

Annual Report 2025

Innovation in the Open

www.linuxfoundation.org

Contents

INTRODUCTION	3
Executive Director Update	4
Board Chair Update	6
Board of Directors List.....	8
Thank You to Our Members.....	9
The Linux Kernel Organization	19
By the Numbers: The LF 2025.....	21

ACCELERATING COMMUNITY ENGAGEMENT	22
Standards and Specification Development	23
Open Source and Regulatory Considerations.....	25
LF	27
Defending Open Source	29
LF Education	30
LF Research	34
LF Events.....	35
Mentorship.....	37

SPHERES OF IMPACT	38
The Open Fork.....	39
Security	41
Cloud	43
AI	45
Data	48
Hardware & Infrastructure	51
Industry.....	53
Trust.....	55

FINANCIAL TRANSPARENCY	57
-------------------------------	-----------

Introduction

Executive Director Update



When we think about the future of open source, it's important to remember where we started. This is especially true when we consider the future of the Linux Foundation.

Today, we are navigating the most significant technology transformation of the past two centuries. Of course, I'm talking about artificial intelligence (AI).

Every day brings news of AI systems advancing across fields and changing how we work. Some argue that AI remains immature and has failed to make significant production inroads. That reminds me of what people said about Linux and mobile devices years ago. In both cases, we saw groundswell adoption, though from different directions. Linux came from developers solving problems. Mobile came from smartphone users adopting devices for business because they crushed existing solutions. Part of our job as an open source foundation is spotting these inflection points and seeing the future while we can still influence it. We want open source to shape the future of technology. The ethos of open source and the non-zero-sum mindset must be injected into any major technological transformation if we are all to benefit.

This past year, we accelerated efforts to build a robust portfolio of critical AI projects. Thanks to your guidance and the work of many LF employees, we are now well-positioned to foster open source AI and build an ecosystem promoting open source values for this critical field. Ben Lorica, who publishes an influential AI newsletter, coined the term “the PARK Stack”: PyTorch, AI models,

Ray, and Kubernetes as the core foundation of AI infrastructure for building, training, and inferencing models. Three of the four PARK Stack components are now Linux Foundation projects.

The PARK Stack echoes the LAMP Stack, perhaps the most important innovation from the Dot Com era and the core architecture that launched the explosion of open source SaaS applications and broad Linux adoption for enterprise production. The PARK Stack could occupy the same influential place. History rhymes in technology cycles.

But the PARK Stack alone is not enough. The open source community lacks a good answer for securing data access at scale for AI builders. This access becomes critical as intellectual property owners wall off content from public crawlers. We also face a gap in state-of-the-art models. While we've seen powerful open source models emerge, there has been no “Linux” of AI models to compete with large proprietary systems. I use the operating system analogy because AI is rapidly taking on an OS-type role for autonomous work.

In that vein, we launched the Agentic AI Foundation (AAIF), a new neutral home for the emerging stack of agentic AI technologies. With founding contributions including Model Context Protocol (MCP), goose, and AGENTS.md, AAIF fills a critical gap by establishing open, community-governed standards for how autonomous systems access tools, data, and workflows. For the Linux Foundation, AAIF represents the same kind of moment we saw with Kubernetes and the cloud-native movement: a chance to convene industry leaders early, set shared standards, and ensure the next phase of AI is built on openness and interoperability rather than fragmentation. This is an important milestone for our organization and a signal that the LF will help guide AI's evolution

at the same scale and significance as prior shifts in modern computing.

We're also thinking about AI in terms of how the LF is moving into the AI era. This past year we repositioned our operations and infrastructure to embrace an AI-first future. Many manual elements of running open source projects and foundations can and should be automated. Our member companies have already taken these steps, using AI to accelerate code refactoring, manage networks, triage security alerts, and generate marketing content. A significant percentage of our maintainers use AI coding tools to ship faster and focus on higher value work requiring judgment and insight. We must be as efficient as our members and learn from them. Very shortly, AI will become just another component in the technology landscape, much as containers went from novelty to Kubernetes and the dominant form of running infrastructure today.

We can focus so heavily on AI because the rest of the Linux Foundation continues to grow and thrive. We've seen another record year in memberships and event attendance. We're seeing massive grassroots growth and taking great leaps in our fast-growing education business. LF Training is transitioning from a certification program toward becoming a comprehensive technology education subscription focused on open source. Our flagship ecosystem, the CNCF, is now enjoying a second wave of highly successful projects such as OpenTelemetry and Backstage. They are approaching the same scale as the world's leading open source juggernauts. The bedrock of the LF, Linux

itself, is progressing through upgrades to Rust, eBPF, and other modern elements, ensuring the world's most popular enterprise operating system remains the most advanced and secure. In industry after industry, from utilities to media and entertainment to finance to telecommunications, capable project leads and member companies are cooperating to build shared technology and advancing at a healthy pace. I am astounded at the ability of so many projects and foundations to fight the gravity of time and size while innovating at breakneck speed.

Significant challenges lie ahead. The world is more complicated economically and politically than five years ago. From deepfakes and AI-powered cyber attacks to the ongoing onslaught of ransomware and supply chain compromises, we face critical challenges to our technology infrastructure, our privacy, and at times our very identity. I remain an optimist. The best path to solving these problems is working together through collective non-zero-sum efforts, prioritizing the common good. We must continue convening like-minded technologists and organizations from around the world to build a shared technology foundation and heritage for all of us and for our children. "People, projects, progress" is our value stack, and it is a clear reminder of how together we have built a multi-trillion dollar technology ecosystem through collaboration and respect.

Thank you for your partnership. Together, let's build the next era of open source and create the future we all want to see.

Jim Zemlin, Executive Director

Board Chair Update



As I enter my sixth year as Board Chair of the Linux Foundation (was elected in 2019), I feel privileged to continue to witness and guide what I believe to be one of the most important movements in human history.

The open source ethos of sharing knowledge, teaching self-reliance, practicing transparency, and placing collective good over individual benefit represents the best of what makes us a great and compassionate civilization. It also represents our ingenuity and our ability to envision a grander future for all. We are now deeply entrenched in the third decade of open source, a pivotal moment that demands both celebration of our achievements and an urgent call to action to address the challenges ahead.

Our journey began with the foundational efforts of the free software movement. Richard Stallman developed the GPL license, defining the four freedoms that would become central to our work, and Linus Torvalds used that GPL to share the code that became Linux, now the default operating system globally. The license was one major factor in the flourishing of open source. Communities of people collaborating on a common mission was another reason for open source's success. Organizations like the LF were established around 2000 to provide a nonprofit, neutral home where companies could collaborate and protect this revolutionary idea from legal, governance, and financial threats.

The second decade saw open source flourish and become ubiquitous. Web-scale companies built massive data centers on Linux and the LAMP Stack. Enterprises, governments, and critical

sectors integrated open source into their core infrastructure. The United Nations affirmed this level of acceptance when it published its open source principles earlier this year. The UN sees the power of open source for solving big human challenges and empowering countries to use digital public goods for serving their citizenry. This year, I had the honor of participating again in the UN Open Source Week, a highlight that underscores how far we have come and how we have become a truly global movement.

In 2025, the LF continued to defy gravity, accelerating growth in contributors, members, and conference attendance—as it has, year after year. A virtuous critical mass continues to gather gravitational pull. More and more, industries and organizations come to the LF to establish collaborative projects with neutral governance. This has meant that the digital fabric of the world's power plants, factories, hospitals, schools, governments, transportation networks, communications infrastructure, and mercantile backbone rests upon LF projects. And yet, we formalized our presence in India, the most populous nation on Earth, only in December 2024. This tells us that despite decades of rapid growth, significant opportunities remain ahead.

There is no better barometer of the growth of open source than the remarkable growth of Open Source Program Offices, or OSPOs. They are the key indicator of future contributions and commitments of organizations to the cause of open source. While we see important contributions from solo developers, it's the enterprises that carry the most water; and the more that join the queue, the more water we can carry. This year saw a threefold increase (from 15% in 2024 to 45% in 2025) in organizations planning OSPOs within two years, all citing improved developer experience as a key driver. Organizations with OSPOs are 2.5 times more likely to enable upstream contributions and nearly twice as likely to encourage open source participation. Even countries and universities see OSPOs as the

galvanizing center for accelerating their innovation and collaboration.

OSPOs demonstrate an organization's seriousness about open source strategy and their willingness to collaborate. They navigate complex realities: abrupt license changes; new government regulations, such as the Cyber Resilience Act (CRA); open source AI debates; and supply chain security. The community has also become more willing to put skin in the game, with well-supported projects like OpenSearch exemplifying how the LF provides neutral homes for true open source collaboration rather than "stadiumware" dominated by single organizations.

With ubiquity comes immense responsibility. We must focus intensely on several key challenges. Open source supports critical infrastructure yet relies on a fragile supply chain maintained in critical places by volunteers who may lack security training. Regulations such as the CRA are forcing maintainers and consumers of open source to increase accountability, know what's in their software, and commit to fixing critical vulnerabilities. We must all be more proactive in tracking and monitoring our software components.

As we enter the next major technology paradigm shift, the LF is stepping up in its role as provider of the most important tools for building, training, and running AI systems. We have an opportunity to ensure AI development is ethical, bias-free, and transparent. We must develop open source AI—models, data, tools, and standards—to prevent technological islands and ensure all can benefit from the remarkable capabilities of these novel thinking systems.

While so much is new, the stark reality is that our community is aging rapidly. The edifice of open source requires a steady influx of fresh talent not only to build the next great thing, but also to ensure that the foundation and plumbing laid many years ago continue to function and are frequently upgraded. Open source needs to open up this opportunity to more entrants.

Many young people take open source for granted, unaware of its history, ethos, and culture. We must mentor and educate the next generation. Our training programs, including free offerings and micro-module platforms, have become the largest collective education endeavor ever recorded, helping thousands gain skills to climb the economic ladder. Programs such as the Andela-CNCF Kubernetes African Developer Training, which has trained over 5,600 technologists, demonstrate our commitment to global talent development. It is also encouraging to see new, excited young minds entering into areas considered less dynamic like mainframes and COBOL—areas that still run a huge portion of the world's critical infrastructure. This is another and equally important side of building a collective and a technology commons.

On the bright side, the market is voting with its labor and attention. The number of people contributing to LF projects continues to increase at a healthy clip. Even Linux and Kubernetes, projects that are "old" by technology standards, remain dynamic and continue to attract both new companies and plenty of new contributors. We also see the world voting with their feet, their eyes, and their ears. Our calendar of thousands of events in 2025—from grassroots meetups to KubeCon, the world's largest technology conference—maps the breadth of our ecosystem. These in-person experiences make us more human and more creative, and they bring open source to life in ways that can never be digitized.

Open source is entering a new, crucial era, building on an incredible run as one of the most successful enablers of global innovation in history. As we navigate the complexities of security, sustainability, and AI, I invite all members to join us in sustaining this movement. The digital fabric of tomorrow depends on the work we do today. Thank you for continuing to elevate open source.

With gratitude and optimism,

Nithya Ruff, Board Chair

Board of Directors



Tim Bird
Sony - Gold Director



Erica Brescia
At-Large Director



Kimberly Craven
Red Hat



Eileen Evans
At-Large Director



Frank Fanzilli
At-Large Director / Treasurer



Peixin Hou
Huawei



Takehisa Katayama
Renesas - Gold Director



Sachin Katti
OpenAI



Ken Komiyama
Fujitsu



April Kyle Nassi
Google



David Marr
Qualcomm



Ben Maurer
Meta



Yuichi Nakamura
Hitachi



Shojiro Nakao
Panasonic - Gold Director



Daniel Park
Samsung



Phil Robb
Ericsson



David Rudin
Microsoft



Nithya Ruff
Chair



Emilio Salvador
GitLab



Dan Williams
Intel



Jim Wright
Oracle



Katsuyuki Yamamoto
NEC Corporation



Jim Zemlin
Linux Foundation

Thank You to Our Members

2025 was a defining year for open source and for the Linux Foundation.

Open technologies drove some of the most important advances in AI, infrastructure, and security—reinforcing open source as the backbone of modern innovation.

Our communities delivered major innovations, onboarded influential new projects, and expanded collaboration across every sector. That surge of progress was made possible by the commitment of our members.

As we head into 2026, we're building on this momentum. Our focus is on scaling the infrastructure and open governance models that enable projects at the LF to grow and thrive. Your membership, leadership, and technical contributions fuel this work and strengthen the impact of open source worldwide.

Thank you for your partnership and your dedication. Together, we're shaping the future of the world's most important technologies.

WAHIBA SANDS, OMAN

Platinum members



Gold members

ANTHROPIC



DELL Technologies



HONDA
The Power of Dreams

LY
LY Corporation



Panasonic

RENESAS

SONY

TOSHIBA

TOYOTA

Silver members

#

11ai.co
1NCE GmbH
1Nebula
1Password
23 Technologies GmbH
24x7 Geeks Solutions Pvt Ltd.
2bcloud
3-Shake Inc
321 Gang
42on
45Drives
6WIND S.A.
8gears

A

Agenda d.o.o.
A10 Networks

ABB Switzerland Ltd,
Group Technology
Management
Ac6
ACC ICT
Accenture Global
Solutions Limited
Access Digital
Accuknox
Acend GmbH
ACKSTORM
Acme Gating
Acompany Co., Ltd.
Acornsoft
Ad-ID
Ada Logics
Adaptive Financial
Consulting Limited
Adaptive6
Addresscloud
Adhara Limited

Adobe Inc.
Adtran Holdings, Inc.
Aduna Global LLC
Advanced Micro Devices
(AMD)
Aembit Inc
AENIX INC
aeolabs
Aerospike
Aetheros
Afi Technologies
Agenda d.o.o.
AgileOps
AIA Shared Services (Hong
Kong) Limited
Airbnb
Airwayz
Aisin Corporation
Aiven
Akamai Technologies, Inc.
Akamas S.p.A.

Akara Technologies, Inc
Akenes SA (Exoscale)
Akka
Akuity, Inc.
Alauda, Inc
Alerant Zrt.
Alibaba Cloud (Singapore)
Private LTD
Alif Semiconductor
Almaviva S.p.A.
Alpha Networks Inc.
Alphaus Inc.
alphawave semi
Alter Way
Altinity
Amadeus SAS
amazee.io
Amazon Web Services, Inc.
Amberflo.io
Ambient IT
American Airlines

American Express
AMI US Holding Inc.
Amnic
Ampere Computing
Amphenol Corporation
Amundi Asset
Management
Anaconda, Inc
Analog Devices, Inc.
Andes Digital
Anjuna Security, Inc.
Anodot Inc.
Anonymome Labs, Inc.
Anritsu
Ansys
Ant Group Co., Ltd.
Antmicro
AnyLog
anynines GmbH
AOE
Aokumo Inc.

ape factory
Apica
Apideck
Apiiro
Apollo GraphQL
Apple Inc.
Applied Blockchain
Applied Materials
applied37
Apptio
Aptum
Aqua Security Software,
Inc.
AQUILA CLOUDS
ARAADIGIT
Arcfra
Archera
Archestra.AI
Arcjet
Arcontech Group PLC
Arctera

Arduino
 areti S.p.A.
 Arista Networks, Inc.
 Arm Limited
 ARMO (Cyber Armor)
 ARONETICS
 Arrcus
 ARTELYS
 Artian Inc
 Aruba SpA
 Aryn, Inc.
 Ascensio System SIA
 ASML
 Aspen Technology, Inc.
 ASRock Rack
 Incorporation
 Assa Abloy Group
 Astera Labs, Inc.
 Asterfusion Data
 Technologies
 Astrape
 ASUS Cloud Corporation
 AT&T Services, Inc.
 Aternos GmbH
 ATIX AG
 Atlassian US, Inc
 Atolio, Inc.
 Attribute
 Atym
 Audiokinetic Inc.
 Aureka
 AuriStor Inc.
 Authzed
 Autodesk
 AUTOSOL
 Autovia GmbH

Avanade Inc.
 AVAP
 Avassa
 AVEVA Group
 Aviatix
 Avisi Cloud Services B.V.
 Aviz Networks
 awesome information
 technology
 AXA Group
 Axis Communications
 AXLBIT, Inc.
 AyanWorks Technology
 Solutions Pvt. Ltd.

B
 B1 Systems GmbH
 BACK MARKET, Inc.
 Bai Xin Information
 Technology Co., Ltd.
 Balena Ltd.
 Banco de Crédito BCP
 Bancolombia
 Bank of America
 Corporation
 Bank of Montreal
 basysKom GmbH
 Baumer Management
 Services AG
 BayLibre Inc.
 BE Networks
 BEGASOFT
 Beijing Baolande Software
 Corporation
 Beijing Big Data Co., Ltd.

Beijing Digital China Cloud
 Technology Information
 Technology Co., Ltd
 Beijing Dosec Technology
 Co., Ltd
 Beijing Sup-info
 Information Technology
 Co. Ltd
 Beijing Tenxcloud
 Technology Co. Ltd.
 Beijing Tongtech Co., Ltd.
 Beijing Xiaomi Mobile
 Software Co., Ltd
 BellSoft
 Beningo Embedded
 Group
 Bentley Systems
 Bindplane
 BISDN
 BlackRock, Inc.
 Blacksmith
 BlakYaks
 Blecon Ltd
 Block Majority
 Block, Inc.
 Blockchain Game Partners
 Inc.
 Bloombase, Inc.
 Bloomberg Finance L.P.
 Bloomfilter
 Blue Sentry
 Bluebricks
 Blues Inc.
 BMW
 BNY
 Boeing

Bolt Graphics
 BONbLOC Inc
 Boost Security
 Bootlin
 Booz Allen Hamilton, Inc.
 Border0
 Bosch
 Boston Consulting Group
 BrainGu, LLC
 BrBPO
 Bright Machines, Inc.
 Broadcom Corporation
 Brobridge
 BTS Group
 Buf Technologies, Inc.
 Buildkite Pty Ltd
 Bull SAS
 Buoyant, Inc.
 ByteDance Ltd.
 bytesatwork
 ByteSource Technology
 Consulting GmbH

C
 Cable Television
 Laboratories Inc.
 Cachengo
 Cadence Design Systems,
 Inc.
 Caligra
 Calyptia
 Camptocamp
 CanaryBit
 Canon Inc
 Canonical Group Limited

Capgemini
 Capital One Services LLC
 Carbonated
 CardinalHQ
 CARIAD SE
 CAST
 Cast AI Group, Inc.
 Catalyst Cloud
 Cathay Financial Holding
 Co., Ltd.
 Causely
 CDW Corporation
 CECLOUD Computing
 Technology Co., Ltd
 CelerData
 Celestica
 Certero
 Certizen Limited
 Chainguard
 Chainlink Labs
 Chainloop
 Chaos Software LTD
 chargebyte
 Charter Communications
 Checkly
 Cheesecake Labs
 Chengdu Jingrong
 Lianchuang Technology
 Co., Ltd.
 China Mobile
 Communication
 Company Ltd
 China Systems Holdings
 Limited

China
 Telecommunications
 Corporation
 China Unicom
 Chislitel Lab
 Chkk Inc
 Chronosphere, Inc.
 Cielara
 Ciena Canada, ULC
 Cinema GmbH
 Circle Internet Financial,
 LLC
 Circle Internet Services,
 Inc
 Circulor Ltd.
 Cirrus Logic
 Cisco Systems, Inc.
 Citi
 Civo Ltd.
 claion
 Clastix SRL
 Cleura
 ClickHouse, Inc.
 Clockwork.io
 Cloud 66
 Cloud Ace
 Cloud Computing HK
 Limited
 Cloud ex Machina
 Cloud Software Group,
 Inc.
 Cloud&Heat Technologies
 GmbH
 Cloudbase Solutions S.R.L
 CloudBees, Inc.
 CloudBolt Software

CloudCapital
Cloudchipr
CloudEagle
Cloudera, Inc.
CloudFerro S.A.
CloudFix
CloudGeometry Inc.
CloudGov
CloudHiro
Cloudification
CloudLinux
Cloudmate
CloudNuro Corp
Cloudpi
Cloudsmith Ltd
CloudThrottle
CloudZero
Clounix (Shanghai)
Technology Limited
Cloverleaf Analytics
Clush
Clusys Inc
CLYSO GmbH
CME Group Inc.
CoBank
Code Intelligence
Coder
Codethink
Collabora Ltd.
Comcast Cable
Communications, LLC
CoMira Solutions Inc
Common Tools
Commonwealth Bank of
Australia
ComplianceCow

Component Soft Kft.
con terra
Confident Security
ConfidentialMind
Confluent Consulting LLP
Connect 5G, Inc.
Conoa AB
Consensys AG
Control Plane Corporation
Control-Plane.io
ControlMonkey
ControlTheory
Converge Technology
Solutions Corp
CORE 24/7 LLP
Coredge.io
CoreLogic Solutions, LLC.
CoreStack
CoreWeave, Inc.
Cornelis Networks
Corsha
Cortex
Cosmonic
CPMC
Crayon
Creationline, Inc.
Credo Semiconductor Inc
croit GmbH
Crowd Favorite Ltd
CrowdStrike
Crunchy Data Solutions, Inc.
CTO2B
Ctrlstack
CUE Labs
Cuemby Inc.
CVS Health

Cybertrust Japan Co., Ltd.
Cybozu, Inc.
Cyient Limited
Cyso BV

D

d-Matrix
DABCo Ltd
DACHS IT GMBH
DAEKYO CNS
DaoCloud Network
Technology Co., Ltd.
Dash0
Data Storage Research,
LLC d/b/a DSR
Corporation
Databricks Inc.
DataCore Software
Datadog, Inc
Datafy LTD
DataStax, Inc.
DB Systel GmbH
De Novo LLC
Dedge Security SL
DeeperThanBlue Ltd
Deepshore GmbH
Defense Unicorns
Deloitte Consulting LLP
Delta Electronics, INC.
Dembach Goo Informatik
GmbH & Co. KG
Denodo Technologies
DENSO CORPORATION
Denvr Dataworks

DENX Software
Engineering GmbH
(DENX)
DeployHub, Inc.
Depository Trust and Clea-
ring Corporation (DTCC)
Depot
Design Barn Inc
Desotech srl
Deutsche Bank AG
Deutsche Telekom AG
Development Seed
DevHub LLC
Devtron Inc.
DevZero
Dfns
Dhiway Networks Private
Limited
Diagrid
Diamanti, Inc.
Dianomic
DigiCert, Inc.
DigiSolve Solutions
Digital Asset (Switzerland)
GmbH
Digital Challengers
Digital China Data Cloud
Technology Co.,Ltd.
DigitalEx, Inc.
DigitalFish Inc
DigitalOcean, LLC
Dina IT Solutions SA
Discover Financial Services
Dlocal
dlthub
DMetaSoul

DNEG
Docker, Inc.
DoiT International
Dongobi
Dorado Software
DornerWorks, Ltd.
Doubleword
Doulos
Douyin Vision Co., Ltd.
Dr. Droid
Draftt
Dragonflydb
DreamBig Semiconductor
Inc.
DriveNets
DTEX Systems Inc.
Duckbill
Dynatrace LLC

E

E.ON
EasyStack
EasyStack Inc.
eBay, Inc.
Echo
EcoFlow Inc.
Edera, Inc.
Edge & Node
Edge Delta
Edgecore Network
Corporation
Edgeless Systems
Edgenesis
Effectual Inc
EigenQ

Elasticsearch, Inc.
Elastisys AB
Elastx AB
Electricity Maps
Electronics and
Telecommunications
Research Institute
Elektrobit Automotive
GmbH
Element Labs
Elementl
Eleven Labs Inc
Elgin White
Eliatra
embear
embedd.it
Embeint
Embrace Mobile Inc
embraceable Technology
Emerson
Emilia Capital
emlix GmbH
EMMA
Encora Holdings LLC
Energy Web Foundation
Enfabrica
Ensignia
Enterprise DB Corporation
Entigo OÜ
EntServ UK Limited
enum GmbH
Enveil
Environmental Systems
Research Institute, Inc.
(ESRI)
Envisor

ENVZERO
 EPAM Systems, Inc
 Epic Games, Inc
 Epsio Labs LTD
 EQTY Lab
 Equifax Inc.
 Equinix Services, Inc.
 Eridu AI
 esatus AG
 Escala24x7
 Espeo Software
 Espresso AI
 Etherealize
 eTopus technology
 European IT Consultancy
 EITCO GmbH
 Eve Security
 Evenkeel Inc. d/b/a Densify
 Everest DX Inc
 Evonem LLC
 Excellion Sdn Bhd
 Exein
 Exivity
 Exostellar
 Expedia Group
 Exponential Science
 Foundation
 Extreme Networks, Inc.
 EYGS LLP
 Ezurio

F

F5, Inc.
 Facets Cloud Inc.
 Factory

Fairsquarelab
 Fairwinds Ops, Inc
 Far-Galaxy Networks, Inh.
 Sebastian Fohler
 Fastly, Inc.
 Federal National Mortgage
 Association (Fannie Mae)
 Fermyon Technologies
 Fidelity Investments
 Filecoin Foundation
 FinOpsly Inc.
 Finout
 Firefly
 FiveTwenty Inc.
 FLC Technology Group
 Flexera
 Flox
 Ford Motor Company
 Formal
 Fortified
 FOSSA
 FossilID
 Foundries.io LTD
 FPT Smart Cloud
 Company Limited
 FrOntierX Inc
 Framestore
 Framework
 Frontier
 ftrack AB
 FUJIFILM Corporation
 FullStackS
 FuriosaAI, Inc.
 Fused
 FusionLayer Inc

G

G-Research
 Gaia Information
 Technology
 Garden Technologies Inc.
 Garmin International, Inc.
 gateway
 Gcore
 GDIT General Dynamics
 Information Technology
 GEICO
 Gen Digital
 General Electric Company
 Genesis Global
 Technology Limited
 Getty Images
 Ghana Link Network
 Service
 Giant Swarm GmbH
 GienTech
 Gitbook
 GitCode
 GitHub, Inc.
 GitLab Inc.
 Gitpod GmbH
 Glide Identity
 Global Data Quantum
 Globant LLC
 Goldman Sachs & Co. LLC
 GoLedger
 Golem Cloud
 Goliath
 Grab Holdings Inc.
 Grafbase

GramLabs, Inc. (d/b/a
 StormForge)
 Granica
 Grape Up Sp. z.o.o.
 Graphcore
 Gravitational, Inc
 Green Hills Software LLC
 Greenpixie
 Grepr
 Greptime Inc
 groundcover Ltd.
 Grovf LLC
 Guangdong OPPO Mobile
 Telecommunications
 Corp., Ltd.
 Guida
 Guidewire Software, Inc.

H

HackerOne
 Hammerspace
 Hangzhou EMQ
 Technologies
 Hangzhou Harmony Cloud
 Technology Co., Ltd.
 Hangzhou Xinqi Electronic
 Co., Ltd.
 HAProxy Technologies
 Harness Inc.
 Hashgraph Foundry Inc.
 (DBA Hashgraph)
 HashiCorp Inc
 Hashlock
 Hasura, Inc.
 Hatchet

HawkStack Technologies
 Private Limited
 HCL Technologies Ltd.
 Hedera Hashgraph LLC
 Hedgehog
 Heeddata
 Here
 HERE Global B.V.
 herodevs
 Hewlett Packard Enterprise
 Development LP
 Hexagon AB
 Hidora
 Highlight
 Homestylar (Shanghai)
 Techonology Co., Ltd
 Hon Hai Precision
 Industry Co., Ltd.
 Honor Device Co. Ltd
 hoop.dev
 Hostersi Sp. o.o.
 Hound Technology Inc.
 dba Honeycomb
 HP Inc.
 HSBC
 Hubble Network
 Hud
 Hugging Face Inc
 Hunter Strategy
 Hushmesh
 Hydrolix
 HyperDX
 Hyperglance, Limited
 Hystax
 Hyundai Motor Group
 Hyve Managed Hosting

I

iA cloud Inc.
 IAR
 Icon Business Systems Ltd
 iCubed
 iExec Blockchain Tech
 IF Information Systems
 IFS World Operations AB
 Igalia, S.L.
 IITS Consulting
 ILKI FRANCE
 Imagination Technologies
 Ltd.
 Incloud Limited
 Indeed, Inc.
 induz
 IndyKite Inc.
 Infineon Technologies AG
 Infisical
 InfluxData Inc
 Infoblox Inc.
 Infomaniak Network SA
 Information Data Systems
 Information Security
 Infosys Limited
 Infracloud Technologies INC
 Infracost
 Inframappa
 initializ
 Innogrid
 innogrit Corporation
 inovex GmbH
 Inspur Cloud Information
 Technology Co.
 Instnt Inc.

Instruqt B.V.
 Integrated Computer
 Solutions, Inc.
 INTEGRITY Security
 Services
 IntellectEU
 Intellectual Highway, Corp.
 Intellias Global Limited
 Intelligent Systems Services
 Inter IKEA Systems B.V.
 Intercept
 Internet Initiative Japan
 Interop Labs
 Interop.io
 Intersect
 Intesi Group SpA
 Intuit, Inc.
 Intuitive Technology
 Partners, Inc.
 Invary
 IO Builders Blockchain
 Technologies & Ventures
 IOG Singapore Pte. Ltd
 IONOS SE
 IOTech Systems Limited
 IRNAS
 Isovalent Inc.
 ITGix
 ITQ Consultancy B.V.
 ITV PLC
 IVCISA

J

Japan Securities Clearing
 Corporation (JSCC)

Jetify
 Jetstack Ltd
 JFrog, Inc
 Jiangsu Bocloud Techno-
 logy Co., Ltd. (BoCloud)
 Jingdong Scientific and
 Technology Information
 Technology Co., Ltd.
 Jit
 Joby Aero
 JPMorgan Chase
 Jump Operations, LLC
 Juniper Networks, Inc.
 JUXT Ltd
 JVCKENWOOD Corporation

K

Kaleido
 kapa.ai
 Kapstan
 Kedify
 Keep Alerting LTD
 Kentik
 Kestrel AI
 Key State Capital
 Keyfactor
 Keysight Technologies Inc.
 Kion
 KIOUY
 Kioxia Corporation
 Kiratech SpA
 Kitware
 Kiwimoore
 Kloia Software and
 Consulting Ltd

Kloulfuse, Inc.
 KodeKloud
 Kodem Security
 Komodor Inc.
 Kondux
 Kong Inc.
 Konsulko Group
 Kosli
 KPIT Technologies Limited
 KPMG LLP
 krick.com GmbH + Co. KG
 Kry10 Limited
 Krypt Corporation
 kt cloud Co., Ltd
 KubeDB LLC
 KubeOps GmbH
 Kubermatic GmbH
 Kubernetes Innovation
 Labs LLC (Kubeshop)
 Kubiya Inc
 Kublr
 KUKA Deutschland GmbH
 Kumina B.V.
 Kunlun Technologies
 Kusari Inc
 KylinSoft Coporation
 Kyndryl
 KYXSTART

L

L4B Software GmbH
 Lablup Inc.
 Laika
 Larch Networks
 Last9 Inc

Launchnodes
 LayerZero Labs Ltd.
 Leading Point
 Leaseweb Global B.V.
 Legit Security
 Leica Camera
 Leminnov
 Lenovo (Beijing) Co., Ltd
 LG Electronics Inc.
 Lightmatter
 Lightning AI
 Lightwheel Limited
 Like Minds Consulting
 Linaro Limited
 LINBIT
 Lineo Solutions, Inc.
 LinkData Technology
 (Tianjin) Co., Ltd
 Linutronix GmbH
 Liquid Reply
 Lloyds Banking Group
 LMAX Exchange Ltd
 Lockheed Martin
 Loft Labs, Inc.
 Logitech Europe SA
 Logshero Ltd.
 London Metal Exchange
 London Stock Exchange
 Group
 Loongson Technology
 Corporation Limited
 Lowe's Companies, Inc.
 LPI.org
 LSD OPEN
 Lucid Computing Inc.
 Lucidity.cloud

Lutech Advanced
 Solutions S.p.A.
 Lynx Software
 Technologies

M

MacStadium
 MagicOrange Group
 Limited
 Mainsail Industries
 MangoBoost, Inc.
 ManTech International
 Corporation
 Mantech Solution
 Marvell Technology, Inc.
 Maryville Consulting
 Group
 Massdriver Inc
 MasterCard Incorporated
 Matrix I.T CloudZone LTD
 MATRIX Software
 MatX
 MavenSolutions
 Maxon Computer GmbH
 Mazda Motor Corporation
 MBDA Italia S.p.A
 McKinsey & Company, Inc
 MediaTek USA Inc.
 MediConCen Limited
 MegazoneCloud
 Meinberg Funkhrehn
 GmbH & Co KG
 Memfault Inc
 MemVerge
 Menzel IT GmbH

MEP
 Mercedes-Benz Tech
 Innovation GmbH
 Merly Inc
 MetalBear Tech Ltd
 Metoro
 MetroStar Systems
 Mezmo
 MIA s.r.l.
 Micas Networks Inc.
 Michelin
 Microchip Technology Inc.
 MicroEJ
 Micron Technology
 Micware Co. Ltd.
 Middleware
 Midokura Japan K.K.
 Milligan Partners
 mimik Technology Inc
 mind
 Minimus
 Minio, Inc
 MIPS Tech LLC.
 Mirantis, Inc.
 Miraxia Edge Technology
 Corporation
 MOBILTECH
 mogenius
 MontaVista Software, LLC
 Moody's
 MoonLight Marketing
 Morgan Stanley
 Motorola Solutions
 Moxa Inc.
 Mozilla Corporation
 MyFitnessPal LLC

N

Nadrama
 Napatech
 NatWest
 ndustrial
 Nearby Computing
 NearForm Ltd
 Nearmap Australia Pty Ltd
 Nebius B.V.
 Nebula Matrix
 Neo4j, Inc.
 NEOS
 NetApp, Inc.
 NetBird
 Netdata
 Netflix, Inc.
 Nethermind
 Netris, INC.
 NETSIA Inc.
 NETWAYS Managed
 Services GmbH
 Netweb Technologies
 Network Optix Inc
 NeuReality Ltd
 Neusoft Corporation
 New H3C Technologies
 Co., Ltd
 New Relic, Inc.
 nexB Inc.
 Nexthop AI
 NGINX International
 Limited
 ngrok
 NHN
 NHN Corporation

Nikon Corporation
 NIO
 NIPA
 Nippon Seiki Co. Ltd.
 Nirmata, Inc.
 Nissan Motor Co., Ltd.
 Nokia Corporation
 nops.io
 Nordic Semiconductor ASA
 Northflank Ltd
 Nosana
 Notarize, Inc. d/b/a
 Proof.com
 NovaGlobal Pte Ltd
 Novatus Global Ltd
 Novo Nordisk A/S
 Nscale
 NTT DATA MSE
 CORPORATION
 NTT, Inc.
 Numascale
 Numbers
 NuNet
 Nutanix, Inc.
 Nuvitek
 Nuvotex GmbH
 NVIDIA Corporation
 NXP Semiconductors
 Netherlands B.V.

O

Ochestra Technologies
 Limited
 Octopus Deploy PTY Ltd.
 OdineLabs

OGIS-RI Co., Ltd.
 Okahu Inc
 OKESTRO
 Okta Inc.
 Ollama
 Omatrix Ltd. Co
 Omnistrate
 Oodle AI, Inc.
 Opaque Systems Inc.
 Open iT, Inc.
 Open Source Automation
 Development Lab
 (OSADL) eG
 Open Source Consulting Inc.
 OpenAI Inc
 OPENMARU
 OpenMetal.io
 OpenNebula Systems
 OpenObserve
 OpenOps
 Openpaya
 OpenSource Connections
 OpenTeams
 OpenZeppelin
 Opera Norway AS
 Operant
 OpsLyft
 OpsMx
 OpsNow Inc.
 Orange SA
 Osaka NDS Co., Ltd.
 OSISM GmbH
 OSNEXUS
 OSSO B.V.
 Oteemo Inc.
 Oticon A/S

OTOY, Inc.
 Otterize
 OVH SAS
 Oxide Computer Company

P

Palark GmbH
 PalCNetworks
 Palo Alto Networks
 PANTHEON.tech s.r.o
 Paramount Software
 Solutions Inc.
 Parity
 Parler Cloud Technologies
 Parseable, Inc
 Parsimo Inc.
 Parsolvo
 Partior Pte Ltd
 Patchstack OÜ
 Pay-i Inc
 PBG Consulting
 Pelanor
 Pelotech
 Peloton Interactive
 Penten
 Percepio AB
 Percona
 Peridio
 Permify
 Permit.io
 Perr&Knight
 pgEdge
 Phala
 Phoenix Software
 International
 PHYTEC Technologie
 Holding AG
 Pickford
 Pier Cloud
 PingCAP
 Pionative
 Pioneer Corporation
 Pionix GmbH
 Pipekit Inc
 PlanetHoster Inc.,
 PlanetScale, Inc.
 PlatCo Group
 Platform Engineering Labs
 Platform Engineering
 Masters GmbH
 Platform9 Systems, Inc.
 Platformatic
 plural
 plusserver
 PLVision Corporation
 Point72, L.P.
 PointFive US Inc
 Polar Signals Inc
 Polygon Labs Services
 (Switzerland) AG
 Porch Financial
 Port
 Portainer.io
 Portal26
 Posit
 Postman
 Powerleader Computer
 system co.,ltd
 Precisely Holdings, LLC
 Precision Innovations Inc
 Preferred Networks, Inc.

PremAI
 Prescient Security
 Presidio Inc
 Provider
 Priceline.com LLC
 Print2Block
 Prodigy Education Inc.
 PRODYNA SE
 Profisea
 Proofcraft Pty Ltd
 ProSiebenSat.1 Tech &
 Services GmbH
 ProsperOps
 ProveAI, LLC
 Provectus IT Inc
 Proven Optics
 ProximaOps LLC
 PS Internet Company LLP
 PTC Inc.
 PTV Group
 Publicis Groupe
 PufferSoft
 Pulumi
 Pure Storage
 Puzzle ITC GmbH
 PwC
 Pydantic

Q

QANplatform
 QAware GmbH
 QBO
 Qiming Information
 Technology Co., Ltd.

Qinghong Electronic
(Suzhou) Co., Ltd.
QLAD
QlikTech International AB
Quality Cloud Corp.
QuantStack
QUANTUM C&S
Quesma
Quilyx
Qumulo
QuSecure Inc

R

R3 LLC
Rackner
Rackspace US, Inc.
Radisys Corporation
Rafay Systems, Inc.
Raft
Raintank, Inc. ? Grafana
Labs
Rambus Inc.
Randstad Digital Germany
AG
Rapid Cycle Solutions
RapidFort, Inc.
RAPIFUZZ
Rayls Foundation
Raynault VFX
Raytac
RBC Capital Markets, LLC
re:cinq
RealCloud
Reality Defender

Realtek Semiconductor
Corporation
RealTheory Inc
Rebussions
Recurve
Redeploy
Redocly Inc
Redpanda Data
Redpill Linpro
ReeVo Spa
REGnosys Limited
Relvy AI Labs, Inc.
Renault
reo.dev
Replicated, Inc.
Resolve Technology
ResolveAi
reThought Flood
ReversingLabs
Revolgy
Ricoh Company, Ltd.
RISCstar Solutions
Rivos Inc
RNG Technology
Roadie
Robin Systems, Inc
Robotec.ai sp. z o. o.
Robusta.dev
Rocket Software, Inc.
Rockwell Automation
Roku, Inc.
Root
Rootly, Inc.
Royal Dutch Shell
RSM Collective

RTE (Reseau de Transport
dElectricite)
RTX
Ruantong computer Co.,
Ltd.
Ruijie Networks Co., Ltd
RunWhen Inc
RX-M, LLC

S

S&P Global Inc.
SafeGraph
Safespring
SAIC Motor Corporation Ltd
Salesforce.com, Inc.
Sambanova Systems
Sanborn
SandboxAQ
SanDisk Technologies, Inc.
Sanofi
SAP SE
Sardina Systems
Sartura
SAS Institute Inc.
SASHA
Sateliot
Savoir-faire Linux
Sawmills
Scala Computing
Scalar
ScaleOps
ScaleUP
Scaleway
Schneider Electric
Schrödinger Inc.

Schweitzer Engineering
Laboratories, Inc.
Science Applications Inter-
national Corporation
SCITIX (SGP) TECH PTE. LTD.
Scope3
Scotiabank
Scott Logic Ltd
ScoutAPM
SDAX Exchange Pte Ltd
Seacom srl sb
Seafarix
Seagate Technology LLC
Seal Security
Seal Software (Shenzhen)
Co.,Ltd.
Searce
Second State
Secondfront
Sedai
SEIYAJ TECH
Selective Insurance Group
Semaphore
Senofi
Sense Reply
Sent, Inc.
Sentry Software
Seowon Information
Co., Ltd.
SerNet GmbH
servicememe
ServiceNow
SERVICEWARE
SGDL Innovation SA
Shabodi

Shandong Cvicse
Middleware Co., Ltd.
Shanghai Mandao
Technology Co., LTD
Shanghai UniVista
Industrial Software
Group Co., Ltd.
Shanghai Yunsilicon
Technology Co.,Ltd.
Shanghai Yunzhou
Technology Co.,Ltd.
(ZStack)
Sharktech Inc
SHE BASH
Shenzhen Jaguar
Microsystems Co., Ltd.
Shenzhen Wise2C
Technology Co.,Ltd
SHI International Corp.
Shielded Technologies
ShineSoft Co. Ltd.
Shopify Inc.
Shutterstock, Inc.
SICHUAN HONGXIN
SOFTWARE CO., LTD.
SICHUAN HUACUN ZHIGU
TECHNOLOGY CO., LTD
Sichuan Huakun Zhenyu
Intelligent Technology
Co., Ltd.
SideFX
Sidero Labs
Siemens AG
SiFive
Signadot
SigNoz

Silence Laboratories
Silicon Innovation
Microelectronics Co., Ltd.
Silicon Laboratories Inc.
Silver Lining Information
Technology Co., Ltd.
SK hynix Inc
Skydance Animation, LLC
Skyloud
Skypuzzler
SKYXOPS Corp.
SmartBear Software, Inc.
SmartDV
Smartiful, Inc.
Snowflake Inc.
Snyk Limited
Société Générale
Socionext Inc.
Socket
socradev GmbH
SoftBank Corp.
Software Mind
Software Safety
Technology co., Ltd
SoftwareONE AG
SoKube
Solidigm
Solo.io, Inc.
Sonatus, Inc.
Sonatype, Inc.
SOOHO.IO
SORAMITSU CO., LTD.
Soundpatrol Inc.
Source Allies
Southern California
Edison

Southworks	State Street Bank and Trust Company	SVA System Vertrieb Alexander GmbH	TDT AG	The Constant Company, LLC / Vultr	Trimble Inc.
Spacelift, Inc.	STCLAB	Sweet Security	team.blue	The Foundry	Tripadvisor LLC
SpacemiT (Hangzhou) Technology Co. Ltd	steadybit GmbH	Swisscom	Tech Mahindra Limited	Visionmongers Limited	TripleID Solutions
SPARE Consulting	Steamhaus	Switch, Inc.	Tech Soft 3D Inc	The Guild	TrueFoundry
Sparkfabrik srl	Stellar Development Foundation	Symbotic	technative	the LEGO Group	TrueFullstaq
Sparkgeo	StepSecurity	Symmera Inc	Technology Innovation Institute	The Mastermind Group	Truepic, Inc.
Spatial	StitcherAI	Symphony Communication Services LLC	Tektronix Inc	The Qt Company Oy	TrustAsia Technologies
Speakeasy Development, Inc.	STMicroelectronics International N.V.	Synadia Communications	Telechips, Inc.	The Scale Factory Limited	Trustwise
Spectro Cloud, Inc.	STORDIS GmbH	Synax GmbH	Teledyne LeCroy Inc	The SIMI Group	Tsavorite Scalable Intelligence, inc
SpeedScale	Storm Reply srl.	Synopsys, Inc	Telefonica, S.A.	The Walt Disney Studios	Turk Telekomünikasyon A.S.
Sphinx	Storware	Syntasso	TELUS Corporation	Thoras.ai	turntabl
SpiralDB	Story Foundation	Synx Data Labs	Tempest Labs Inc	Thornstein Groep	Tuxera Inc.
Spirent Communications Inc	Strategic Blue	Synyega	Temporal Technologies Inc	ThoughtWorks, Inc	Txture
Spitzkop	Stratox Cloud Native	Sysdig, Inc.	Tencent Holdings Limited	Thread AI	Tyk Technologies Ltd.
Splunk Inc.	Strava	SysEleven GmbH	Tencent Technology (Shenzhen) Co., Ltd.	Thunder Software Technology Co. Ltd.	
Spotify AB	Structsure, LLC	SYSGO GmbH	TenneT	Tianjin Aojin Network Technology Co., Ltd.	U
Springer Nature	Styra Inc		Tensor9	Tigera, Inc.	U.S. Bank
Sprint Corporation	Succinct		TensorMesh, Inc.	Tigris Data Inc.	Uber Technologies Inc.
Squarespace, Inc.	Sudo Information Technology Co. Ltd.		Tenstorrent	TIM S.p.A	UBS AG
SQUER Holding GmbH	Sumitomo Electric Industry, Ltd.	T	TeraSky	Tinfoil	Ufi Space
SQUILD	Sumo Logic, Inc.	TAC Security Inc	Ternary	TINTRI BY DDN INC	Ultimum Technologies
SSL Inc	Super Micro Computer, Inc.	Taiji Computer Corporation Limited	Terracotta AI	Tokenovate	Ultraviolet Consult DOO
Stability AI Ltd	Super Protocol	taikun.cloud a.s.	Terramate	TomTom International B.V.	Ulyssean
StackGen	SuperOrbital, LLC.	Tailscale Inc	TestifySec	Trace Machina - Simulate Robotic Systems	UMB AG
stackgenie	Surveil	TangentWorks	Testkube	Trace3	Unified Streaming
StackGuardian	SUSE LLC	Tangoe US , Inc.	Tetrate.io	Traceroute42	Unikraft GmbH
StackHPC	Suzhou Centec Communications	Target Corporation	Teuto.net Netzdienste GmbH	TradeHeader	Union.ai
STACKIT GmbH & Co. KG	Suzuki Motor Corporation	Tata Communications Limited	Texas Instruments Incorporated	Traefik Labs SAS	Uniqconn
Stacklet		Tata Consultancy Services Limited	Thales SA	Trail of Bits	Unravel
Stacklok		TCC Consulting Limited	Thavron Solutions	Travelping GmbH	Upbound, Inc.
Stackwatch Inc			The Coca-Cola Company	Trend Micro Incorporated	UpCloud Ltd
Starix Technology, Inc.				Triam Security	Upscale AI, Inc.

Upsider
Uptycs, Inc.
USU GMBH
Utilidata
UtilityAPI
UTMStack
Uturn Data Solutions

V

VA Linux Systems Japan
K.K.
Valor Capital Group
ValueMentor
Valve Corporation
Vantage
Vantor
Vates
Vattenfall Eldistribution AB
Vaxowave
vbrick
VDURA
Veea Inc.
Vega Cloud Inc
Velocity
Ventana Micro Systems
VeriSilicon, Inc.
Veritas Automata
Verizon Corporate Services
VES LLC
VESSL AI Inc
VEXXHOST, Inc.
VHL Technologies
Viavi Solutions Inc.
VictoriaMetrics
Viettel

Virasemi Inc
Virtasant
Visa Inc.
Vivendi
vivo Mobile
Communication Co., Ltd.
VJSAR Consultant Service
Ltd
Vodafone Group Plc.
VSHN AG
VyOS Networks

W

Wētā FX Limited
Walmart Inc.
Wanchain
Wand Cloud
Wargaming.net Limited
Wavecon GmbH
Wavelabs
Wellington Management
Company, LLP
Welotec
WeScale SAS
Western Digital
Corporation
WhaTap Labs Inc
Wherobots, Inc
Whirlpool Corporation
Whitestack LLC
WhizUs GmbH
WIIT AG
Wind River Systems, Inc.

Wing Programming
Language(Monada EBS
LTD)
Wipro Limited
Wiv.ai
Wiz
Workday, Inc.
WorkOS
WorldTech IT LLC
Worley Limited
Worteks
WPP plc
WSO2 LLC.
Wuhan Yangtze
Computing Technology
Co.LTD

X

x-cellent technologies
GmbH
x-ion GmbH
Xero Limited
xFlow Research.com
Xfusion International PTE.
LTD.
Xi'an Tieke Jingwei
Information Technology
Co., Ltd. (CARS)
Xosphere
XPerf Inc.
Xsight Labs


Y

YÄRKEN

Yazaki Corporation
Ybor.AI
Yellowbrick Data
YLD! Limited
Yonder
Yovily

Z

ZConverter
Zed
ZEDEDA, Inc.
Zephr.xyz
Zesty Tech Ltd.
zode
Zoi
Zoniqx Inc.
Zoss Team, LLC
ZTE Corporation
Zuplo Inc
ZutaCore
Zylo Inc



The Linux Kernel Organization

In tandem with security, nurturing open source innovation to create a better world is at the heart of the Linux Foundation's activities.

In 2021, we celebrated the Linux kernel's 30th birthday. Four years later, Linux remains among the top three global open source projects in terms of development velocity, as it has been for many decades. Each release results from the work of thousands of contributors worldwide and hundreds of different organizations. Out of 2,057 developers who contributed to the Linux 6.16 release, 310 contributed their first patch to the kernel. The kernel community actively maintains a steady flow of innovative improvements to expand the footprint of Linux and enhance its capabilities for new hardware and workloads, as well as integrating the memory-safe Rust language into the codebase to ensure that developers wishing to use that language can also participate.


The Linux kernel community formalized the [process](#) for taking remedial measures when Code of Conduct violations take place.

The Linux kernel community has streamlined the process of [assigning CVEs](#) (security notifications) to relevant bug fixes over the past year, providing detailed information for all consumers of Linux to be able to programmatically determine if they are vulnerable to the specific issue or not. This is essential given that the rate of CVEs being issued for the Linux kernel, 13 a day, precludes the possibility of reviewing them and is key in ensuring that Linux will be able to meet the upcoming requirements of the E.U. with the recently enacted Cyber Security Act, which will be required in the next years by all manufacturers.




Linux kernel maintainers and developers will get together in Japan in December for the annual Maintainer Summit, Kernel Summit, and Linux Plumbers conferences to discuss a range of technical and procedural topics, such as the role of AI in the kernel development process. The kernel community has been using homegrown AI tools for automating selecting fixes for stable releases for over a decade and as a signal for potential CVE designations of kernel bug fixes. The discussions on AI will be about regulating patches containing AI-assisted code among other important topics.

Serving Nearly 1,500 Open Source Project Communities



Project technical segment

- Cloud, Containers, & Virtualization **23%**
- Networking & Edge **14%**
- AI, ML, Data & Analytics **12%**
- Cross-Technology **12%**
- Web & Application Development **10%**
- Privacy & Security **5%**
- Blockchain **4%**
- IoT & Embedded **4%**
- DevOps, CI/CD, & Site Reliability **3%**
- Open Source & Compliance Best Practices **3%**
- System Administration **2%**
- Storage **2%**
- Linux Kernel **2%**
- System Engineering **2%**
- Open Hardware **1%**
- Safety-Critical Systems **1%**
- Visual Effects **1%**




Project type

- Open Source Software **72%**
- Open Standard / Specification **17%**
- Community Initiative **5%**
- Open Data **2%**
- Open Hardware **2%**
- Open Governance Network **1%**
- Peer Network **1%**

By the numbers: The Linux Foundation 2025

21,624 organizations contributed to Linux Foundation open source projects.




LF Education enrolled **329K unique learners**




2,126 members supported the Linux Foundation.




402k+ LF YouTube subscribers with **1.73M views** of open, global, always accessible event content



100,182 developers actively contributing




7 Mentorship programs with **15k mentee applications**



LF Education issued **45.8k certifications**



17k repos in 2025 over 11k in 2024



The Linux Foundation hosted **229 events** with **118K+ attendees** from **10k organizations**, across **110 countries**.





Accelerating Community Engagement

Standards as Substrate: Cultivating Innovation Through Standards and Specifications

It's no surprise to us that open collaboration wins, and this year, our standards and specifications projects have been winning big: [Alliance for Open Media](#) took home a prestigious Technology and Engineering Emmy® Award for the AV1 Video Codec, and [Overture Maps](#) was named to [Fast Company's 2025 Next Big Things in Tech](#) list.

These projects exemplify how innovation is accelerating faster than ever, especially where open collaboration on software and standards are happening together. Technologies like AI and distributed trust are reshaping how we exchange information. While the pace of change can feel overwhelming, it is creating immense value—benefiting both industry solution providers and customers.

As our latest [State of Open Standards Report](#) shows, standards and specifications play a key role in innovation, balancing new ideas and market opportunities with consensus-driven stakeholder feedback, interoperability, and regulatory requirements. Over three-quarters of respondents said standardization helped their markets mature, fostered competition and innovation, and eased regulatory and business compliance. Moreover, most respondents also viewed heavy adoption of open source software and open specifications as a far better indicator of innovation than traditional GDP measures (e.g., patent counts)..

The Linux Foundation and Joint Development Foundation (JDF) created innovative approaches to standards development that have provided a

faster-paced standards development process. Our standards project enables communities to develop standards in collaboration with open source projects, creating a multiplier effect. This is



readily apparent in our standards communities focused on AI. This year, several important specification releases from LF and the JDF projects are laying the foundation for a safe, secure, innovative, and interoperable AI ecosystem:

- ▶ [Ultra Ethernet Consortium v1.0 specification](#) supports AI/HPC workloads over Ethernet, reducing the total cost of ownership while enabling faster training and shorter time-to-model.
- ▶ [SPDX 3.0](#) extends beyond software to systems, standardizing metadata for model and dataset discovery, reuse, and compliance audits.
- ▶ C2PA specifications operationalize provenance and add AI-specific features, including companion [Guidance for Artificial Intelligence and Machine Learning](#), lays the groundwork for detecting the tampering of media files and streams, and compliance with “deepfake” regulations.
- ▶ [GraphQL’s September 2025 specification](#) release introduces features that strengthen its role as a foundation for AI assistants, codegen tools, and autonomous agents—such as descriptions on executable documents, Schema Coordinates, and OneOf input objects.
- ▶ Open source projects are also advancing AI specifications: Agent2Agent is building a protocol for autonomous AI agent interoperability; AGNTCY is enabling agent discovery and connection; Docling is developing tools to convert unstructured documents into structured AI-ready formats; and BeeAI is delivering a production-ready framework for multi-agent, multi-platform systems.

Building globally relevant standards requires broad collaboration. LF and JDF projects are partnering with European standards organizations such as ETSI on regulatory frameworks, publishing international standards with ISO, co-developing open source projects with IEC, and voicing developer perspectives on major legislation such as the CRA (more on that below). Recent strategy papers from the U.S. and EU emphasize open collaboration as the engine of pro-competitive innovation—and our projects are answering that call.

To support this momentum, the LF is also investing in new resources and training materials for standards and specifications development. These offerings are designed to help contributors quickly gain the skills and context they need to participate in specification development and standards work—lowering the barrier to entry and enabling more voices to shape the future of open innovation. To learn more about these resources and other high-profile standards news, [subscribe to our quarterly newsletter](#).

Securing the Commons: Open Source Policy, Standards, and Global Coordination

With the ubiquity of open source comes greater responsibility on the projects, communities, and their stewards.

Open source ecosystems operate in an environment of increasing complexity. Shifting international dynamics and evolving security considerations challenge traditional models of global collaboration. In 2025, the Linux Foundation's Board of Directors established a Policy, Security, and AI committee to help coordinate and navigate proactively on strategic challenges that require a community response. This committee is also coordinating with project-level and regional policy collaboration efforts to help develop a more consistent collective response across communities.

Yet, these pressures create opportunities to demonstrate open source's enduring value: fostering innovation through inclusive participation across geographical and institutional boundaries and establishing trust through transparency and shared responsibility. Open source creates an inclusive, shared commons of research and value that anyone can benefit from. The LF's response centers on strengthening collaborative frameworks that transcend regional differences, positioning openness and technical excellence as unifying principles in an increasingly fragmented landscape.

Recognizing these increasing challenges, the LF has systematically expanded its regulatory engagement capacity across regions, deliberately framing emerging requirements as catalysts for ecosystem advancement. ETSI membership and EU MSP-ICT participation enable our communities to provide substantive input into European standards development, from CRA implementation to AI Act



frameworks, not merely for compliance, but to elevate collaborative practices across the entire open source landscape. The OpenSSF's leadership role in CRA implementation exemplifies this philosophy: Regulatory change becomes an opportunity to strengthen security practices, improve project governance, and demonstrate open source's capacity for self-improvement consistent with the community's core principles. This distributed expertise model ensures that regulatory responses enhance rather than constrain collaborative innovation.

Strategic alliances enable the LF to support and enhance the messaging of other like-minded organizations, modeling the same collaboration that drives our technical project communities. Participation in the Geneva Dialogue on Responsible Behavior in Cyberspace and partnerships with CEPS and OpenForum Europe reflect its deep engagement in multilateral discussions. Through the Open Source Congress, open source foundations orchestrate unified messaging toward policymakers and global stakeholders. The LF's central role in UN Open Source Week 2025 underscores the ecosystem's evolution into a recognized force in global digital governance. This alliance strategy extends internally to our members, as we support and inform their legal counsel and public policy teams about the broad impact of open collaboration across the technological ecosystem.

The LF engages selectively in providing feedback, education, and best practices guidance to advance open collaboration and open source development policy goals. AI governance engagement demonstrates this strategic focus: Partnering with Hugging Face, Mozilla, and others to develop collaborative guidance on EU AI Act implementation ensures open source methodologies remain viable pathways for responsible AI development. Similarly, the LF's contributions to the review of the EU standardization system advocate for processes that accommodate distributed innovation models open to input from the wider open source community and other nontraditional participants in standards development. These targeted interventions aim to establish transparent, inclusive development practices as regulatory defaults rather than exceptions, creating frameworks that enhance rather than constrain collaborative innovation.

LFX: Understanding Code Contributors and Velocity Through Data

We know that the strength of any open source project depends on its contributors and speed of work.

Over the past year, we've focused on supporting them by leveraging data. Our platform has enabled 100,182 contributing developers and 21,624 contributing companies to actively improve 17,023 repositories, collectively making 1.07M commits.

Listening to What Maintainers Need

Open source maintainers want:

- ▶ More time to work on open source projects
- ▶ Better tools and processes to streamline their work and bring on new contributors

Many developers work on open source projects in their off-hours—late at night or on weekends—indicating insufficient workday time for these contributions. With data, we've opened conversations with employers to provide more dedicated time.

We've also improved tools for maintainers to develop projects efficiently. Each week, 999M lines of code are added and 188M lines removed, ensuring infrastructure and compliance. Through LFX Crowdfunding, we raised \$857,899, supported 15,128 mentees, added 304 corporate CLA signatures and 6,126 individual CLA signatures, and hosted 20,728 meetings. These improvements help maintainers focus on building better open source projects.



LFX Insights: Understanding the World's Most Critical Projects

Rapid growth in 2025 established LFX Insights as the leading platform for understanding project health across the open source ecosystem. Following its relaunch in May, adoption accelerated quickly, with monthly active users rising from 1,693 in April to 13,860 in November—an 8.8x increase in just seven months.

LFX Insights also expanded significantly beyond Linux Foundation projects. The platform now covers more than 1,000 LF projects and more than 3,000 non-LF projects, with many leading communities relying on it to understand trends in activity, contributions, and collaboration.

In 2025, LFX Insights tracked 17,023 repositories containing 1,070,892 commits. Developers added 999M lines of code and removed 187M lines, giving maintainers clear visibility into how their projects evolve. The platform also monitored 1,082,773 logged issues and 655,239 pull requests. Since launch, Insights has recorded 25.9M commits, 10.3M logged issues, and 4.6M pull requests across supported projects.

A major focus this year was improving data quality. Insights removed 7.4M duplicate records, enriched 4.1M records, and manually reviewed 55,000 contributor profiles to improve identity accuracy. These enhancements power advanced analytics capabilities such as the LF Health Score, the LF OSS Index, and detailed contributor and organization-level metrics.

Altogether, the platform now tracks 557M activities, maintains 3.5M contributor profiles and 630K organization profiles, and integrates data from more than 20 sources. This breadth and depth of data gives maintainers and OSS users a reliable and data-driven foundation to navigate and build the future of open source.

Taking Action with Data

LFX Insights provides actionable insights. For example, if code velocity slows down, maintainers can identify bottlenecks.

Proactive Support

At the LF, we actively engage with maintainers and communities. When we see a problem, we step in to help—whether it's fixing a security issue, improving documentation, or advocating for more developer support.

Protection Against Patent Aggression by NPEs

The escalating threat of patent aggression posed by Non-Practicing Entities (NPEs, also commonly referred to as “patent trolls”) continued to grow in 2025, continuing a pattern observed in recent years.

NPE lawsuits targeting open source adopters surged to their highest level this year. This surge coincided with significant policy headwinds: Recent changes in *inter partes review* (IPR) discretionary denial rules resulted in a sharp decline in IPR institution rates at the United States Patent and Trademark Office (USPTO). Since IPRs are a strongly preferred, low-cost tool for affected parties to challenge low-quality patents, this shift underscores the critical importance of the Open Source Zone (OS Zone).

One of The Linux Foundation’s core defensive measures is the Open Source Zone, a partnership with Unified Patents launched in 2019 and expanded in 2024. The OS Zone employs a comprehensive deterrence strategy—utilizing administrative challenges, crowdsourced prior art contests, and intelligence sharing—to protect open innovation from NPE aggression. LF members (more than 1,400 companies) receive access to extensive tools and benefits, based on membership tier, for deterring NPE aggression. In response to both the increased NPE aggression and policy challenges, the OS Zone dramatically ramped up its deterrence activities, tripling the number of crowdsourced prior art contests and almost doubling the number of administrative challenges to NPE-held patents compared to 2024.



Late in the year, further changes to IPR rules—which would make challenging low-quality NPE-held patents more difficult and expensive—were proposed by the USPTO. The LF collaborated with organizations such as the Electronic Frontier Foundation (EFF) and GitHub, filed comments opposing these changes, and [mobilized its community to do the same](#), reinforcing our commitment to safeguarding open innovation.

AI Runs on Open Source... and Real Humans

AI is capturing imaginations and budgets, but beneath every high-performing AI system lies an intricate, often overlooked truth: AI runs on open source and real people.


Open source frameworks, cloud native infrastructure, orchestration layers, and deployment tooling form the spine of modern AI. Without them, no model trains, no inference scales, no system deploys reliably. But just as critical are the people who understand that stack, who guide these technologies with judgment, and who build the bridges between innovation and impact.

Too often, organizations focus narrowly on recruiting AI specialists, assuming that a few well-placed data scientists or model developers will unlock competitive advantage. The reality is more complex and more human. AI succeeds only when embedded in a broader ecosystem of skilled professionals: cloud architects who ensure scale and reliability, DevOps teams who integrate models into pipelines, and engineers fluent in Linux, Kubernetes, and PyTorch who make the system actually run.

FIGURE 8: ORGANIZATIONS ARE UNDERSTAFFED IN KEY TECHNOLOGICAL AREAS

% of understaffed organizations:

Source: 2025 Tech Talent, Q19, Sample Size = 556, DKNS excluded (3% to 15%)



The Gaps Behind the Hype

Just as no AI stack is complete without open frameworks and cloud capacity, no AI strategy succeeds without human oversight. From prompt engineers to AI governance leads, new roles are emerging that underscore how much discretion, ethics, and contextual awareness still matter. Even the most powerful model needs someone to validate, fine-tune, and decide when it's wrong.

The [2025 State of Tech Talent Report](#) shows that while 94% of organizations expect AI to deliver significant value, most are unprepared to support it at scale. They're understaffed not only in AI/ML but also across the infrastructure that supports production-grade deployment: cloud computing (59%), platform engineering

(56%), cybersecurity (65%), and more. Successful innovators increasingly recognize that judgment, real-world experience, ethical decision-making, and strategic reasoning cannot be automated away.

Open source plays a crucial role here. Community-driven tools like Kubernetes, TensorFlow, and vLLM evolve faster and more transparently than closed alternatives. Upskilling across these open ecosystems isn't just good practice; it's essential to remaining competitive.

The path forward isn't AI-only. It's AI-plus: plus infrastructure, plus people, plus open collaboration. That's how organizations move from experimentation to execution.

In the age of AI, innovation doesn't come from machines alone. It comes from the skilled people and open systems working behind the scenes. AI runs on open. And us.

Scholarships

In 2025, The Linux Foundation awarded more than 500 scholarships for training and certification to individuals from 117 countries in memory of Shubhra Kar, LF's CTO. Hundreds more scholarships were awarded via partnerships with nonprofits, including Blacks in Technology, Living Open Source, TransTech Social Enterprises, and Women Blessing Women.

New Products Ignite Open Source Innovation

In addition to our core certification and course offerings, this year we made significant investments in new educational resources to help the open source community grow and learn.

Cybersecurity Skills Framework

In May we launched the [Cybersecurity Skills Framework](#), a global reference guide that helps organizations identify and address critical cybersecurity competencies across a broad range of IT job families that extends well beyond cybersecurity specialists.

Produced in collaboration with the Open Source Security Foundation (OpenSSF) the framework delivers actionable guidance to enterprise leaders looking to systematically reduce cyber risk.

Executive Education

Today's digital landscape requires executives to make faster, more informed decisions about complex technology issues, often without a technical background. To meet this challenge, we launched the [Executive Education](#) program, a new learning experience designed for senior leaders.

Each series brings together a select cohort for four intensive, discussion-driven sessions. Guided by expert facilitators, participants learn alongside peers facing the same high-level pressures. The emphasis is on real-world strategy, actionable insights, and direct peer exchange, not technical execution. Our first series focused on cybersecurity strategy and risk management.

Microlearning

Our new microlearning products address a number of needs. First, they significantly expand the learning content we offer at no cost. Second, they enable us to provide learning resources across a range of formats, including: microcourses, videos, tutorials, interactive infographics, and more. Lastly, and perhaps most importantly, they deliver dynamic learning at the speed of tech, with new content launching weekly.

This year alone, we launched more than 200 new microlearning resources, resulting in more than 13,000 registrations. The most in-demand topics included: scanning container images, shifting security left, AI accelerators, WebAssembly, and cybersecurity.

In addition to the free microlearning resources, we also launched a range of premium microlearning resources through our THRIVE annual subscription offerings.



SERENGETI NATIONAL PARK, TANZANIA

2025 Course & Certification Launches

Instructor-Led Courses

- ▶ Cilium in Action: Kubernetes Security & Insight
- ▶ WebAssembly in Action: Develop, Deploy, and Debug
- ▶ Zephyr RTOS Programming

Express Learning

- ▶ Authentication and Authorization for Web/AI
- ▶ Secure AI/ML-Driven Software Development
- ▶ Understanding the EU Cyber Resilience Act (CRA)

E-Learning Courses

- ▶ A Beginner's Guide to Linux Kernel Development—Japan
- ▶ Feature Flagging with OpenFeature
- ▶ GitHub for Open Standards Development
- ▶ Introduction to JavaScript Security
- ▶ Introduction to Kubernetes—Japan
- ▶ IT Specification Basics for New Developers
- ▶ Open Source Essentials for Developers
- ▶ Open Source RT-Thread RTOS on RISC-V
- ▶ Porting Code from Arm to RISC-V
- ▶ Security for Software Development Managers

Certifications



LF Research

In 2025, LF Research produced new, world-class insights to inform, engage, and strengthen open source communities worldwide.

Combining rigorous data analysis with global stakeholder insights, we published more than two dozen studies at the intersection of open source and industry, government, and society. Here are a few highlights:

AI

Unsurprisingly, AI was the focus of numerous studies this year. We assessed the [Economic and Workforce Impact of Open Source AI](#) globally, in Latin America, Africa, the Middle East, Türkiye, Canada, and across Asia-Pacific economies and explored concepts of sovereignty in [The State of Sovereign AI](#). These studies reveal AI's powerful role as a general purpose technology, driving innovation and economic value, with open source ecosystems being essential to realizing its promise.

Technical Talent

With the release of our [Tech Talent 2025](#) report, we provided fresh evidence to illustrate how AI is reshaping the technology workforce. Contrary to headlines, our study revealed a global net hiring effect, requiring organizations to upskill and cross-skill in order to meet their AI and innovation agendas while simultaneously retaining talented teams.



Hilary Carter and Anna Hermansen of LF Research, and Frank Nagle, Chief Economist, discuss the workforce and economic impact of open source AI in Washington, D.C.

Cybersecurity and the CRA

Legislation like Europe's CRA inspired two reports this year, [Unaware and Uncertain](#) measuring CRA awareness, and [Pathways to Cybersecurity Best Practices](#). These studies support manufacturers, stewards, and contributors alike as they work to strengthen their security practices to meet regulatory requirements.

Commercial Open Source and OSPOs

2025 marked the inaugural [State of Commercial Open Source](#) report, illuminating the growing role that open source plays in global venture markets, the superior financial performance of commercial open source companies, and that business and community value is mutually reinforcing. Our [2025 State of the OSPO](#) report illustrates the benefits of a well-defined open source strategy, documenting the OSPO's role in enabling culture, innovation, and collaboration.

Our research community enables us to create accessible content that strengthens project, enterprise, and government decision-making on all matters open. We extend our sincere gratitude to the thousands of contributors who make this work possible—those who participate in our surveys, lend their perspectives during interviews, and sponsor our studies.


LF Events

In 2025, the Linux Foundation's events continued to serve as the global meeting place for open source collaboration, drawing together developers, maintainers, end users, enterprise leaders, and policymakers across every major region.

These gatherings once again demonstrated the vital role that in-person events play in advancing open source technology, strengthening project ecosystems, and ensuring communities have the space to innovate, align, and build together. While virtual-only events are no longer a major focus, the availability of high-quality session recordings continues to expand global access to open source education, creating a lasting knowledge base for developers everywhere.

The cloud native community remained exceptionally strong in 2025, with **KubeCon + CloudNativeCon Europe** marking a major milestone as the largest KubeCon event ever, welcoming **12,418 attendees from 2,808 organizations** to London in April. Energy was high throughout the week, with attendees eager to learn from **665 speakers, 325 sessions**, and deep dives into this year's dominant themes—**platform engineering, AI, and observability**. The event also celebrated a historic community achievement: **ten years of the Cloud Native Computing Foundation (CNCF)**.

Across the rest of the KubeCon portfolio, the momentum continued. **North America 2025** in Salt Lake City gathered **150 more attendees than in 2024**, despite the government shutdown causing geopolitical challenges and widespread flight cancellations. **KubeCon + CloudNativeCon India**, now in its second year,



expanded significantly, bringing together **4,017 attendees from 1,256 companies**, with 92% coming from India's fast-growing cloud native ecosystem. And in a major new milestone, **KubeCon + CloudNativeCon Japan** launched for the first time in 2025, selling out both registration and sponsorships. With **1,502 attendees from 471 companies**—74% based in Japan—the debut event showcased the rapid growth of open source infrastructure and cloud native adoption across Asia.

AI continued to cut across nearly every event series in 2025, driving both content and attendance. **AI_dev Europe**, co-located with **Open Source Summit Europe** in Amsterdam, brought together builders working in open source AI, GenAI, and ML. The **PyTorch Conference 2025** saw another dramatic surge in participation, reaching **3,432 registrants from 1,026 organizations**, with top job functions including AI engineers, ML engineers, and executive technical leaders—a testament to the accelerating demand for open, developer-driven AI innovation.

Open Source Summits continued to anchor the LF's regional programming. **OSS North America** made its way to Denver, **OSS Europe** to Amsterdam, and **OSS Japan** to Tokyo, where interest was so strong that **OSS Japan + Automotive Linux Summit + AI_dev Japan** (December 2025) surpassed **all previously recorded attendance levels**, underscoring Japan's growing leadership in AI, automotive, and embedded open source technologies. New regional expansions also took shape; **Open Source Summit India** debuted with **984 attendees**, and **Open Source Summit Korea** launched in Seoul with **663 attendees**, further strengthening the LF's footprint across Asia.

While global socioeconomic factors—economic uncertainty, travel restrictions, geopolitical concerns, and reduced corporate travel budgets—affected attendance and sponsorship levels in some regions, the overall picture remained resilient. Certain communities, most notably in Japan and across AI-focused

technologies, experienced strong growth, demonstrating that open source continues to be a strategic priority even in challenging economic conditions.

Despite the full return to physical events, virtual access remains essential. Across the LF's YouTube channels, session recordings continued to reach a global audience. In 2025, LF-affiliated channels grew to **402,366 subscribers** and generated **1.73 million views**, with the majority of viewer engagement coming from openly available conference content. These platforms play a crucial role in democratizing access to open source education, allowing developers who cannot travel to stay connected with cutting-edge technical insight and community dialogue.

Together, the LF's 2025 events reaffirmed the importance of in-person collaboration while simultaneously expanding global access to knowledge. Despite economic headwinds, the passion, innovation, and cross-project collaboration seen across these gatherings highlight the continued growth and resilience of open source communities worldwide.



Mentorship

Growing new talent and attracting new developers is challenging for open source communities. Yet, it is vital to reach out to train the next generation of developers to keep the open source communities healthy and sustainable.

Equitable access to learning resources is a barrier for a significant number of new developers. It isn't easy for new developers to get a start in open source, connect with open source communities and contribute to them. It is equally challenging for employers to find new developers to add to their technical projects.

Since its inception in 2019, the Linux Foundation's mentorship program has gained popularity across the globe. Each year, our mentorship programs train, educate, and mentor new developers strengthening open source communities. Several of our graduates have gone on to become maintainers of open source projects they contributed to during their mentorship program work. Select quotes from our graduates:

"An incredibly valuable experience that has broadened my technical horizons and enhanced my practical skills."

- CNCF Mentorship Graduate

"The Kernel was this magical thing I only knew how to build but couldn't begin to imagine how it worked. The program bridged that gap and showed that I could, even with limited experience, contribute to the Kernel in meaningful ways."

- Linux Kernel Mentorship Graduate

"Profoundly transformative and immensely rewarding – transitioned from being a novice in the open-source community to a confident contributor. Enriched my professional development and instilled a strong sense of perseverance and continuous learning."


- CNCF Mentorship Graduate

"Collaborative effort led to significant personal and professional growth that equips me with the skills to tackle similar challenges in the future and leaves me with a profound sense of accomplishment."

- LF Decentralized Trust Mentorship Graduate

Our graduates say that learning from experts and experienced developers in open source communities and collaborating with them has given the opportunity to contribute to real-world projects. They say that the mentorships helped them grow from novices to confident contributors, providing valuable hands-on experience. These experiences contributed to their personal, professional, and technical growth.

We at the LF believe today's new developers are tomorrow's leaders, maintainers, and experts. We are proud of our six-year journey to provide learning resources for new open source developers, opportunities to experts in open source communities to train and mentor the next generation, and make newly trained talent available to prospective employers.



Spheres of Impact

The Open Fork Forges Ahead

Following on the heels of successful forks of formerly open source projects last year, 2025 proved that forking is not merely a technical response to licensing disputes, but also a transformative force that fundamentally reshapes the economics of enterprise software.

Thanks to the hard work of several LF projects and their contributing members and organizations, the “Right to Fork” is now well-ensconced in the options of any open source community to protect their commons from license changes and other acts to reduce and remove community control.

Together, the Linux Foundation and our member organizations are building a better understanding of what makes a project fork succeed. At the first CHAOSS Data Science Hackathon in Denver, researchers analyzed open source project behaviors and identified critical failure and success patterns. The study found that forked projects under neutral foundation governance consistently attract more diverse contributor bases than their proprietary predecessors. Some patterns were obvious. Having a blog, newsletter, and public repo matters in success. Other success factors were less obvious, such as language differentials affecting project progress and growth. Across more recent projects that joined the LF, all experienced increases in contribution diversity, with more individuals and organizations contributing. This diversity drives innovation while reducing the single points of failure that plague vendor-controlled development. The emergence of well-orchestrated



and well-supported forks transforms communities to become more open to contributors of all stripes, absent the pure profit motive that previously limited participation. The proof of progress, however, can be easily found in the project numbers.

When Redis changed to a “dual source-available” license, the open source community responded with [Valkey](#), which has since demonstrated the superior innovation velocity possible under neutral governance. In the year since the fork, Valkey has developed a thriving community and solidified its reputation as a drop-in replacement. With 50 contributing companies, 1,300 commits, 150 contributors, over 30 million container pulls, and 13 releases, Valkey reached version 9.0 while maintaining exceptional performance. A Valkey 9.0 cluster can support over 1 billion requests per second, hitting a coveted hyperscale metric that would meet the needs of the majority of enterprise and large organizational users. The project’s success stems not from copying existing functionality but from unleashing collaborative development on new features that single-company control prevented.

[OpenTofu](#)’s evolution after HashiCorp’s change to the Business Source License (BSL) tells a similar story of liberation through forking. The project’s 1.8 release delivered enhanced state encryption and improved provider compatibility that exceeded the capabilities available in the BSL alternative. By version 1.10.0, OpenTofu had crossed 10 million downloads and rolled out critical enterprise capabilities, including object store usage, telemetry, and OCI

Registry support. OpenTofu enjoyed participation from dozens of contributing companies after the fork, and that participation continues to grow, alongside the diligence of the core maintainers. The LF's provision of neutral governance proved decisive, enabling contributions from organizations that had been reluctant to contribute to vendor-controlled infrastructure.

Perhaps most striking is [OpenSearch](#)'s trajectory since its fork from Elasticsearch — exceeding 900 million downloads this year. The project's 3.0 release—the first major release in three years—introduced AI-powered semantic highlighting and disk-based vector search, innovations driven by community needs rather than commercial priorities. The release integrated Apache Lucene 10 to boost search efficiency, added new vector database features supporting AI workloads, and enhanced cluster stability with improved observability dashboards. GPU-accelerated vector search functionality improved indexing speed by 9.3x, cut costs by 3.75x, and reduced CPU utilization by 2.5x compared to baseline performance. Companies like KRUU, Europe's largest photo booth rental company, implemented complete observability solutions in a single day. Freshworks, a provider of enterprise AI-assisted service software, chose OpenSearch specifically for its AI capabilities and vendor neutrality. The project demonstrates how forks can evolve beyond their origins to become superior platforms by incorporating the requirements, features, and continuous innovation from an open community.

Consequently, a deeper understanding of the Right to Fork is making its way beyond mainstream tech and open source communities into traditionally siloed industries. For instance, when Camunda — a widely used BPM framework — changed its open source license and long-term-support policies, five international financial institutions (Fidelity, Natwest, Capital One, Deutsche Bank, and Bank of Montreal) came together to launch [Fluxnova](#), an open source business process management and orchestration framework. As the only openly governed fork of

Camunda, the project immediately attracted strategic interest from enterprises within and beyond financial services, arriving just as workflow orchestration enters a renaissance, with AI agents increasingly acting across intra- and inter-firm workflows. Crucially, Fluxnova, under the auspices of FINOS, the Fintech Open Source Foundation, has also cemented open source “co-build” as a strategic third option for executive leadership, relieving the constant build-vs-buy tension in this heavily locked-in industry.

The power of successful forking has had immediate market effects. Some companies that have switched licenses on communities have reverted to OSI-approved licenses in direct response to customer feedback and competitive pressure from their respective forks. Often, however, these vendors retain single corporate control without neutral governance, meaning they could still change licenses — or restrict features contributed by the community — at any time. Redmonk, an influential open source research firm, performed an [economic analysis](#) of the post-fork results of these companies and found no clear lift, and possibly a diminution of revenues, after relicensing. This raises questions about the economic benefits of changing licenses without community consultation or approval.

The upshot? Today, the possibility of successful forking now empowers and shows communities they have the collective power to exercise a Right to Fork. If a project owner is not acting as a community steward, open governance can be an alternative path forward. Open source is a community's covenant, with both legal and unspoken expectations. Multiple companies have abandoned planned license modifications after recognizing that sustainable business models require genuine partnership with open source communities rather than extraction from them. The fork thus functions as both a remedy and a deterrent, creating powerful incentives for collaborative, rather than exploitative, approaches to open source governance.

Supply Chain and Code Security

In 2025, open source faced an array of security challenges and heightened regulatory scrutiny, alongside project innovation and global education on the strength of open source security.

Echoing the past Log4J and XZ Utils incidents, supply chain compromise of open source continues to gather momentum as a recurring attack vector, with attacks growing in both frequency and complexity. Nation-state actors sought to prey on key weaknesses in the open source supply chain to compromise core components of the open source fabric. AI is likely being used already by bad actors to facilitate these attacks, which meld novel malware design with obfuscation and social engineering to trick communities into accepting evil payloads.

On the side of regulation, governments are asking for greater security guarantees and accountability from open source projects. The European Union's Cyber Resilience Act and AI Act, along with directives from U.S. government entities and new regulations from China, the United Kingdom, and others, have increased the security complexities for open source maintainers and contributors. In the realm of open source AI models, overzealous security constraints pose a potent risk to innovation and research by smaller players.

In 2025, under the leadership of new General Manager, Steve Fernandez, the [OpenSSF](#) renewed its momentum and began to meet these challenges on multiple fronts — community events, education and guidance, software innovation, and critical services. Other foundations and projects, including [LF Europe](#), the [CNCF](#), [ELISA](#), and [Zephyr](#), all put time and effort into security. The result

was a very strong year of engagement, education and innovation.

A busy events slate kicked off in February 2025 with FOSDEM 2025, where LF Europe and OpenSSF announced a global CRA preparedness initiative to help open source projects and adopters understand obligations and prioritize practical controls. The program focused on actionable guidance, templates, and training, and it became a recurrent educational and event theme throughout the calendar year. In March, OpenSSF held a Policy Summit in Washington, D.C., under Chatham House rule, allowing attendees from government and technology organizations to speak freely about how to align regulations and laws with the realities of open source security capabilities, policies, and initiatives. In April, OpenSSF sponsored and played a key role at the second VulnCon, an event designed to foster collaboration among various vulnerability management and cybersecurity professionals to develop forward-leaning ideas to benefit the global cybersecurity ecosystem.

Over the course of the year, OpenSSF hosted and sponsored multiple community days in the U.S. Japan, India, the Netherlands, and Korea to build regional awareness of key initiatives in open source security. These events focused on helping participants and the community move the needle on open source initiatives like Sigstore, SBOMs, GUAC, and other elements of code signing and provenance that will become essential to ensuring continued trust in open source. The OpenSSF crew also participated in the UN's Open Source Week in June. Throughout 2025, OpenSSF provided mentorship and guidance to participants in DARPA's AIxCC competition, a \$29 million contest to build new open source security projects that guard against AI threats or leverage AI to improve cybersecurity processes and outcomes. This was nerd nirvana for White Hats, who got to strut their stuff at DefCon 2025 in Las Vegas, the Big Event for cutting-edge cybersecurity practitioners.

In tandem with events, the OpenSSF pushed through multiple educational initiatives. The CRA Preparedness event generated an on-demand webinar, training course, and a set of guidelines for open source contributors, maintainers, OSPOs, and others. OpenSSF partnered with [LF Education](#) to launch multiple new security courses covering JavaScript, Kubernetes, and the fundamentals of secure software development. The new Open Source Project Security (OSPS) Baseline won a strong reception, providing a much-needed roadmap for open source security practitioners wishing to implement relevant frameworks and approaches with clear mappings to the frameworks about which developers are often asked.

Of course, shipping innovation remains a critical signal. OpenSSF delivered. In April, contributors from across the AI/ML working group cut Model Signing v1.0, a spec, library, and CLI to sign and verify ML artefacts. In February 2025, OpenSSF released the OSPS Baseline in a tiered framework of security practices that evolved with project maturity. The baseline synthesized existing guidance from OpenSSF and other expert groups, outlining tasks, processes, artifacts, and configurations for enhanced software development security. GUAC 1.0, from the [Graph for Understanding Artifact Composition](#) project, reached its 1.0 milestone in June 2025, providing a stable core for aggregating software supply chain metadata to enhance security. The project created a knowledge graph connecting SBOMs, vulnerability data, and build metadata, enabling queries about software composition and risk that would have been impossible with isolated data sources. Package registries tightened provenance further. npm continued rolling out publication paths that link releases to source and to the exact build process used, with signing backed by Sigstore. Over the course of the year, [Alpha-Omega](#) made 11 grants totalling more than \$4 million for different projects to improve open source security across a wide spectrum, including Javascript, Rust, Ruby, Python and open source cryptography libraries. Alpha-Omega also provided expert guidance and content on the nascent but crucial area of open source AI security.

Educating the community on how to improve supply chain security was a key focus. The [Continuous Delivery Foundation](#) published a comprehensive resource for implementing open source security tooling into CI/CD pipelines in October 2025, helping DevOps engineers build security-compliant workflows. OpenSSF published a maintainers guide for securing CI/CD pipelines in June 2025, responding to several high-profile attacks targeting CI/CD infrastructure in open source projects. The foundation hosted a tech talk on securing the AI lifecycle in October 2025, covering tooling, model signing, provenance in Trusted Execution Environments, GPU-based integrity verification, AI/ML SBOMs using SPDX 3.0, and concrete roles and controls with demonstrations of signed models on Kaggle and NVIDIA Model Hub.

Equally important is the chopping wood and carrying water of security work. Third-party code reviews continued apace. OSTIF completed audits this year of Karmada, Notary Project (cryptography), HickoryDNS, Linkerd, Logback, RSTUF, Scorecard, PHP, Istio ztunnel, NATS, nghttp3 and ngtcp2, Apache Log4Net and Log4CXX, Ruby on Rails, Volcano, conda-forge, PowSyBI, OpenEXR, and MaterialX, covering cloud native infrastructure, languages, build and packaging systems, and the media stack relied on by film and design. Looming ahead is more work on security for the open source AI protocols MCP, A2A and AGNTCY and the daunting but critical task of securing predictive GenAI systems as they become more commonly used and deployed. Thanks to Steve's leadership and the growing participation and commitment from a wide range of contributors, we see 2026 as being a breakout year for OpenSSF — responding to the ever more urgent need to secure the open source foundations that underpin our economies, governments, and societies.

Cloud in 2025: Scale, Velocity, Innovation, AI

Modern applications are built in the cloud. This is no longer a new or emerging concept. Yet cloud, and Kubernetes in particular, are enjoying a revival thanks to AI.

A decade after the [CNCF](#)'s was founded, and cloud native went mainstream, the Kubernetes community is quickly accelerating from ML pilots to full-scale platforms that treat AI as just another production service. Kubernetes has become the default platform for training and serving AI applications. The evidence shows in Kubernetes itself: Two releases in 2025 made accelerators easier to share and govern. The steadily evolving Gateway API is making Kubernetes simpler to manage, scale and observe, which matters greatly for latency-sensitive AI applications.

The broader CNCF community continues to grow and mature, building a solid foundation for the AI future while delivering new capabilities for more traditional workloads. The transition has been remarkably smooth, thanks to dozens of key maintainers and module owners and a whole new cohort of AI companies and ML engineers bringing expertise to the CNCF. Like a well-oiled open source software machine, the CNCF can read the room, react, pivot, and evolve to meet the needs of the ecosystem, while still serving the requirements of more traditional applications. This is the hallmark of a successful ecosystem, populated by dedicated maintainers and contributors and supported by member organizations with strong commitment to building together.

As in previous years, the midyear data from CNCF showed steady velocity increases — even for older projects. Kubernetes still draws the largest contributor base and maintains one of the highest project velocities in the Linux Foundation. A testament to the health of the ecosystem, the CNCF has spawned multiple mega-projects, some serving Kubernetes directly, others leaping beyond cloud native into broader adoption. [OpenTelemetry](#), the second engine of activity, is becoming the default way to gather metrics, logs, traces and profiles across all application types. [Backstage](#) leads internal developer portals, turning platform work into visible, self-service tools that developers trust. [Crossplane](#) grew strongly, reflecting demand for a cleaner way to model infrastructure as simple, reusable APIs across clouds. [Kubeflow](#) moved up the activity ranks, a sign that ML pipelines are settling into everyday operations rather than languishing in labs.

The AI story connects these dots. Dynamic Resource Allocation, a priority area of improvement in Kubernetes releases, makes GPUs and other types of hardware for AI more manageable and efficient at scale. Gateway API improvements not only supported better handling of traditional Kubernetes traffic but also set up



a way to treat AI model traffic like any other service. KServe and related projects give teams a practical path to deploy and update models alongside the rest of their stack. These are just some of the AI developments, that are too numerous to recount in this space. Collectively, these changes let platform teams run AI using processes they already know. The benefits are faster delivery, clearer guardrails and higher utilization. The upshot? CNCF is not only cloud native but AI native, as well.

In 2025, CNCF's events reached a new and impressive scale. All told, CNCF ran five flagship KubeCons in 2025 across Europe, North America, Japan, India and China, with co-located events accessible via all-access passes. KubeCon + CloudNativeCon Europe in London set records. By the numbers, the E.U. event boasted more than 12,500 attendees, 2,808 companies, 665 speakers, and 325 sessions. More and more a gathering place beyond Kubernetes, the confab also anchored 16 co-located events, drawing 6,149 registrants. Asia events for CNCF continue to expand. China's Hong Kong edition announced 65-plus sessions with 11 maintainer tracks. Thank you for making the 2nd annual KubeCon + CloudNativeCon India another sold-out success! Registration and sponsorships for KubeCon India sold out for the second year in a row, underscoring the incredible growth and momentum of cloud native and open source technologies in this region. Also this year, Japan hosted its first KubeCon + CloudNativeCon and sold out sponsorships and registrations, signaling deep regional demand.

Grassroots programs continue firing on all cylinders. There will be 30 Kubernetes Community Days in 2025 across the Americas, Europe, APAC, and Africa, with returning stops to Helsinki, New York, Bengaluru, Budapest, Warsaw, Edinburgh, and Colombo and new debuts such as the inaugural KCD SF Bay Area and KCD Porto. These more intimate events make Kubernetes and cloud native

more accessible to a broader community. Building on the success of the CNCF Ambassadors program, the Kubestronaut skills program passed 2,000 Kubestronauts and the Golden Kubestronaut program is nearing the 200 mark.

Not all growth was organic. In 2025, the [OpenInfra Foundation](#) joined the LF, further expanding our reach in the world of open source software broadly and cloud computing more specifically. The unification brings the [OpenStack](#) and OpenInfra communities alongside Linux and Kubernetes while preserving project identities and governance. We welcome Jonathan Bryce, Executive Director for the OpenInfra Foundation and the CNCF and his talented team, as well as their valuable slate of projects to the LF.

The LF sits in a unique position in the world of cloud computing. It is safe to say that LF projects have become the de facto standards for large swaths of computing. This alone is an enormous accomplishment and one that speaks to the incredible commitment and quality of LF members and maintainers. Only three years into the AI revolution, it is also safe to say that LF projects have become the de facto standard for orchestration and delivery of AI in the cloud.

With a decade of history, Kubernetes is defying the traditional maturation curve by demonstrating continued velocity growth, echoing the footsteps of Linux. The growing array of mega-projects within the CNCF points toward an even faster-growing and brighter future as AI provides the rocket fuel where traditional cloud begins to slow down. The continued relevance and the growing importance of cloud computing projects and foundations inside the LF are classic examples of community development and common innovation that benefit all and build the future while supporting the present.

The Linux Foundation: The Trusted Home for Open Source AI

Over the past year, the LF has emerged as the center of gravity for open source AI. The PyTorch Foundation and LF AI & Data efforts have grown the number of AI projects inside the LF with major wins in critical areas, such as vLLM, Ray, Deepspeed, Newton, and Docling. The CNCF is shifting into the next gear of its evolution with AI—cloud native deployments becoming the default for AI infrastructure and tooling.

Kubernetes and its surrounding ecosystem has become the preferred scalable infrastructure for deploying and managing AI training applications across their entire lifecycle. The PyTorch transition to an umbrella foundation has opened the door for additions in the increasingly important realm of AI tooling and MLOps components. And the formation of the Agentic AI Foundation (AAIF) signaled the next leap forward in open source AI.

The most popular AI application stack is now termed the PARK Stack: PyTorch, AI models, Ray, and Kubernetes. The LF hosts three of those components and is gaining traction in AI models. MCP, contributed as a founding project



WESTERN CAPE, SOUTH AFRICA

of the AAIF, is becoming the de facto standard for agentic AI communications, often compared to HTTP. The AAIF includes founding Platinum members from all leading frontier model companies, making the LF the neutral host for critical determinations of how AI systems connect and communicate.

This builds on previous work. In the past year, the LF also became host to three protocols that together define how agents discover, communicate, and coordinate. The A2A) protocol, contributed in June 2025 with backing from over 100 organizations, establishes standards for agent-to-agent communication. Building on MCP—which connects agents to external data sources—A2A provides the language for agents to communicate with each other. The AGNTCY protocol followed in July, creating an “Internet of agents” for discovery and management, while September brought agentgateway for secure routing across different environments.

Together with MCP, these protocols create an interoperability stack preventing fragmentation that would occur if agents from different vendors couldn’t collaborate. Under LF governance, these standards evolve through open collaboration rather than competitive positioning.

In addition to the AAIF, LF AI & Data and PyTorch Foundation now include 66 projects. In the past year, critical projects have been added, including vLLM, Docling, DeepSpeed, and Newton. Membership continues to grow, with over 100 members across the combined foundations including most leading AI companies. Today, virtually every major AI company now participates in an LF project.

A key watershed moment was the PyTorch Foundation’s transition into an umbrella organization for leading AI projects. The addition of vLLM and DeepSpeed as first-hosted projects signaled that the AI ecosystem recognizes the benefits of neutral governance and tighter coupling of related projects. As evidenced in the CNCF, integrated toolchains working together on an open source canvas

enable faster innovation than any single entity can achieve alone.

This consolidation addresses a fundamental challenge in AI development. Building production AI systems requires a constellation of specialized tools—frameworks for training models, engines for efficient inference, platforms for data preparation, and systems for deployment at scale. When these tools exist in isolation, developers face integration nightmares and duplicated effort. Bringing them under shared governance within the PyTorch ecosystem creates seamless pathways from experimentation through production deployment.

vLLM exemplifies why this matters. As a high-throughput inference and serving engine, it tackles one of the most persistent bottlenecks in deploying AI systems: efficiently running large language models in production. Through innovations in memory management and batching, vLLM enables organizations to serve AI applications up to 24 times faster than traditional approaches on the same hardware. vLLM is a critical ecosystem player, partnering with DeepSeek and Hugging Face on new model launches, among others.

DeepSpeed specializes in making it feasible to train increasingly large models by optimizing memory usage and computation across multiple GPUs. Together under the PyTorch umbrella, these projects



forge an integrated stack handling the full lifecycle of AI development, from training massive models efficiently to serving them at scale in production.

The AI tidal wave also raised other boats at the CNCF. Because AI is container-native, Kubernetes has become the dominant orchestration and deployment modality for large AI applications, training, and inference. With this lift, the CNCF has regained growth momentum around AI-centric projects joining the foundation, including Kubeflow, Dapr, K8sGPT, KServe, Fluid, and Volcano.

While the PyTorch ecosystem addresses model development and serving and CNCF evolves into the premiere innovation engine for AI orchestration and infrastructure, the Open Platform for Enterprise AI (OPEA) tackles a different challenge: helping organizations build complete AI applications quickly and reliably. OPEA solves the challenge of integrating the growing array of open models with organizational infrastructure by providing standardized, composable building blocks for enterprise AI systems and yielding a retrieval-augmented generation offering to boot.

Rather than starting from scratch, organizations can assemble proven components—retrieval systems, prompt engines, language models, and evaluation frameworks—into complete solutions. By establishing common patterns and interfaces, the OPEA enables a marketplace of compatible components. Organizations can mix and match the best tools for their specific needs rather than being locked into any single vendor's stack. This modularity and interoperability represent the same principles that made Linux and Kubernetes successful, now applied to AI.

On the standards side, the Model Openness Framework (MOF) was solidified this year as a joint project of LF AI & Data and PyTorch Foundation, establishing the critical process of defining new standards for the AI Age.

This framework addresses a fundamental crisis of trust in the AI ecosystem. Many models marketed as “open” employ restrictive licenses, provide sparse documentation, or omit critical components like datasets and training code. This “openwashing” threatens to undermine the collaborative ethos that made open source software transformative. The MOF, developed through the Generative AI Commons initiative, offers both a north star to aspire toward and a practical roadmap for getting there.

While much AI attention focuses on frontier models with hundreds of billions of parameters, the LF ecosystem continues supporting development and deployment of smaller, more efficient models. vLLM's efficiency improvements make it practical to run capable models on more modest hardware. The MOF applies equally to small and large models, encouraging transparency regardless of scale.

This is crucial because the future of AI likely involves a distribution of model sizes. Massive models will cover the most complex tasks; smaller specialized models will work for specific applications where latency, cost, and local deployment matter. By providing infrastructure and standards that work across this spectrum, the LF enables innovation at every scale. We also have a strong point of view: We are in the Mainframe Era of AI, but the Unix Era of small tools for discrete purposes that all work together is emerging fast, largely driven by the open source community.

Like any major technology shift, AI is a journey, not a destination. Open source is coming on strong. Aside from AI tools and models, the LF contributes through critical projects for orchestrating AI (Kubernetes and related projects), managing AI data structures (OpenSearch, Delta Lake, Vortex), and providing the open data required to keep AI accessible to all. We are determined to make sure the LF is a key player in this once-in-a-generation technology revolution.

Building the Data Foundations for an AI-Powered World

Data is the lifeblood of AI. Every model, every prediction, every breakthrough depends on the quality and accessibility of the information feeding it. Yet, the challenge facing organizations today is not just collecting more data. It is making that data work together.

The Linux Foundation recognizes this reality and has continued to expand efforts in data, both supporting efforts to make large, broadly useful pools of data openly available, and to build out tooling and systems that can facilitate the collection, transformation, distribution, and use of data. These efforts span multiple foundations and multiple industry verticals, as well as important tools for the data landscape.

After joining LF Charities in October 2024, the [Jupyter Foundation](#) became one of the LF's flagship projects. Among the world's largest ecosystems for data science and AI practitioners and researchers, Jupyter had nearly 260 million package downloads from May through September 2025. The project spent 2025 shoring up its own foundations with new leadership, improved security practices, and many new features shipped. The LF hosts and supports numerous other projects that play key roles in managing or storing the world's data including Valkey, OpenSearch, Marquez, Delta Lake, JanusGraph, Marquez, Milvus, Ceph, and Vortex, all of which continue to progress.

A second area of emphasis is open data and open standards for data exchange. The [Overture Maps Foundation](#) reached a watershed

moment in 2025. The organization launched its Global Entity Reference System in June, fundamentally changing how geographic data connects across platforms. GERS assigns unique and persistent identifiers to every building, road, and place on Earth, eliminating the costly conflation work that has plagued the geospatial industry for decades. Organizations spend up to 80 percent of their resources on data preparation rather than analysis, with geographic data presenting particularly acute challenges. GERS solves this challenge by providing what amounts to a universal adapter for location data. What once required weeks of complex geospatial conflation now happens with simple column joins in minutes.



A wide community effort among dozens of organizations makes this possible. Overture's reference map delivers monthly validated datasets connecting GERS IDs to real world entities. Data changelogs track every change between releases, allowing efficient updates rather than complete reprocessing. The GERS Registry maintains the complete lifecycle of every ID ever created. Bridge files provide prebuilt mappings to existing systems like OpenStreetMap and Esri Community Maps, making adoption immediate rather than eventual. Real world adoption validates the approach. Meta integrates Overture's data directly into Facebook and Instagram maps, reaching billions of users. TomTom incorporates it into Orbis Maps for developers building everything from navigation systems to logistics tools. Esri embeds it in ArcGIS, putting contributed business listings onto maps used daily by government agencies and businesses worldwide. The [Open Travel Alliance](#) is using Overture and GERS to provide embedded place data for its members, allowing, for example, hotels to share precise location information with ride sharing apps using common GERS identifiers. The integration also will allow hotels, airlines, transport companies, restaurant reservation platforms, and other travel sector players to coordinate to build better and more up to date agentic AI travel capabilities.

In short, GERS and Overture have fundamentally changed the dynamics of geospatial data and democratized what was once a scarce and expensive resource into a common data holding any organization can leverage. The technical foundation continues strengthening. The foundation's May member summit brought together representatives from 30 organizations, marking the shift from building datasets to enabling an ecosystem. Monthly data releases throughout 2025 delivered steady improvements, with the transportation theme reaching general availability and places data expanding to more than 60 million points of interest.

Amy Rose became Overture's first CTO in March and was appointed to the federal government's National Geospatial Advisory Committee

in late 2024. Her dual role positions open data at the heart of national spatial data infrastructure planning. The foundation also welcomed a new Executive Director. Will Mortensen, who most recently served as the director of the Office of Geography at the National Geospatial-Intelligence Agency (NGA), joined the project in October, replacing Overture's inaugural executive director, Marc Prioleau, who announced that he would be retiring at the end of the year.

As demand for energy surges, driven by AI data centers and continued societal growth, the energy sector is moving to a shared commons of data and data standards to improve efficiency, availability, and sustainability. [LF Energy](#)'s portfolio of projects addresses multiple layers of this challenge, from AI-powered grid optimization to standardized data exchange for decarbonization and vehicle charging.

[GridFM](#) aims to create foundation data and AI models specifically for power grids, applying the breakthrough approach that transformed natural language processing to electrical infrastructure. The first version, GridFM v0, trained on more than 300,000 models solved optimal power flow problems. The model promises computation speedups of three to four orders of magnitude, enabling real-time analysis that current systems cannot achieve. The [Battery Data Alliance](#) released the largest open source battery dataset ever in August 2025. The dataset contains cycling data from 199 coin cell batteries tested for 1,000 cycles under fully automated workflows, all formatted in the Battery Data Format. SC Decarbonisation Hub, contributed by Shell, creates visibility into Scope 3 emissions across supply chains. LF Energy members made significant progress on the TROLIE standard, with additional language. NODE Collective builds comprehensive databases of residential electrification incentives across the United States. Previously an industry plagued by information silos and poor compatibility, the energy sector is shaping up to be a primary exemplar of how open source can provide a sound data foundation for broad sector improvements.

In finance, the [FINOS](#) and [FinOps](#) Foundations both continue to build out the common data foundation of an industry that is already largely digital but has traditionally struggled to build shared technology. FINOS and FinOps continue to make strong progress. In 2025, FINOS advanced practical interoperability and data standards with FDC3 2.2 for trading desktop app messaging and a major CDM update v6.0 to make trade, product, and lifecycle data machine readable and executable across firms. It also broadened compliance guardrails, shipping new Common Cloud Controls catalogs including an RDBMS set and unveiling Common Controls for AI Services to standardize governance of AI in regulated environments. For its part, the FinOps team shipped

a new version of the FOCUS spec for interoperable cloud cost and resource consumption comparisons while expanding to SaaS and PaaS products. With all of the major hyperscalers now integrating FOCUS capabilities, it is rapidly becoming the de facto method for measuring and controlling cloud costs and resource consumption.

Lastly, the LF is not just stewards of standards and datasets but also makers of data products. The [LFX Insights](#) platform is now generating automated reports on key health, velocity, and participation metrics for all LF projects. This is already driving engagement and improving understanding of the value of participating in open source and the dynamics of specific projects and foundations.

TUSCANY, ITALY

Hardware & Infrastructure: Nascent Ecosystems, AI Standards, and Vibrant Standard Bearers

The [Joint Development Foundation \(JDF\)](#) marked its 10-year anniversary as a key catalyst for open standards in key technology areas.

Supported by more than 500 member organizations and 3,500 participants, the JDF marked this milestone with the launch of the [OpenSTX Foundation](#), the JDF's 36th standards collaboration, and the publication of the 3MF File Format Specification Suite, an ISO/IEC JTC 1 standard that is the sixtieth such success for the organization. The Linux Foundation is proud to have played a role in the growth and success of this organization.

2025 was the year [RISC-V](#) matured from upstart to truly production-ready. The global RISC-V Summits drew thousands of attendees, venture investment accelerated, and some of the world's top chip designers launched new efforts around an open ISA. For companies building atop a fully open, neutrally governed standard, RISC-V emerged as the path forward. NVIDIA, Red Hat, Alibaba, Huawei, Canonical, and dozens more shipped RISC-V products and integrations across CPUs, SoCs, maker boards, and specialty silicon. Alibaba introduced the C930, its first server-grade RISC-V CPU. Enterprise readiness advanced as SiFive and Red Hat launched a developer preview of Red Hat Enterprise Linux on SiFive's HiFive Premier platform, and Android 15 was successfully demonstrated on a high-performance RISC-V device.

The work on the project is moving quickly. RISC-V repos have been forked more than 3,000 times, and the project sees more than 1,600 pull requests each month. In 2025, the RISC-V Foundation

approved multiple key specs, including RVA23 profiles and work on Vector-C intrinsics that simplify C programming on RISC-V platforms. The project was also approved as a recognized publicly available specification (PAS) submitter by the ISO/IEC Joint Technical Committee (JTC 1). Linux 6.18 included improvements that allow RISC-V chips to run better on RISC-V silicon. With a neutrally governed, global standard gaining momentum amid geopolitical uncertainty, all indicators point to RISC-V entering the mainstream within the next year or two and providing a growing challenge to Arm and x86 in the competition for global mindshare and presence on silicon.

As AI models grew, the network became the bottleneck, which pushed an industry-wide rethink of the data fabric around open standards. The [Ultra Ethernet Consortium \(UEC\)](#) published its 1.0 specification in June 2025. UEC delivered an Ethernet-based fabric engineered for large-scale AI traffic patterns and provided an open alternative to proprietary interconnects. Inside the



server, two complementary open efforts advanced quickly. The [Open Programmable Infrastructure](#) project gained traction standardizing DPUs and IPUs as common offload engines, highlighted by the first OPI Summit in San Jose in October and steady progress toward a 1.0 specification.

Open networking matured in parallel. The [SONiC Foundation](#) community continued to expand deployments and contributors. Operators demonstrated SONiC as the open NOS underpinning performant, vendor diverse leaf-spine fabrics for AI clusters with shared control planes, open telemetry, and rapid pipeline iteration. The growth of SONiC complements UEC by pairing an open network operating system with an open, AI-aware transport fabric.

Upstream in open silicon, [CHIPS Alliance](#) members advanced open IP and tooling. Communities improved open simulators and verification flows, contributed updates to embedded RISC-V security blocks such as Caliptra, and continued work that ties open architectures to manufacturable designs. The result is a clearer bridge from specification to silicon that benefits RISC-V cores, heterogeneous accelerators, and domain-specific SoCs.

Telecom followed suit. Backed by U.S. FutureG initiatives, [OCUDU](#), short for Open Source Centralized Unit and Distributed Unit, advanced efforts to open the 5G and 6G RAN CU and DU stack, with a clear path to neutral governance at the LF. Award announcements recognized leadership from DeepSig and SRS on the initial implementation track. OCUDU brings the last major proprietary stronghold in mobile networking into the same open, modular orbit as cores and cloud. This shift matters for secure, resilient, and rapidly evolvable national infrastructure across military, private, and commercial applications.

On the HPC side, the [High Performance Software Foundation](#) (HPSF) held its inaugural conference in Chicago in May, which also hosted the Kokkos User Group. The community showcased

steady progress toward Kokkos 5.0, which targets performance portability across heterogeneous, massively parallel architectures. The HPSF also welcomed OpenCHAMI, broadening the ecosystem of production grade open HPC software that maps cleanly onto emerging accelerators and fabrics. This represents the essential link between open hardware and the scientists and engineers who rely on stable, portable programming models.

The energy that powers it all is infrastructure, too. [LF Energy](#) projects saw growing adoption as utilities and operators modernized grids for AI era loads, flexibility markets, and data center interconnection. These efforts connect open standards in power systems with the compute, cooling, and networking standards that define the modern data center, which helps ensure reliability while accommodating new classes of high density workloads.

The [Open Mainframe Project](#) capped a banner year. Zowe advanced through four LTS releases in 2025, with v3.1 in February, v3.2 in May, v3.3 in September, and v3.4 slated for GA in November. IBM made Zowe 3.0.x a fully supported, generally available distribution for enterprise customers on May 2. These are clear signals of production readiness across the mainframe ecosystem. Community programs expanded impressively. The Open Mainframe blog remained one of the liveliest at the LF, filled with stories of how mentees learned, how professionals built careers on mainframes, and what is new across the platform. The Summer 2025 Mentorship drew roughly 1,100 applicants for 13 projects, and a new cohort of Ambassadors launched in April, which underscores a healthy talent pipeline. In AI, Open Mainframe supported multiple efforts, including Zorse, a dataset tool for training large language model code assistants for COBOL and other mainframe-dominant languages. The project celebrated its tenth anniversary and demonstrated how an established technology base can stay vital, crucial, and innovative.

Climbing the Technology Stack in Industry

Open source is climbing the value chain. Over the past decade, industry foundations, SIGs, and working groups inside the Linux Foundation have helped move open source from pure code to industry-focused and enterprise-grade building blocks that drop into real-world systems in finance, telecom, mobility, aerospace, energy, industrials, healthcare, and media.

Multiple publications reported on the growing acceptance of open source collaboration in the highly competitive financial sector. The LF has both accelerated and benefited from this shift. [FINOS](#) crossed the 100-member mark in January and kept adding blue-chip firms through the year. CDM, now a FINOS graduated project, shipped 6.0 with more than 100 model improvements, further standardizing programmatic representations of financial products. The Financial Desktop Connectivity and Collaboration Consortium (FDC3) shipped version 2.2 for simplified desktop interoperability across research and trading stacks. A new Common Controls for AI Services project gave CISOs and compliance leaders an auditable baseline for safe model use, while the [Fluxnova](#) project marked the first time major financial institutions exercised their Right to Fork into an openly governed workflow orchestration framework. [OSFF](#) London and NYC, combined with global hackathons and meetups, drew over 5,000 attendees from scores of institutions.

FinOps matured from cloud cost management to cloud-plus value management. The [FinOps Foundation](#)'s FinOps Framework update widened scopes beyond public cloud to SaaS, on-prem, and licensing. FOCUS 1.2 made invoice-level reconciliation and multi-currency normalization routine. All of the major hyperscale clouds added or expanded FOCUS support so buyers could trust a common data schema for cloud spend and resource usage. FinOps X drew over 2,000 attendees. Regional X Day events took best practices to Tokyo, São Paulo, Amsterdam, and Washington, D.C. A new FinOps for AI learning path taught teams how to govern token and usage-based model spend. The practice is becoming a cross-functional operating model rather than a niche cost-cutting project.

Telecom's open stack shifted from virtualized networks to AI-assisted automation. [LFN](#) launched Essedum 1.0 as a core platform to carry data, models, and applications into live network operations. [ONAP](#) delivered two major releases. Oslo in January and Paris in July pushed modular architectures, stronger security, and intent-driven workflows that make GenAI a tool rather than a demo. Anuket's Quinpiac release enhanced support to project [Sylva](#) to keep telco cloud substrates standard across vendors. On the service layer, [CAMARA](#) grew fast. The Spring meta-release offered 38 mature network APIs. The Fall meta-release pushed the catalog to 60 with additions in SIM-swap detection, quality of service, and edge.

A notable inflection arrived with [OCUDU](#). The U.S. DoD FutureG Office selected DeepSig and SRS to lead an open CU and DU software effort for 5G and 6G RAN. The work aims to deliver a Linux of RAN-style reference for carrier-grade stacks and become the "Linux" of mobile communications. Beyond the core, the [Aether Project](#) documented live deployments and research integrations, showing an open 5G stack that is ready for production blueprints, not just proofs of concept.

Software-defined vehicles have moved into the fast lane, with [Automotive Grade Linux](#) (AGL) occupying pole position for infotainment systems. AGL convened the ecosystem twice for conferences in Tokyo in February and Berlin in July. OEMs such as Honda and Toyota and suppliers like Panasonic shared concrete progress on SDV architecture, compliance, and supply chain alignment. [ELISA](#) added new insights on building safe systems for cars. [Zephyr](#) RTOS featured in automotive SoC discussions and continued the quiet work required for functional safety certification. For electric vehicles, LF Energy's EVERest project delivered numerous advances in 2025, including advancing charging standards, the implementation of ISO 15118-20, and continued development of OCPP 2.1. On the roadmap are new technology integrations like the Megawatt Charging System (MCS), EEBus for bidirectional charging, and the Matter smart home standard.

Up in the air, ELISA's [Space Grade Linux](#) initiative stepped into the spotlight at Open Source Summit Europe. The team outlined a safety-aware Linux foundation for space systems along with a Yocto-based reference approach. The Aerospace Working Group paired with the Safety Architecture group to define kernel requirements through a cabin-lighting use case. The discipline of safety engineering met the cadence of open source.

The [Dronecode Foundation](#) delivered a banner year, shipping PX4 v1.16 shipped after a year-long cycle with many new features. QGroundControl 5.0 brought a major UI and UX refresh. MAVSDK advanced through the 3.10 series. Membership expanded with companies across sensing, platforms, and UTM. Real-world deployments are growing and include wildland fire and avalanche mitigation, high-altitude anti-icing, logistics, and disaster response. The ecosystem counted tens of thousands of contributors and support for more than a million vehicles across dozens of board variants.

[LF Energy](#) projects crossed the chasm from trial to grid impact. RTE documented SEAPATH in production as the virtualization

substrate for digital substations with additional sites planned. In the U.K., National Grid and GE Vernova highlighted collaboration that fed into the first official SEAPATH release. PowSyBI completed a third-party audit, added fuzzing in OSS-Fuzz, and resolved issues in the open. In health, the HealthyPi Move, an open source, watch-form-factor biometric monitor running on Zephyr, was added to the project's portfolio in July. The device monitors ECG, heart rate, SpO₂, and other vital signs.

As always, the [Academy Software Foundation](#) (ASWF) delivered blockbuster results. The Academy honored technologies foundational to the ASWF ecosystem with scientific and technical awards in recent seasons. [OpenVDB](#) earned a Scientific and Engineering Award for its lasting impact on volumetric imagery. [OpenUSD](#) was recognized for advances in scene description. [OpenEXR](#), [OpenColorIO](#), [OpenTimelineIO](#), [OSL](#), and [MaterialX](#) powered major 2025 winners, including Dune: Part Two and The Wild Robot. In codecs, the [Alliance for Open Media](#) added hardware credibility with a Technology and Engineering Emmy for AV1 and an IBC Innovation Award for SVT-AV1. To cap it off, the Academy of Motion Pictures itself contributed the [Academy Color Encoding System](#)—a global standard for color management, to the ASWF. In an era of cost squeezes for entertainment companies, open technologies are proving more and more technically competitive and commercially viable. ASWF mentorships and summer programs remained among the hottest education tickets in the LF.

The LF's user members cannot be thanked enough for making the progress in industry verticals possible. Projects in these arenas require deep sector expertise that may not be quite as broadly applicable as contributions to an operating system or a container orchestration solution. That said, the industry communities have come to recognize and embrace that, even when the solution is industry specific, sharing the burden of technology development can benefit all.


Building Trust Infrastructure for the Digital Age

Trust is under attack across the digital landscape, from deepfakes to identity theft and email compromise. Projects and communities supported by the Linux Foundation are building trust systems in three places that matter – content, code and currency.

The common thread is verifiability and traceability through the life of an artifact, whether that artifact is a photo, a binary, or a payment event. We focused on open standards, neutral governance, and production adoption. The results are tangible and repeatable.

Content authenticity advanced dramatically in 2025. Google enabled [C2PA's Content Credentials](#), a collaborative standard hosted by Joint Development Foundation Projects (JDFP), across the Pixel 10 camera workflow and documented Assurance Level 2 for the Pixel Camera app. Photos captured on those devices carry cryptographic provenance out of the box. Alterations become easier to detect. Claims are easier to test. Distribution systems can choose to honor the signal and label content accordingly. That puts a real cost on manipulation at the moment where it used to be cheap.

The ecosystem followed. TikTok began attaching C2PA credentials to clips created by its AI Alive feature. This comes on the heels in 2024 of YouTube integrating C2PA and setting up expectations of integration with content creators who embedded C2PA credentials in their work prior to upload. Sony introduced a camcorder that writes C2PA authenticity data at record time and worked with BBC teams to run end to end newsroom tests. Nikon, Canon and other



leading camera makers also added C2PA to devices. A chain of custody is taking shape that starts at the lens and spans broadcast, publication, and distribution. Interestingly, security video has started to align on C2PA, too. ONVIF, a leading provider of security camera technology, opened work to bring C2PA-compatible credentials into its interface profiles. The upshot? C2PA is well down the path toward broad adoption and standardization. That said, challenges remain, as C2PA is not yet widely used by consumers or creators and is a default setting in the largest platforms. This is the heavy lift in front of us.

[LF Decentralized Trust](#) continued pushing forward with new trustless and zero-trust systems that are paving the way for a more efficient and more trustworthy future. The foundation's projects logged a busy 2025. Hedera's open codebase now operates as the [Hiero project](#) with vendor-neutral governance and an actively running network, while [Hyperledger Iroha](#) reached 2.0 with role-based permissions, multisig, a richer query model, smart contracts, and WASM support that backed live central-bank pilots in Papua New Guinea and the Solomon Islands. [Hyperledger Fabric](#) advanced the 3.x line with batching and SmartBFT improvements that lift throughput for credential and registry backends. [Hyperledger Besu](#) posted a marquee production case with the World Bank's FundsChain for end-to-end tracking of development funds, strengthening the public-sector audit story. [Hyperledger Cacti](#) matured as the interoperability layer with new connectors that let verifiers, wallets, and registries span heterogeneous ledgers. [Web3j](#) 5 gave enterprises a maintained Java pathway for Ethereum integration that simplifies build pipelines and production operations, making these stacks practical substrates for trust registries, privacy-preserving revocation, and verifier integrations now moving from pilot to platform.

Trust efforts did not stop at code and content. Payments are a crucial trust layer for every modern economy; however, the best anti-fraud systems have traditionally been pricey and out of reach

for financial institutions in the developing world. [Tazama](#) fills this gap. Part of LF Charities, in 2025, Tazama moved from launch to concrete partnerships and deployments that raise the fraud floor for instant payments. It is an open platform for real-time monitoring and fraud detection that ingests transaction streams, applies rules and models, and emits actionable alerts before funds clear. It was built for national switches and mobile money, with an emphasis on privacy and data sovereignty. In 2025, the program announced a strategic partnership with FNA to bring network analytics to anti-scam operations and reported "go lives" that support providers in Mongolia and Nigeria.


Trust in digital identity systems also advanced through the [Decentralized Identity Foundation \(DIF\)](#), which celebrated its first anniversary at the LF. The foundation is playing a decisive role in global efforts to build foundational, open source components for decentralized identity systems. The DIF's Trust Registry Query Protocol reached version 2.0, and DIF Labs hosted two impressive problem-solving cohorts tasked with solving real-world challenges including: ZKMPA for zero-knowledge multi-signature credentials, privacy-preserving revocation with a taxonomy, benchmarks, and reference implementations; and QVC binding W3C verifiable credentials to eIDAS-grade qualified electronic signatures.

These initiatives interconnected across technical and organizational boundaries. C2PA content credentials leveraged decentralized identity infrastructure for creator verification. Software supply chain security depended on similar cryptographic primitives as financial fraud detection. Building trust at scale required tamper-evident logs, cryptographic signatures, transparent verification, and neutral governance. The LF and its supported foundations, such as OpenSSF, JDFP, LF Decentralized Trust, and DIF, provided the venues where these patterns developed as open standards rather than proprietary solutions. The work accelerated throughout 2025 as each domain contributed to a broader architecture for trustworthy digital systems operating across organizational and national boundaries.


Financial Transparency

The Linux Foundation’s revenue is derived from four main sources: Memberships and Donations, Project Support, Training and Certifications, and Event Registration and Sponsorship.

In 2025, we are forecasting revenues of over \$310M. In 2025, the LF is forecasting to spend nearly \$285M supporting our mission.



Membership & Donations	\$133,345,419
Project Services	\$83,571,956
Training & Certification	\$29,615,980
Event Sponsorships & Registrations	\$58,633,555
Other	\$6,176,111



Corporate Operations	\$15,834,749
Linux Kernel Project	\$8,410,114
Event Services	\$16,813,013
Community Tooling	\$15,726,845
Project Infrastructure	\$17,733,121
Training	\$21,637,925
Project Support	\$181,889,435
International Operations	\$6,750,480

TRAVEL
FUNDING

In our continued efforts to broaden accessibility to in-person events, we provided over:

\$1.5M

in travel funding and registration scholarships



\$900,000

Travel Funding



Almost

1100

Registration Scholarships

60%

of funding went to diverse community members



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 including Linux, Kubernetes, Model Context Protocol (MCP), OpenChain, OpenSearch, OpenSSF, OpenStack, PyTorch, Ray, RISC-V, SPDX, and Zephyr provide the foundation for global infrastructure. The Linux Foundation is focused 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. 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

2810 N Church St
PMB 57274
Wilmington, Delaware, 19802-4447 US

info@linuxfoundation.org
www.linuxfoundation.org