


 LF AI & DATA

2024

LF AI & Data Foundation Annual Report

lfai.data.foundation



Contents

Reflections on 2024 and Strategic Vision for 2025	3
2024 By the Numbers	4
Project Collaboration Report	5
Shaping the Future of Generative AI Survey	12
Generative AI Commons Committee.....	12
From the TAC Chair	13
2024 Event Highlights	13
Community Engagement.....	19
LF AI & Data Committees	20
Blog & Resources.....	22
Looking Ahead to 2025.....	23

2024 Key Highlights

11 NEW PROJECTS

Welcomed 11 new projects in incubation supporting their growth and sustainability. (total of **67 projects** in incubation)



11 NEW MEMBERS

Welcomed 11 new members including board members Neural Magic, Nvidia, Microsoft & Databricks. (total of **75 members**)

MOF

Published the Model Openness Framework and launched the Model Openness Tool. [<https://isitopen.ai/>]



LF EVENTS

Hosted three AI_dev Conferences and supported dozens of project and community meetups

REGIONAL USER GROUPS

Launched the APAC Regional UserGroup (RUG) and Japan RUG supporting our local communities and geographies.



Reflections on 2024 and Strategic Vision for 2025



By **Ibrahim Haddad, PhD.**,
General Manager, LF AI & Data

As we close the chapter on 2024, I am proud to reflect on the significant progress LF AI & Data has made in advancing the open source AI and data ecosystem. Our organization continues to serve as a neutral host and collaboration facilitator, providing the infrastructure, credibility, and expertise that empower a global community of developers, organizations, and end users to innovate together.

This year marked a key milestone with the publication of our report, [Shaping the Future of Generative AI](#). This comprehensive study explored the transformative potential of generative AI, outlining the critical challenges and opportunities ahead. As a trusted leader in open source innovation, LF AI & Data tackles these challenges through collaborative, transparent, and community-driven initiatives.

Our projects have addressed key issues across industries, from enhancing interoperability in edge computing to advancing ethical AI frameworks. These efforts reflect our ongoing commitment to open collaboration and delivering impactful solutions that benefit all participants.

Open source plays a pivotal role in expanding the capabilities of Generative AI by fostering global collaboration, sparking innovation, and making advanced technology more accessible. With open source, developers worldwide can share their expertise, driving rapid progress and continuous improvements

in GenAI frameworks. At LF AI & Data, we have over **100,000** developers contributing to **67** hosted projects from more than **3,000** organizations.

We remain dedicated to open source principles, interoperability, and building trust within the AI ecosystem. Our rich contributor programs, neutral IP zone, and collaborative partnerships position us to lead purposefully and deliver meaningful impact across sectors.

The insights from [Shaping the Future of Generative AI](#) will continue to guide our efforts as we expand tools for generative AI, strengthen edge computing capabilities, advance global ethical AI standards, and champion sustainable AI practices. This report serves as both a roadmap and a call to action for all AI and data ecosystem stakeholders.

As we look to the future, our collective efforts will shape the next phase of AI and data innovation while ensuring its inclusive and far-reaching benefits.

Thank you for being part of this journey. Let's make 2025 another year of transformative progress for LF AI & Data and the broader AI community.

“Organizations with higher levels of GenAI adoption are helping to shape next-generation frameworks and models, aligning them more closely with advanced, real-world use cases.”

2024 By the Numbers



3 Projects Graduated in 2024

docarray Bitol kedro

LF AI & DATA

Project Collaboration Report

The LF AI & Data Foundation is pleased to present the [LF AI & Data Project Collaboration Report](#), offering an overview of the synergies and integrations across our hosted projects. As a neutral host and collaboration facilitator, LF AI & Data brings together top developers, organizations, end users, and members to address key AI and data technology challenges.

This approach drives impactful solutions and ensures that the entire ecosystem benefits from a shared commitment to advancing AI and data technologies.


We invite you to explore how LF AI & Data is shaping the future of AI through impactful, community-driven collaboration.

“Intel is at the forefront of incubating open source development to build trusted, scalable open infrastructure that enables heterogeneity and provides a platform for developer innovation. Generative AI is at this moment; OPEA, with the support of the broader community, will address critical pain points of RAG adoption and scale today. It will also define a platform for the next phases of developer innovation that harnesses the potential value generative AI can bring to enterprises and all our lives.”

- MELISSA EVERS, VICE PRESIDENT OF SOFTWARE ENGINEERING GROUP AND GENERAL MANAGER OF STRATEGY TO EXECUTION, *INTEL*

The graphic features a white background with a blue and cyan header. At the top left, the word 'PROJECTS' is written in blue. Below it is the Egeria logo, which consists of a stylized profile of a face with a large eye, flanked by two blue chevrons pointing outwards, followed by the word 'EGERIA' in a bold, black, sans-serif font. Below the logo, there is a paragraph of text describing Egeria as the world's first open-source metadata standard, offering open APIs, event formats, types, and integration logic. It also mentions that Egeria was originally contributed to ODPi by IBM and ING Bank in August 2018 and became a Graduate Project in November 2020. Below this text, there is a section titled 'Egeria actively collaborates with Open Lineage, Unity Catalog, and Marquez to enhance data management and governance in the following ways:'. This section is divided into three columns, each featuring a logo and a brief description of the integration. The first column features the Open Lineage logo (a blue arrow pointing right) and describes how Egeria supports sending, receiving, storing, augmenting, and distributing open lineage events. The second column features the Unity Catalog logo (a colorful hexagon) and describes how Egeria surveys Unity Catalog content to create resource summaries and catalogs data resources. The third column features the Marquez logo (a blue and red network diagram) and describes how Egeria integrates with Marquez through Open Lineage. At the bottom of the graphic, there is a footer with the text 'LF AI & DATA PROJECT COLLABORATION' on the left and the number '2' on the right.


PROJECTS


 EGERIA


Egeria is the world's first open-source metadata standard. It offers open APIs, event formats, types, and integration logic, enabling enterprise-wide data management and governance without reformatting or restricting data to a specific format, platform, or vendor.

Originally contributed to ODPi by IBM and ING Bank in August 2018, Egeria became a Graduate Project in November 2020.

Egeria actively collaborates with Open Lineage, Unity Catalog, and Marquez to enhance data management and governance in the following ways:

 **OpenLineage**
Egeria supports sending, receiving, storing, augmenting, and distributing open lineage events. When running governance processes, it generates events to enhance lineage visibility.

 **Unity Catalog**
Egeria surveys Unity Catalog content to create resource summaries, catalogs data resources (enabling search and dependency lineage across instances), and shares this information with other tools. It can also provision new resources into Unity Catalog as part of governance workflows.

 **MARQUEZ**
Egeria integrates with Marquez through Open Lineage, enhancing metadata tracking and interoperability.

These integrations are maintained within the Egeria project, leveraging Egeria's integration frameworks to streamline interoperability with these systems and provide a unified approach to data governance.

LF AI & DATA PROJECT COLLABORATION 2


PROJECTS



Elyra is an open-source low-code/no-code framework for building reproducible, scalable, component-based data science pipelines. It enables senior data scientists to create reusable components, allows citizen data scientists to reuse code without programming, and provides MLOps engineers with tested, scalable deliverables on platforms like Kubeflow and Airflow.

Contributed by: IBM in October 2022

Elyra integrates with Claimed and the Machine Learning Xchange (MLX) to provide data scientists with easy access to curated, reusable components, streamlining the development and deployment of machine learning workflows



Claimed

Elyra collaborates actively with Claimed components, allowing users to incorporate pre-built, reusable elements into their data science pipelines. These components are available through the Machine Learning Exchange (MLX), a repository of curated machine learning assets.



MLX

Elyra simplifies workflow integration by discovering claimed components from MLX, streamlining project development, promoting collaboration, and accelerating machine learning pipeline creation and deployment.

By leveraging claimed components, Elyra not only promotes reusability and efficiency but also ensures that teams can stay up-to-date with the latest advancements in machine learning through access to a dynamic exchange of resources.

PROJECTS



Flyte is a scalable, cloud-native workflow orchestration platform for ML and data pipelines in Python and Java. It manages data flow, parallelization, scaling, and orchestration on Kubernetes with Docker. Flyte is a graduated project of the LF AI & Data Foundation.

Contributed by: Lyft in February 2021 as an incubation-stage project and graduated in December 2021.

These integrations help Flyte users streamline machine learning and data engineering pipelines, ensuring scalability, reliability, and ease of use across various environments.



MLflow

Flyte can integrate with MLflow to manage model tracking, experiment logging, and versioning within workflows, supporting the complete ML lifecycle.



Kubernetes

Flyte itself runs on Kubernetes and can manage workflows within containers, allowing for containerized, scalable, and reproducible ML pipelines.

Flyte integrates with a range of open-source projects, such as Apache Spark to streamline machine learning and data engineering pipelines, supporting scalable workflows on Kubernetes, that enable distributed data processing, and managing the entire ML lifecycle with tracking and observability tools like Prometheus and Grafana.

PROJECTS



Horovod, developed by Uber, simplifies distributed deep learning by reducing model training from days to hours. Horovod improves the speed, scale, and resource utilization of deep learning training. Horovod works with on-premises and cloud platforms like AWS, Azure, Databricks, and integrates with Apache Spark to unify data processing and model training.

Horovod is a graduation-stage project of the LF AI & Data Foundation. Contributed by: Uber in December 2018 as an incubation-stage project and graduated in September 2020.

Horovod integrates with Kubernetes and PyTorch to streamline distributed training, enabling scalable deployment and efficient orchestration of multi-node, multi-GPU machine learning models in the following ways:



kubernetes

Kubernetes

Kubernetes integration enables scalable, containerized deployment of Horovod in distributed clusters, making it easier to manage multi-node, multi-GPU training workloads in cloud and hybrid environments.



PyTorch

Works seamlessly with the PyTorch framework in addition to TensorFlow and MXNet, allowing users to scale their training across multiple GPUs and nodes with minimal code changes.

These integrations make Horovod versatile, enabling users to build robust, scalable, and flexible deep learning pipelines across various platforms and environments.

PROJECTS



kedro

Kedro is an open-source Python framework for creating reproducible, maintainable and modular data science code. It borrows concepts from software engineering best-practice and applies them to machine-learning code; applied concepts include modularity, separation of concerns and versioning.

Kedro is an incubation-stage project of the LF AI & Data Foundation. Contributed by: McKinsey and QuantumBlack in August 2021

Kedro integrates with Delta Lake and ONNX by enabling data pipelines that leverage ONNX models for deployment and inference while using Delta Lake for reliable, scalable, and version-controlled data storage, ensuring consistent and efficient workflows for machine learning and data processing.



Delta Lake

Kedro integrates with Delta Lake by enabling seamless data pipeline orchestration with Delta Lake's robust storage, ensuring reliable, scalable, and version-controlled data workflows while leveraging Delta Lake's ACID transactions for data consistency and quality.



ONNX

Kedro integrates with ONNX by allowing data scientists and engineers to incorporate ONNX models into Kedro pipelines, enabling smooth deployment and inference of machine learning models across various platforms and runtimes.

PROJECTS



ONNX is an open format for representing deep learning models, allowing AI developers to easily move models between tools and choose the best combination for their needs. Created by Facebook and Microsoft, ONNX enables interoperability between frameworks, eliminating the lock-in to a single ecosystem.

The project aims to help developers select the right tools for their projects, accelerating AI from research to reality. Contributed by the ONNX community and its partners in October 2019.

ONNX serves as an intermediary format that enables interoperability between machine learning frameworks like PyTorch and tools like the Adversarial Robustness Toolbox (ART)



PyTorch

PyTorch: ONNX is natively supported by PyTorch, enabling seamless export of PyTorch models to the ONNX format for cross-platform compatibility.



Adversarial Robustness Toolbox

ART supports models in ONNX format, making it easier to evaluate robustness and conduct adversarial training across frameworks that export to ONNX.

ONNX's integrations across these tools and frameworks enable interoperability, making it easier to build, train, and deploy models across diverse hardware and software environments.

PROJECTS



OpenLineage proposes an open standard and API for lineage collection that data processing engines can implement to publish at run time details of the data sources that it is reading, the types of processing it is performing and the destination of the results. OpenLineage is a graduation-stage project of the LF AI & Data Foundation.

Contributed by: Datakin in May 2021 as an incubation-stage project and graduated in July 2023.

OpenLineage integrates with Marquez and Kubernetes to track, monitor, and visualize data lineage within Kubernetes-managed workflows, enhancing data pipeline transparency and reliability



Marquez

OpenLineage and Marquez work together to capture, store, and visualize data lineage. OpenLineage collects lineage data from tools like Airflow and Spark, while Marquez stores and displays this metadata, offering visualization tools for exploring dependencies, performing impact analysis, and improving data observability.



Kubernetes

OpenLineage can be integrated with Kubernetes, often through orchestration tools like Airflow or Spark, to track lineage across containerized workloads. This is especially useful in cloud and hybrid environments where data workflows span multiple clusters and environments.

These integrations make OpenLineage a versatile and powerful tool for tracking, visualizing, and maintaining lineage in modern data engineering environments, enhancing observability and governance across data workflows.

PROJECTS



Pyro is a universal probabilistic programming language (PPL) written in Python and supported by PyTorch on the backend. Pyro enables flexible and expressive deep probabilistic modeling, unifying the best of modern deep learning and Bayesian modeling. Pyro is a graduation-stage project of the LF AI & Data Foundation.

Contributed by Uber in January 2019 as an incubation-stage project and graduated in February 2021.

Pyro integrates with various open-source projects to enhance its capabilities in probabilistic modeling and machine learning, making it a flexible tool for building complex models while benefiting from the strengths of other libraries.



PyTorch

Pyro is built on top of PyTorch, utilizing its deep learning and tensor computation features for model building, optimization, and training.

Pyro integrates with projects like PyTorch, NumPyro, TensorFlow Probability, Stan, MCMC libraries, Matplotlib, Seaborn, and Apache Spark to enhance probabilistic modeling, inference, and scalability across diverse machine learning and statistical applications.

PROJECTS



vLLM is an open-source library for fast LLM inference and serving. vLLM utilizes PagedAttention, an attention algorithm that effectively manages attention keys and values. vLLM equipped with PagedAttention redefines the new state of the art in LLM serving: it delivers up to 24x higher throughput than HuggingFace Transformers, without requiring any model architecture changes.

The University of California – Berkeley donated vLLM to LF AI & Data Foundation as an incubation-stage project in July 2024.

vLLM integrates with both PyTorch and ONNX by optimizing the inference of models trained in PyTorch and exported to the ONNX format, ensuring efficient and scalable deployment across different platforms and runtimes.



PyTorch

vLLM integrates with PyTorch by optimizing the inference of large language models, enabling efficient execution and scaling of PyTorch-based models in production environments, while leveraging vLLM's advanced performance features for faster and more cost-effective deployment.



ONNX

vLLM integrates with ONNX by enabling the efficient inference of models that are exported to the ONNX format, allowing users to leverage vLLM's optimized engine for faster, scalable deployment across various platforms and runtimes, regardless of the original training framework.

vLLM integrates with projects like Hugging Face Transformers, PyTorch, TensorFlow, ONNX, DeepSpeed, and CUDA to optimize large language model inference, ensuring efficient, scalable, and high-performance deployments.

2024 Growth & Membership

LF AI & Data experienced significant growth in 2024, welcoming 11 new members including:

Premier Members



General Members



LF AI & Data offers a range of membership levels to accommodate diverse organizations and ensure their active participation in shaping the future of AI and data innovation. See right for a categorized list of current members organized by membership level: Premier, General, and Associate.

Becoming a member of LF AI & Data allows organizations to actively contribute to and benefit from a dynamic open-source data and AI ecosystem. Members enjoy the opportunity to:

- ▶ Influence the development of cutting-edge data and AI technologies.
- ▶ Access a global network of innovators and industry leaders.
- ▶ Support the adoption and sustainability of open-source initiatives.

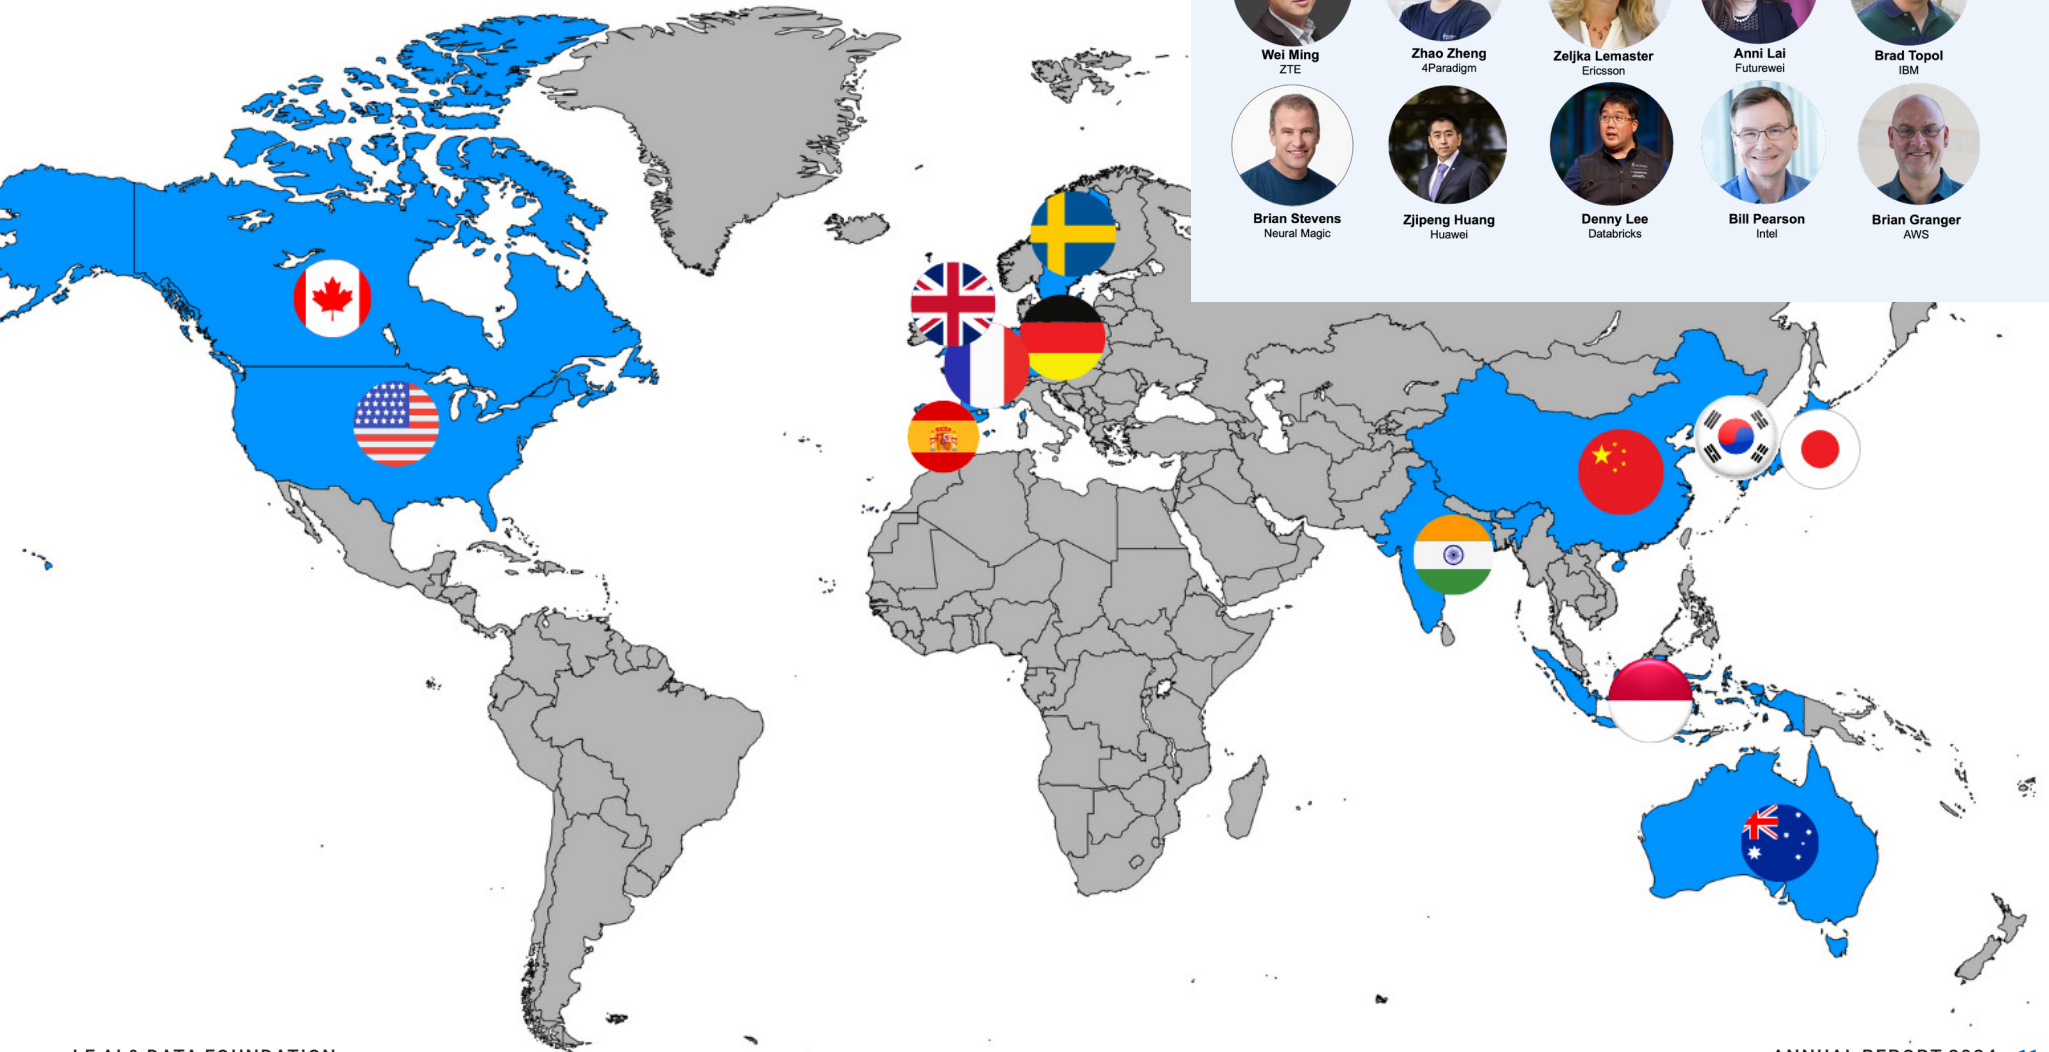
We nurture open-source AI and data projects, growing them from seed to fruition with full support and resources.



To join, please visit our member [enrollment page](#).

Geographical Distribution of Members

LF AI & Data's members span the globe, reflecting a truly international commitment to advancing open source data and AI technologies. The map below illustrates the geographic distribution of our member organizations, showcasing the breadth and diversity of our global community.



2024 LF AI & Data Governing Board

 Junping Du Datastrato	 Jorge Silva SAS	 Liam Zheng OPPO	 Ali Dallou Microsoft	 Vini Jaiswal TikTok
 Wei Ming ZTE	 Zhao Zheng 4Paradigm	 Zeljka Lemaster Eriksen	 Anni Lai Futurewei	 Brad Topol IBM
 Brian Stevens Neural Magic	 Zijpeng Huang Huawei	 Denny Lee Databricks	 Bill Pearson Intel	 Brian Granger AWS

Shaping the Future of Generative AI Survey



Published November 2024

The report highlights the pivotal role of open source in Generative AI (GenAI) adoption, underscoring its strategic importance in driving innovation and fostering a community-driven ecosystem. It confirms open source's key role in GenAI growth, offering businesses actionable insights as they navigate this

rapidly evolving field. The report emphasizes open collaboration, ethical AI development, and interoperable standards for responsible usage and broad adoption.

It also points to the underrepresentation of open source in GenAI compared to proprietary models, highlighting licensing complexities and the need for

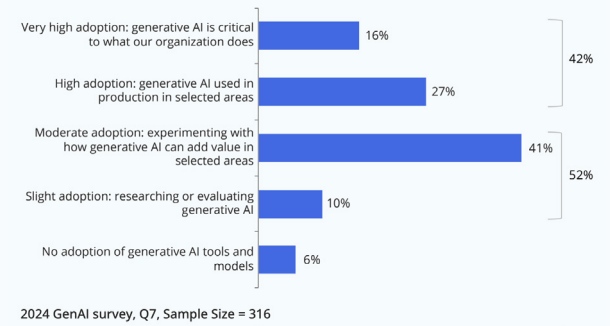
workforce readiness.

Despite this, open source makes GenAI more accessible, with diverse applications across sectors. The report stresses the community's shared responsibility to ensure ethical AI development.

Open source enables widespread GenAI adoption by providing accessible, flexible, and cost-effective tools. It promotes rapid innovation, community-driven advancements, and interoperability while offering businesses greater control and avoiding vendor lock-in. The transparency of open source AI models builds trust, especially in regulated industries, and supports workforce upskilling.

FIGURE 1: HOW ORGANIZATIONS ARE ADOPTING GEN AI

To what extent has your organization adopt generative AI? (select one)



Generative AI Commons SIG



Anni Lai, Generative AI Commons Chair, LF AI & Data

In 2024, the Generative AI Commons community achieved key milestones, driving innovation and collaboration in the open source GenAI space. Our efforts focused on four core workstreams: Framework, Model-Data-Application, Education & Outreach, and Responsible AI.

This year's important accomplishment was the release of the [Model Openness Framework Specification 1.0](#) and the accompanying [Model Openness Tool](#), both designed to address misrepresentation of openness in GenAI development.

Additionally, the community developed the [Responsible AI Framework](#), which provides a consensus-based definition of Responsible AI and outlines its nine critical dimensions, spotlighting projects that advance Responsible AI systems.

The community also enriched the GenAI ecosystem with a comprehensive glossary, educational resources, blogs, and webinars.

Looking ahead to 2025, we aim to expand these efforts by developing a detailed GenAI landscape, reference architecture, and best practices. These initiatives will accelerate open source innovation and enhance the practical application of GenAI technologies across industries.

We're always excited to welcome new community members! Join us by [subscribing](#) to the Generative AI Commons mailing list:

Email: gen-ai-commons+subscribe@lists.lfai.foundation

From the TAC Chair

The LF AI & Data Foundation has witnessed an extraordinary year in 2024, marked by transformative growth and impactful contributions to the open-source AI and data ecosystem. The Technical Advisory Committee (TAC) took massive strides towards expanding the Foundation's project portfolio and prioritized project visibility.

This year, 11 new projects were welcomed into the LF AI & Data ecosystem. Each project brings fresh perspectives, addressing critical challenges and advancing what is possible in AI and data technologies. From scalable frameworks to innovative tools, these projects showcase the diversity, creativity, and energy of the community.

An exciting initiative launched this year is TAC Talks, a live series designed to spotlight the innovative work happening within our

projects. In these sessions, project maintainers guide the community through technical deep dives into their work creating a space for meaningful dialogue, knowledge exchange, and collaboration.

This report captures the progress made in 2024 and the foundation laid for continued success. It is also an invitation to all stakeholders to engage with the LF AI & Data Foundation's mission and contribute to shaping the future of open-source AI and data innovation.

Thank you for your continued support.



Vini Jaiswal
TAC Chair, LF AI & Data

2024 Event Highlights



TAC Talks



In 2024, we launched *TAC Talks* to spotlight our projects and their impact. Hosted by Vini Jaiswal, this biweekly series delves into the cutting-edge initiatives shaping the future of AI and data within the LF AI & Data

ecosystem. Each episode features project maintainers and core contributors who share technical insights, showcase their work, and explore innovative solutions.

Virtual Events

Virtual events play a critical role in promoting connection, collaboration, and growth within the LF AI & Data community and the project ecosystem. In 2024, LF AI & Data hosted a series of impactful virtual events, including webinars, live demos, and community meetups.

Top sessions include:



[TAC Talks: Unity Catalog](#)

Contributed by Databricks, Unity Catalog is a Sandbox-stage project transforming data management and AI technologies. It strengthens data governance and security, offering valuable insights into managing modern data and AI workflows.



[TAC Talks: Open Platform for Enterprise AI \(OPEA\)](#)

This *TAC Talks* session introduced OPEA, a Sandbox-stage project simplifying GenAI app development and adoption for enterprises, exploring its origins, objectives, and the technical challenges it addresses.




[TAC Talks: Monocle](#)

This *TAC Talks* session featured Monocle, a Sandbox-stage project by Okahu that simplifies GenAI application tracing. Monocle enables developers to trace complex workflows, from code to vector databases, with minimal changes, enhancing visibility and streamlining workflows.

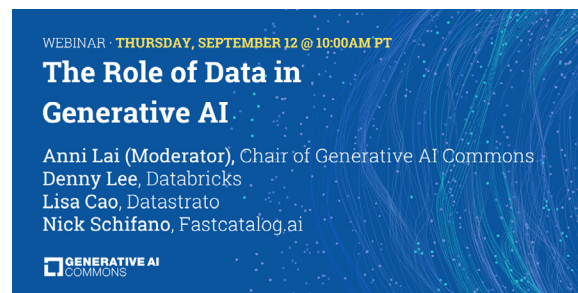
LF AI & Data 2024 Webinars and Events

[Generative AI Commons](#) hosted virtual events offering valuable education and expert insights on generative AI, fostering open dialogue, knowledge sharing, and actionable takeaways.



[The Importance of Openness in Generative AI](#)

This webinar highlighted the role of openness in shaping new technologies, emphasizing how it promotes transparency and accountability in AI while addressing risks of proprietary approaches.



[The Role of Data in Generative AI](#)

In this panel, experts Denny Lee, Lisa Cao, and Nick Schifano, moderated by Anni Lai, explored data's critical role in generative AI. Topics included data collection, preprocessing, governance, and strategies to address bias, ensure fairness, and maintain data quality and diversity.



[Demystifying Generative AI](#)

This webinar introduced generative AI and its real-world applications. Jeremiah Long, a member of the Generative AI Commons Education & Outreach Committee, showcased practical use cases and offered insights on integrating AI tools into work and daily life.

Unity Catalog

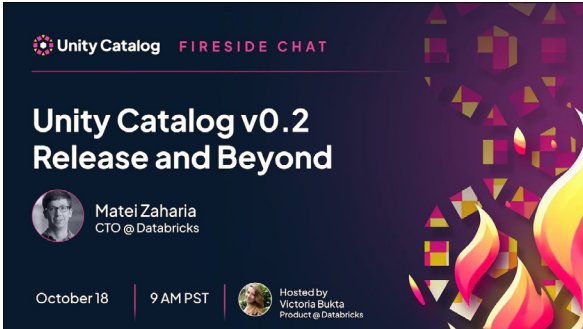
Unity Catalog hosted virtual events to foster community connection, learning, and collaboration.

- ▶ **Community Meetups** showcased demos, project updates, and discussions on Unity Catalog’s evolving capabilities, encouraging dialogue among contributors.
- ▶ **Fireside Chats** featured industry leaders and maintainers sharing insights on Unity Catalog’s impact on data governance, security, and AI integration.
- ▶ **Zero to Hero** offered step-by-step sessions, guiding participants from setup to real-world applications.



DuckLake + Unity Catalog

On December 5, the Unity Catalog Community Meetup featured DuckLake, a new integration between DuckDB and Unity Catalog. The session highlighted how DuckLake simplifies data lake and governance workflows, included a live demo, and offered attendees a chance to engage with contributors and ask questions.



Fireside Chat: Unity Catalog v0.2 Release and Beyond with Matei Zaharia and Victoria Bukta

Matei Zaharia, creator of Apache Spark™, joined the fireside chat to discuss Unity Catalog’s vision and the importance of open-sourcing the project. The conversation highlighted the need for a catalog in data analytics and AI, emphasizing Unity Catalog’s role in advancing open-source data governance and management.



Zero to Hero: Setting Up Cloud Storage

Victoria Bukta and Unity Catalog contributor Bijan Houle live-coded the setup of Unity Catalog from scratch, demonstrating how to run it locally and connect it to cloud storage. The interactive session offered real-world insights, highlighting challenges and solutions while answering questions in real time.

“As a Premier Member of LF AI & Data, we strengthen our commitment to supporting diverse open-source AI and data communities. We believe that open source technologies are foundational to accelerating innovation, enhancing transparency, and ensuring that technological progress is accessible to a broad audience from the outset.”

- OLAF HUBEL, DIRECTOR OF DEVELOPER RELATIONS, DATABRICKS

Delta Lake

Delta Lake hosted interactive live events, including Ask Me Anything sessions and technical deep dives, strengthening its community and promoting best practices for data lakehouse implementations.




[Thinking Fast and Thinking Slow: Adopting the Streaming-First Mindset](#)

Attendees learned to shift from a batch-oriented to a streaming-first approach, unlocking new possibilities for dynamic data use and empowering data engineers to lead their data strategies.

OPEA

OPEA's virtual events provide practical insights and technical guidance to help community members bring GenAI solutions to market efficiently. Initiatives like demopaloozas and hackathons offer valuable knowledge and create opportunities for active contributions to the project's growth.



[Ask the Authors](#)

The Delta Lake community launched the Ask the Authors series, a live event with the authors of *Delta Lake: The Definitive Guide*. The session offered insights into optimizing data lakes, tackling data challenges, and exploring Delta Lake's advanced capabilities.



[Delta AMA: Meet Dominique Brezinski \(our newest TSC\)](#)

This event introduced the newest Delta Lake TSC member. Attendees explored Delta Lake's origins and learned how it was deployed into production in under a year, processing hundreds of petabytes reliably and efficiently.

[OPEA Demopalooza: GenAI RAG Workflows in Action](#)

The OPEA Demopalooza featured two demos, showcasing OPEA in action and inspiring participants to kick-start their next GenAI projects.



[OPEA Hackathon - AMA - Open Source Contribution & Mentorship Event](#)

Held during the OPEA Hackathon, this event spotlighted contributions in code and content, featuring contributors who closed issues, partners working on pull requests, and updates on mentor-mentee progress.



Check out what events are coming in 2025 by visiting the [LF AI & Data Community Calendar](#).



2024 Linux Foundation AI Global Events



In 2024, LF AI & Data expanded its global reach with AI_dev conferences in Paris, Japan, and Hong Kong. AI_dev is a hub for developers exploring open-source generative AI and machine learning. Centered on the belief that open source drives AI innovation, the event brings together top developers to foster discussions, collaborations, and shape the future of open-source AI. Attendees will participate in hands-on sessions, technical discussions, and leverage collective expertise to push the boundaries of AI with open source.

AI_dev
Open Source GenAI & ML Summit

19-20, JUNE 2024 | PARIS, FRANCE


DEMOGRAPHIC DATA

THE LINUX FOUNDATION
#AIDev

640 ATTENDEES

61% OF ATTENDEES IN TECHNICAL POSITIONS

405 TOTAL ORGANIZATIONS REPRESENTED



KubeCon | **CloudNativeCon**

THE LINUX FOUNDATION

OPEN SOURCE SUMMIT | **AI_dev**

China 2024

THE LINUX FOUNDATION

OPEN SOURCE SUMMIT
JAPAN

AI_dev
Open Source GenAI & ML Summit

2024 OCTOBER 28-29 | TOKYO, JAPAN

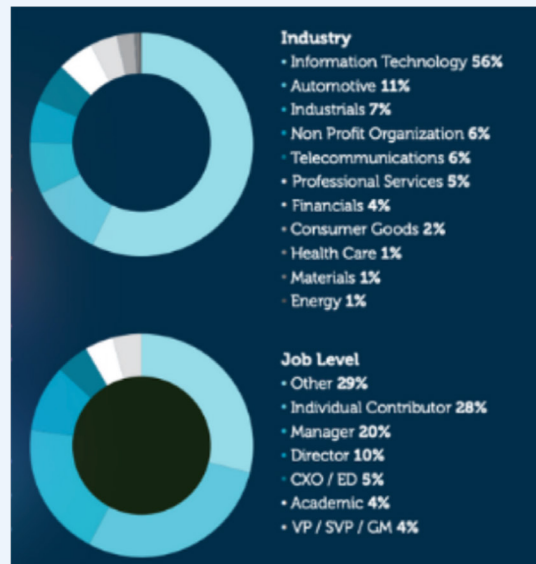

DEMOGRAPHIC DATA

#OSSUMMIT & #AIDEV

878 IN-PERSON ATTENDEES

55% OF ATTENDEES IN TECHNICAL POSITIONS

560 TOTAL ORGANIZATIONS REPRESENTED



54%
First-time attendees

527
CFPs submitted

TOP 3 JOB FUNCTIONS

DEVELOPER

272

ARCHITECT

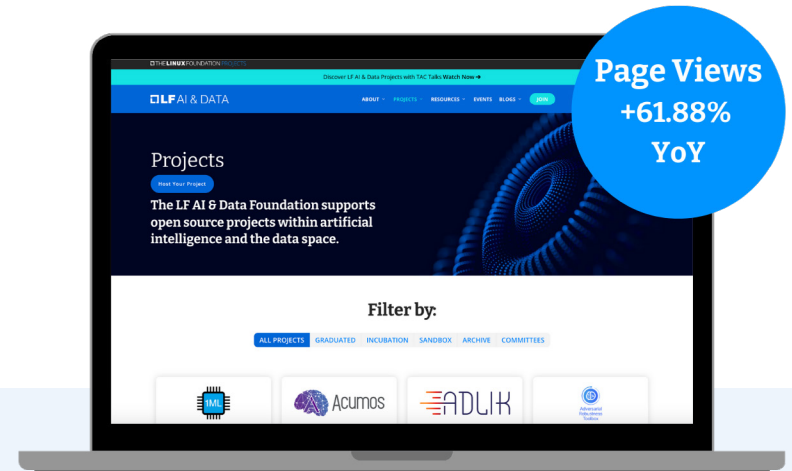
246

DEVOPS/SRE/SYSADMIN

217

Community Engagement

Our projects pages account for one third of our website traffic. Here are our top project pages:



YouTube

The LF AI & Data YouTube channel continued to grow as a vital resource for our community in 2024. It serves as a comprehensive repository for knowledge-sharing, hosting webinars, tutorials, project demos, and event recordings. Check out the [LF AI & Data YouTube channel](#).



LinkedIn

Our presence on [LinkedIn](#) expanded significantly, allowing us to connect with a broader audience. Key achievements include:

- ▶ **Follower Growth:** Our LinkedIn page saw a **52.7% increase** in followers year over year.
- ▶ **Content Engagement:** We published **323 posts** throughout the year, achieving over **113,000 impressions** and generating **2,272 interactions**, including likes, comments, and shares.



Slack

[Slack](#) serves as the central hub for real-time collaboration and communication within our community. With nearly **3,000 members**, our Slack workspace is where community members connect, share ideas, and collaborate. We invite you to join the conversation and stay engaged with our vibrant community.

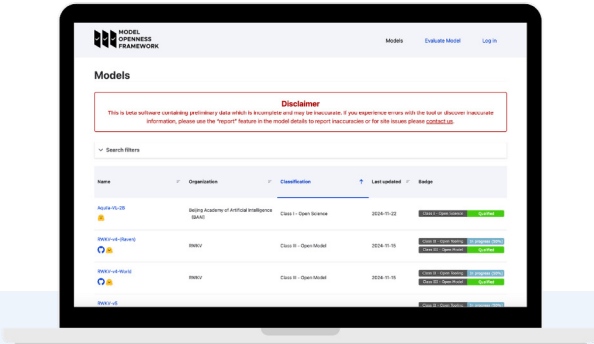
LF AI & Data Initiatives

Generative AI Commons



Since its launch in 2023, Generative AI Commons has made tremendous strides in promoting and advancing Generative AI open source innovations through neutral, open and transparent collaboration. This year, the committee achieved significant milestones across its four focused workstreams:

- ▶ Models, Applications & Data (MAD)
- ▶ Frameworks
- ▶ Education and Outreach
- ▶ Responsible AI



Model Openness Framework & Model Openness Tool



Recognizing the need for greater transparency in AI, Generative AI Commons created the [Model Openness Framework](#) (MOF) as a comprehensive framework to evaluate and classify the openness of AI models. The MOF provides a structured approach to assess models across 16 critical components, setting high standards for openness and accessibility. By defining three tiers of openness—Open Model, Open Tooling, and Open Science—the framework guides model producers and

enables consumers to make informed decisions about the transparency and usability of AI systems.

To facilitate the practical implementation of the MOF, the committee launched the [Model Openness Tool](#) (MOT). MOT enables developers to evaluate their models against MOF criteria and generate actionable transparency scores.

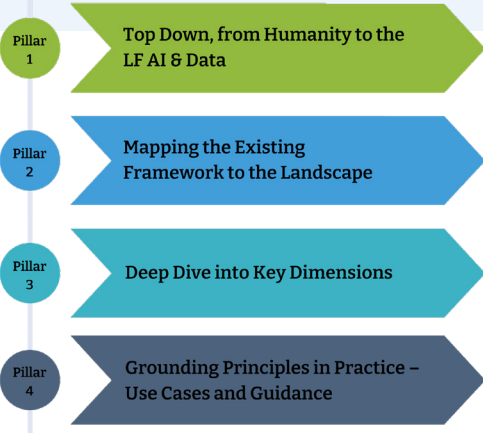
The [BAAI Aquila-VL-2B model](#), developed by the Beijing Academy of Artificial Intelligence (BAAI) research team, became the first model to earn the prestigious Class I “Open Science” rating under the MOF.

Responsible AI Pathways

Recognizing the transformative potential of AI and the accompanying risks—such as bias, safety concerns, and misuse—the committee’s Responsible AI workstream launched the [Responsible AI Pathways Initiative](#) to ensure ethical AI development and deployment.

This initiative is structured around four core pillars that ensure a holistic and actionable approach to Responsible AI. These principles will form the

foundation of the upcoming Responsible AI Framework, scheduled for release in 2025. The framework will provide actionable guidelines to address the challenges of Responsible AI across key dimensions, serving as a cornerstone for fostering ethical innovation in AI.



Regional User Group (RUG)



In 2024, LF AI & Data launched two pioneering RUGs: RUG China and RUG Japan. These groups were established with a shared mission to enhance regional collaboration, knowledge exchange, and innovation within the open source AI

ecosystem. RUGs under LF AI & Data are designed to serve as local outreach hubs, aligning with the Foundation's overarching vision of expanding the global open source AI and data community. As RUGs grow and evolve, they will remain instrumental in driving the adoption and development of open source AI technologies, strengthening the Foundation's global impact.

Generative AI Glossary

Generative AI Commons introduced the [Generative AI Glossary](#), a dynamic and ever-evolving resource designed to make generative AI terminology easier to understand. As a living document, it evolves with community contributions to stay current with the latest advancements in the field.

“Our vision at Neural Magic is The Future of AI is Open, so it would seem only natural to be joining the LF AI & Data as a Premier Member. We look forward to contributing within the LF AI & Data community to help make our vision become a reality”.

- BRIAN STEVENS, CEO, NEURAL MAGIC

BI & AI Committee

The BI & AI Committee promotes the integration of Business Intelligence (BI) and Artificial Intelligence (AI) to create Cognitive Intelligence (CI), combining the speed and efficiency of AI with the intuitive insights of BI. Established to address the growing need for interoperability, best practices, and practical guidelines, the committee focuses on aligning these disciplines for smarter, data-driven decision-making.

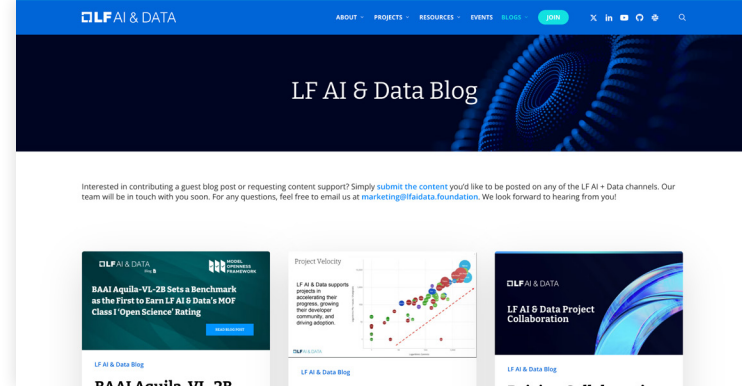
In 2024, the committee launched the white paper [The Alchemy of Intelligence: How Generative AI can Revolutionize Business Intelligence and Analytics in Modern Enterprises](#). This publication highlights how Generative AI is transforming BI and analytics by enabling predictive insights, fostering efficiency, and empowering businesses to make smarter decisions. Through real-world scenarios, the white paper explores opportunities and challenges in adopting Generative AI, offering actionable strategies for implementation within enterprises.




Blog & Resources

We published **82 blog posts** across the LF AI & Data Blog and the Community Blog, covering diverse topics that reflect our ecosystem's passion and innovation. The blog section drove **20% of website traffic**.

The LF AI & Data Blog highlights official news, updates, and milestones, offering insights into project launches, key initiatives, and event recaps.




Launched this year, the Community Blog shares insights and thought leadership from our community and industry experts, featuring the latest trends, analyses, and member-contributed articles.




Top 5 Blogs of 2024

[Introducing the Model Openness Framework: Promoting Completeness and Openness for Reproducibility, Transparency and Usability in AI](#)



[LF AI & Data Foundation Launches Open Platform for Enterprise AI \(OPEA\) for Groundbreaking Enterprise AI Collaboration](#)



[Translation Augmented Generation Breaking Language Barriers in LLM Ecosystem](#)



[LF AI & Data Announces vLLM as its Latest Incubation Project](#)

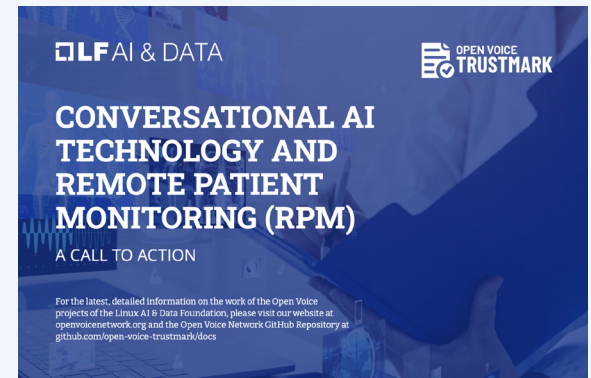


[Embracing the Future of AI with Open Source and Open Science Models](#)

Conversational AI and Remote Patient Monitoring (RPM)

In March 2024, Open Voice TrustMark published the white paper [Conversational AI Technology and Remote Patient Monitoring \(RPM\)](#),

exploring the transformative role of conversational AI in the future of healthcare. The paper highlights how advancements in natural language processing, understanding, generation, and generative AI are reshaping RPM by enabling lower costs, increased adoption, broader utility, and greater inclusivity. It envisions a future where conversational AI enhances patient care by identifying early indicators of physical diseases and mental health disorders, paving the way for more efficient and accessible healthcare worldwide.



Looking Ahead to 2025

Open Source AI Trends

Open source AI will drive innovation, reduce costs, and enable organizations of all sizes to adopt cutting-edge AI technologies, unlocking new industry opportunities. It will transform sectors such as healthcare, financial services, retail, manufacturing, autonomous vehicles, agriculture, energy, education, legal, and creative industries by providing accessible, customizable, cost-effective building blocks for enterprise solutions.

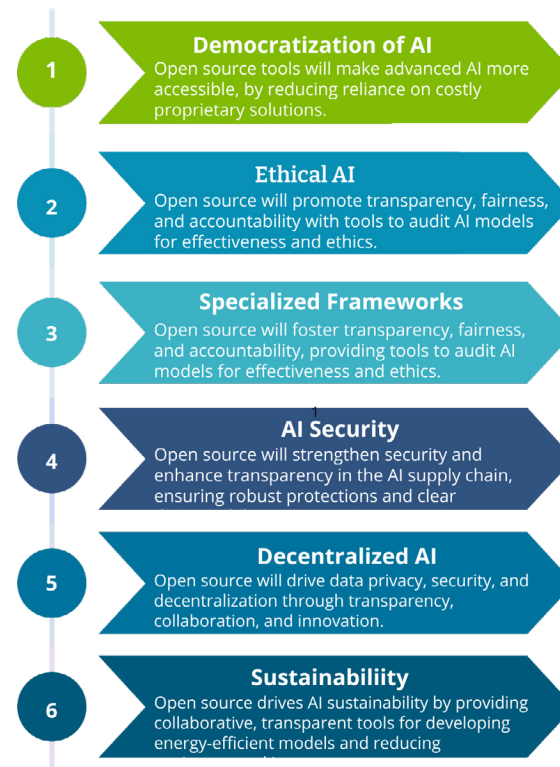
The Evolving Role of Corporations in Open Source AI

- ▶ Corporations will play a significant role in advancing open source AI through increased investment. As AI's value continues to grow, companies will recognize the importance of supporting open source projects by providing financial backing and resources for development.
- ▶ Greater collaboration between corporations, academia, and independent developers will also emerge. As AI models become more complex, organizations will need to work closely with external contributors, fostering stronger partnerships and joint research efforts.
- ▶ As AI regulations evolve, open source AI communities will align with new rules, particularly regarding data privacy and ethical AI. This alignment will ensure transparency and accountability in AI development, fostering greater public trust in the technology.

Regulatory Considerations

EU regulations will significantly shape open source AI development. The EU AI Act will promote ethical AI practices, driving transparency, fairness, and accountability in open source projects and encouraging collaboration on bias detection and ethical auditing tools.

The GDPR will continue to influence how open-source AI handles sensitive data, fostering privacy-preserving techniques like differential privacy and federated learning and positioning open-source AI as a leader in privacy-focused solutions.



Intellectual property and AI-generated content regulations will require updates to open source licensing models, clarifying data ownership and usage. The EU's push for AI standardization will also rely on open source AI to ensure interoperability across sectors like healthcare and smart cities.

While these regulations pose challenges, they also offer open source AI the opportunity to lead in responsible, secure, and innovative AI development.

Get Involved

Explore Membership:
Become a member of LF AI & Data and help shape the future of open source artificial intelligence and data innovation. Discover membership opportunities here.

Subscribe to Our Mailing List: Receive updates on groundbreaking projects, upcoming events, and community initiatives directly in your inbox. [Subscribe here.](#)

Follow Us on Social Media:







About The Linux Foundation

The Linux Foundation is the world's leading home for collaboration on open source software, hardware, standards, and data. Linux Foundation projects are critical to the world's infrastructure including Linux, Kubernetes, Node.js, ONAP, OpenChain, OpenSSF, PyTorch, RISC-V, SPDX, Zephyr, and more. The Linux Foundation focuses on leveraging best practices and addressing the needs of contributors, users, and solution providers to create sustainable models for open collaboration. For more information, please visit us at linuxfoundation.org. The Linux Foundation has registered trademarks and uses trademarks. For a list of trademarks of The Linux Foundation, please see its trademark usage page: www.linuxfoundation.org/trademark-usage. Linux is a registered trademark of Linus Torvalds.

Connect with us

 twitter.com/linuxfoundation

 facebook.com/TheLinuxFoundation

 linkedin.com/company/the-linux-foundation

 youtube.com/user/TheLinuxFoundation

 github.com/LF-Engineering

548 Market St
PMB 57274
San Francisco, California 94104-5401 US

info@linuxfoundation.org
www.linuxfoundation.org

