Meeting of the LF AI & Data Technical Advisory Council (TAC)

August 26, 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

LF AI & DATA

Agenda

- Roll Call (2 mins)
- Approval of Minutes from previous meetings (5 mins)
- Kedro Project (40 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.

Approval of July 29th, 2021 Minutes

Draft minutes from the July 29th TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

- That the minutes of the July 29th meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

Approval of August 12th, 2021 Minutes

Draft minutes from the August 12 TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

- That the minutes of the August 12th meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

Introducing Kedro

Kedro at LF AI + Data

26th of August 2021

QUANTUMBLACK
A MCKINSEY COMPANY

Agenda

1. Why would we like to donate Kedro to LF?
2. Meet the team
3. What problem are we trying to solve?
4. What is Kedro?
5. Where does Kedro fit in the ecosystem?
6. What's coming next?
7. What is our impact to date?
8. A Kedro example

Why would we like to donate Kedro to LF?

Commitment to Open-Source

- A signal to the open-source community that the Kedro project is working towards a long-term maintenance model
- McKinsey and QuantumBlack would like to leverage the initial marketing announcements to build credibility in their technical and product-related capability
- Open governance and open roadmap allows more users to participate in the development of Kedro
- Neutral holding ground for Kedro's assets

Building a Standard

- We aim to create an open, de-facto standard around Machine-Learning Engineering and use Kedro as the vehicle for that, therefore we would like to:
  - Increase the Kedro user and contributor bases
  - Collaborate with complementary projects
- Become a thought-leader and established product in the MLOPs or Machine-Learning Engineering spaces

Meet the team

Framework

Ivan

Tech lead

Lorena

SWE

Merel

SWE

Antony

DS SWE

Ignacio

DS SWE

Jiri

DE SWE

Design

Gabriel

Visual Designer

Hamza

Design Research

Frontend

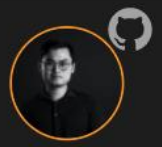

Lim

Viz tech lead

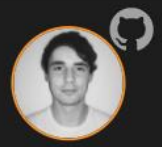

Liam

Front end

Susanna

Front end

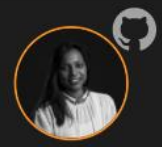

Rashida

Front end

Product

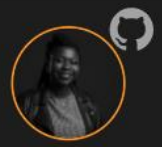

Yetu

Product Lead

Joel

Product

What problem are we trying to solve?

03

What are you trying to build?

Insights

Data science code that no one will use after your project is complete

What are you trying to build?

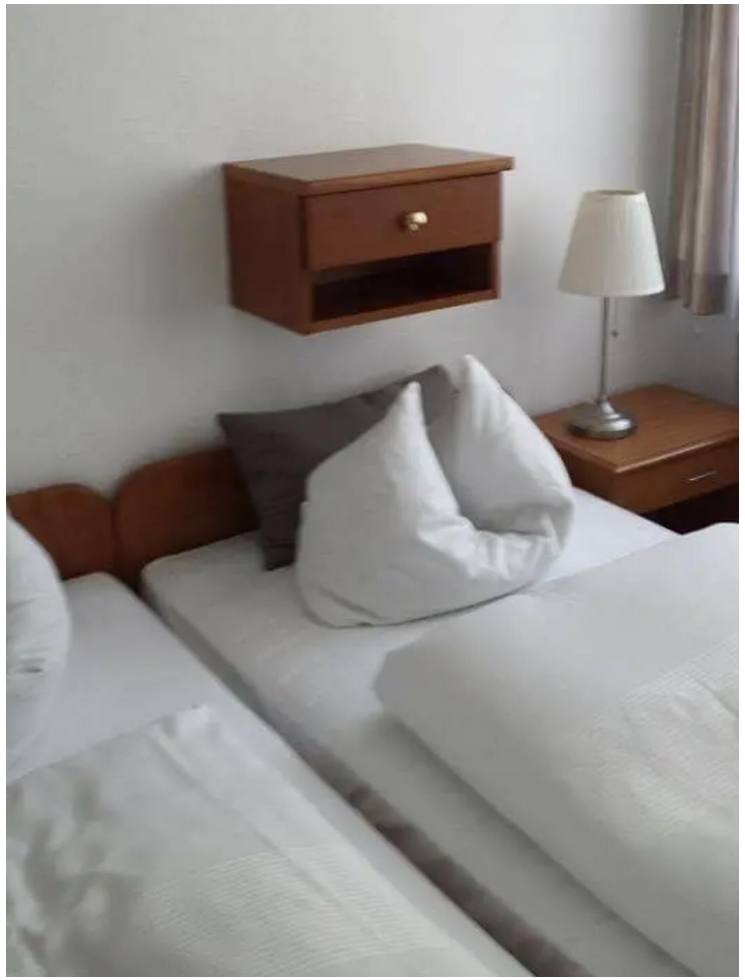

1. This clever bedroom storage option

2. This timely workaround

3. This energizing electrical hookup

Source: https://www.buzzfeed.com/philippjahner/craftsmanship-fails

What are you trying to build?

Machine Learning Product

Data science code that needs to be re-run and maintained

The challenges of creating machine learning products

The Jupyter notebook workflow has 5Cs of challenges

Challenge 1
Collaboration

Multi-user collaboration in a notebook is challenging to do because of the recommended one-person/one-notebook workflow.

Challenge 2
Code Reviews

Code reviews, the act of checking each other's code for mistakes, requires extensions of notebook capabilities. Often meaning, reviews are not done for code written in notebooks.

Challenge 3
Code Quality

Writing unit tests, documentation for the codebase and linting (like a grammar check for code) is not something that can be easily done in a notebook.

Challenge 4
Caching

The convenience of caching in a notebook sacrifices an accurate notebook execution flow leading you to believe that your code runs without errors.

Challenge 5
Consistency

Reproducibility in notebooks is challenge. A 2019 NYU study[1] executed 860k Notebooks found in 264k GitHub repositories. 24% of the notebooks completed without error; 4% produced the same results.

Source: 1. Pimentel, J., Murta, L., Braganholo, V. and Freire, J. (n.d.). *A Large-scale Study about Quality and Reproducibility of Jupyter Notebooks*. [online] Available at: http://www.ic.uff.br/~leomurta/papers/pimentel2019a.pdf [Accessed 23 Sep. 2020].

The challenges of creating machine learning products

A workflow beyond notebooks still has challenges

"Data scientists have to learn so many tools to create high-quality code."

"I have to think about Sphinx, flake8, isort, black, Cookiecutter Data Science, Docker, Python Logging, virtual environments, Pytest, configuration and more ."

"I spend a lot of time trying to understand a codebase that I didn't write."

"It's tedious to always setup documentation and code quality tooling my project."

"Everyone works in different ways."

"It takes really long to put code in production and we have to rewrite and restructure large parts of it."

"My code will not run on another person's machine."

"No one wants to use the framework I created."

"We all have different levels of exposure to software engineering best-practice."

What is Kedro?

04

What is Kedro?

It is developed and maintained by QuantumBlack; and, is McKinsey's first open-source product

Reproducible, maintainable and modular data science ***solved***

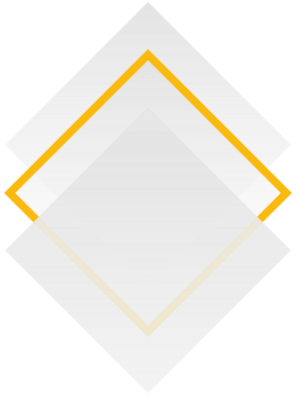

What is it?

- An open-source Python framework created for data scientists, data engineers and machine-learning engineers
- It borrows concepts from software engineering and applies them to machine-learning code; applied concepts include modularity, separation of concerns and versioning

Why do we use it?

- Addresses the main shortcomings of Jupyter notebooks, one-off scripts, and glue-code because there is a focus on creating maintainable data science code
- Increases the efficiency of an analytics team
- We use it to build reusable code stores like how React is used to build Design Systems

Impact on MLOPs?

- It won Best Technical Tool or Framework for AI in 2019 (Awards AI) and merit award for the Technical Documentation, is listed on the 2020 ThoughtWorks Technology Radar and the 2020 Data & AI Landscape
- Used at start-ups, major enterprises and in academia

Concepts in Kedro

Ships with a CLI and UI for visualizing data and ML pipelines

Project Template

A series of files and folders derived from Cookiecutter Data Science. Project setup consistency makes it easier for team members to collaborate with each other.

Configuration

Remove hard-coded variables from ML code so that it runs locally, in cloud or in production without major changes. Applies to data, parameters, credentials and logging.

The Catalog

An extensible collection of data, model or image connectors, available with a YAML or Code API, that borrow arguments from Pandas, Spark API and more.

Nodes & Pipelines

A pure Python function that has an input and an output. A pipeline is a directed acyclic graph, it is a collection of nodes with defined relationships and dependencies.

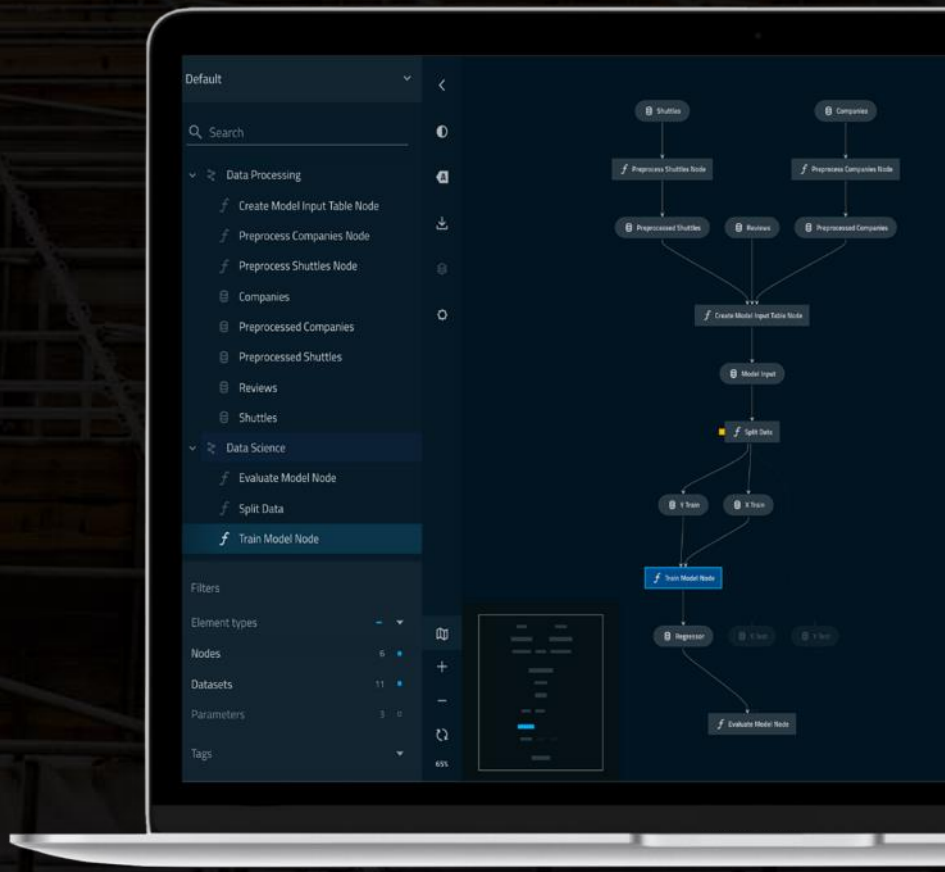

Project template

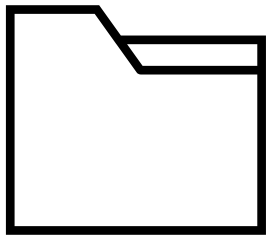

Configuration

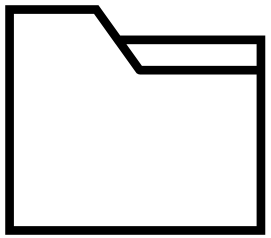

Notebooks

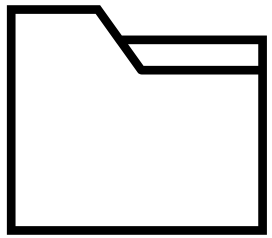

Logs

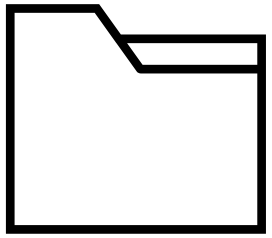

Python Script

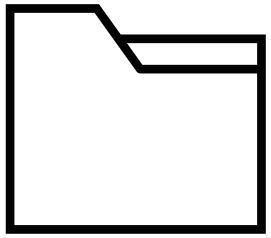

Tests

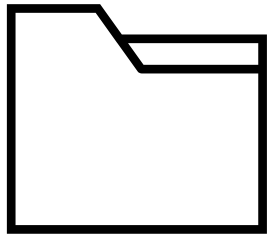

Project Documentation

What is the project template?

- A modifiable series of files and folders
- Built-in support for Python logging, Pytest for unit tests and Sphinx for documentation

What does the project template help you do?

- Spend time on documenting your ML approach and not on worrying how your project is structured
- You spend less time digging around in previous projects for useful code
- Makes it easier for collaborators to work with you

Configuration

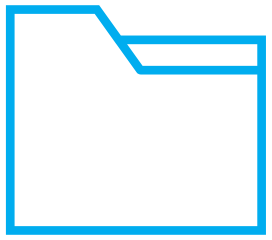

Configuration

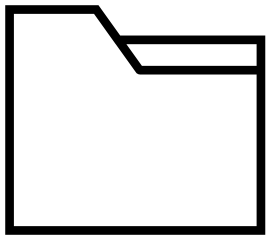

Notebooks

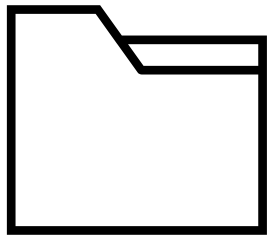

Logs

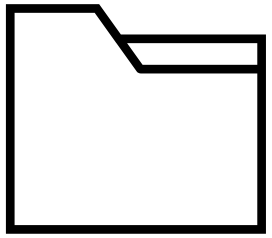

Python Script

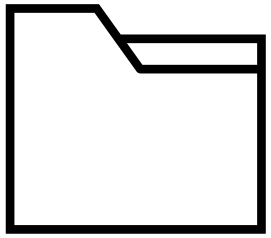

Tests

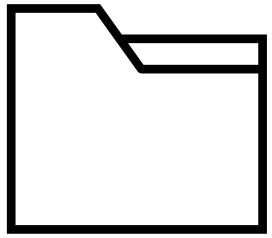

Project Documentation

What is configuration?

- “Settings” for your machine-learning code
- A way to define requirements for data, logging and parameters in different environments
- Helps keep credentials out of your code base
- Keep all parameters in one place

What does configuration help you do?

- Machine learning code that transitions from prototype to production with little effort
- Makes it possible to write generalizable and reusable analytics code that does not require significant modification to be used

The Catalog

Integrations in the Catalog

pandas	Pandas
Spark	Spark
DASK	Dask
SQLAlchemy	SQLAlchemy
NetworkX	NetworkX
matplotlib	MatplotLib
Google BigQuery	Google BigQuery
Google Cloud Storage	Google Cloud Storage
amazon REDSHIFT	AWS Redshift
amazon S3	AWS S3
Microsoft Azure Blob Storage	Azure Blob Storage
hadoop	Hadoop File System

What is the Catalog?

- Manages the loading and saving of your data
- Available as a code or YAML API
- Versioning is available for file-based systems every time the pipeline runs
- It's extensible, and we accept new data connectors

What does the Catalog help you do?

- Never write a single line of code that would read or write to a file, database or storage system
- Makes it possible to write generalizable and reusable analytics code that does not require significant modification to be used
- Access data without leaking credentials

Nodes & Pipelines

What is a node?

- A pure Python function that has an input and an output.
- Node definition supports multiple inputs for things like table joins and multiple outputs for things like producing a train/test split.

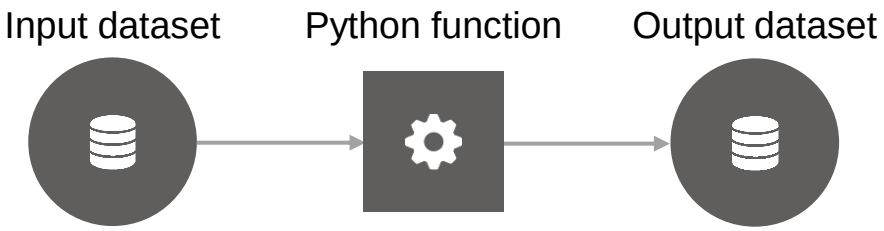

What is a pipeline?

- It is a directed acyclic graph.
- A collection of nodes with defined relationships and dependencies.

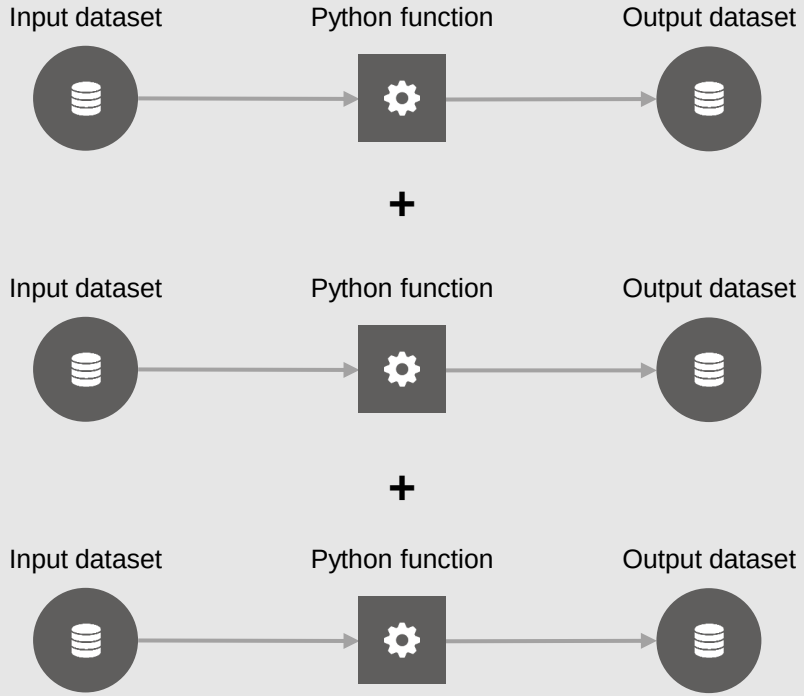

Pipeline visualization

Gives you x-ray vision into your project. You can see exactly how data flows through your pipelines. It is fully automated and focuses on your code base.

PLUGIN

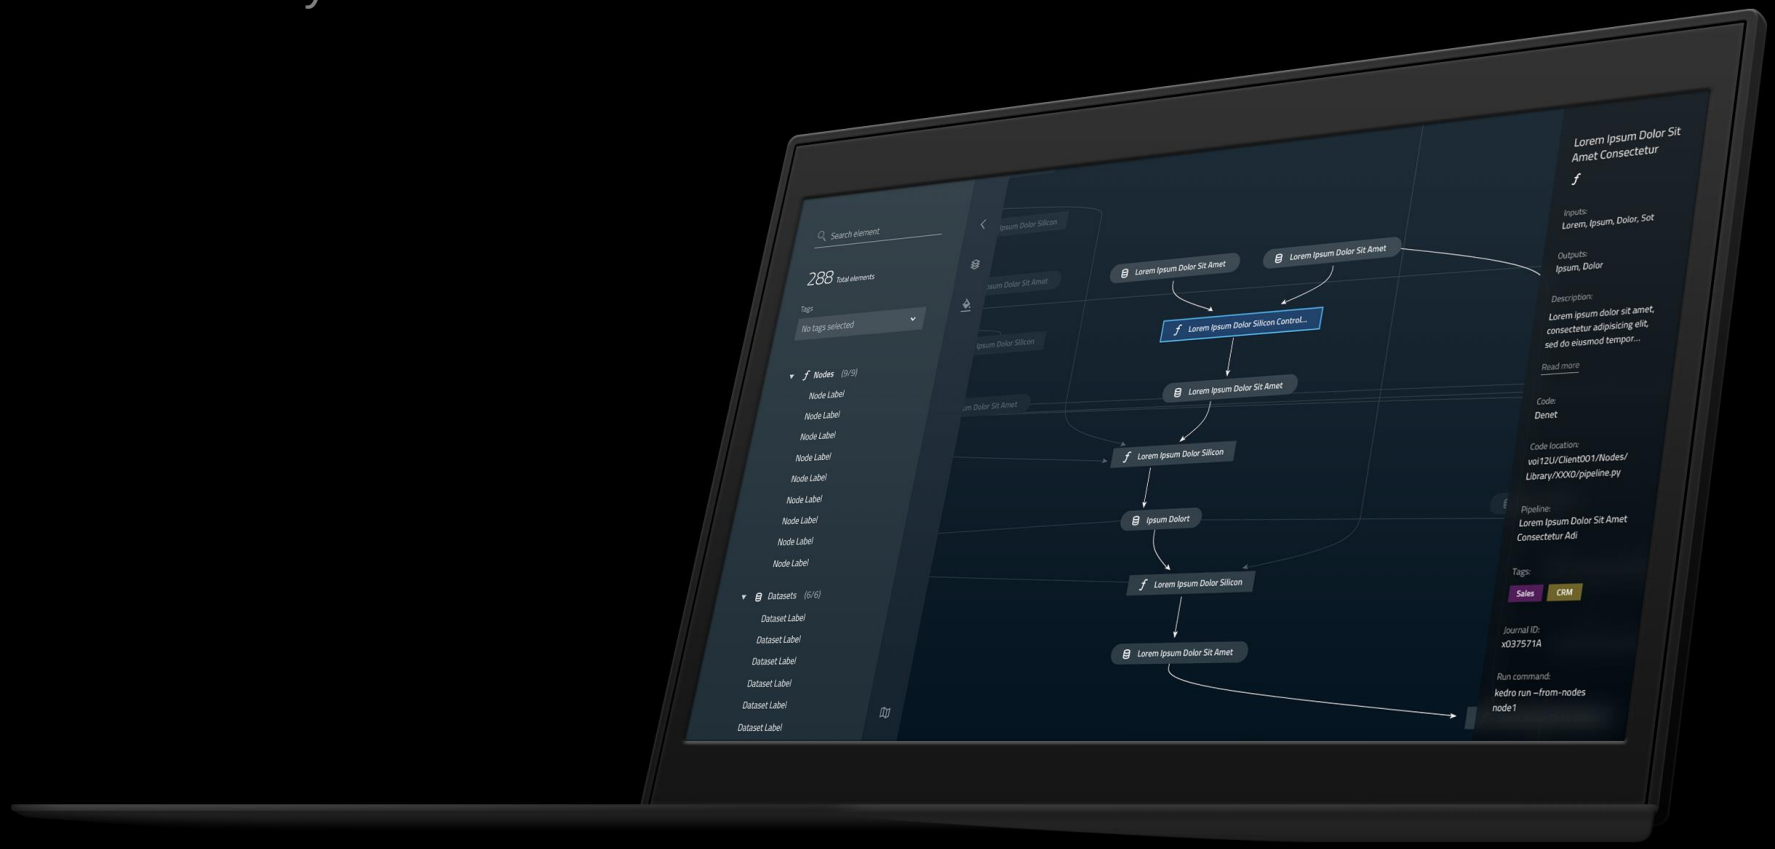

Demo: quantumblacklabs.github.io/kedro-viz/

Flexible deployment

Kedro supports multiple deployment modes

Deployment modes

Kedro currently supports:

- Single-machine deployment on a production server using:
  - A container based using Kedro-Docker
  - Packaging a pipeline using `kedro package`
  - A CLI-based approach using the Kedro CLI
- Distributed application deployment allowing a Kedro pipeline to run on multiple computers within a network at the same time

Supporting a range of tools

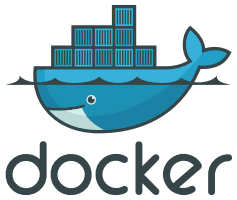

AWS Batch

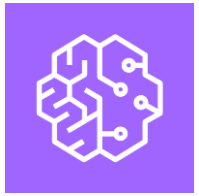

Amazon SageMaker

Kedro is actively maintained by QuantumBlack

We are committed to growing community and making sure that our users are supported for their standard and advanced use cases.

The Kedro community is active on:
https://discord.gg/7sTm3y5kKu

Documentation is available on Kedro's Read The Docs: https://kedro.readthedocs.io/

OUR SUPPORT CHANNELS

The Kedro community is active on:
https://github.com/quantumblacklabs/kedro/
The team and contributors actively maintain raised feature requests, bug reports and pull requests.

Where does Kedro fit in the ecosystem?

05

Where does Kedro fit in the ecosystem?

Kedro is the *scaffolding* that helps you develop a data and machine-learning pipeline that can be deployed

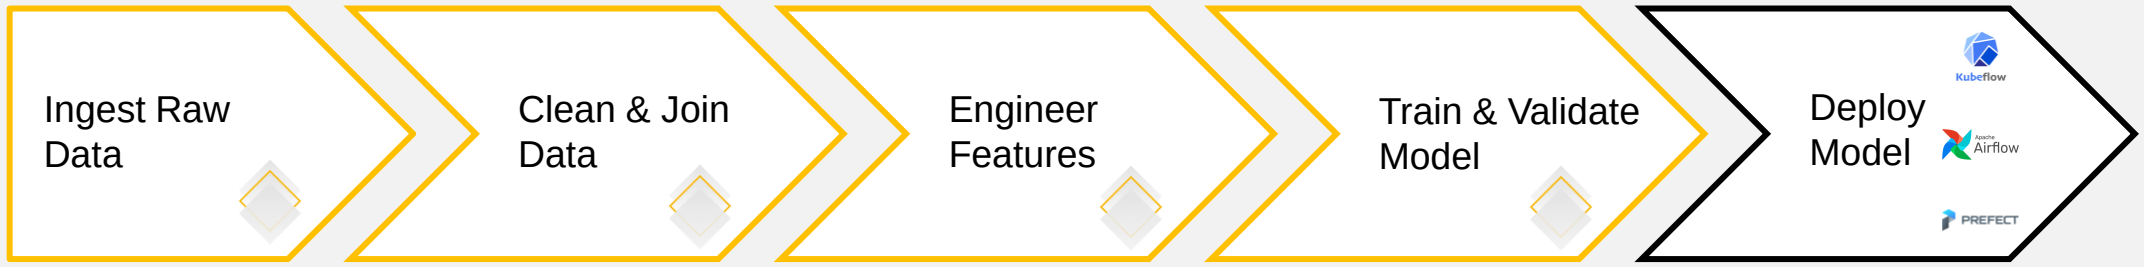

Philosophy of Kedro

- Kedro focuses on how you work while writing standardized, modular, maintainable and reproducible data science code and does not focus on how you would like to run it in production
- The responsibility of *"What time will this pipeline run?"* and *"How will I know if it failed?"* is left to tools called orchestrators like Apache Airflow, Kubeflow and Prefect
- Orchestrators do not focus on the process of producing something that could be deployed, which is what Kedro does

Competitive Space

Tool	Focus	Project Template	Data Catalog	DAG Workflow	DAG UI	Experiment Tracking	Data Versioning	Scheduler	Monitoring
Kedro	“Kedro is an open-source Python framework for creating reproducible, maintainable and modular data science code.”	✓	✓	✓ * Datacentric DAG	✓	* Coming soon!	✓ * Basic feature	✓ * With integration	✓ * With integration
ZenML	“ZenML is an extensible, open-source MLOps framework for using production-ready Machine Learning pipelines, in a simple way."			✓ * Taskcentric DAG	✓	✓	✓	✓	
COOKIECUTTER	“A logical, reasonably standardized, but flexible project structure for doing and sharing data science work.”	✓							
mlflow	“MLflow is an open source platform to manage the ML lifecycle, including experimentation, reproducibility, deployment, and a central model registry.”			✓ * Basic feature		✓	✓ * Models		✓
INTAKE	“Data catalogs provide an abstraction that allows you to externally define, and optionally share, descriptions of datasets, called catalog entries.”		✓						
	Orchestration platforms allow users to author, schedule and monitor workflows task-centric data workflows.			✓ * Taskcentric DAG	✓			✓	✓
DVC	“DVC is built to make ML models shareable and reproducible. Designed to handle data sets, machine learning models, and metrics as well as code.”						✓		
Pachyderm version-controlled data science	“Pachyderm is the data layer that powers your machine learning lifecycle. Automate and unify your MLOps tool chain. With automatic data versioning and data driven pipelines.”			✓	✓		✓		

What’s coming next?

06

Current and upcoming focuses

Current focus

- Experiment tracking
- Micro-packaging workflow

Upcoming focus

- Scaling configuration
- Interactive Jupyter notebook workflow
- Deployment plugin

What is our impact to date?

07

Who is using Kedro?

There are Kedro users across the world, who work at start-ups, major enterprises and academic institutions

6,500+

Monthly Active Users of the Kedro Documentation

4,000+

GitHub stars

100+

GitHub Contributors

Impact tracking

Impact Tracking; count

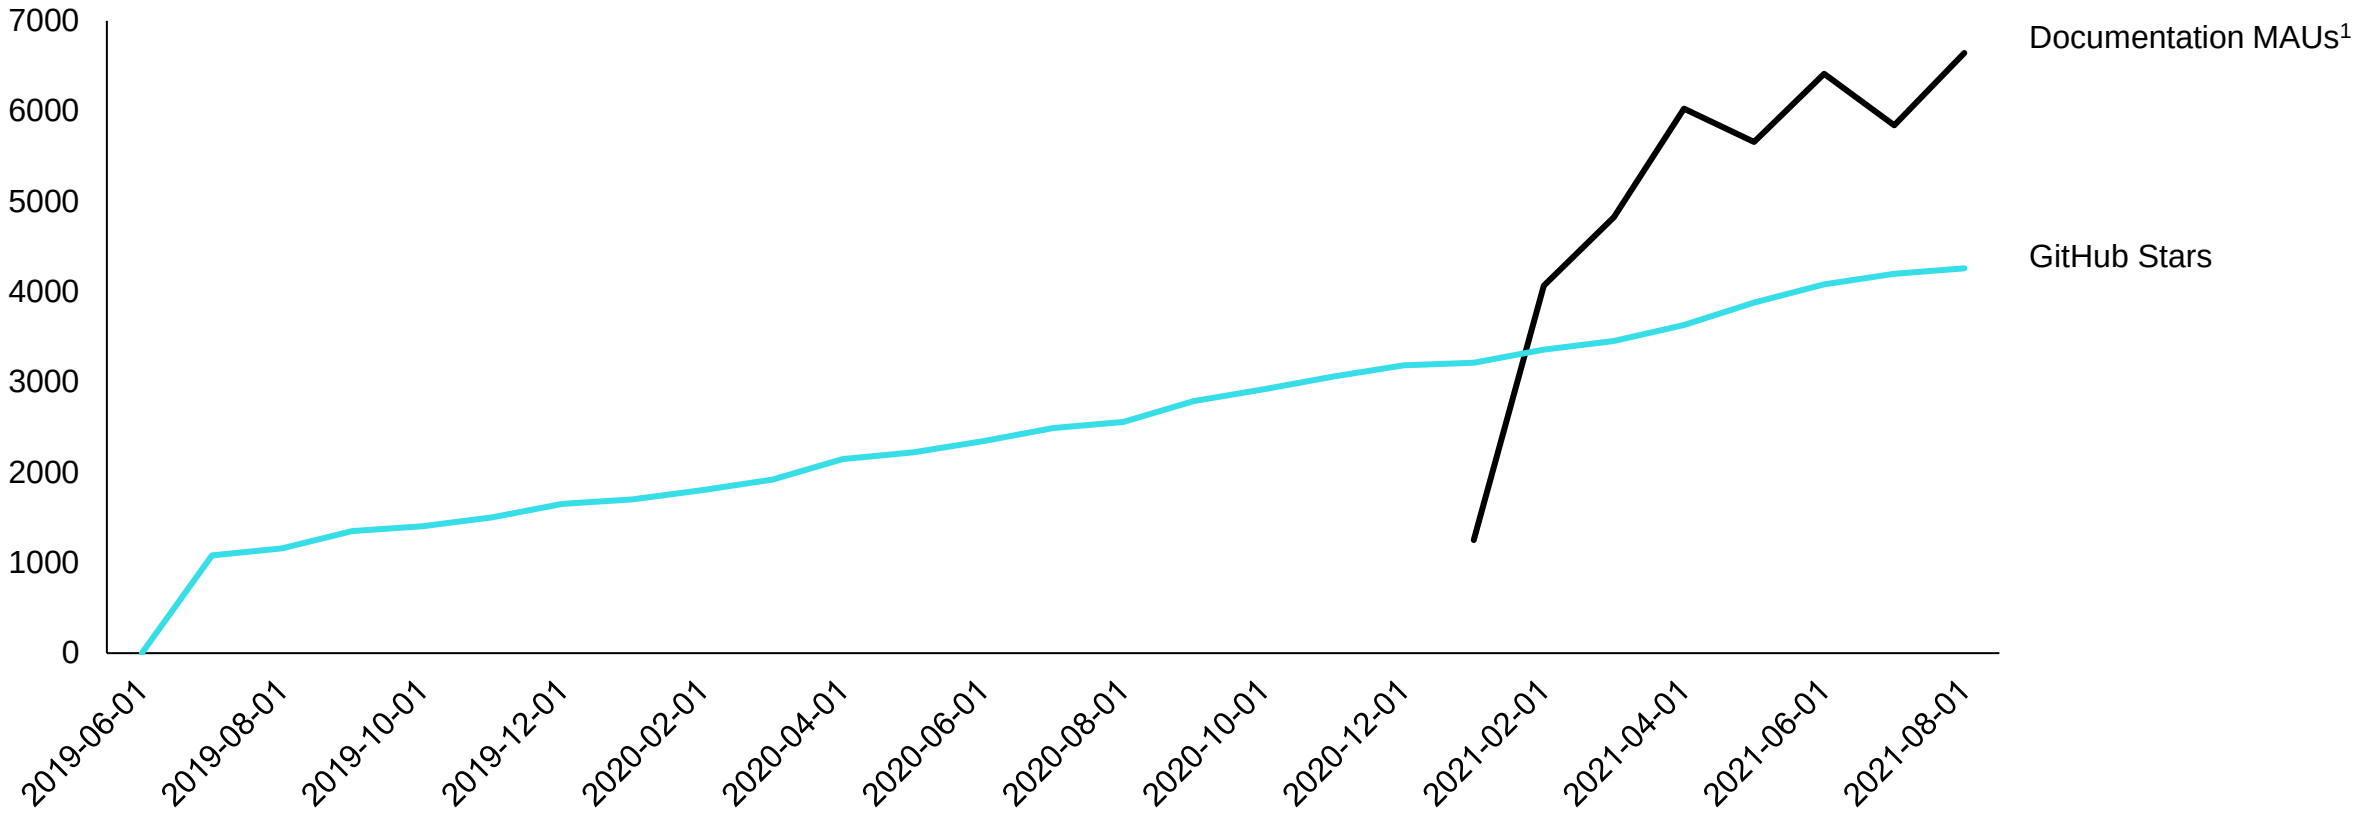

1. Analytics on the Kedro documentation was implemented in December 2020

A Kedro example

An excerpt from the Spaceflight tutorial in the Kedro documentation

08

The Blueprint view

The Data Processing Pipeline of Kedro-Viz

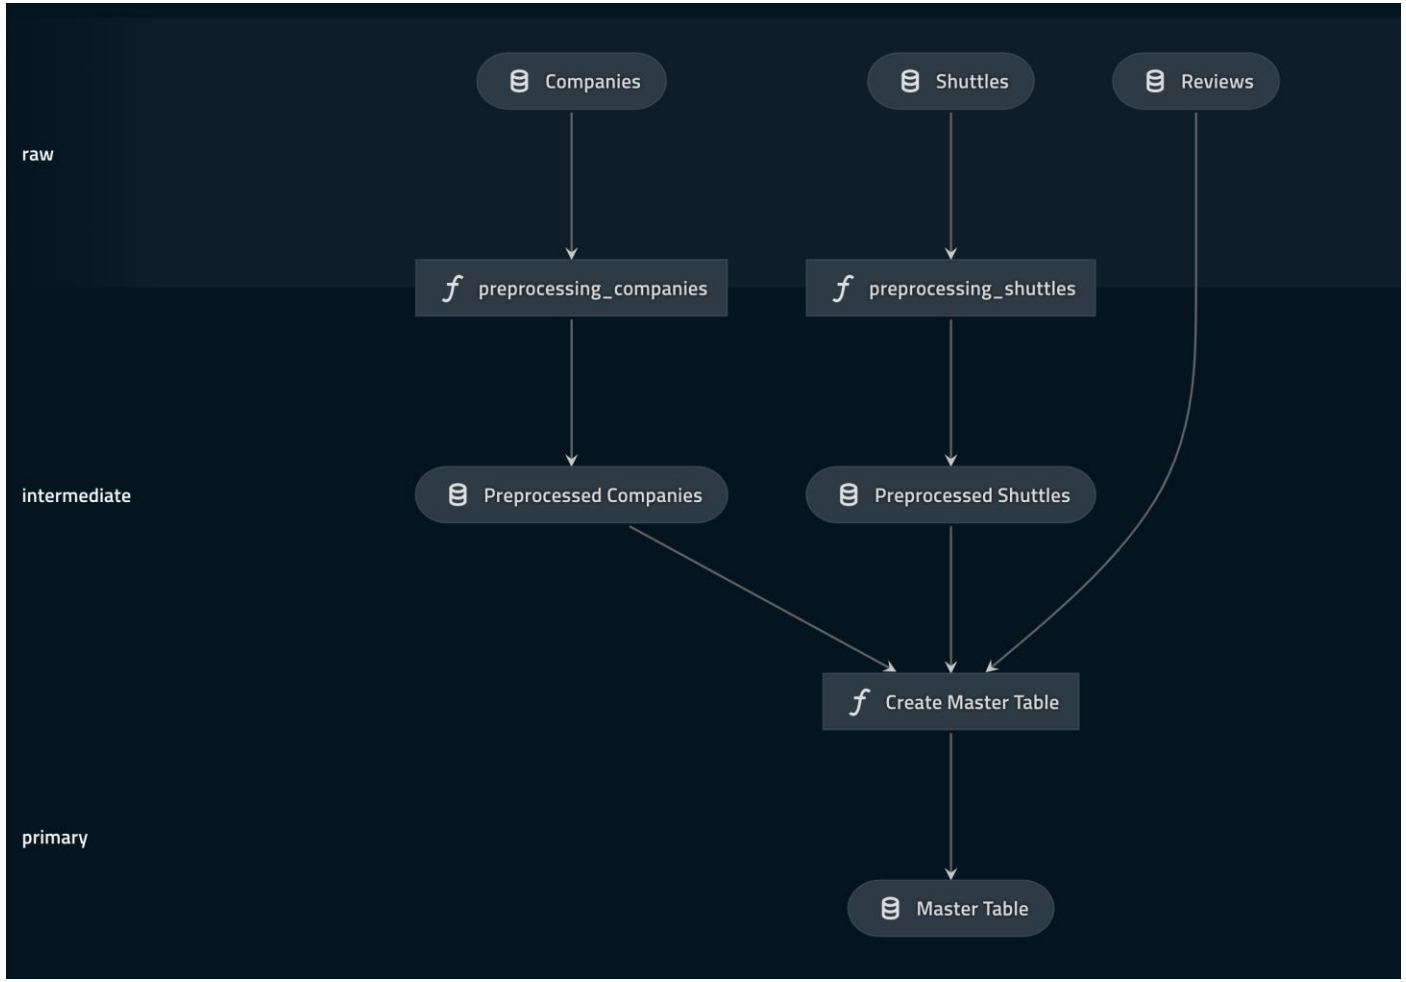


```
kedro viz
```

```
2020-10-12 13:47:00,906 - werkzeug - INFO
-  * Running on http://127.0.0.1:4141/
(Press CTRL+C to quit)


2020-10-12 13:47:01,296 - werkzeug - INFO
- 127.0.0.1 - - [12/Oct/2020 13:47:01]
"GET /api/main HTTP/1.1" 200 -


2020-10-12 13:47:01,478 - werkzeug - INFO
- 127.0.0.1 - - [12/Oct/2020 13:47:01]
"GET /manifest.json HTTP/1.1" 304 -


2020-10-12 13:47:03,111 - werkzeug - INFO
- 127.0.0.1 - - [12/Oct/2020 13:47:03]
"GET /service-worker.js HTTP/1.1" 304 -
```

The Blueprint view

The Data Processing Pipeline of Kedro-Viz

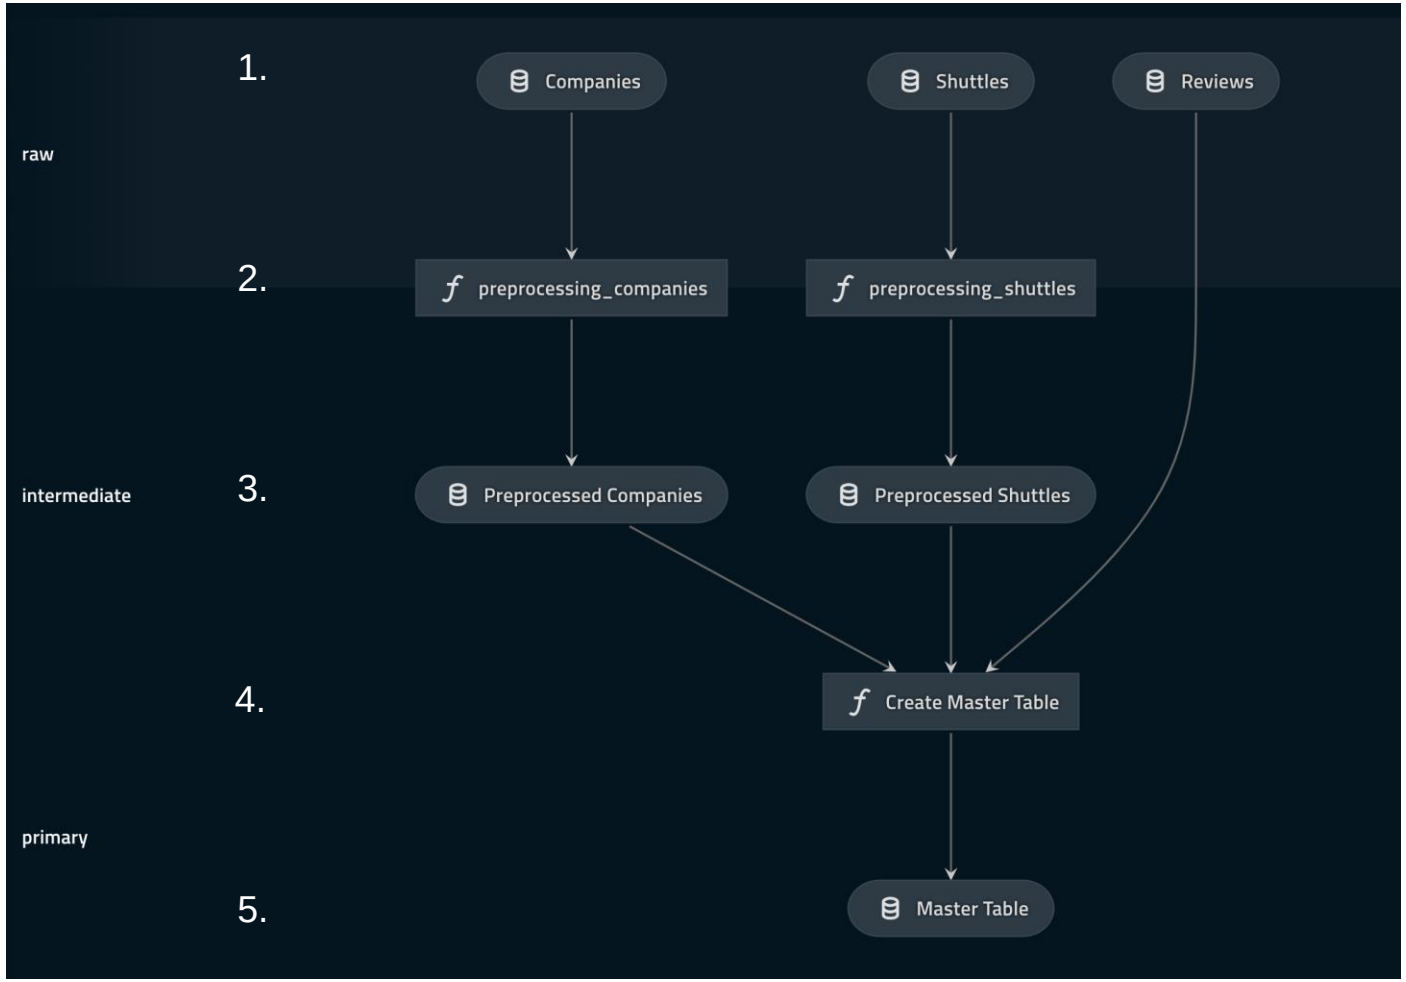

What happens in this data pipeline?

1. Input three tables which have data from *Companies*, *Space Shuttles* and *Customer Reviews*
2. Pre-process the *Companies* and *Shuttles* data
3. Output the processed the *Companies* and *Space Shuttles* tables
4. Join the tables to form a master table, using the processed *Companies* and *Space Shuttles* tables and the *Reviews* table
5. Output a master table

Adding entries to the Catalog

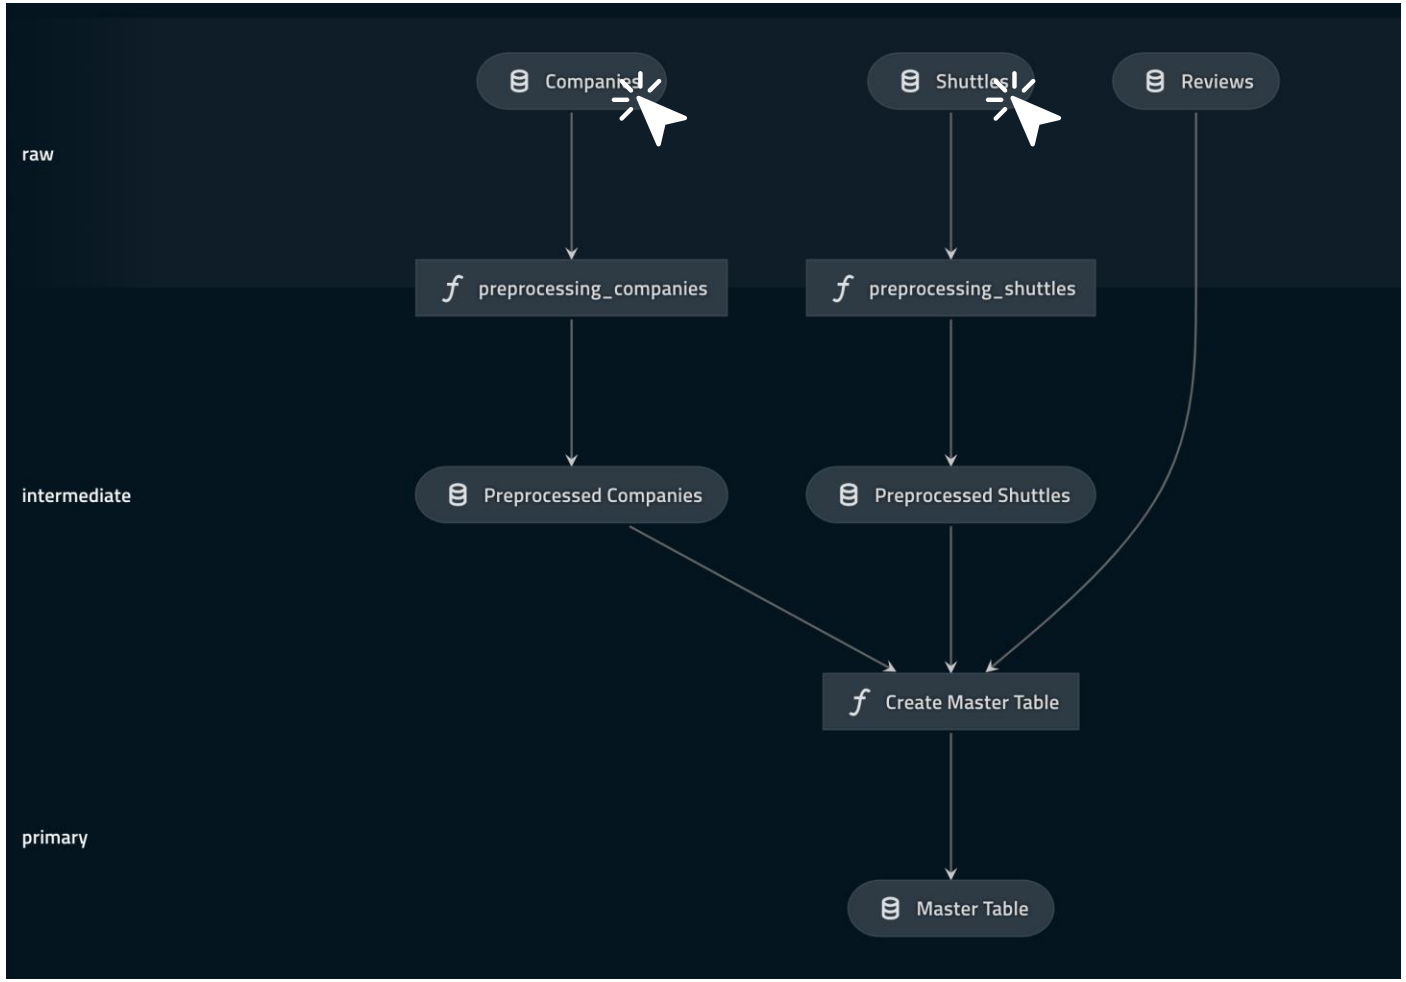


```
# catalog.yml

companies:
  type: pandas.CSVDataSet
  filepath: data/01_raw/companies.csv
  layer: raw

shuttles:
  type: pandas.ExcelDataSet
  filepath: data/01_raw/shuttles.xlsx
  layer: raw
```

Creating a node

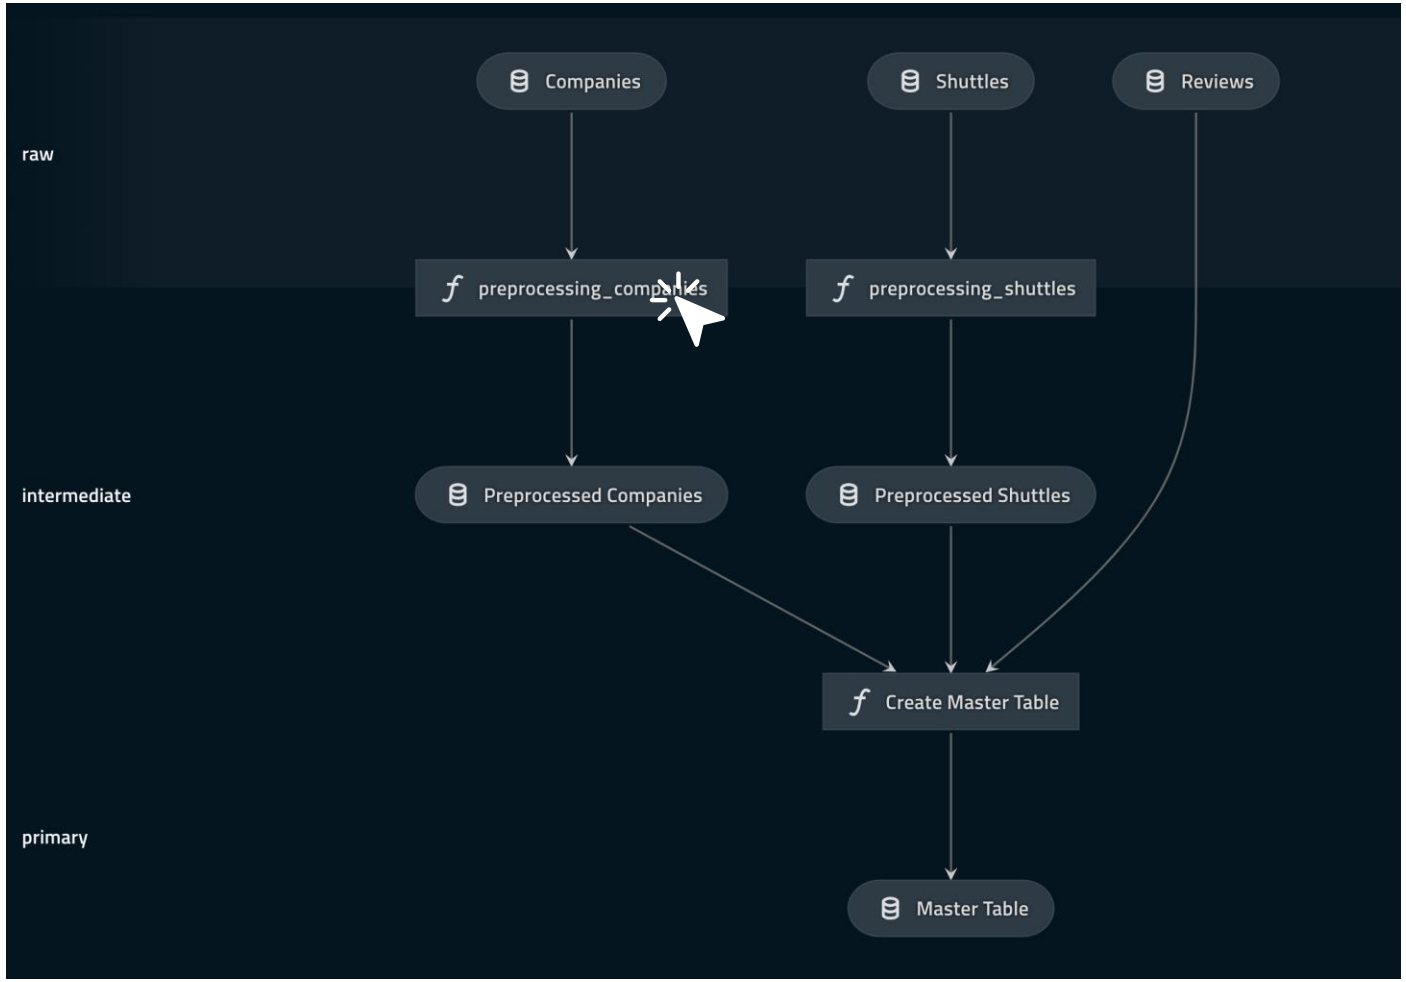


```python
# nodes.py

def preprocess_companies(df):

    df["company_rating"] =
df["company_rating"].apply(_parse_percentage)

    return df
```

Creating a pipeline

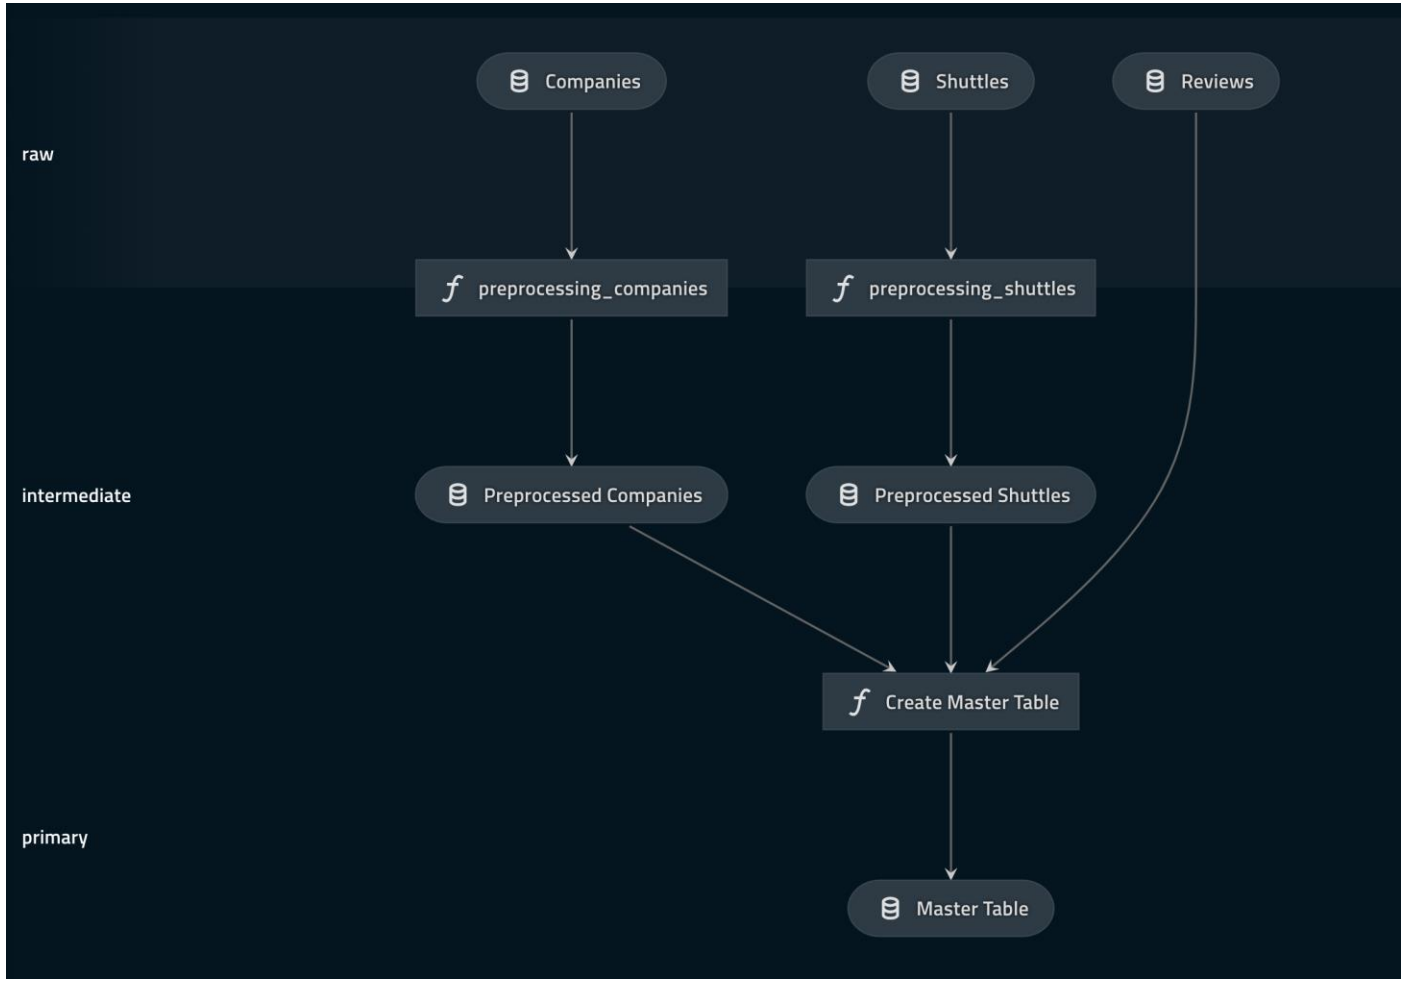


```
# pipelines.py

def create_pipeline(**kwargs):
    return Pipeline([
            node(
                func=preprocess_companies,
                inputs="companies",
                outputs="preprocessed_companies"),
            node(
                func=preprocess_shuttles,
                inputs="shuttles",
                outputs="preprocessed_shuttles"),
            node(
                func=create_master_table,
                inputs=["preprocessed_shuttles",
"preprocessed_companies", "reviews"],
                outputs="master_table"),
        ])
```

Running a pipeline

```
kedro run
```

```
2020-10-12 14:58:02,801 - kedro.io.data_catalog - INFO - Loading data from `shuttles` (ExcelDataSet)...
2020-10-12 14:58:11,592 - kedro.pipeline.node - INFO - Running node: preprocessing_shuttles: preprocess_shuttles([shuttles]) -
> [preprocessed_shuttles]
2020-10-12 14:58:11,663 - kedro.io.data_catalog - INFO - Saving data to `preprocessed_shuttles` (CSVDataSet)...
2020-10-12 14:58:12,085 - kedro.runner.sequential_runner - INFO - Completed 1 out of 3 tasks
2020-10-12 14:58:12,085 - kedro.io.data_catalog - INFO - Loading data from `companies` (CSVDataSet)...
2020-10-12 14:58:12,118 - kedro.pipeline.node - INFO - Running node: preprocessing_companies:
preprocess_companies([companies]) -> [preprocessed_companies]
2020-10-12 14:58:12,162 - kedro.io.data_catalog - INFO - Saving data to `preprocessed_companies` (CSVDataSet)...
2020-10-12 14:58:12,387 - kedro.runner.sequential_runner - INFO - Completed 2 out of 3 tasks
2020-10-12 14:58:12,388 - kedro.io.data_catalog - INFO - Loading data from `preprocessed_shuttles` (CSVDataSet)...
2020-10-12 14:58:12,489 - kedro.io.data_catalog - INFO - Loading data from `preprocessed_companies` (CSVDataSet)...
2020-10-12 14:58:12,521 - kedro.io.data_catalog - INFO - Loading data from `reviews` (CSVDataSet)...
2020-10-12 14:58:12,577 - kedro.pipeline.node - INFO - Running node:
create_master_table([preprocessed_companies,preprocessed_shuttles,reviews]) -> [master_table]
2020-10-12 14:58:14,530 - kedro.io.data_catalog - INFO - Saving data to `master_table` (CSVDataSet)...
2020-10-12 14:58:23,737 - kedro.runner.sequential_runner - INFO - Completed 3 out of 3 tasks
```

TAC Open Discussion

LF AI & DATA

LF AI & Data - Ongoing Annual Project Reviews

LF AI & DATA

Date	Project	Presenter
April 6, 2021	Egeria	Mandy Chessell (slack) - TAC recording / deck
April 6, 2021	OpenDS4all	Andre de Waal (slack) - TAC recording / deck
May 20, 2021	ONNX	Jim Spohrer (slack) - TAC recording / deck
July 15, 2021	EDL	Ti Zhou (slack) deck
July 29, 2021	Angel	Bruce Tao (slack) (confirmed) deck
July 29, 2021	Adlik	Meng Wei (slack) (confirmed) deck
Aug 12, 2021 (potentially Aug 12)	Sparklyr	Sigrid Keydana Yitao Li (slack) (confirmed) deck
Aug 12, 2021	Milvus	Jun Gu (slack) (confirmed)
Aug 26, 2021	Kendro new project into incubation	Yetunde Dada <yetunde_dada@mckinsey.com>
Sept 9, 2021	Marquez	Julien le Dem (slack) (confirmed)
Sept 9, 2021	Acumos	Amit Kumar (slack) (tentative)
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)
Jan 6, 2021	Datashim	Yiannis Gkoufas (to be asked)
Jan 6, 2022	Flyte	Ketan Umare (to be asked)
Jan 20, 2022	RosaeNLG	Ludan Stoeckle (to be asked)
Jan 20, 2022	SubstraFramework	Camille Marini (to be asked)
	Machine Learning Exchange	Animesh Singh (to be asked)
	VulcanKompute	Alejandro Saucedo (to be asked)
	OpenLineage	Julien le Dem (to be asked)
	MARS	Chris Qin (to be asked)

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

LF AI & Data - New Updates

LF AI & DATA

Upcoming TAC Meetings

LF AI & DATA

Upcoming TAC Meetings (Tentative)

- Sept 9: Marquez Annual Project review; Acumos Annual project review; LF Outreach Committee update and review
- Sept 23: NNStreamer Annual Project review; ForestFlow Annual project review

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