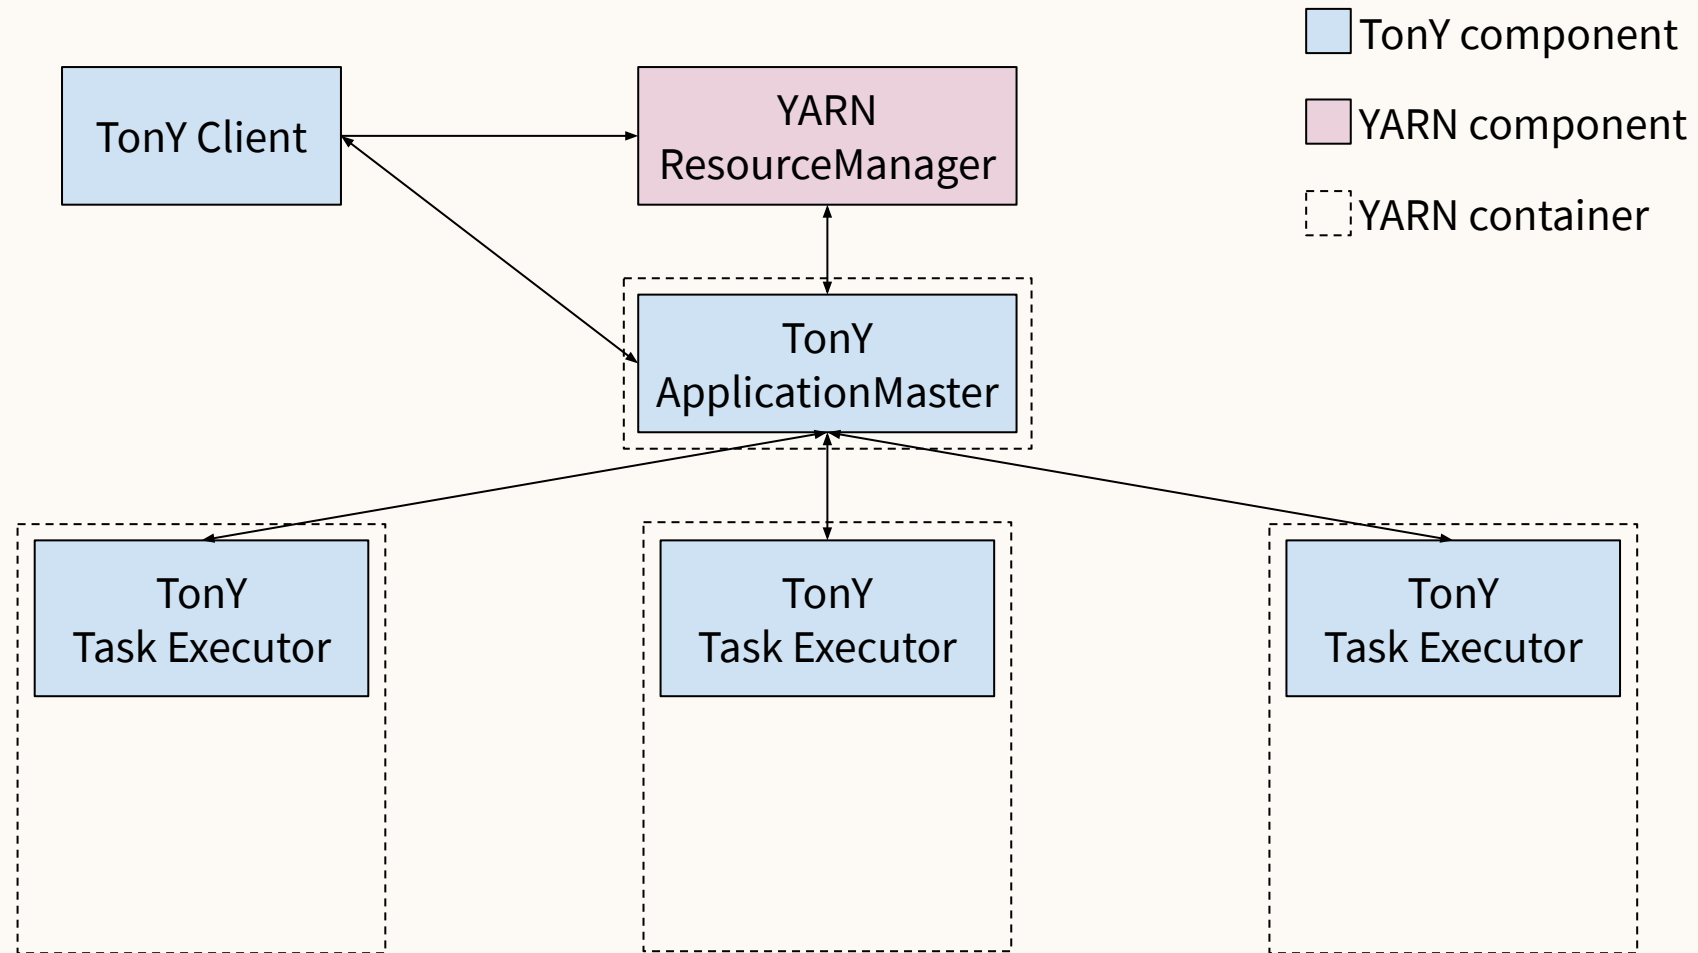
□ TonY component

□ YARN component

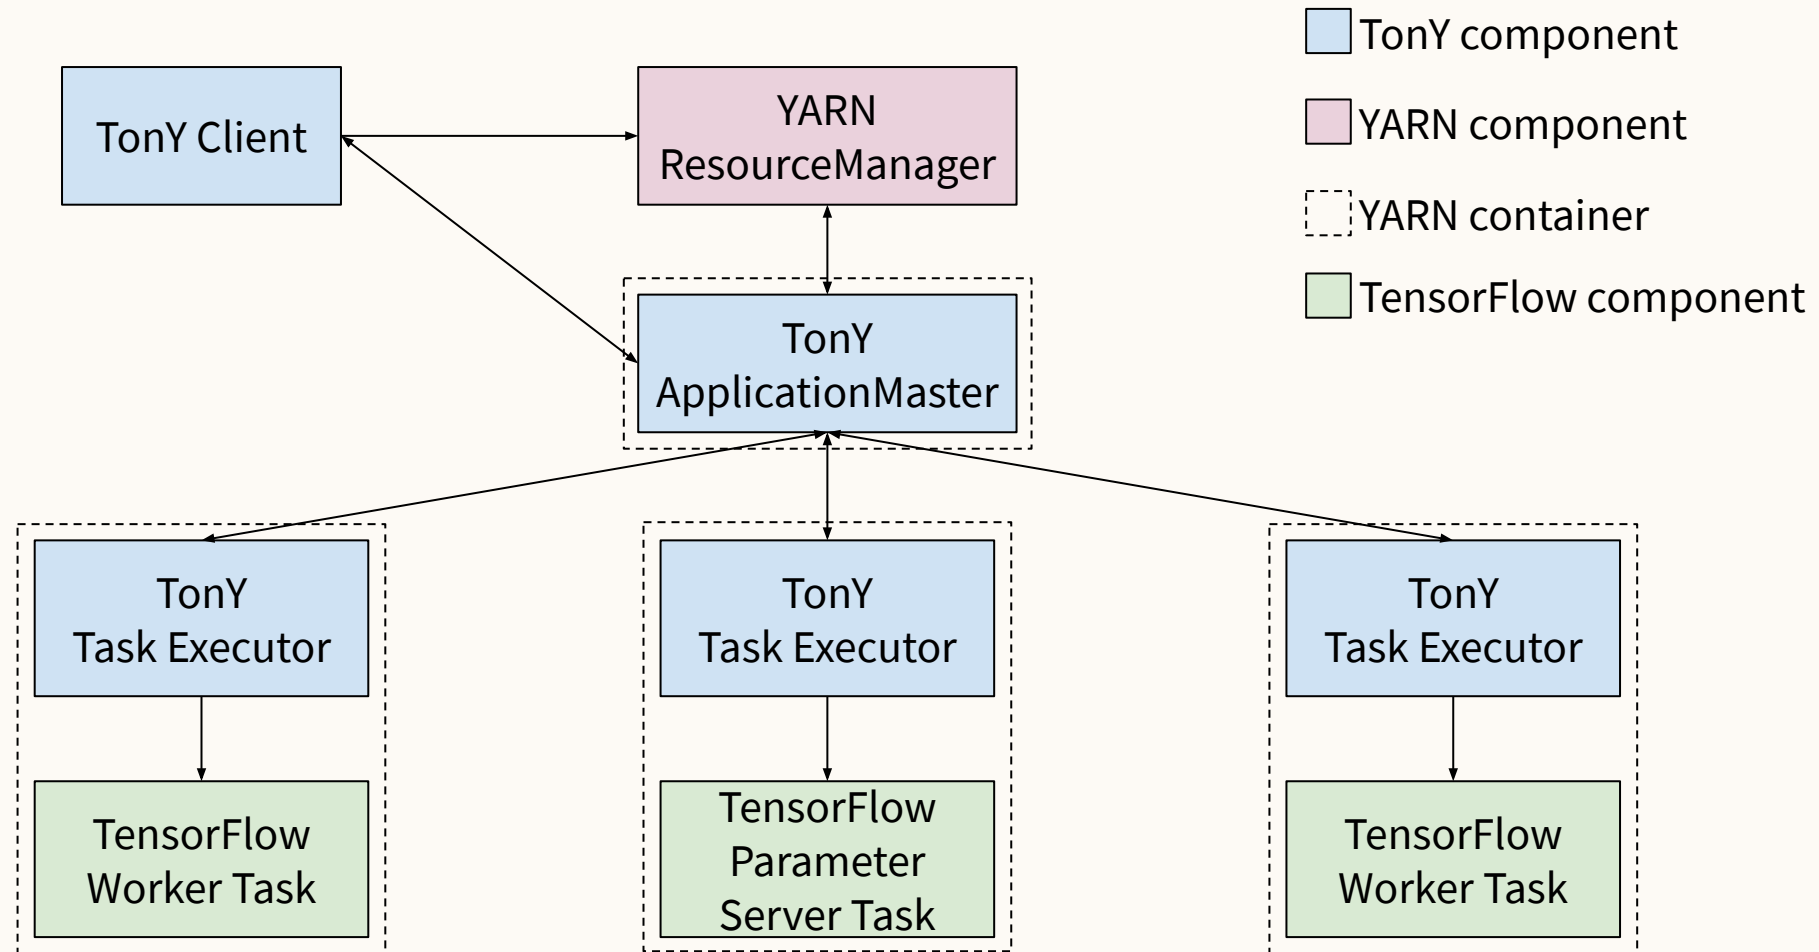
# TonY Under the Hood



# TonY Under the Hood

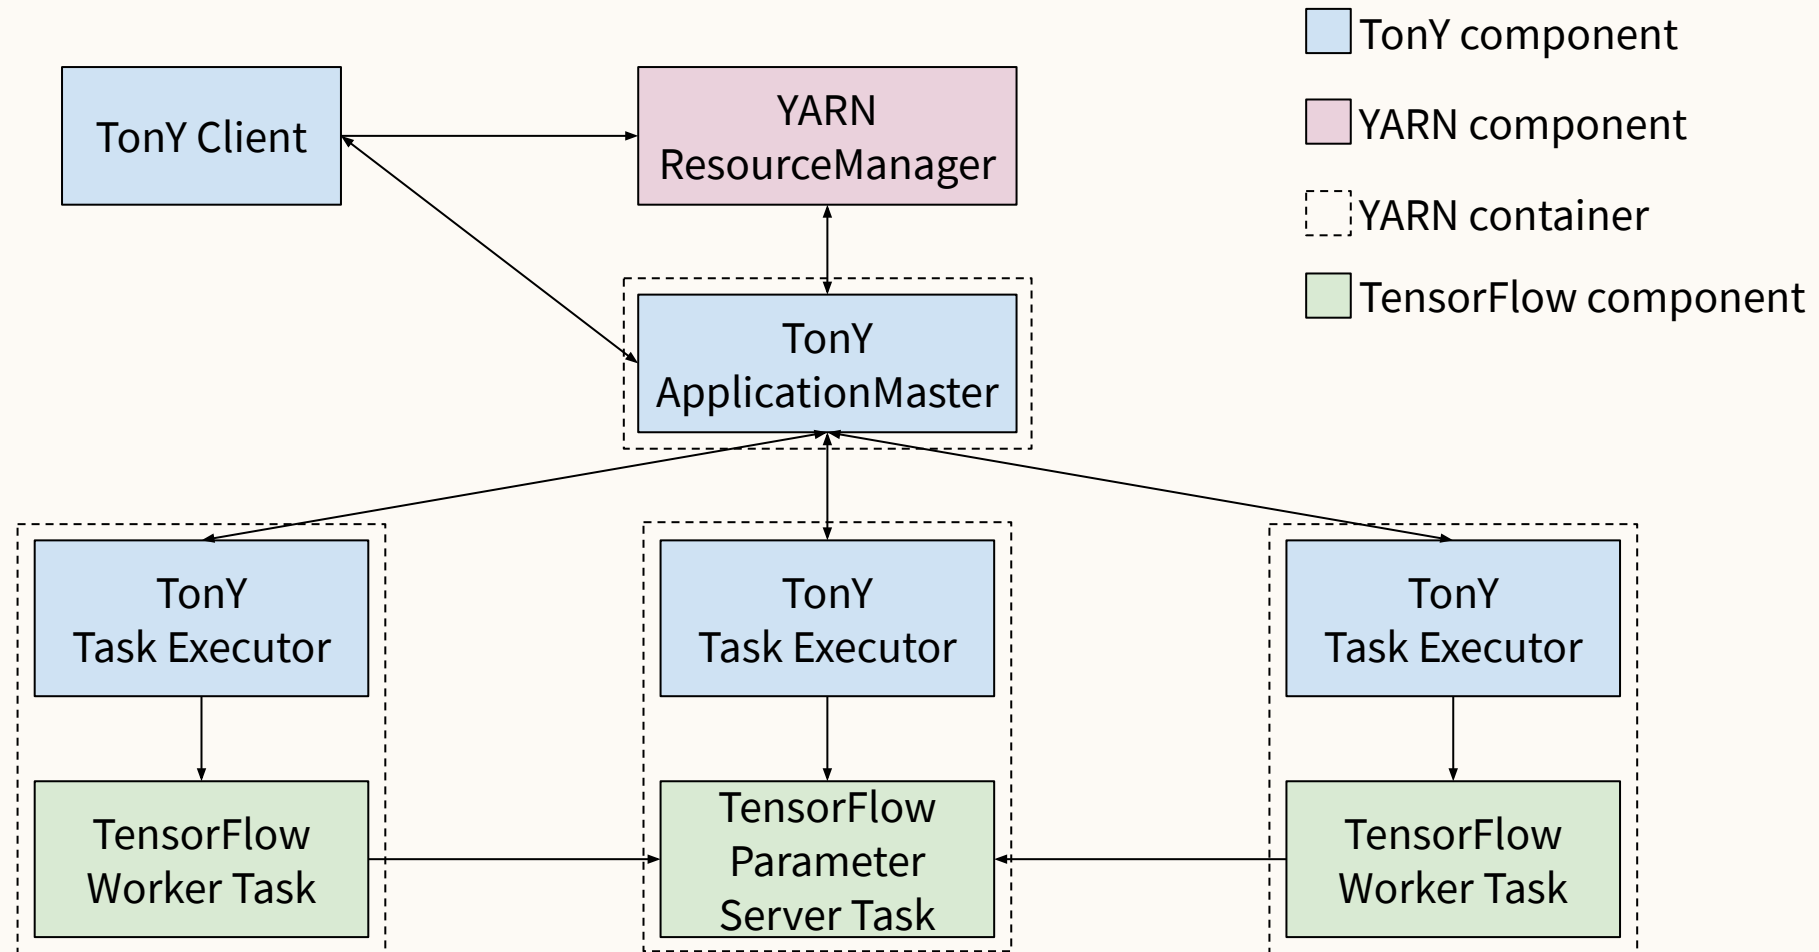


# TonY Under the Hood





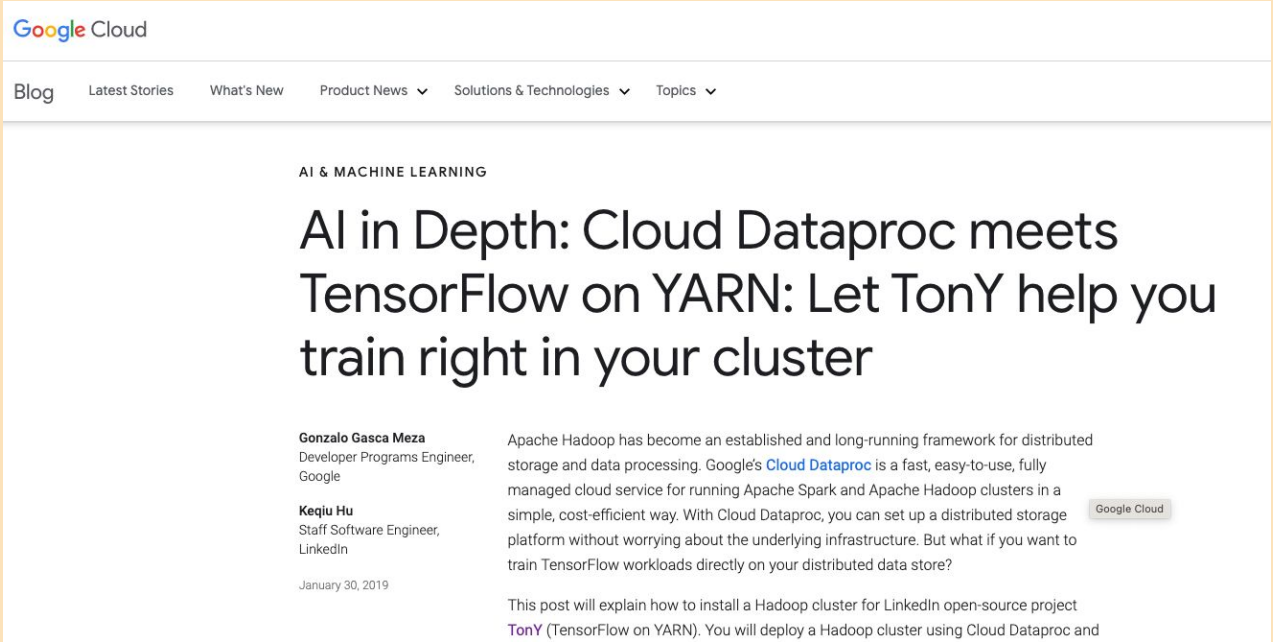
# TonY Under the Hood



# TonY in Community

- Open Sourced:  
<https://github.com/linkedin/TonY>
- 640 Stars, 150 Forks  


- 29 contributors
- Adopted at Google Cloud, iQiyi.com, NetEase, and other companies.
- [Horovod integration](#) contribution from iQiyi.



The screenshot shows a Google Cloud blog page. The header includes the Google Cloud logo and navigation links for 'Blog', 'Latest Stories', 'What's New', 'Product News', 'Solutions & Technologies', and 'Topics'. The main content area is titled 'AI & MACHINE LEARNING' and features the article 'AI in Depth: Cloud Dataproc meets TensorFlow on YARN: Let TonY help you train right in your cluster'. The author is identified as Gonzalo Gasca Meza, a Developer Programs Engineer at Google. The article text discusses Apache Hadoop's role in distributed storage and data processing, and how Google's Cloud Dataproc simplifies running Apache Spark and Hadoop clusters. It also mentions that the post will explain how to install a Hadoop cluster for the LinkedIn open-source project TonY (TensorFlow on YARN) using Cloud Dataproc.

# TonY in Community

- [TonY @ Strata NY '18](#)
- [TonY @ AI Conf London '19](#)
- [Scaling TensorFlow at LinkedIn @ TensorFlow World '19](#)
- [TonY @ OpML '19](#)











## Next Step

- Evolve TonY as part of generic AI training platform on Hadoop and Kubernetes.
- History server support.
- Cloud integration.

Thank you

# Summary of open-source solutions

Open-source solution	Pros	Cons
<p><a href="#">Kubeflow</a> / <a href="#">Kubernetes</a> (Google)</p>  	<ul style="list-style-type: none"> <li>• Large marketplace of libraries and plugins</li> <li>• Active community</li> </ul>	<ul style="list-style-type: none"> <li>• Does not run on Hadoop</li> <li>• May not scale to very large clusters</li> </ul>
<p><a href="#">TensorFlow on Spark</a> (Yahoo!)  <a href="#">Spark Deep Learning</a> (Databricks)</p> 	<ul style="list-style-type: none"> <li>• Integrates with Spark</li> </ul>	<ul style="list-style-type: none"> <li>• No GPU resource support until Spark 3.0 (<a href="#">SPARK-20327</a>)</li> <li>• No heterogeneous resource support</li> </ul>
<p><a href="#">TOY: TensorFlow on YARN</a> (Intel)  <a href="#">XLearning</a> (Qihoo)</p> 	<ul style="list-style-type: none"> <li>• YARN native, works out-of-the-box</li> </ul>	<ul style="list-style-type: none"> <li>• No GPU resource support</li> </ul>
<p><a href="#">Horovod</a> (Uber)</p> 	<ul style="list-style-type: none"> <li>• Supports synchronous distributed training</li> </ul>	<ul style="list-style-type: none"> <li>• MPI on YARN requires Docker</li> </ul>
<p><a href="#">YARN Native Service</a></p> 	<ul style="list-style-type: none"> <li>• YARN native</li> </ul>	<ul style="list-style-type: none"> <li>• Distributed TensorFlow requires YARN DNS Registry and ZooKeeper</li> </ul>

# TAC Vote on Project Incubation Proposal: TonY

## **Proposed Resolution:**

The TAC approves the TonY as an Incubation project of the LF AI & Data Foundation

## Next Steps

LF AI & Data staff will work with the TonY project team to onboard the project leading to the announcement of the project joining LF AI & Data

Explore potential integrations between the project and other LF AI & Data projects

Integrate the project with LF AI & Data operations



# Annual Project Review - EDL

<https://github.com/lfai/proposing-projects/blob/master/proposals/edl.adoc>  
(contact: Ti Zhou - zhouti@baidu.com)

# EDL

## Brief Description:

Computing resources on cloud such as Amazon AWS, Baidu Cloud have multi-tenancy. Deep learning model training and inference with elastic resources will be common on cloud. We propose Elastic Deep Learning (EDL) that makes training and inference of deep learning models on cloud easier and more efficient.

TSC Chairperson: Ti Zhou

## Contributed by:

Baidu contributed EDL to LF AI in August 2018, and now EDL is the incubating project.




## Key Links:

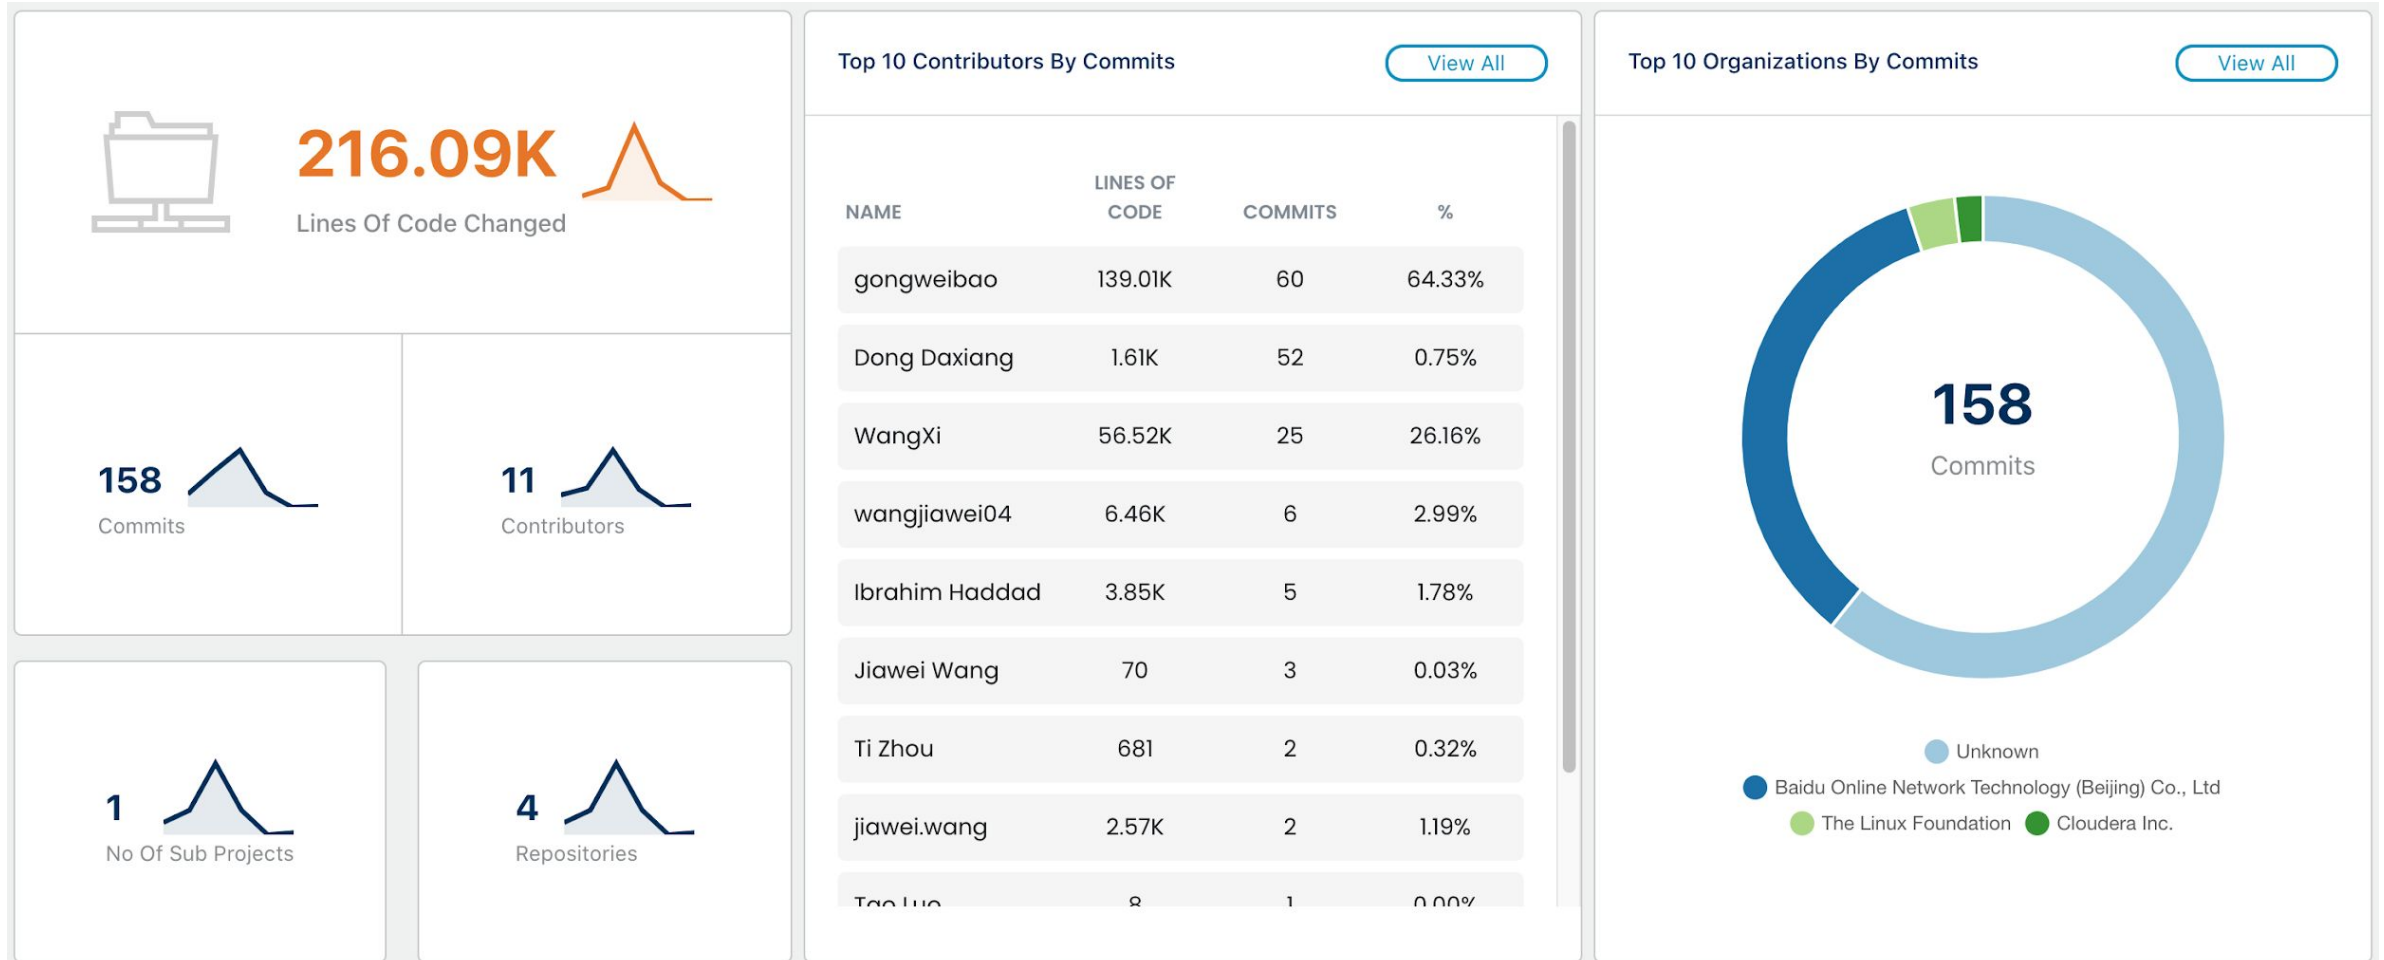
Github: <https://github.com/elasticdeeplearning/edl>

Homepage: <https://elasticdeeplearning.ai/>

Artwork: <https://artwork.lfaidata.foundation/projects/edl/>



# Contributions (over last two years)



# Key Achievements in the past year

## Two Releases

- 0.2.0 fault tolerance
- 0.3.0 knowledge distillation

## Enhanced Roadmap

- 0.4.0 live fault tolerance
- 0.5.0 auto scaling

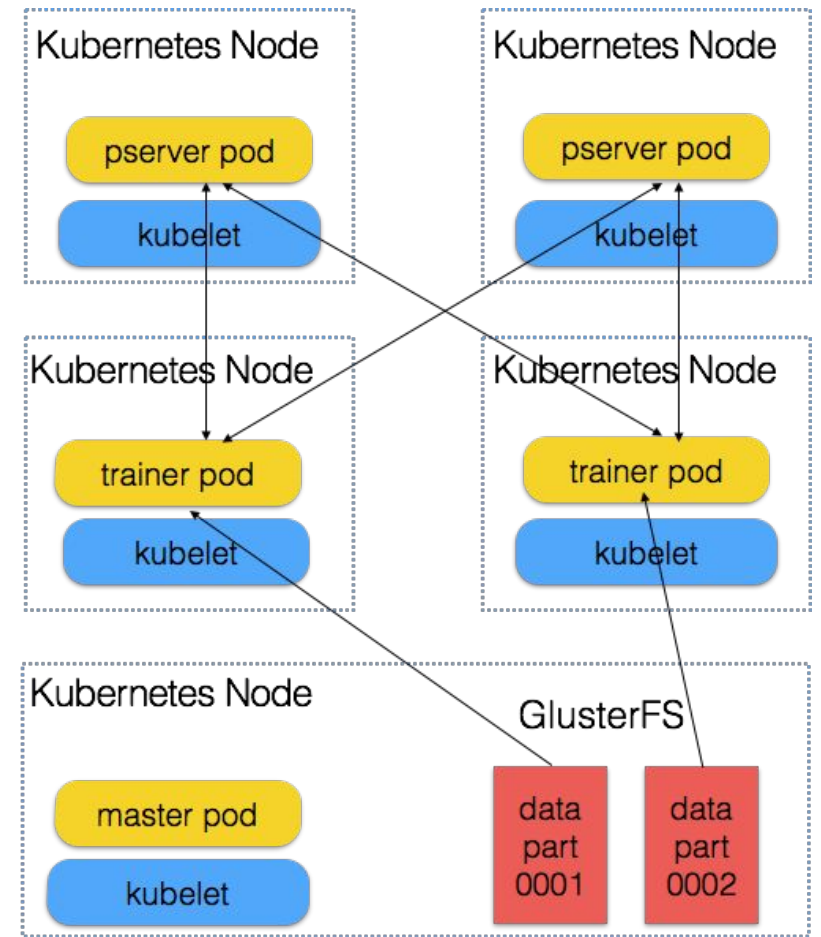
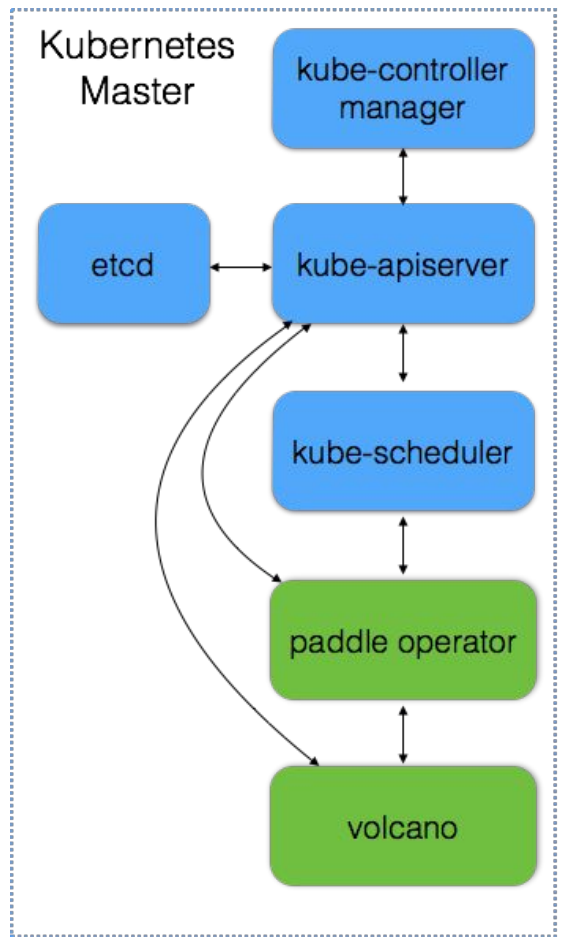
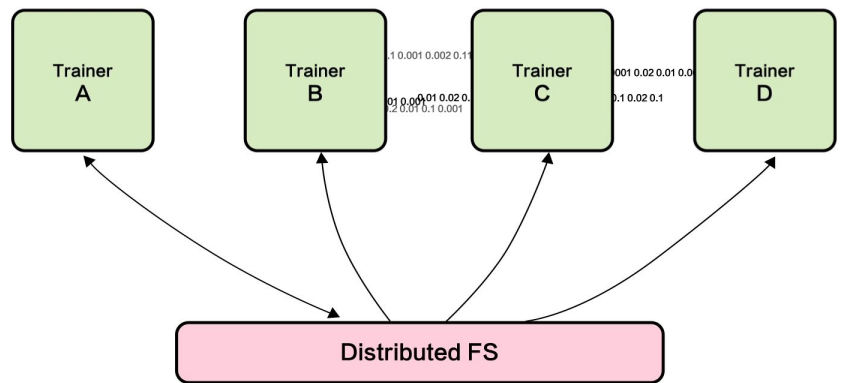
# Release 0.2.0 – Fault Tolerance

## Job recovery with EDL

- Snapshots based elastic deep learning
- Training Job can be killed any time
- Job recovery is based on last snapshots

## User Interface Design


- Non-intrusive API for PaddlePaddle users
- Minimum user awareness of elastic training



# Release 0.3.0 – Knowledge Distillation


## Save GPU memory resources

- Heterogeneous resources between teacher and student
- Flexible ratio between teacher and student



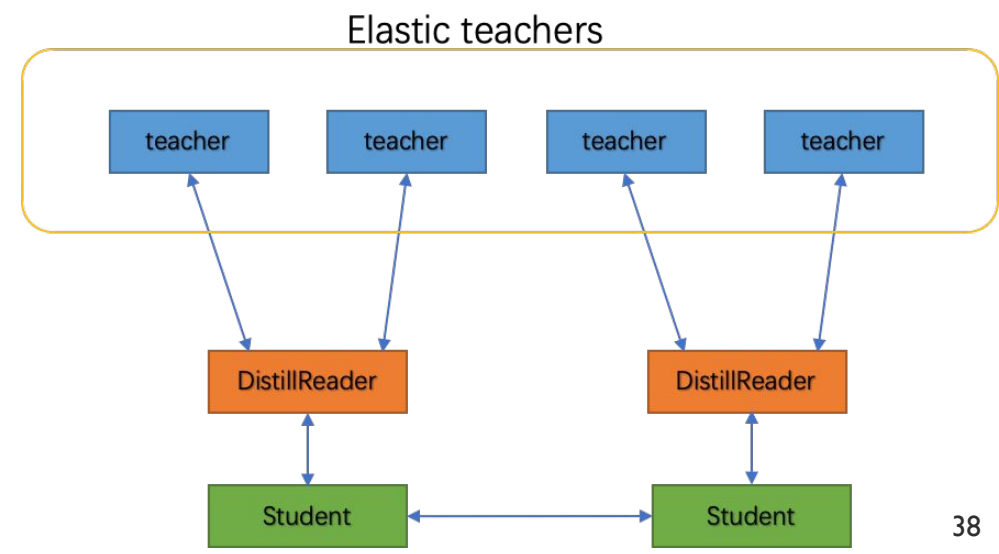
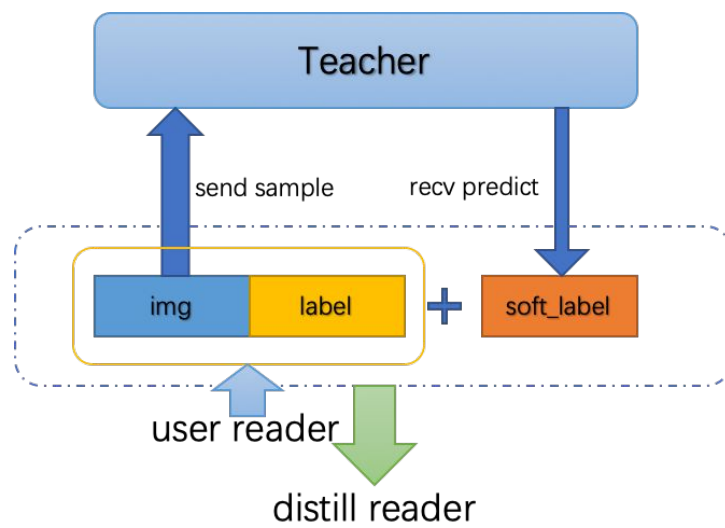
## Improve training speed

- Larger batch size
- Heterogeneous pipelines



## Improve utilization of resources

- Teacher using inference card
- No task failures caused by preemption of online resources
- Make full use of k40 cluster





## Areas the project could use help on

Would like to cooperate with Horovod's elastic training feature, to provide users the state-of-the-art elastic training on both collective mode and parameter server mode.



# Feedback on working with LF AI & Data

Benefit from LF AI & Data Slack, Zoom, Mailing Lists for WG meetings, etc.

Value opportunities to work with related projects (ONNX/Horovod/etc).

# LF AI & Data - Ongoing Annual Project Reviews

 LF AI & DATA

## Annual Review schedule

Date	Project
April 6, 2021	<a href="#">Egeria</a>
April 6, 2021	OpenDS4all
May 20, 2021	<a href="#">ONNX</a>
July 15, 2021	<a href="#">Acumos</a>
July 29, 2021	Angel
July 29, 2021	<a href="#">Adlik</a>
Aug 26, 2021	<a href="#">EDL</a>
Aug 26, 2021	<a href="#">Sparklyr</a>
Sept 9, 2021	Marquez
Sept 9, 2021	<a href="#">Milvus</a>
Sept 23, 2021	<a href="#">NNStreamer</a>
Sept 23, 2021	<a href="#">ForestFlow</a>
Oct 7, 2021	Ludwig
Oct 7, 2021	Amundsen

Oct 21, 2021	AI Fairness 360
Oct 21, 2021	AI <a href="#">Explainability</a> 360
Oct 21, 2021	Adversarial Robustness Toolbox
Nov 4, 2021	<a href="#">Horovod</a>
Nov 4, 2021	FEAST
Nov 18, 2021	<a href="#">SOAJS</a>
Nov 18, 2021	Delta
Dec 2, 2021	<a href="#">DataPractices.org</a>
Dec 2, 2021	<a href="#">JanusGraph</a>
Dec 16, 2021	<a href="#">Pyro</a>
Jan 6, 2021	<a href="#">Datashim</a>
Jan 6, 2022	<a href="#">Flyte</a>
Jan 20, 2022	<a href="#">RosaeNLG</a>
Jan 20, 2022	<a href="#">SubstraFramework</a>
	<a href="#">MLX</a>
	<a href="#">VulcanKompute</a>

[Schedule:](https://wiki.lfaidata.foundation/pages/editpage.action?pageId=43286684) <https://wiki.lfaidata.foundation/pages/editpage.action?pageId=43286684>

# LF AI & Data - General Updates

 LF AI & DATA

Linux Foundation AI & Data Landscape  
2021-06-26T04:29:53Z 35c7a79

See the interactive landscape at [lfaidata.foundation](https://lfaidata.foundation)

Greyed logos are not open source

**309 projects**  
**2.2M+ GitHub Stars**  
**450M+ LoC**  
**40K+ developers**

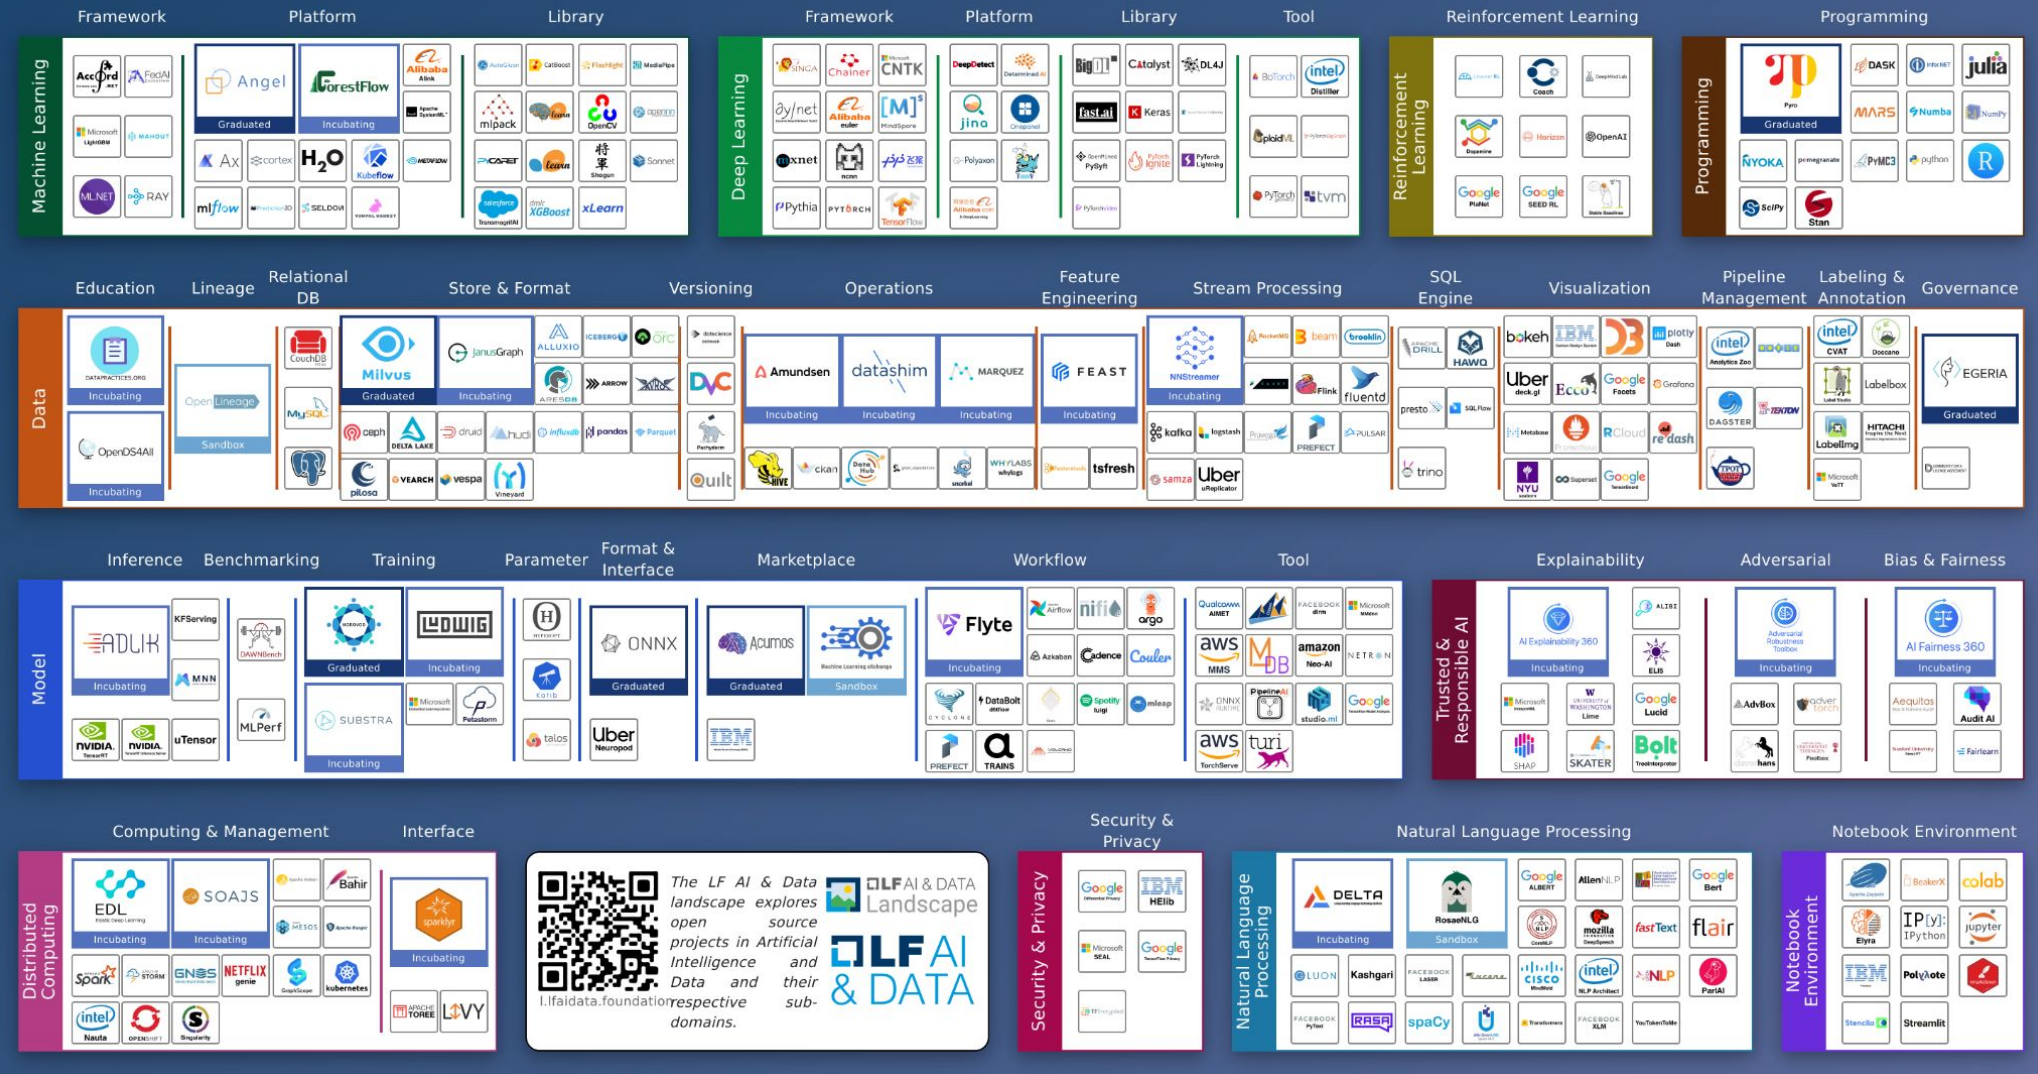
**140+ founding orgs**

**10+ universities**

**13 licenses**

**>450M LoC**

**>1M new LoC weekly**



Help us improve the landscape:  
<https://github.com/lfai/lfaidata-landscape/>

Machine Learning	Framework	Platform	Library	Framework	Platform	Library	Tool	Reinforcement Learning	Programming
		 Graduated	 Incubating						 Graduated

Data	Education	Lineage	Relational DB	Store & Format	Versioning	Operations	Feature Engineering	Stream Processing	SQL Engine	Visualization	Pipeline Management	Labeling & Annotation	Governance
	 Incubating	 Sandbox	 Graduated	 Incubating		 Incubating	 Incubating	 Incubating	 Incubating	 Incubating			 Graduated

Model	Inference	Benchmarking	Training	Parameter	Format & Interface	Marketplace	Workflow	Tool	Explainability	Adversarial	Bias & Fairness
	 Incubating		 Graduated	 Incubating	 Graduated	 Graduated	 Sandbox	 Incubating	 Incubating	 Incubating	 Incubating

Distributed Computing	Computing & Management	Interface	Security & Privacy	Natural Language Processing	Notebook Environment
	 Incubating	 Incubating	 Incubating	 Incubating	 Sandbox



The LF AI & Data landscape explores open source projects in Artificial Intelligence and Data and their respective sub-domains.







lfaidata.foundation























# Projects

<https://landscape.lfai.foundation/card-mode?project=company>




Graduated LF AI & Data Projects (6)

 Acumos LF AI & Data Foundation ★ 16	 Angel-ML LF AI & Data Foundation ★ 6,247	 Egeria LF AI & Data Foundation ★ 509	 Horovod LF AI & Data Foundation ★ 11,246	 ONNX LF AI & Data Foundation ★ 10,681	 Pyro LF AI & Data Foundation ★ 6,939
---	--	--	--	---	--

Incubating LF AI & Data Projects (20)

 ADLIK LF AI & Data Foundation ★ 256	 Adversarial Robustness Toolbox LF AI & Data Foundation ★ 2,250	 AI Explainability 360 Toolkit LF AI & Data Foundation ★ 840	 AI Fairness 360 Toolkit LF AI & Data Foundation ★ 1,397	 Amundsen LF AI & Data Foundation ★ 2,164	 DataPractices.ORG LF AI & Data Foundation ★ 10	 datashim LF AI & Data Foundation ★ 101	 DELTA LF AI & Data Foundation ★ 1,413	 Elastic Deep Learning (EDL) LF AI & Data Foundation ★ 132	 Feast LF AI & Data Foundation ★ 1,881
 Flyte LF AI & Data Foundation ★ 1,421	 ForestFlow LF AI & Data Foundation ★ 46	 JanusGraph LF AI & Data Foundation ★ 3,972	 Ludwig LF AI & Data Foundation ★ 7,678	 Marquez LF AI & Data Foundation ★ 630	 Milvus LF AI & Data Foundation ★ 6,260	 NNStreamer LF AI & Data Foundation ★ 362	 OpenDS4All LF AI & Data Foundation ★ 310	 SOAJS LF AI & Data Foundation ★ 46	 sparklyr LF AI & Data Foundation ★ 785

Sandbox LF AI & Data Projects (3)

 Machine Learning eXchange LF AI & Data Foundation ★ 11	 OpenLineage LF AI & Data Foundation ★ 272	 RosaeNLG LF AI & Data Foundation ★ 20
--	---	---

\* Missing Substra Framework (pending logo redesign)











# New projects in 2021

→ Track incoming proposals via: <https://github.com/lfai/proposing-projects>

1. **Datashim:** Open source enablement and acceleration of data access for Kubernetes/OpenShift workloads in a transparent and declarative way
2. **Flyte:** Production-grade, declarative, structured and highly scalable cloud-native workflow orchestration platform
3. **RosaeNLG:** Open source project, template-based Natural Language Generation (NLG) automating the production of relatively repetitive texts based on structured input data and textual templates, run by a NLG engine
4. **Substra Framework:** Low-layer framework, offering secure, traceable, distributed orchestration of machine learning tasks among partners.
5. **ML eXchange:** Data and AI Assets Catalog and Execution Engine
6. **Kompute: Blazing fast, mobile-enabled, asynchronous, and optimized for advanced GPU processing usecases.**
7. **Open Lineage:** Open standard for metadata and lineage collection designed to instrument jobs as they are running



# Active and growing developer community

 <b>LF AI &amp; Data Foundation</b> LF AI & Data Foundation drives Open Source Innovation in Artificial Intelligence, Machine Learning, Deep Learning, and Data <a href="#">Go To Overview</a>				 <b>10.06M</b> Lines Of Code	 <b>138.90K</b> Pull Requests / Changesets	 <b>No Data</b> Builds	
				 <b>10.13K</b> Emails	 <b>13.31K</b> Total Mentions	 <b>59.85K</b> Total Issues	
				 <b>1.65K</b> Pages	 <b>6.90K</b> Messages	 <b>15.28K</b> 50th Percentile Of Median Pulls By Image	
<b>Contributors</b> <b>16.87K</b>	<b>Contributions</b> <b>282.21K</b>	<b>Commits</b> <b>128.09K</b>	<b>Repositories</b> <b>429</b>				

# Looking to host a project with LF AI & Data

- › Hosted project stages and life cycle:

<https://lfaidata.foundation/project-stages-and-lifecycle/>

- › Offered services for hosted projects:

<https://lfaidata.foundation/services-for-projects/>

- › Contact:

Jim Spohrer (TAC Chair) and Ibrahim Haddad (ED, LF AI & Data)

# Promoting Upcoming Project Releases

We promote project releases via a blog post and on LF AI & Data [Twitter](#) and/or [LinkedIn](#) social channels

For links to details on upcoming releases for LF AI & Data hosted projects visit the [Technical Project Releases wiki](#)

If you are an LF AI & Data hosted project and would like LF AI & Data to promote your release, reach out to [pr@lfai.foundation](mailto:pr@lfai.foundation) to coordinate in advance (min 2 wks) of your expected release date.

# Note on quorum

As LF AI & Data is growing, we now have 18 voting members on the TAC.

TAC representative - please ensure you attend the bi-weekly calls or email Jacqueline/Ibrahim to designate an alternate representative when you can not make it.

We need to ensure quorum on the calls especially when we have items to vote on.

# Upcoming TAC Meetings

# Upcoming TAC Meetings (Tentative)

- › July 29: TBD - Annual project review + TBD
- › Aug 12: TBD - Annual project review + TBD

Please send agenda topic requests to  
[tac-general@lists.lfaidata.foundation](mailto:tac-general@lists.lfaidata.foundation)

# TAC Meeting Details

- › To subscribe to the TAC Group Calendar, visit the wiki: <https://wiki.lfaidata.foundation/x/cQB2>
- › 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 Notice

- › 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](mailto:legal@linuxfoundation.org) with any questions about The Linux Foundation's policies or the notices set forth on this slide.