Meeting of the Technical Advisory Council (TAC)

July 29th, 2021

LF 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 Updegrove 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
- Wiki: wiki.lfaidata.foundation
- GitHub: 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

- 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 July 15th (3 mins)
- Adlik Annual Project Review (20 minutes)
- Angel Annual Project Review (20 minutes)
- LF AI General Updates
- Open Discussion

TAC Voting Members

* = still need backup specified on wiki

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

Approval of July 15th, 2021 Minutes

Draft minutes from the July 15th TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

- That the minutes of the July 15th meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

Adlik Annual Report

Yuan Liya

07/29/2021

THE LINUX FOUNDATION

LF AI

What's Adlik

Adlik [ædlik], a toolkit for accelerating deep learning inference on specific hardware.

- Support several kinds of hardwares.
- Collaborate with existing inference solutions with unified entrance.

An open source project of LFAI and code hosted on GitHub. https://github.com/Adlik

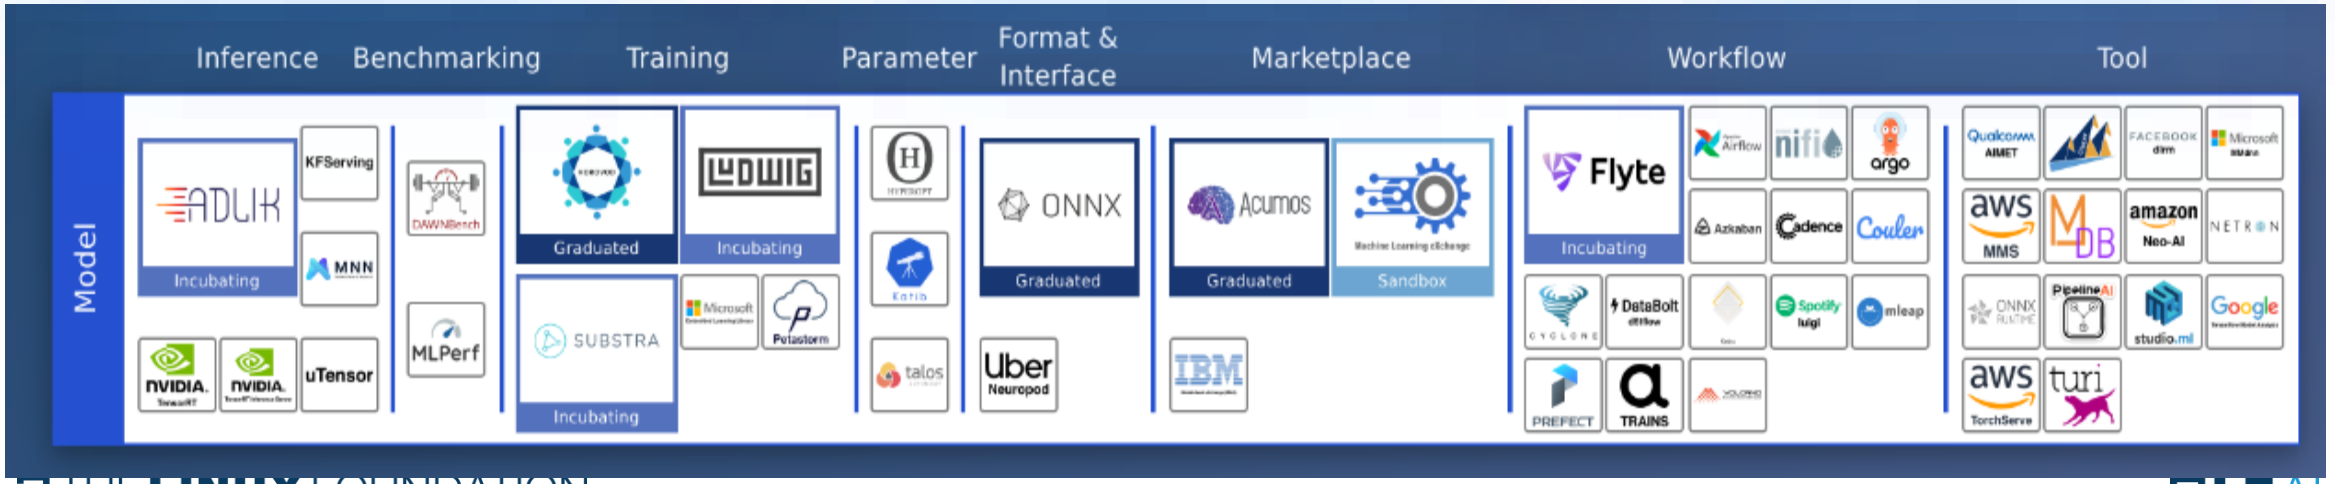

Adlik Architecture

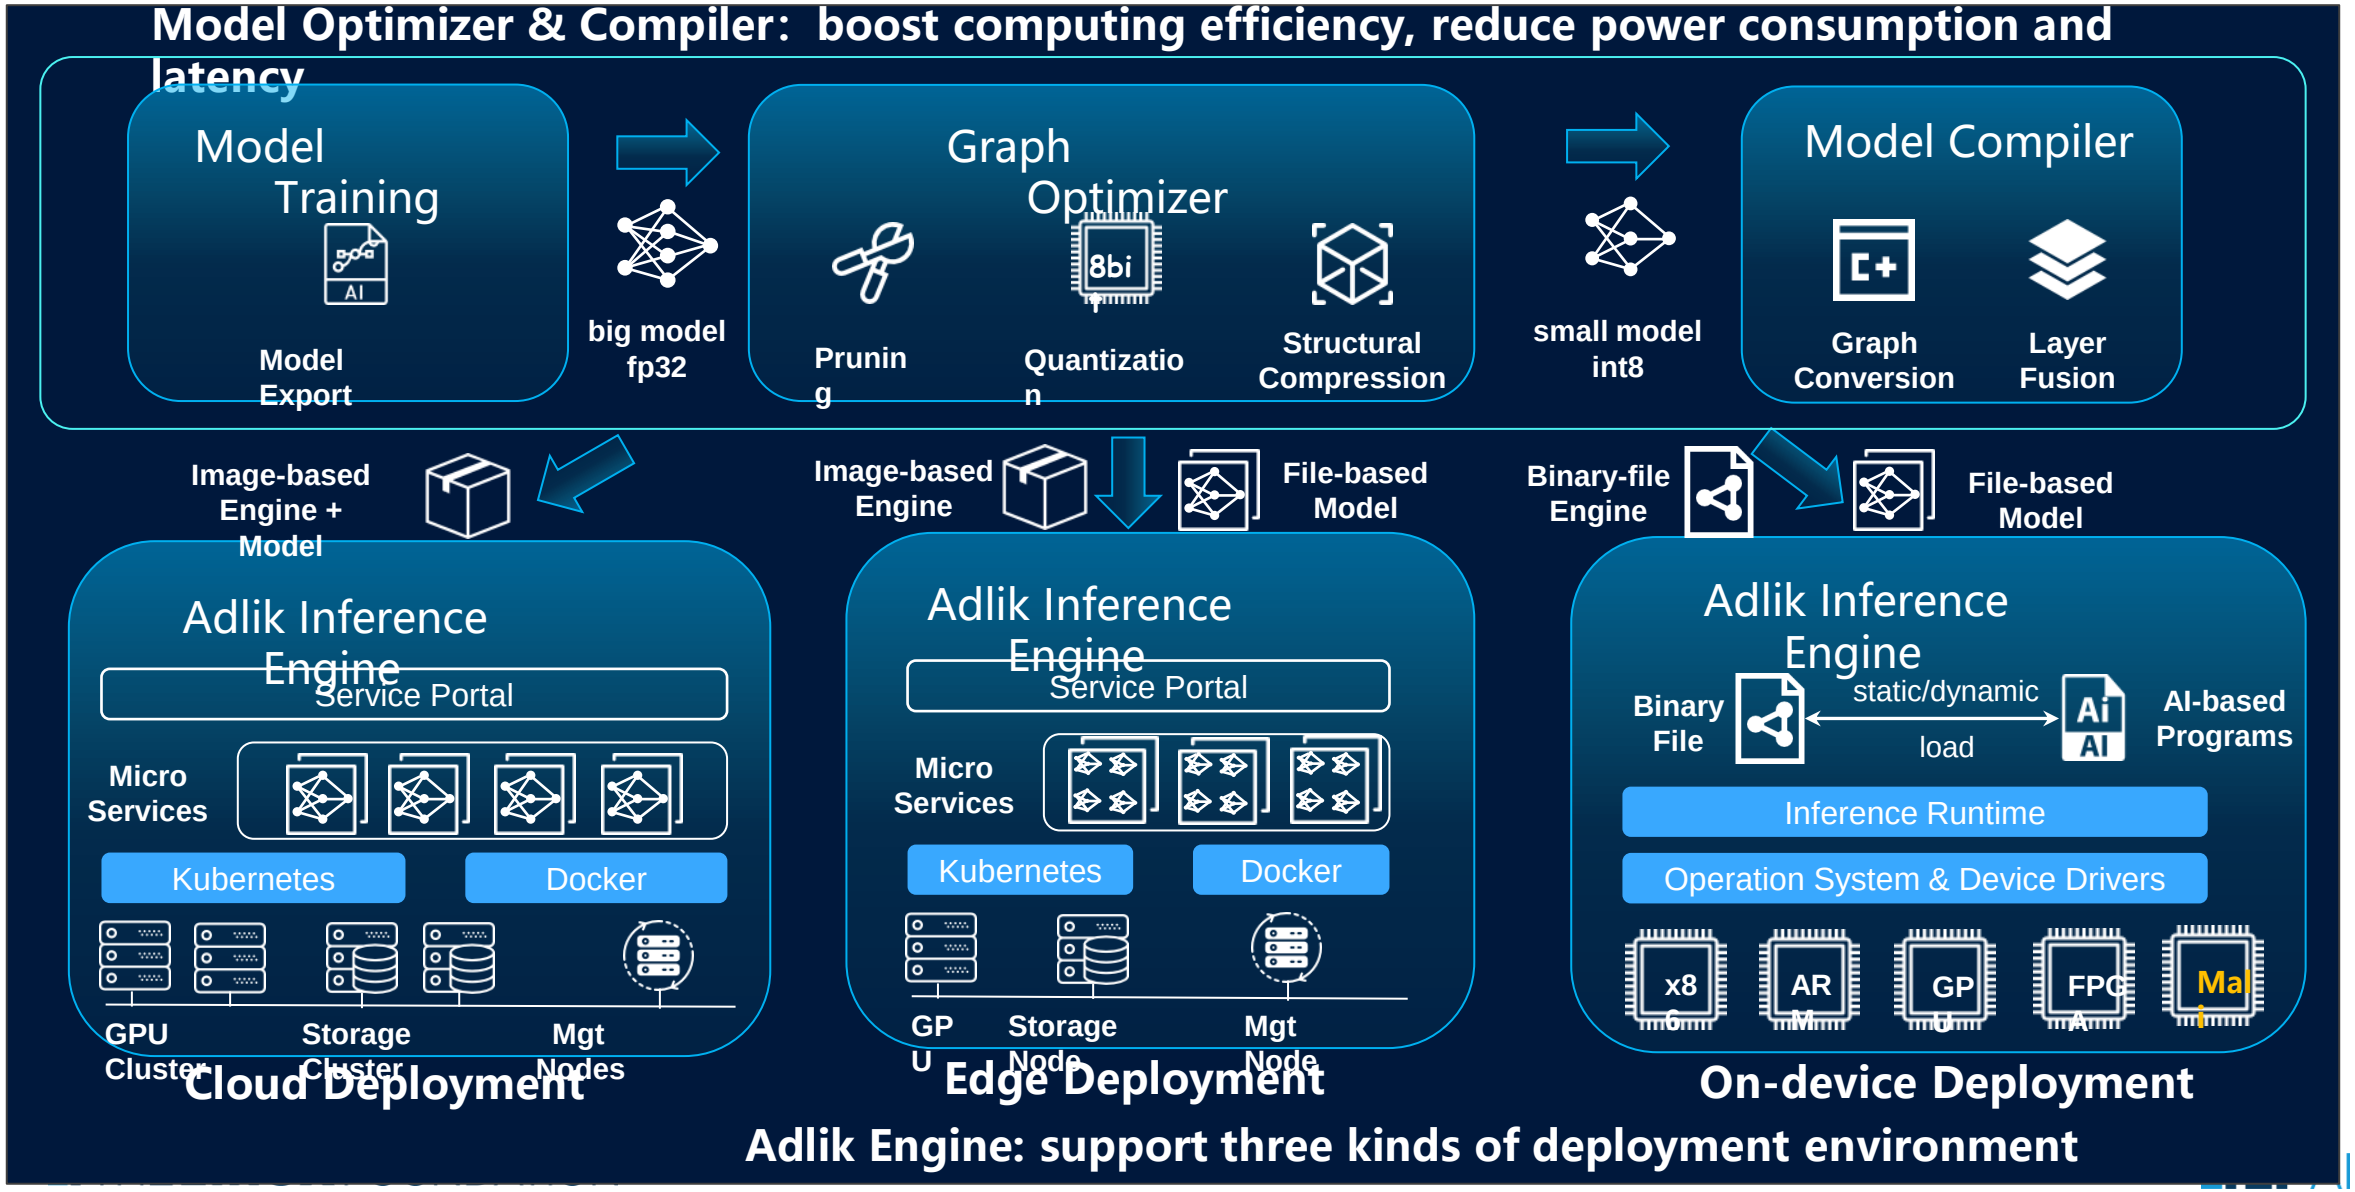

Insights of Adlik (over the last year)

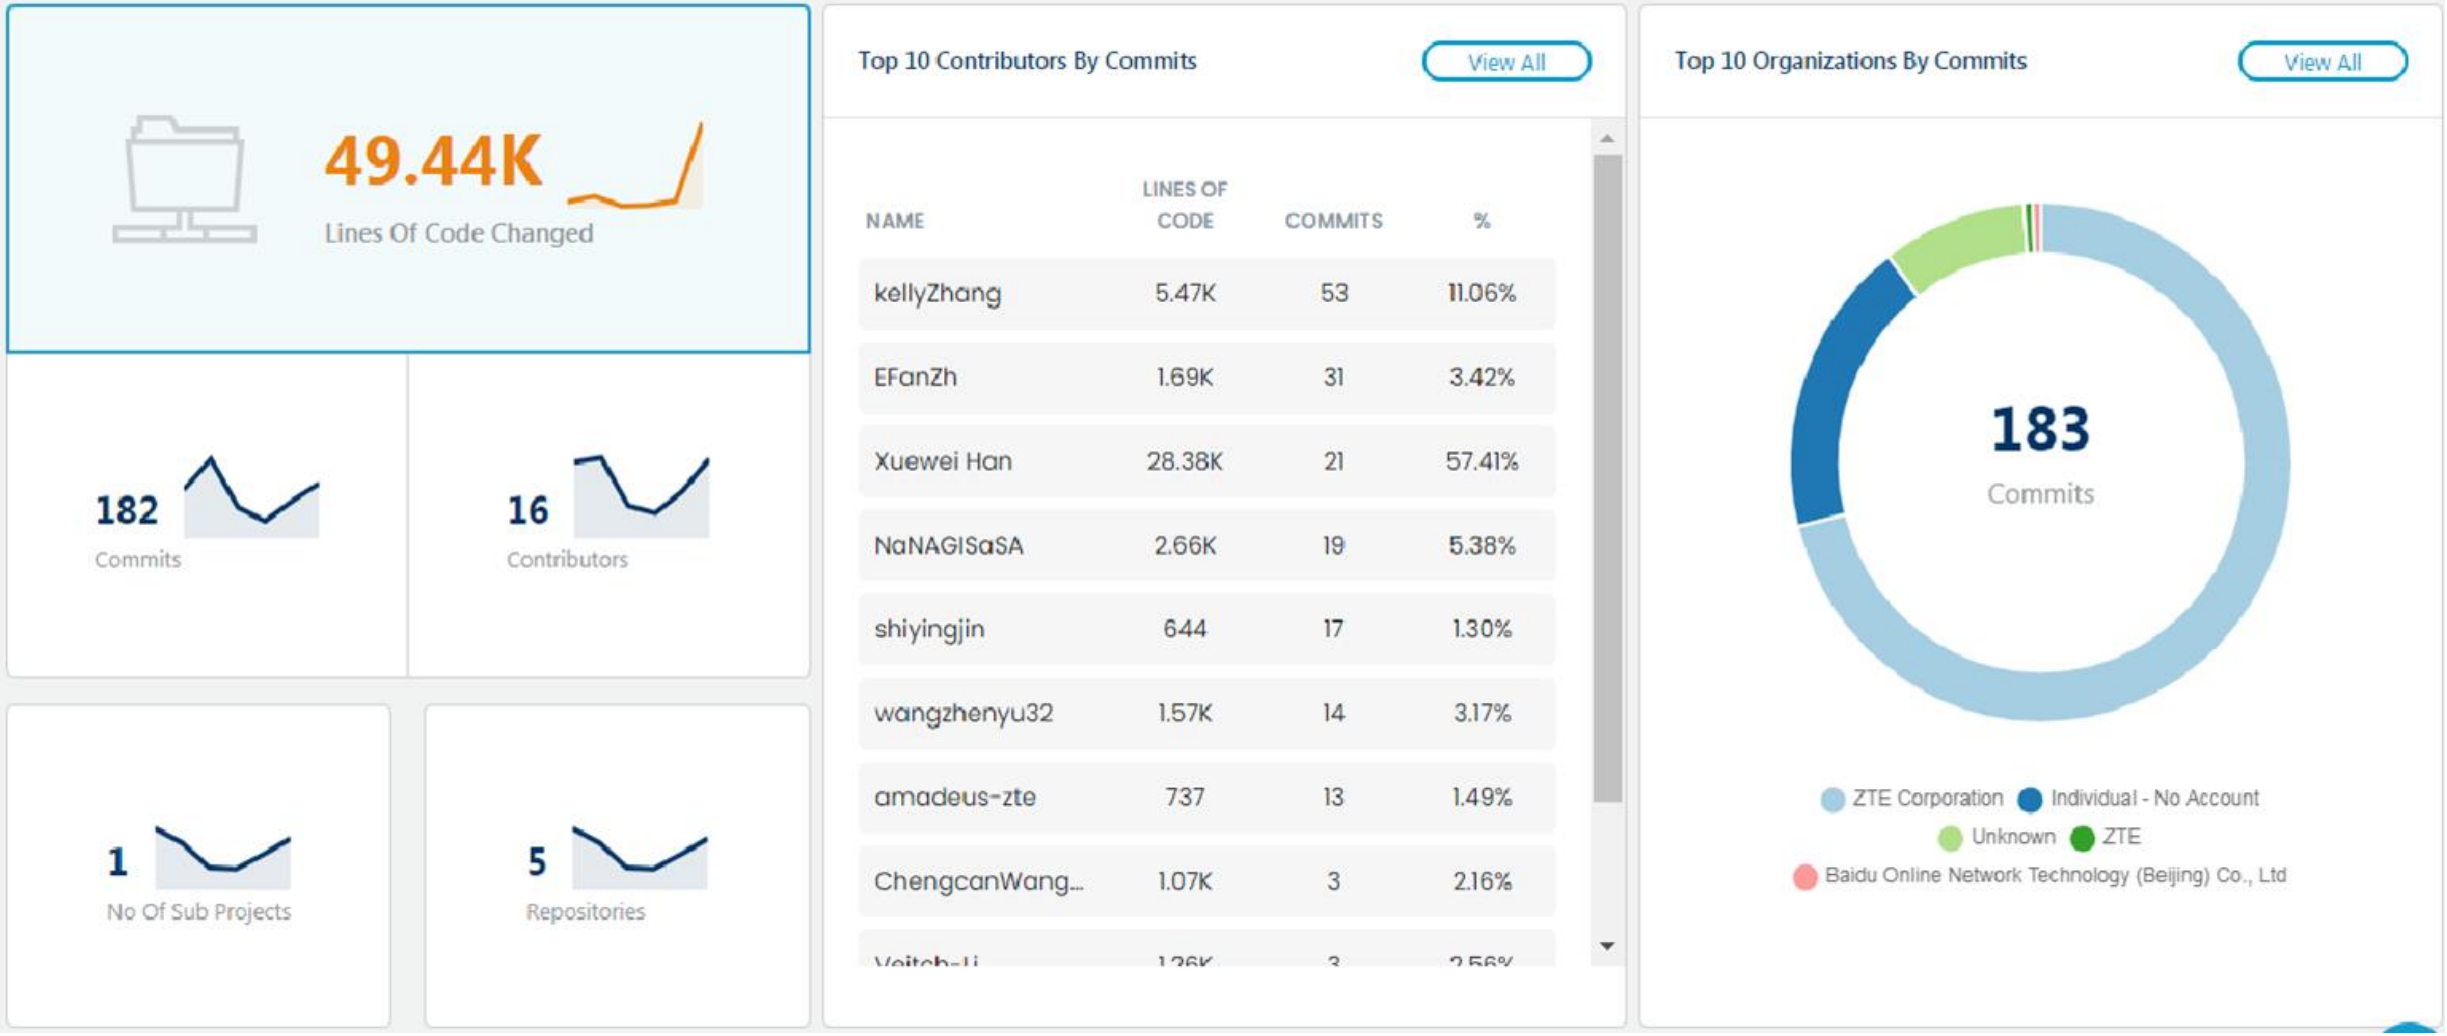

Key achievements

- *Released Version 0.2.0 (Bear): 2020.11*
  - *Provide new compiler framework.*
  - *Support hybrid scheduling of ML and DL inference jobs.*
  - *Support image based deployment of Adlik compiler and inference engine in cloud native environment.*
  - *Support Adlik running in RaspberryPi and JetsonNano*
  - *Benchmark test for ResNet-50, Inception V3, Yolo V3 and Bert.*

- *Released Version 0.3.0 (Cheetah): 2021.6*
  - *Specific optimization for YOLO v4 and Resnet50 v1/v2*
  - *Model compiler with PaddlePaddle/MXNet/Caffe supported*
  - *TVM auto scheduling*
  - *TVM/OpenVINO/TFLite/TensorRT/TensorFlow runtime integrated*
  - *Paddle models supported in benchmark test framework*

Adlik optimizer

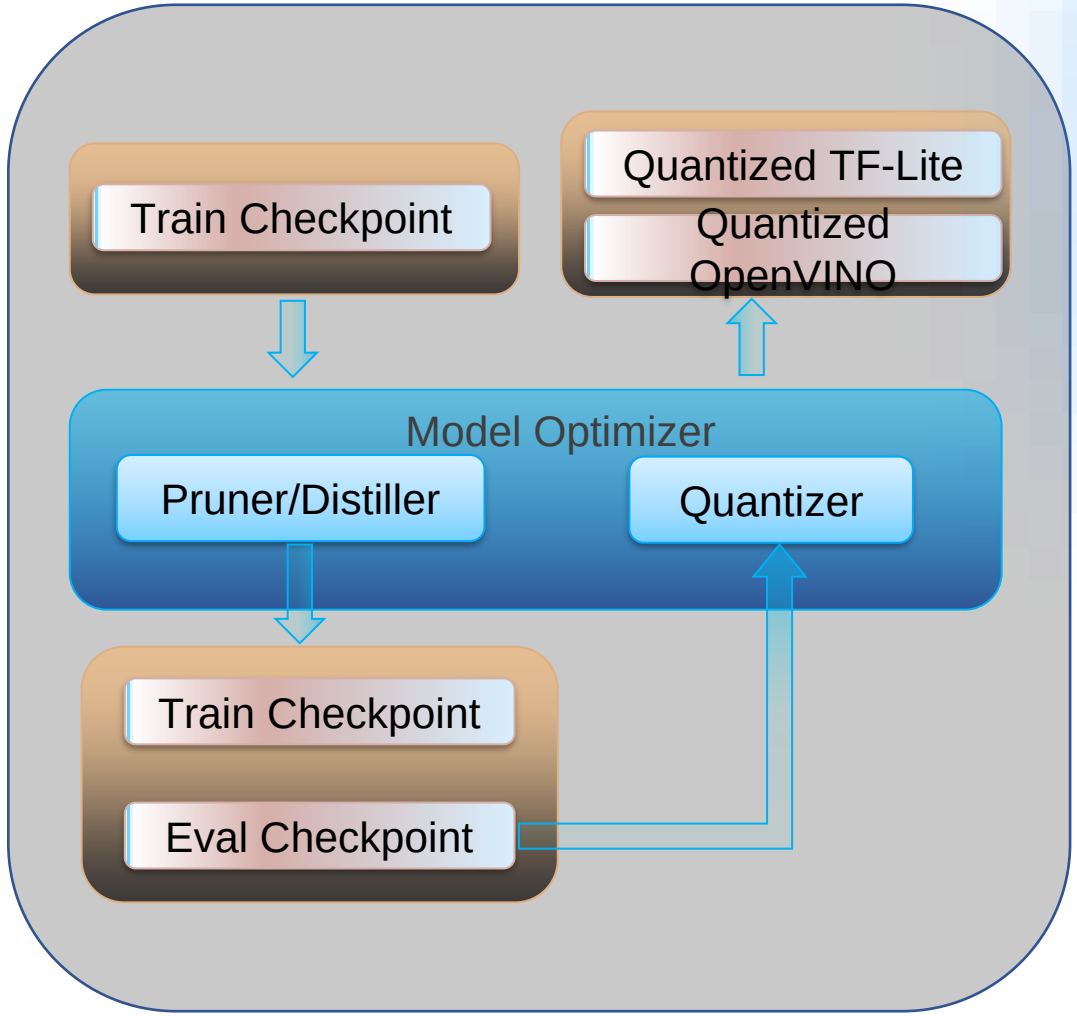

- Supporting combined distillation, which greatly improves the accuracy of the model
- Supporting 8-bit Calibration Quantization. Quantizing process needs only a small batch of datasets and few minutes.

	Params	Flops	Accuracy	Size
ResNet-50	25610152	3899M	76.174%	99M
+ pruned(72.8%)	6954152	1075M	72.28%	27M
+ distill	6954152	1075M	76.39%	27M
+ quantize			75.938%	7.1M

Model Optimizer Result: 7.1/99 = 7.2%

Adlik Cloud Native use case

Docker Environment

```
docker run -it --rm -v /media/B/work/keras:/model 10.233.170.2:5000/adlik/model-compiler:7.0_10.0 bash
root@ecaf2fd16421:/# cd model/
root@ecaf2fd16421:/model# python3 compile_model.py
Source type: ONNXModelFile.
Target type: OpenvinoModel.
Compile_path: ONNXModelFile -> OpenvinoModel.
{'status': 'success', 'path': 'model_tf_yolov3_608_128/yolov3_1.zip'}
docker run -it --rm -v /home/t630/zkl:/model -p 31000:8500 10.233.170.2:31000/00253486/adlik_serving-openvino:latest bash
/# adlik-serving --model_base_path=/model/yolov3_repos/ --grpc_port=8500 --http_port=8501
I adlik_serving/server/core/server_core.cc:54] Adlik serving is running...
I adlik_serving/server/grpc/grpc_options.cc:88] grpc server port: 8500
I adlik_serving/server/grpc/grpc_server.cc:24] grpc server is serving...
I adlik_serving/server/http/http_options.cc:35] http server port: 8501
python3 yolov3_client.py -n yolo416 -b 1 dog.jpg
```

Kubernetes Environment

```
kubectl create -f compiler.yaml
pod/model-compiler created
kubectl get pod | grep compiler
model-compiler                          1/1     Running           0          24s
ls
yolov3  yolov3_1.zip
kubectl create -f openvino-serving.yaml
kubectl get pod | grep openvino-serving
openvino-serving                         1/1     Running           0          24s
kubectl create -f openvino-svc.yaml
kubectl get pod | grep openvino-serving
openvino-service          NodePort    10.254.255.197   <none>          8500:31501/TCP                79s
python3 yolov3_client.py -b 1 dog.jpg
```

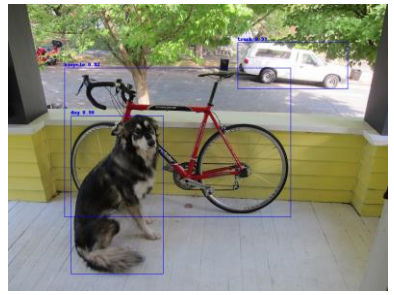

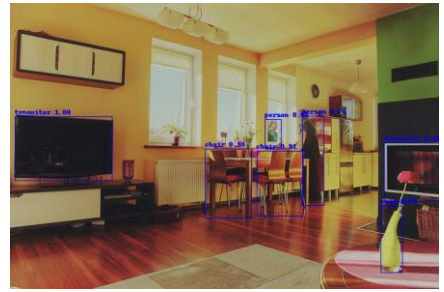

Adlik Compiler Practice: Operator Schedule Optimization

Step1: Schedule parameter optimization for single op

Step 2：Schedule parameter optimization in graph view

Schedule Parameters Optimization Algorithm

Schedule Template

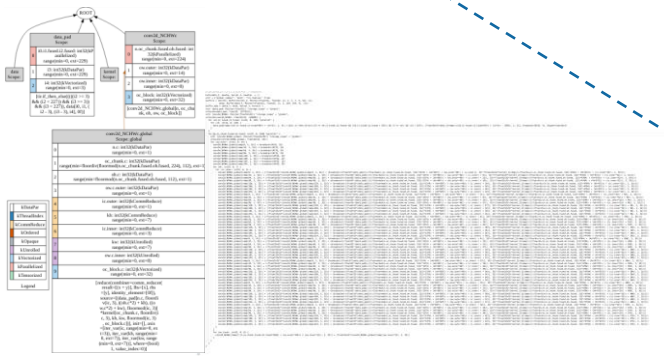

Operator Schedule Template

Optimized schedule parameter

调度参数搜索

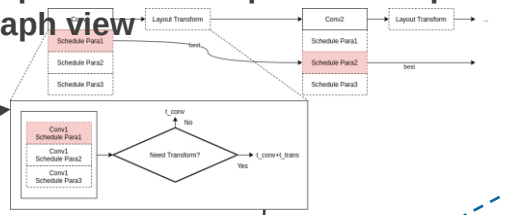

- Generate schedule based on schedule parameters and template
- Execute the program and collect performance information using VTune (Hostpot,/Memory Consumption/Memory Access)

调度模板设计

调度性能分

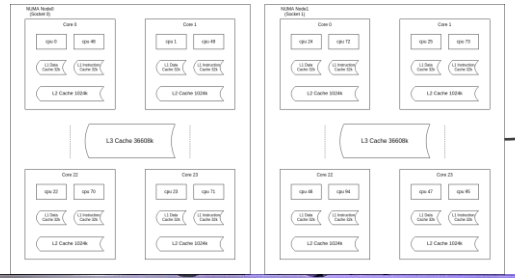

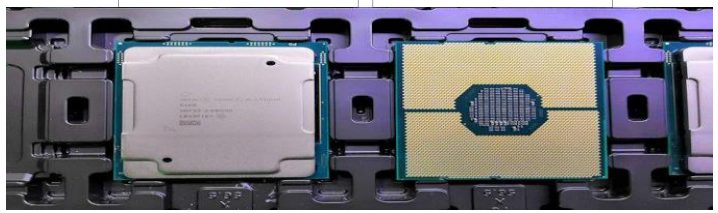

CPU Micro Architecture

1×3×224×224

Conv

W ⟨64×3×7×7⟩
kernel_shape = 7, 7
pads = 3, 3, 3, 3
strides = 2, 2

operator workload

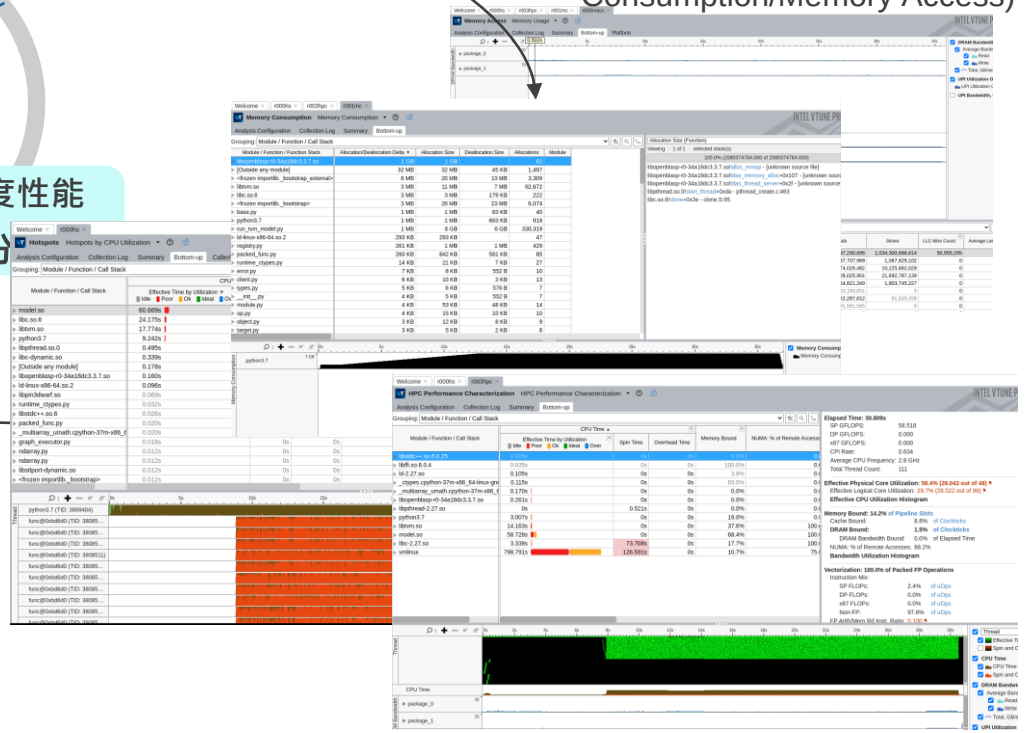

THE LINUX FOUNDATION

Next: RL-based AutoScheduler

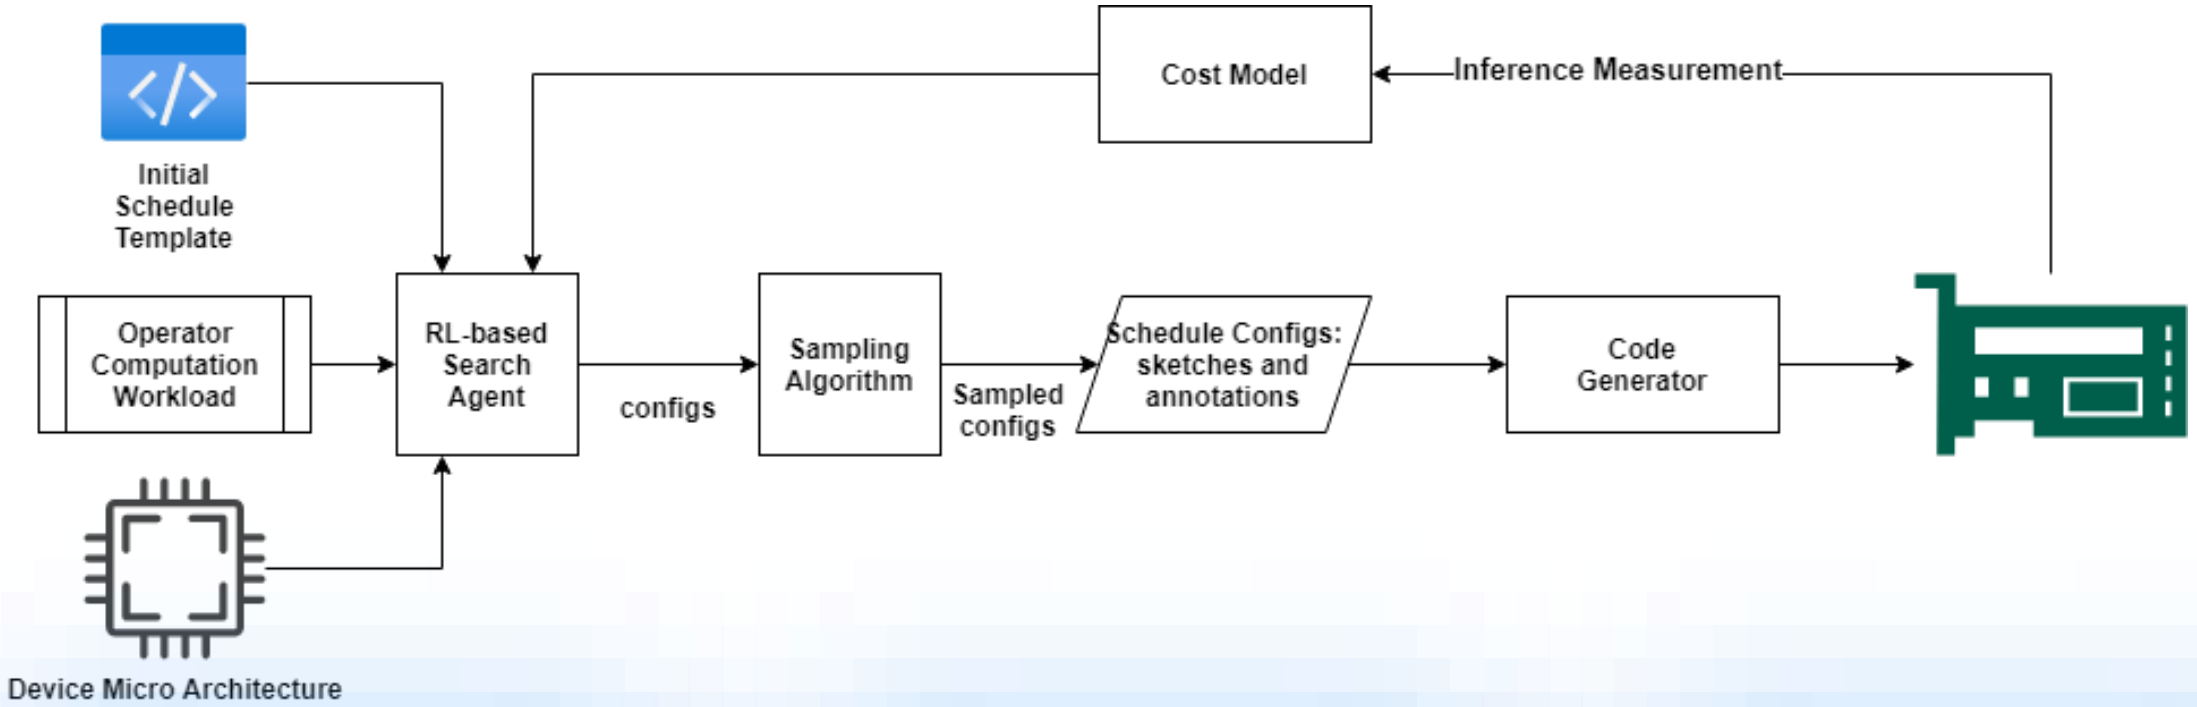

Based on Ansor (a.k.a TVM auto scheduler)

Agent generates new config based on:
1. Micro-architecture of device
2. Init/last episode configs
3. Operator workload

A cost model is trained to accelerate config evaluation

Adlik Community

Activities :

- Routine TSC meetings.
- Stable cooperation with CMCC, Unicom, AIIA.
- Submit CR in ORAN community, introduce Adlik into ORAN framework.
- Cooperation intention with PaddlePaddle community.

Presentations

CSDN ProCon 2020

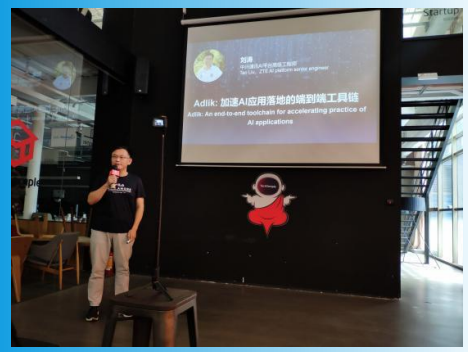

AI OpenDay in Shenzhen, 2020

COPU Open China Open World Summit, 2021

Open source AI summit，2020

AIIA annual meeting in Xiamen, 2020

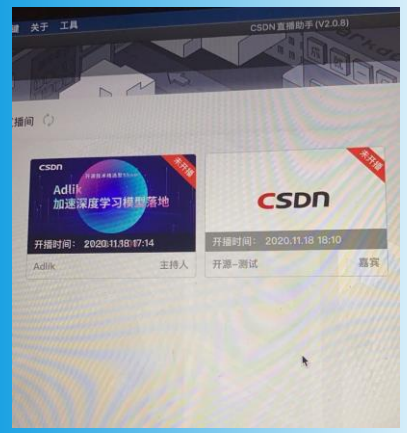

CSDN Show, 2020

Wave summit, 2020

Areas the project will use help on

Would like to explore coorperation with ONNX project.

Would like to have a community meetup on September 10th.

Feedback on working with LF AI & Data

- Benefit from LF AI & Data resources including website, Zoom, mailing lists, wiki space, etc.
- Great support for Adlik Community Meetings
- Great support for Adlik Release Blogs and Social Media
- Great opportunities for Adlik to reach out

谢 谢

Thank You

Annual Review for Angel

Bruce Tao

07/29/2021

LF AI & DATA

Angel

Brief Description:

Angel is a high-performance distributed machine learning platform based on the philosophy of Parameter Server. It is tuned for performance with big data from Tencent and has a wide range of applicability and stability, demonstrating increasing advantage in handling higher dimension model.

Contributed by:

Tencent in August 2018 as an Incubation Project
Angel graduated in December 2019

Key Links:

Github

- https://github.com/Angel-ML/angel

Artwork

- https://github.com/lfai/artwork/tree/master/projects/angel

Mail Lists

- angel-announce
- angel-technical-discuss
- angel-tsc

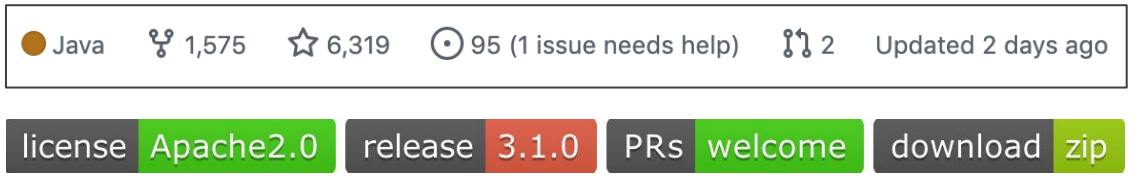

Organizations contributing

Contributions from increasing number of organizations and thanks to

Commits in the past year

132.77K
Lines Of Code Changed

101
Commits

12
Contributors

1
No Of Sub Projects

3
Repositories

Top 10 Contributors By Commits View All

NAME	LINES OF CODE	COMMITS	%
Ouyang Wen	35.06K	26	26.41%
paynie	78.93K	17	59.45%
xs-li	211	14	0.16%
joyjxu	4.16K	11	3.13%
rachelsunrh	4.29K	8	3.23%
howiehywang	2.99K	7	2.25%
Ryantaocer	446	6	0.34%
jiangyasong	4.02K	6	3.02%
shichenggao	2.61K	3	1.97%

Top 10 Organizations By Commits View All

LF AI & DATA

Pull Requests in the past year

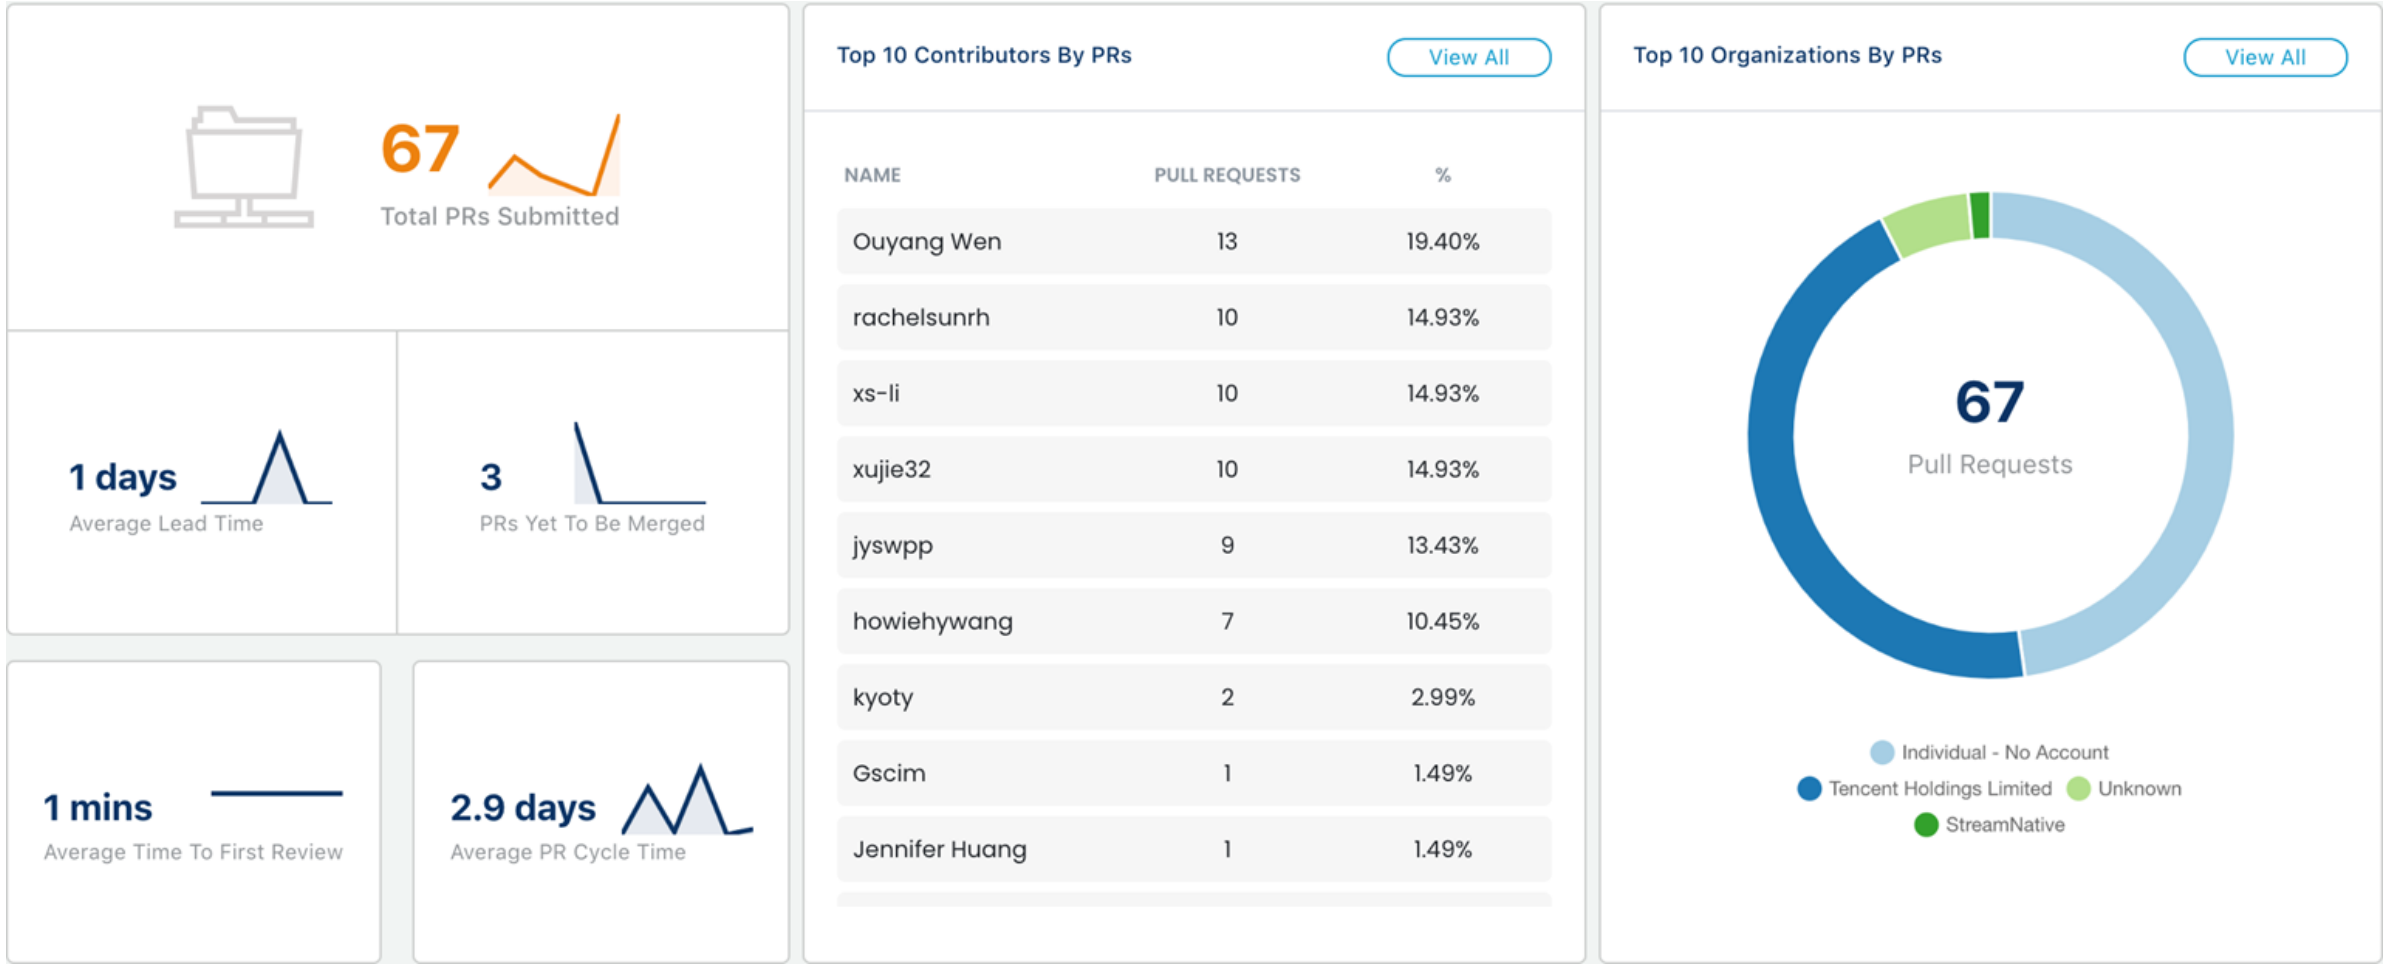

Issues in the past year

Top 10 Contributors By Issues Submitted

View All

NAME	ISSUE	%
paynie	11	11%
howiehywang	9	9%
rachelsunrh	9	9%
xujie32	9	9%
graphv	7	7%
jyswpp	6	6%
dongxuej	5	5%
Gscim	3	3%
Ouyang Wen	3	3%

Top 10 Organizations By Issues Submitted

View All

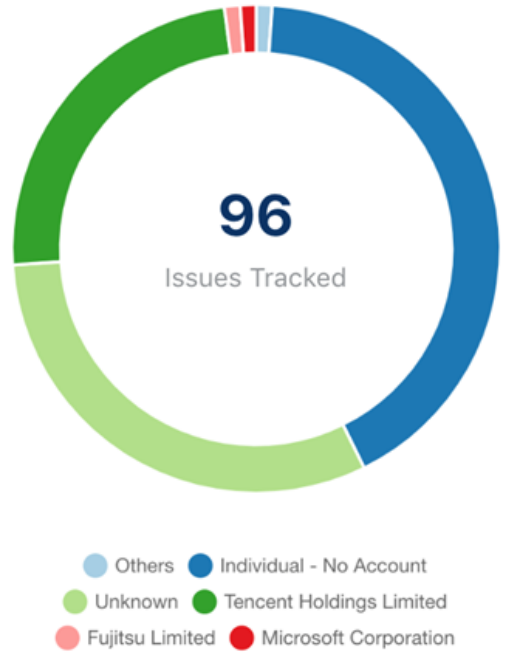

LF AI & DATA

Key Achievements in the past year

- Enterprise and academic user cases grow by ~100%
- Contributions from outside Tencent is increasing and diverse

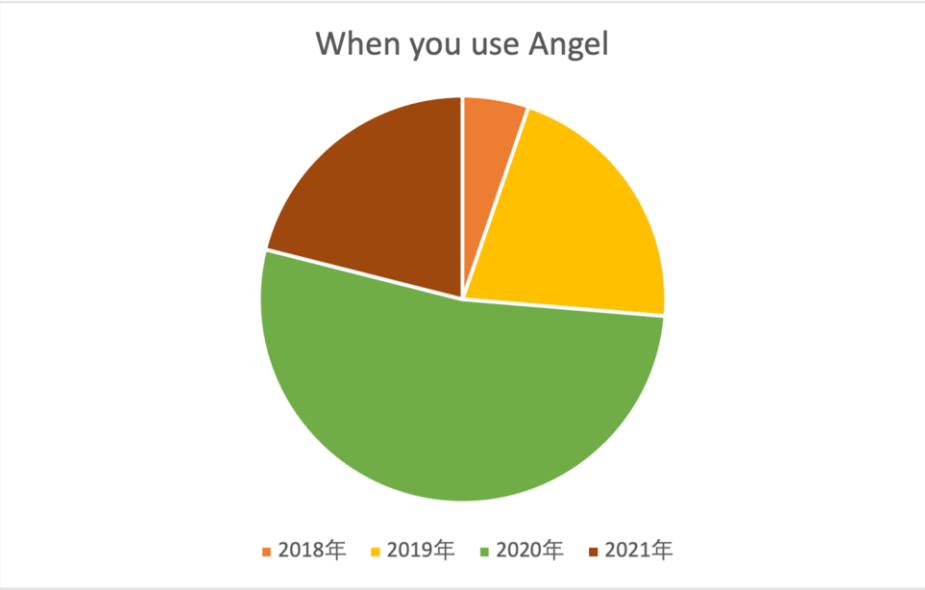

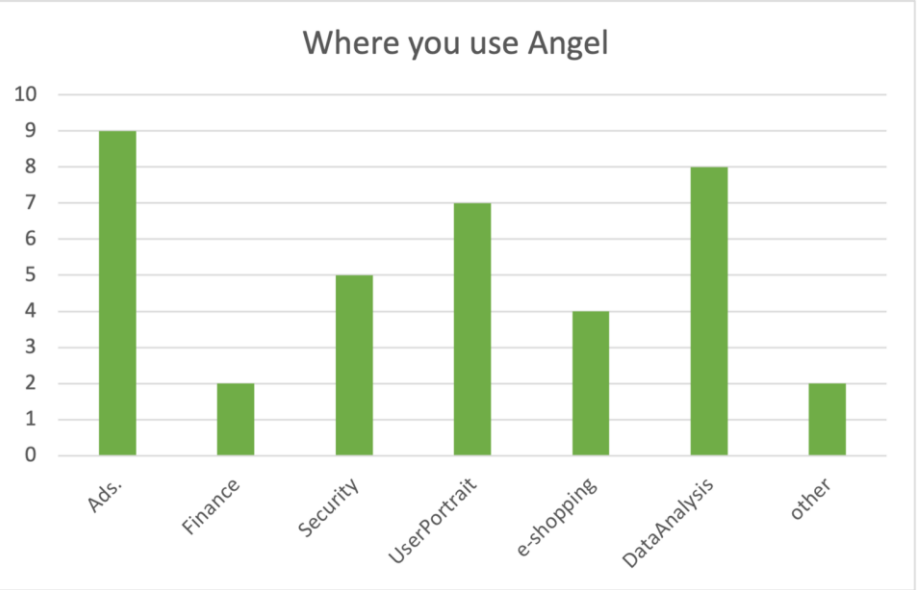

Key Achievements in the past year

- Enterprise and academic user cases grow by ~100%
- Contributions from outside Tencent is increasing and diverse
- More collaborations with Acumos, TRAINS projects

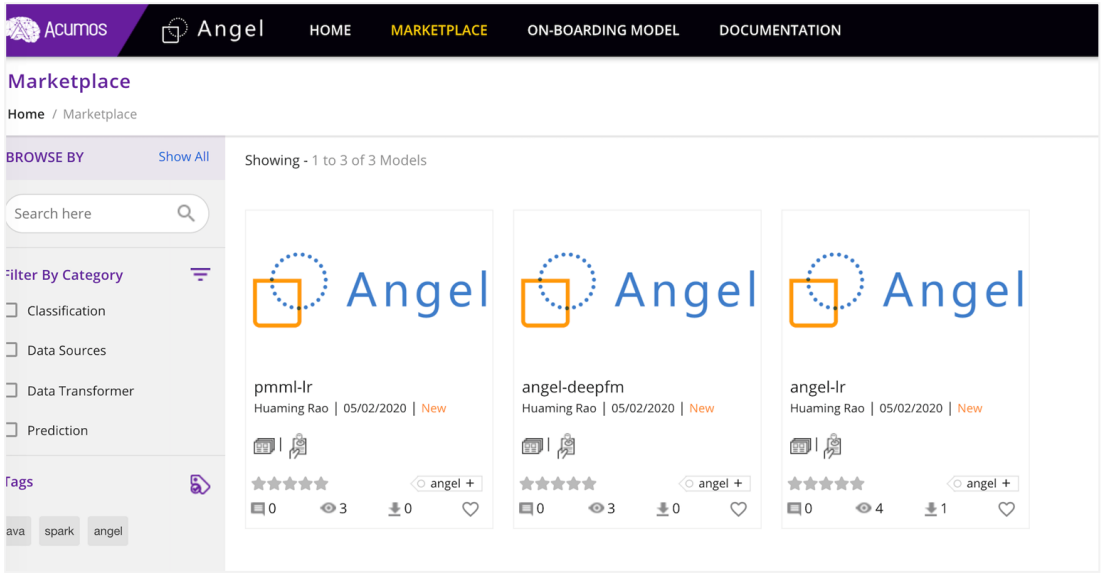

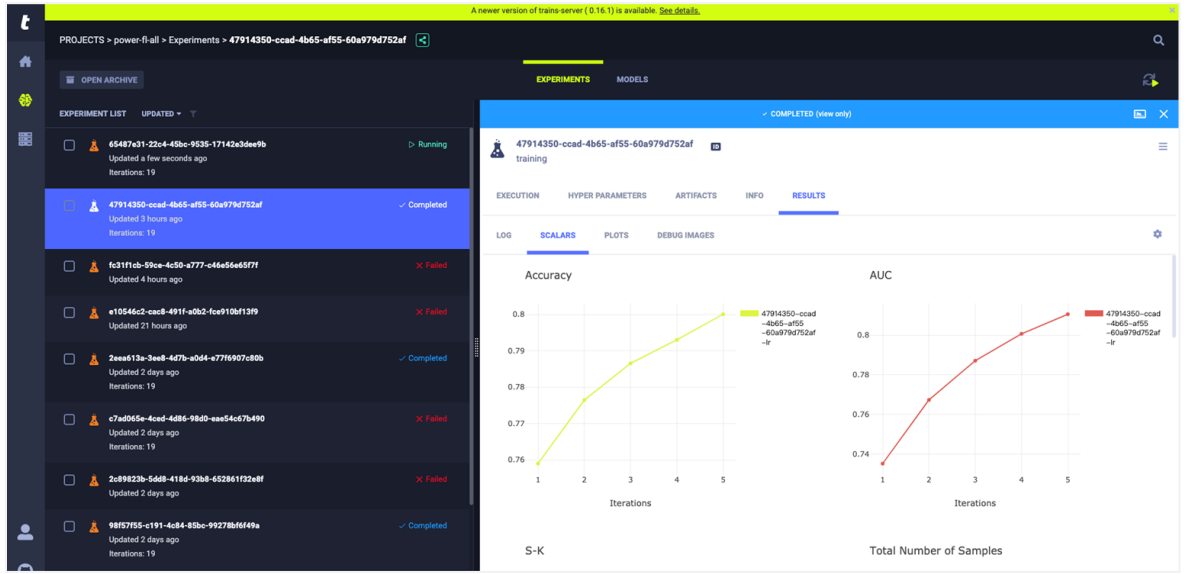

Key Achievements in the past year

- Enterprise and academic user cases grow by ~100%
- Contributions from outside Tencent is increasing and diverse
- More collaborations with Acumos, TRAINS projects
- Active community events
  - LF AI & Data mini summit on OSS 2020 Oct.
  - 10 meetup activities
  - talk on DTCC conference

Key Achievements in the past year

- Enterprise and academic user cases grow by ~100%
- Contributions from outside Tencent is increasing and diverse
- More collaborations with Acumos, TRAINS projects
- Active community events
  - LF AI & Data mini summit on OSS 2020 Oct.
  - 10 meetup activities
  - talk on DTCC conference
- two major releases: v3.1, v3.2
  - Graph Neural Network model support
  - Performance improved by 3 times

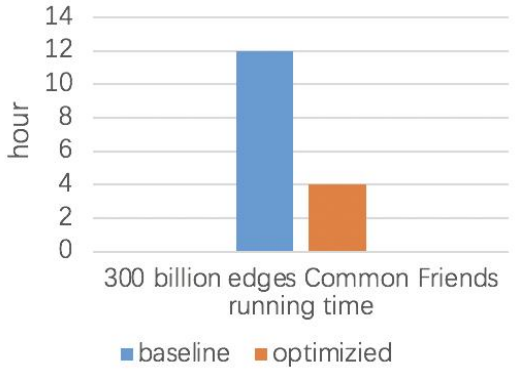

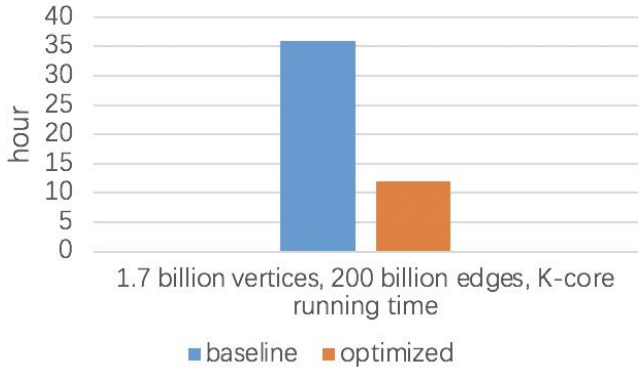

Areas the project could use help on

- Diverse contributions from more organizations
- Community operations like code review, issue management, documentations etc.
- Cross collaborations with other open source projects
- Outreach promotions with Linux Foundation channels
- Joint events both online and offline

Feedback on working with LF AI & Data

- Many thanks to the foundation for the guidance and help
- More and more professional in the project staging process like sandbox/incubation/graduation
- Congratulations to the fast growth of the project family in the past year

- It would be better if there is more coach for the incubating projects and more strategy planning for the graduation projects
- There could be more correlations and collaborations among the projects in and outside the foundation.

Open Discussion

LF AI & DATA

LF AI & Data - Ongoing Annual Project Reviews

LF AI & DATA

July 29, 2021	Adlik	Meng Wei (slack) (confirmed)
Aug 26, 2021	Acumos	Amit Kumar (slack) (confirmed)
Aug 26, 2021	Sparklyr	Sigrid Keydana Yitao Li (slack) (confirmed)
Sept 9, 2021	Marquez	Julien le Dem (slack) (confirmed)
Sept 9, 2021	Milvus	Jun Gu (slack) (confirmed)
Sept 23, 2021	NNStreamer	MyungJoo Ham (slack) (confirmed)
Sept 23, 2021	ForestFlow	Ahmad Alkilani (slack) (confirmed)
Oct 7, 2021	Ludwig	Piero Molino (slack) (confirmed)
Oct 7, 2021	Amundsen	Mark Grover (slack) (confirmed)
Oct 21, 2021	AI Fairness 360	Animesh Singh (to be asked)
Oct 21, 2021	AI Explainability 360	Animesh Singh (to be asked)
Oct 21, 2021	Adversarial Robustness Toolbox	Animesh Singh (to be asked)
Nov 4, 2021	Horovod	Travis Addair (to be asked)
Nov 4, 2021	FEAST	Willem Pienaar (to be asked)
Nov 18, 2021	SOAJS	Antoine Hage (to be asked)
Nov 18, 2021	Delta	Kun Han (to be asked)
Dec 2, 2021	DataPractices.org	Patrick McGarry (to be asked)
Dec 2, 2021	JanusGraph	Jason Plurad (to be asked)
Dec 16, 2021	Pyro	Fritz Obermeyer (to be asked)

Schedule: https://wiki.lfaidata.foundation/pages/editpage.action?pageId=43286684

LF AI & Data - New Updates

LF AI & DATA

MLOps Committee

The Governing Board approved establishing a formal MLOps Committee.

Subscribe to the calendar:

https://lists.lfaidata.foundation/g/mlops-committee/calendar

Subscribe to the mailing list:

https://lists.lfaidata.foundation/g/mlops-committee/editsub

Wiki space (to be setup):

https://lists.lfaidata.foundation/g/mlops-committee/wiki

Associate Members Monthly Call

A monthly call dedicated to enable collaboration between Associate Members of LF AI & Data (Universities, Non-profit organizations, R&D labs, etc)

Subscribe to mailing list, calendar:

https://lists.lfaidata.foundation/g/associatemembers/editsub

Upcoming TAC Meetings

LF AI & DATA

Upcoming TAC Meetings (Tentative)

- Aug 12: TBD
- Aug 26: Annual Project Review: Acumos, Sparklyr

Please send agenda topic requests to 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

LF AI & DATA

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 with any questions about The Linux Foundation's policies or the notices set forth on this slide.

LF AI & DATA