LF NETWORKING

2025 Year In Review

Powering the Open AI-Native Networking Revolution

Table of Contents

Introduction

The Linux Foundation Networking Mission

"Drive an open source ecosystem that revolutionizes the movement or communication of data on a network—including its data plane, control plane, analytic, orchestration, and automation technologies—for enterprise, cloud, and carrier network constituents."

LF Networking (LFN) continues to serve as the world's largest and most collaborative open source networking ecosystem—uniting service providers, cloud operators, enterprises, vendors, and system integrators under one shared mission: to accelerate the adoption of open, interoperable, and AI-native networking technologies.

Now entering its eighth year of community growth and market maturity, LFN's portfolio has become a cornerstone for digital transformation across domains,from the data center to the network edge. In 2025, LFN drove the next phase of industry evolution by enabling intelligent, automated, and sustainable networks built on open innovation.

Highlights of 2025 progress include:

- AI-Native Networking Momentum: Advancing domain-specific AI capabilities across projects to optimize network performance, operations, and orchestration. This includes the induction of the Essedum project, an open source AI-native networking initiative, designed to optimize network performance, operations, and orchestration.

Introduction

- Cloud Native Telco Initiative (CNTi): The joint effort with CNCF has built out its beta certification program and continues to expand ecosystem adoption through industry collaboration and operator validation.

- Cloud Native Automation Evolution: Continued alignment on cloud-native automation frameworks supporting multi-vendor, multi-cloud orchestration.

- Cross-Industry Deployments: Real-world implementations showcased at ONE Summit 2025 demonstrated measurable business impact—from hyperscalers to public sector networks—proving open source networking's role in powering the world's digital infrastructure.

As we look ahead, LFN remains committed to uniting the networking and cloud communities in pursuit of scalable, AI-driven, and interoperable open source solutions. We invite you to explore our projects and join us in shaping the future of connectivity through open collaboration.

The AI Native Networking Era

LFN continues to stand at the forefront of open innovation—serving as the central hub where industry collaboration transforms into deployable, interoperable, and AI-native networking solutions. With a diverse ecosystem spanning service providers, cloud operators, enterprises, vendors, and system integrators, LFN's projects power the foundations of global connectivity and digital transformation.

Through our flagship initiatives such as ONAP (network service automation), Anuket (cloud-native infrastructure architecture), L3AF (eBPF lifecycle management), and Nephio (Kubernetes-based automation)-as well as newer, evolving projects like Duranta and Essedum that enable the full stack from RAN to AI, LFN's collaborative model continues to bridge open-source technology with real-world deployment. In 2025, that model reached new heights, driven by the growth of domain-specific AI, multi-cloud networking, and sustainable automation frameworks designed for the next generation of communications.

From Cloud Native to AI-Native: 2025 Marks a Turning Point

LFN's 2025 journey reflects the industry's pivot from cloud-native architectures toward AI-native networks. This evolution took a major step forward with the introduction of Essedum 1.0, a groundbreaking open-source initiative that provides a foundation for domain-specific AI in networking. Essedum enables communities to collaboratively develop and validate AI applications and agents that improve network performance, energy efficiency, and service assurance, laying the groundwork for intelligent, self-optimizing infrastructure.

LFN also strengthened cross-foundation synergies to accelerate this AI-native vision. Projects such as Anuket's Thoth advanced work on curated data sets for telecom AI, while CNTi (Cloud Native Telco Initiative) matured its certification and

best-practice framework to support production-ready deployments across operators and vendors. Together, these efforts reflect a decisive shift from experimentation to implementation.

Open Collaboration Expands: Building the AI-RAN and Beyond

2025 saw a surge in high-impact partnerships that extended LFN's influence across adjacent domains:

- Duranta, based on seed code contributed by OpenAirInterface Software Alliance (OSA), became a flagship project for advancing open RAN innovation. It enables researchers and developers to integrate cloud-native automation into RAN and core networks, fostering faster innovation cycles for 5G/6G evolution.
- A formal alliance with ATIS deepened collaboration on Open RAN interoperability and open source test frameworks, aligning global standards with open development.
- A strategic collaboration with the AI-RAN Alliance reinforced the role of open source software in applying artificial intelligence to optimizing RAN performance, spectrum management, and energy efficiency.
- More collaboration with premier member, Infosys, strengthened responsible AI governance through the donation of the seed code to the Salus project, ensuring transparency and fairness in the deployment of AI across communications networks.

Together, these alliances underscore LFN's role as the collaboration fabric that interconnects the open networking ecosystem—from AI and automation to next-generation wireless.

The AI Native Networking Era

Ecosystem Growth and Real-World Impact

LFN's collaborative approach has continued to produce tangible results. The 2025 Global Open Source Networking Survey, conducted in partnership with LF Research, revealed that over 92 percent of organizations now rely on open source technologies for cloud-native networking, Open RAN, and AI integration. This growing adoption validates the maturity of open projects as production-grade enablers of digital transformation.

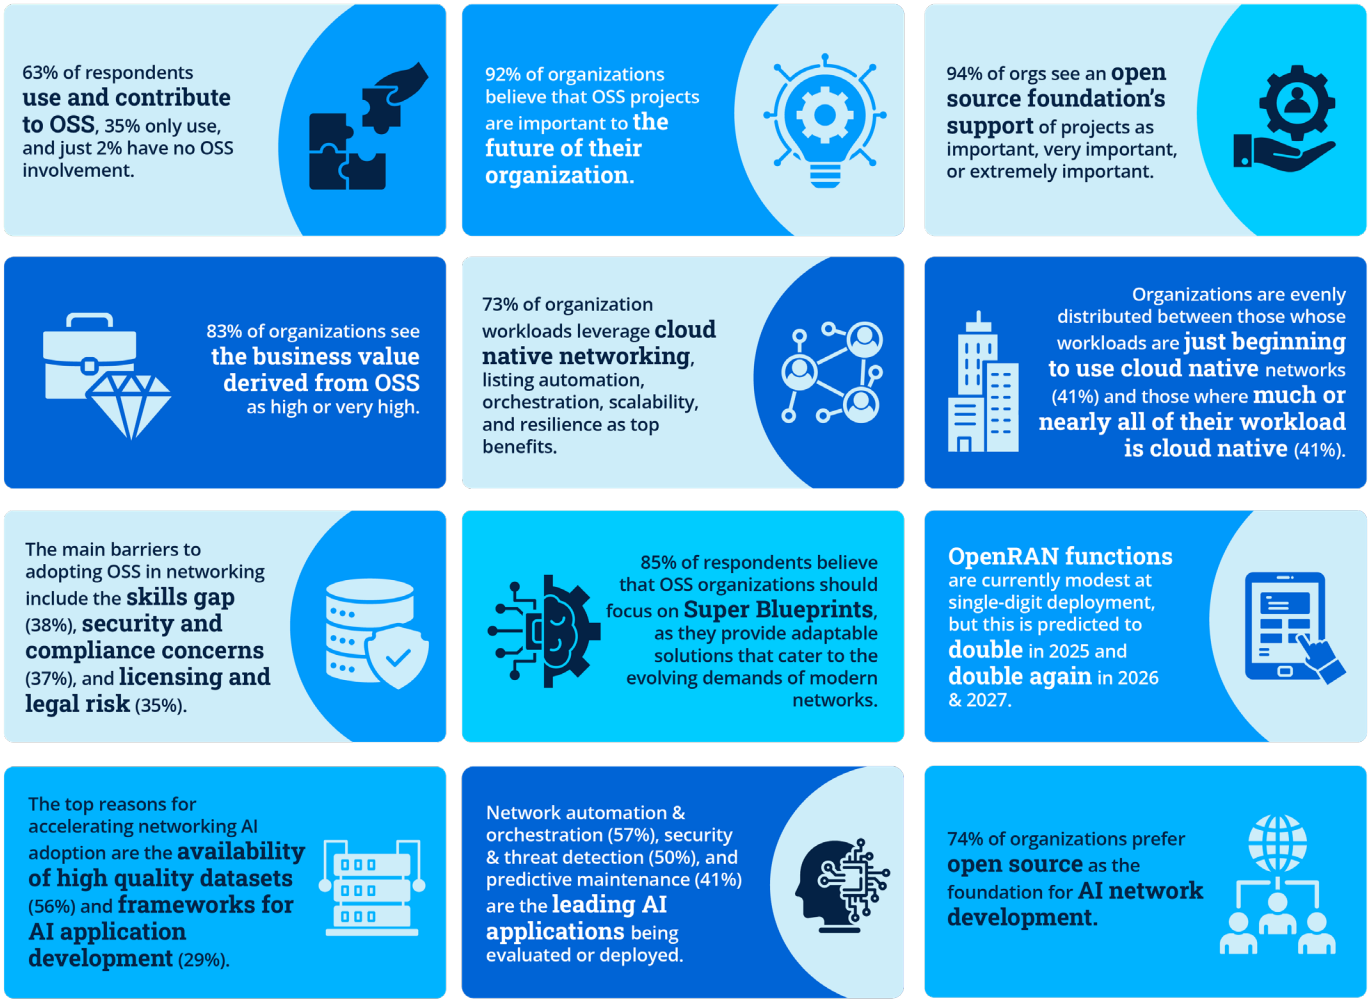

The AI Native Networking Era

At our flagship Open Networking & Edge Summit 2025 event, held in London alongside KubeCon + CloudNativeCon Europe, the community showcased these results in action. The event's sessions explored AI-driven automation, network observability, and edge computing deployments that have moved from prototypes to production. Case studies highlighted how telecom operators, enterprises, and hyperscalers are leveraging open source to drive business outcomes—improving operational efficiency, accelerating service rollout, and enhancing resilience across distributed infrastructures.

Security, Interoperability, and "Next-G" Readiness

As adoption accelerates, LFN continues to address challenges central to enterprise and carrier-grade deployment. The TAC Security & Quality Workstream, established in 2025, set measurable goals for code quality and secure software development across all projects, aligning with OpenSSF best practices. This initiative ensures that security remains integral to every layer of the LFN ecosystem.

At the same time, LFN's Next-G initiatives - including reference architectures for 5G and 6G through the 5G Super Blueprint - provide researchers and developers with open, reproducible frameworks for testing and innovation. These blueprints integrate modular components from multiple projects such as ONAP, Nephio, and OAI, offering a reference architecture for both current and future wireless generations.

Navigating the AI-Native Future

Looking ahead, LFN is uniquely positioned to lead the convergence of AI, cloud, and connectivity to enable the autonomous networks of the future . Predictive and generative AI use cases will expand - from automated incident postmortems to adaptive service orchestration - powered by shared data models and community-driven frameworks. LF Networking's open source projects are poised to deliver the essential building blocks for developing sophisticated AI agents that will autonomously

manage the complex networks of tomorrow. These AI agents are task-oriented, AI-powered applications granted the autonomy and intelligence necessary to independently operate, optimize, and maintain the networking infrastructure.

As agentic-AI workloads continue to distribute across the device-edge-cloud continuum, LFN's projects will provide the connectivity, orchestration, and observability layers that enable true multi-cloud portability. Open APIs, like those developed under sister project **CAMARA** and LFN's Nephio will remain key to monetizing network capabilities and creating new service-provider revenue streams.

The year ahead will demand continued collaboration across open-source communities, standards bodies, and industry consortia. LFN's governance model - rooted in neutrality, transparency, and interoperability - positions it as the anchor for this evolution, ensuring that open innovation continues to thrive at the intersection of AI, networking, and the cloud.

Message from the LFN Board

SEBASTIAN ZECHLIN
Chair, LFN Governing Board, and Lead Architect Technology Architecture & Innovation, Deutsche Telekom

Agentic AI Has Come to Stay

The major shift shaping our industry in 2025 is the rise of Agentic AI—AI that doesn't just respond, but plans, reasons, and acts. What seemed like an experiment a year ago is now running in production networks around the world.

Self-healing has moved from concept to reality.

Modern networks now auto-remediate faults in minutes using live telemetry and policy engines. Intent-based systems are deploying and scaling network functions without a single CLI command. Security operations are becoming proactive, with autonomous agents detecting and neutralizing threats faster than any human SOC could. Even customer experience is transforming—agents are resolving most inquiries before a ticket is opened.

The numbers tell the story: the agentic AI tools market is expected to reach $10.4 billion this year with a 56 % CAGR. According to analysts, a quarter of telecom operators are already running pilots, and half will follow within the next few years. Yes, there's plenty of hype—and even talk of an AI bubble—but this time, much of the excitement is backed by tangible results.

In the foreseeable future, we'll see networks evolve into multi-agent systems, where specialized agents coordinate everything from optimization to billing intelligence. Human operators will move from firefighting to strategic oversight. Transparency, openness, and trust will become essential, as regulators and boards demand explainable decisions.

The challenge—and opportunity—for the open source community is to make these systems reliable, interoperable, and safe. LFN and its projects are uniquely positioned to lead this effort. Our history of collaboration, open innovation, and creating de-facto standards gives us the foundation to shape how agentic AI becomes part of the network fabric.

As 2026 unfolds, here are some aspirations for the year ahead:

- That agentic AI frees us from repetitive work, giving us more time for creative problem-solving.
- That LFN projects harness AI responsibly to drive automation, efficiency, and innovation.
- That we strengthen collaboration across the community to share knowledge, tools, and data safely.
- And that LFN continues to serve as the open, trusted hub for innovation in intelligent networking.

A heartfelt thank-you to every contributor, member, and staff who made 2025 such a transformative year.

Here's to an exciting 2026—where networks don't just run, they think.

LFN Board Members

Christian Olrog
Ericsson

Chuanyu Chen
Huawei

Douglas Knisely
Qualcomm

Frank Brockners
Cisco

Frederic Desjarlais
Walmart Global Technology

Hanen Garcia
Red Hat

Jonne Soininen
Nokia

Junlan Feng
China Mobile

Olaf Renner
Nokia

Praveen Atreya
Verizon Global Network And Technology

Ryuji Wakikawa
Keio University

Sebastian Scheele
Kubermatic

Sebastian Zechlin
Deutsche Telekom

Shankar Malik
NgKore

Sreekanth Sasidharan
Infosys

Sun Qiong
China Telecom

LFN Members

Platinum Members


Silver

Associate




Message from the General Manager

BY ARPIT JOSHIPURA
SVP and General Manager, Networking, Edge, & IoT, the Linux Foundation

As we close another year of open innovation and collaboration, I want to reflect on the meaningful progress our community has made in 2025—and look ahead to what's next.

This year, our Collaboration Hub focused aggressively on enabling the ecosystem to lead digital transformation, especially around domain-specific AI and cloud native automation. The community's momentum has been impressive. For example, the induction of Essedum and its subsequent release of Essedum 1.0 provided a core platform for AI-powered networking applications. And the survey we released in March confirmed that more than 92% of respondents are relying on open source projects for cloud-native adoption, Open RAN and domain-specific AI priorities.

On the collaboration front, we announced key partnerships to advance open networking innovation:

- A formal collaboration with AI-RAN Alliance to jointly advance "AI-RAN" capabilities in open networks.
- A new collaboration with OpenAirInterface Software Alliance (OSA) to launch the new Duranta project, designed to accelerate open-source RAN innovation.
- A partnership with ATIS to further develop and integrate open-source and Open RAN technologies.
- A collaboration with premier member organization, Infosys, on strengthening responsible AI for global networks.

We also brought together the growing ecosystem for the 2025 Open Networking & Edge Summit in London (alongside KubeCon + CloudNativeCon Europe) where thought leaders explored the future of AI-native, open end-to-end networks.

As we look toward 2026 and beyond, we acknowledge the increasingly complex geopolitical, economic and technological terrain. Yet one thing remains unchanged: open source thrives in uncertainty. Through collaboration, transparency, and shared purpose, we're building resilient, interoperable networks designed to power the future. I invite you to dive deeper, explore our projects and join us in shaping the next chapter of open innovation.

Message from the CTO

BY RANNY HAIBY
CTO, Networking, Edge, Access at the Linux Foundation

2025 was characterized by the successful culmination of our industry's shift to cloud-native architectures. Our diverse open source communities, spanning operators, vendors, and developers, fully embraced the promise of containerization and, crucially, standardized on GitOps-centric approaches. Projects like Nephio matured rapidly, providing carrier-grade intent automation that simplifies the deployment and lifecycle management of cloud-native network functions (CNFs). Simultaneously, ONAP pivoted to a modular, GitOps-ready framework, while the CNTi (Cloud Native Telecom Initiative) provided the essential testing and best practices needed to ensure confidence and interoperability across the ecosystem. This collective commitment to declarative, code-driven infrastructure has reduced operational complexity and accelerated feature velocity, cementing LFN's role as the nexus for modern network automation.

As the network infrastructure becomes cloud-native, the next logical frontier was integrating embedded intelligence. Throughout 2025, AI ceased to be an adjunct technology and became foundational to network operations. Recognizing the need for a standardized, ethical, and open platform to manage this complexity, we proudly launched the Essedum project. Essedum provides the community with a comprehensive framework for applying AI to networking use cases, spanning the entire end-to-end pipeline—from secure data sharing and ingestion to model training and multi-platform deployment. This platform ensures that our members can build, deploy, and govern intelligent network applications efficiently, moving beyond fragmented approaches toward truly data-driven, closed-loop automation within the network itself.

Looking forward to 2026, our sights are set on the next paradigm shift: the creation of truly autonomous networks powered by Agentic AI. This transition moves us beyond static automation scripts to systems capable of reasoning, planning, and executing multi-step goals without continuous human intervention. LFN will be at the forefront of this evolution, focusing on the open frameworks and governance models required for these intelligent agents to interact safely and efficiently. We will dedicate significant community resources to defining how agentic systems can interpret high-level human intents and translate them into reliable, self-governing network actions, paving the way for the cognitive, zero-touch network of the future.

Message from the Technical Advisory Council

BY FATIH NAR
Distinguished Chief Architect, Red Hat – Chair, LFN Technical Advisory Council

In 2025, the Technical Advisory Council has adapted the focus from governance mechanics to strategic outcomes that matter; production-ready implementations, measurable business value, and sustainable open-source collaboration. Our work centered on three interconnected themes that define networking's transformation.

Applied AI moved from experimentation to production operations.

The community demonstrated that AI's value lies not in model size or hype cycles, but in solving real operational challenges, autonomous network management, intelligent root cause analysis, and predictive service assurance. Projects like Essedum aim to deliver the infrastructure for domain-specific AI pipelines, while Salus focusses on responsible AI guardrails addressing bias, privacy, and transparency concerns that enterprises rightfully demand. LFN community work in applied AI showcased practical implementations spanning revenue assurance, network operations automation, and intent-driven orchestration, proving that value-driven AI adoption outperforms market-driven Fear of Missing Out (FOMO).

Infrastructure modernization enabled distributed intelligence at scale.

Production networks require more than powerful AI models, they need the transmission systems to deliver intelligence where decisions happen. The TAC encouraged hybrid cloud architectures, distributed AI serving capabilities, and standardized observability frameworks that make network behavior transparent and actionable. Community contributions in OpenTelemetry-based observability, disaggregated inference architectures, and edge-to-core intelligence distribution established patterns for scalable, vendor-neutral deployments that prioritize price-performance and operational security.

Open source collaboration accelerated real-world adoption.

Cross-project integration matured significantly. Nephio's Kubernetes-based automation connected with Cloud Native Telecom Initiative (CNTi)'s cloud-native best practices, Anuket's infrastructure specifications, and the AI Task Force's data sharing initiatives. The 5G Super Blueprint expanded its use case library, demonstrating how LFN projects compose into production solutions rather than isolated components. This interoperability, technical and organizational, represents the compounding advantage of open-source collaboration over proprietary fragmentation.

Message from the Technical Advisory Council

Looking towards 2026

The TAC encourages continued focus on openness, collaboration, and measurable value creation. We need production reference architectures with performance data, not just conceptual frameworks. We need observability standards that work across vendor boundaries, enabling operators to actually understand and optimize their networks. Most critically, we need AI implementations justified by operational improvements and cost reductions, not technology enthusiasm alone.

To our maintainers, contributors, and adopters; your commitment to solving real problems with open technologies drives networking's evolution.

The future belongs to organizations that deploy intelligence thoughtfully, observe systems comprehensively, and collaborate openly principles that LFN projects continue to exemplify.

LFX Insights

Open source, like any other industry vertical, tracks its performance and growth through the use of a set of objective performance indicators. LFN does this through the use of our LFX Insights analytics platform that provides an overview on the health and emerging trends regarding each LFN project.

One of the biggest challenges for open source projects is being able to define, track, and review key project metrics. The LFX Insights analytics tool provides insights on the health of LF Networking projects and what trends are happening regarding the number of contributors, number of commits, etc.

Both Stable & Growing Communities

Over time, the LFN project community has ebbed and flowed as is the natural course of development. Some of the more mature projects (such as ONAP and Nephio) have held steady with contributions over the past few years (with peaks at milestone releases, for example):

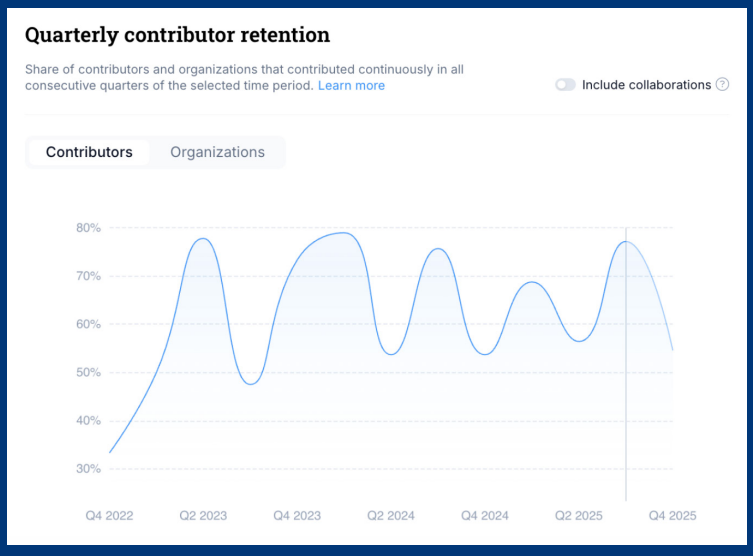

Nephio contributor retention, Q4 2022 - Q4 2025

LFX Insights

ONAP contributor retention, Q4 2022 - Q4 2025

Meanwhile, our newer projects are still in growth mode and in the process of establishing strong contributor bases.

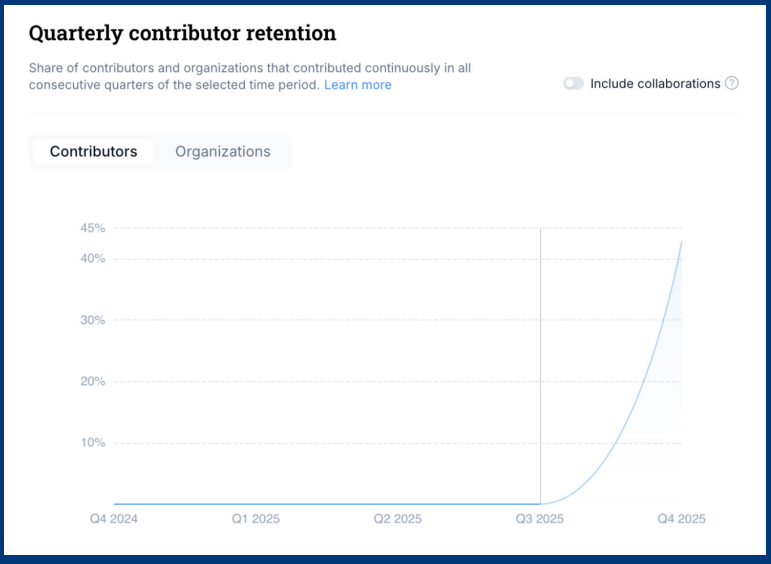

Essedum contributor retention, from inception (Q2025) to present.

Currently, LFN:

Hosts **12** projects | Engages **12,600+** contributors | Has a monetary value of close to **$2B**

LFX Insights

The following snapshot provides examples of the contributions and participation seen across several LFN projects this year:

CNTi

The mission of the Project is to develop and maintain cloud native best practices and related tooling.

Contributors	206
Organizations	72
Software value	$12M

FD.io

FD.io (Fast Data Input/Output) is an open source project focused on data plane packet processing and network I/O...

Contributors	2,167
Organizations	464
Software value	$287M

L3AF

L3AF is an eBPF based platform enabling network functions a service. L3AF provides a platform to launch and manage eB...

Contributors	913
Organizations	156
Software value	$756K

Nephio Project

The mission of the Project is delivers a carrier-grade, simple, open, Kubernetes-based cloud native automation solution...

Contributors	347
Organizations	95
Software value	$24M

ONAP

ONAP (Open Network Automation Platform) is an open source networking project that provides a comprehensive platform f...

Contributors	4,685
Organizations	377
Software value	$1.2B

OpenDaylight

OpenDaylight is an open source platform for software-defined networking (SDN) that provides a modular, pluggable, and...

Contributors	2,057
Organizations	353
Software value	$360M

Check back in with LFX Insights often, as improvements and new functionalities are underway.

5G Super Blueprint: Centerpiece of Open Integration

The 5G Super Blueprint is a LF Networking-led year-round initiative that provides opportunities for open-source communities to build upon some of the most important technologies across the globe through cross-community collaboration. Bringing these pieces together is often challenging, not just for the developers, but for the real-world end-users who need to stitch together multiple open source and proprietary solutions.

The 5G Super Blueprint addresses open-source solutions for industry challenges bringing together multiple projects, communities, and companies to blueprint, prototype and integrate real-world use cases to showcase the capabilities of open-source projects in the context of the full ecosystem in which they exist.

LFN, Collaboration Hub for Open Source Networking, Edge and Access

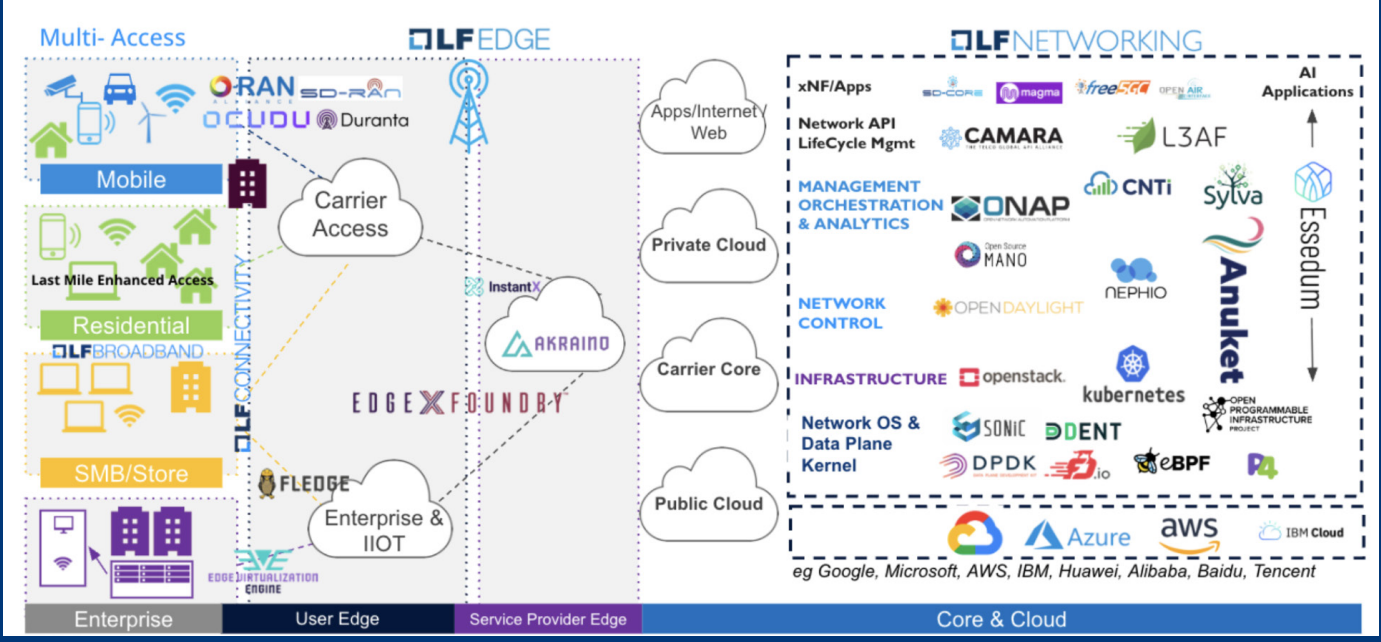

5G Super Blueprint: Centerpiece of Open Integration

5G Super Blueprints Library

The 5G Super Blueprints Library is the central reference point for completed 5G Super Blueprints integration projects. Similar to a public or school library; content of the 5G Super Blueprints library varies in topic/technology, project scope, type and breathe.

5G Super Blueprints integration projects range from complete end-to-end solutions to technology building blocks that can be consumed, evolved, and expanded upon. Our goal is to create individual blueprints for network and industry blueprints, and establish a library of blueprints.

There are currently 9 entries in the 5G SBP Library, with more on the horizon.

Cutting-Edge 5G Super Blueprint Projects Underway: Expanding the blueprints - Going beyond 5G

This year, the community had the opportunity to look further away than just 5G, with deployments of this generation of wireless communications already deep in production around the globe. The focus is naturally shifting towards the next generations of wireless communications, with AI becoming more natively integrated with networks. End users are no longer limited to one type of connectivity, be it public cellular network, private networks, WiFi and non-terrestrial network. The focus of blueprints is shifting towards the needs of modern applications, and making seamless connections over autonomously controlled networks. Some of the notable blueprint activity this year focused around:

- Leveraging Open Source Edge and Networking projects to provide “application centric” connectivity

Harnessing the power of generative AI to augment and enrich the control plane of networks, taking them a step closer to autonomous networks guided by the intent of the operators, expressed in natural human language.

LFN Project Updates

Anuket

Anuket, strengthened its role in 2025 as LF Networking's reference blueprint for telco-grade cloud infrastructure, while formally achieving "Graduated" project status within LFN. The project delivered the Quinnipiac release on March 20, 2025, and planned the Red River release for January, 2026, underscoring its continued focus on providing stable, standards-aligned reference models and architectures rather than fast-moving feature code. Quinnipiac was highlighted by the TSC and LF Networking as another proof point of Anuket's unique position in the industry: integrating operator requirements, normalization, and open source implementation guidance under one initiative.

Technically, Quinnipiac brought significant updates across Anuket's documentation and test frameworks. The Reference Model (RM) refined requirements around simultaneous multithreading (SMT), updated virtual DU (vDU) real-time and PTP timing constraints to align with O-RAN, refreshed references, and completed automation of GSMA's NG.126 specification output—cementing Anuket's role as an upstream source for industry standards.

The RA2 Kubernetes reference architecture advanced to newer Kubernetes baselines, added specifications for a CaaS manager interface and operations, strengthened SMT and real-time vDU requirements, and aligned closely with the Sylva project by synchronizing VDU requirements and merging RC2 components into RA2. In parallel, the Functest-Kubernetes sub-project kept pace with the Kubernetes ecosystem, releasing a 1.32 test framework for automated cluster verification in the Quinnipiac timeframe.

2025 also saw Anuket rationalize and modernize its project portfolio. Legacy implementation tracks such as RI1, RI2, Kuberef, and Barometer were archived to focus community effort on the reference specifications and test frameworks that are most widely reused today. At the same time, work began to refresh RA1 to current OpenStack releases, ensuring that both VM-based and container-based infrastructures remain first-class citizens in the Anuket model. This streamlining reflects how Anuket's artifacts are increasingly consumed by other initiatives—most notably the Sylva project and external efforts such as the EU Telco Cloud Reference Architecture, which implement Anuket architectures and conformance suites as part of their own blueprints.

Community-wise, Anuket maintained an active TSC, continued regular TSC meetings, and deepened its collaborations with projects like Sylva and the O-RAN Alliance, even as it acknowledged the challenges of sustaining a mature, documentation-centric project in a fast-moving industry. Collectively, the Quinnipiac release, the Graduated status, the RM/RA2 and Functest updates, and ongoing alignment with standards bodies position Anuket as the enduring telco cloud reference compass for LF Networking—providing operators, vendors, and adjacent projects with a stable, widely recognized foundation for building and validating telco-grade cloud infrastructure.

CNTi

The Cloud Native Telecom Initiative (CNTi) continued to mature in 2025 as LF Networking's focal point for cloud-native best practices and conformance in telecom. Building on its CNF Test Suite—a vendor-neutral, open source framework that validates Cloud-Native Network Functions (CNFs) against best practices—CNTi delivered 8 releases offering a steady stream of prioritized enhancements, bug fixes, new tests, and maintenance updates, resulting in a more robust and user-friendly validation experience.

From January 1 onward, the community produced over 100+ commits from 8 contributors across 4 companies, backed by more than 35,000 CI jobs executed to validate code and releases, underscoring the project's emphasis on reliability and repeatability in CNF testing.

CNTi's top focus area was technical hardening of the CNF Test Suite. Community efforts focused on code cleanup, refactoring, and test hardening—including work to address "unhealthy" GitHub Actions pipelines. Now the CNF Test Suite benefits from increased stability through use of deterministic Github Actions.

In parallel, the community launched a dedicated effort to stand up an open-source reference CNF, beginning with free5GC, that can exercise the full breadth of tests and serve as a living benchmark for real-world network functions. This push goes hand-in-hand with evolving the project's best practices (for example, revisiting expectations around logging to stdout/stderr to ensure they reflect real telco operations) and adding new tests that cover emerging patterns such as operators, advanced security policies, and observability.

Collaboration was another defining theme. CNTi deepened its partnership with Project Sylva, where the CNTi Test Suite is integrated into the Sylva Validation Center to verify that workloads on Sylva-compliant stacks meet cloud-native standards. Early validations of products such as Open-Xchange PowerDNS and Titan.ium DNS surfaced real-world feedback that fed directly into Test Suite improvements and were published in a subsequent release. Both communities shared progress in a jointly issued blog: Advancing Cloud Native Telco - A Collaboration Between Sylva and CNTi

Discussions also advanced with additional stakeholders—including CSPs, vendors, and communities like NGMN and major cloud providers—to expand both the pool of contributors and the number of CNFs being tested against CNTi's vendor-neutral criteria.

Community visibility and outreach remained strong throughout 2025. CNTi community members were active at KubeCon, Cloud Native Telco Day, and Sylva Days, where they staffed the LFN booth, ran live CNF Test Suite demos, and featured CNTi in multiple talks as a key enabler of cloud-native adoption for networking. CNTi closed Q4 visibility activities strong with a podcast (LFN's new "Layer 8: Voices of Open Networking), featuring viewpoints and thought leadership from two prominent members of the CNTi community, Philippe Ensarguet of Orange and Marc Price of MATRIXX Software.

Combined, these technical, collaborative, and outreach milestones position CNTi as a cornerstone of LF Networking's strategy to give operators and vendors a trusted, open, and extensible way to validate CNF readiness—and to accelerate confident cloud-native deployment in telecom networks.

Duranta

Duranta joined LF Networking in 2025 as a new candidate project from the OpenAirInterface community, giving LFN's open end-to-end stack a dedicated, open source O-RAN CU/DU + UE reference implementation. Built on OpenAirInterface (OAI) software and hosted with neutral governance under LFN, Duranta is designed as a research-grade RAN and UE stack that can be used for lab work, integration, and early deployments, while benefiting from LFN's security processes, governing board structure, and industry-scale collaboration model.

With Duranta, LFN's 5G Super Blueprint now gains a critical RAN building block that complements earlier integrations with Nephio and CNF automation, allowing the community to demonstrate and iterate on open, cloud-native RAN architectures more holistically.

The project was formally reviewed by the LFN TAC and announced as an LFN Candidate project in mid-2025, with a public launch at Open Source Summit Europe in Amsterdam alongside a joint LF Networking / OpenAirInterface press announcement. As part of this launch, Duranta also adopted the new Collaborative Standards Software License (CSSL), aligning its licensing model with the needs of both research communities and standards-driven ecosystems and tying into broader OAI outreach about CSSL and its relationship to Duranta.

Within LFN, Duranta is scoped as both a RAN & UE project—developing CU/DU and UE software plus test assets and documentation—and an operations project, focused on CI/CD pipelines, packaging, lab topologies, and reproducible test harnesses for end-to-end validation.

Looking ahead, establishment of the Technical Steering Committee (TSC) is underway, which will formalize the project's technical governance and roadmap. Early priorities include deepening Duranta's role in LFN's open RAN blueprints, expanding integration with Nephio and other automation projects, and using LFN's shared labs to push toward deployment-ready, standards-aligned open RAN software. Together, these steps position Duranta as the RAN/UE anchor for LFN's AI- and cloud-native networking vision—bridging academic research, open source development, and real-world 5G/6G deployment needs.

Essedum

Essedum strong entrance into the LF Networking in 2025 as its flagship AI application and agent framework for networking. The project became an LFN Candidate project in March 2025, with a charter to accelerate the integration of AI data, models, and applications across telecom and networking use cases. Seed code contributed by Infosys provided a full-stack AI app–building framework (backend and UI) and collaboration with Thoth from the Anuket community is underway, to deliver built-in data anonymization and privacy-preserving preprocessing. TSC governance and regular community meetings were quickly established, giving Essedum a clear roadmap and an active home for contributors across vendors, operators, and academia.

A major technical milestone with Essedum Release 1.0 (R1) seed code, published on GitHub as a complete backend-and-UI platform for building AI-powered networking applications. This first release delivers a modular framework spanning data ingestion, pipeline orchestration, and model deployment, including capabilities for defining secure connections between systems, managing datasets from storage buckets, databases, and REST APIs, orchestrating training and inference pipelines, integrating with cloud AI services (such as AWS SageMaker, Azure ML, and GCP Vertex AI), and managing endpoints and adapters from a centralized interface.

The Responsible AI toolkit from project Salus adds guardrails for governance, transparency, and ethical use of AI in networking, with R2 brings better developer experience, security framework, deeper test automation and tighter Salus integration, followed by an R3 roadmap focused on agentic capabilities.

Community visibility ramped quickly as well. Essedum was highlighted in keynotes at Open Source Summit (OSS) Hyderabad and featured in LFN's broader narrative around AI-native networking and AI security, including coverage of Essedum as a reusable, domain-specific AI framework for telecom networks.

A live demo of the framework during the June 10 community call showcased end-to-end workflows—from connecting data sources, to building pipelines, to deploying models—helping anchor regular, well-attended community sessions that continue to onboard new users and contributors. With R1 and R2 in place, and a clear roadmap into agentic capabilities and richer AI governance, Essedum is quickly becoming the AI "app and agent factory" for networking, unifying LFN's agentic AI, data-sharing, anonymization, and responsible-AI efforts into a single, extensible open platform.

BY DAVE WALLACE , FD.IO TSC CHAIR AND PRINCIPAL SOFTWARE ENGINEER, CISCO

The FD.io release train continued to deliver in 2025, sustaining its reputation as the world's fastest and most secure open source networking data plane. The community shipped three Vector Packet Processing (VPP) releases—25.02 in February, 25.06 in June, and 25.10 in October—each paired with comprehensive CSIT performance and regression reports (CSIT-25.02, 25.06, and additional CSIT reporting aligned with the 25.10 release window) to validate functionality and track trends over time. VPP 25.02 alone included more than 269 commits and 21 new features, with roughly 100 issues fixed, while VPP 25.06 added 28 new features and over 260 commits, including 92 fixes, underscoring a continued focus on both innovation and quality.

Testing and performance engineering remained a core strength. The CSIT team delivered detailed 25.02 and 25.06 release reports in March and July, followed by additional CSIT work in November aligned with the 25.10 release, each documenting extensive functional and performance coverage across VPP, DPDK, and TRex, along with trending analysis to guard against regressions. In parallel, FD.io advanced its strategic goal of expanding testing to the latest-generation silicon. Work progressed to bring Intel's Granite Rapids and Sierra Forest Xeon platforms into the FD.io lab, alongside next-generation NVIDIA ConnectX-7 200 Gbps NICs and Marvell Octeon 10 DPUs—devices that pair closely with VPP and can deliver up to multi-fold packet-processing performance improvements. By Q4, these latest-generation testbeds were fully integrated into the CSIT pipelines.

This broadened hardware coverage helps ensure FD.io remains relevant for both traditional network workloads and emerging data center, edge, and AI use cases. Operationally, the project made important investments in sustainability and efficiency. The migration from Jira to GitHub Issues was completed, simplifying the developer workflow and reducing tooling costs, while planning and proof-of-concept work to transition CI from Jenkins to GitHub Actions matured into a full migration completed in Q4—modernizing the automation stack that underpins VPP and CSIT's high-velocity release cadence and delivering additional operational cost savings.

These improvements, combined with ongoing refinement of infrastructure and lab management, position FD.io for more scalable and maintainable operations in the years ahead. Community engagement also remained strong. FD.io maintained visibility at industry and open source events, including multiple talks and an in-person meetup at FOSDEM 2025, and continued to highlight real-world deployments such as platforms built on Marvell Octeon 10 DPUs and other VPP-based solutions across routers, gateways, and edge systems.

Collectively, these efforts illustrate a mature, production-grade project that continues to evolve: expanding hardware coverage, modernizing its CI and tooling, and delivering predictable, high-quality releases. FD.io enters 2026 well-positioned to remain the universal dataplane of choice for high-performance, programmable networking.

LFN Project Updates

L3AF advanced significantly in 2025 as LF Networking's eBPF-based networking and observability platform, anchored by the L3AF 2.1.0 release early in the year. This release extended L3AF's role as a complete life-cycle manager for eBPF programs—supporting sophisticated traffic management, security, and visibility use cases in large-scale production environments such as Walmart's global retail and e-commerce infrastructure.

With 2.1, L3AF sharpened its focus on cloud-native deployment and zero-downtime operations, reinforcing its value as a vendor- and cloud-agnostic way to orchestrate kernel-resident functionality.

From a technical standpoint, L3AF 2.1 delivered several key capabilities that raise the bar for eBPF program management. The release introduced a graceful restart for l3afd, allowing control-plane upgrades without interrupting running data-plane programs, and enabled l3afd to run cleanly inside containers for modern cloud-native environments. L3AF images are now published on Docker Hub, making it easier to integrate into Kubernetes-based CI/CD pipelines.

On the data-plane side, the eBPF Package Repository was enhanced with BPF CO-RE support to improve portability across Linux kernels, while new support for kprobes, uprobes and tracepoints expanded observability and troubleshooting options. Additional enhancements included the ability to attach programs dynamically to interfaces created on the fly. Together. we're making L3AF more robust in complex, multi-interface and multi-cloud environments.

Looking ahead, the community translated these foundations into a focused roadmap. Work is underway to deepen support for observability use cases and eBPF program metrics, to load TC programs using tcx, and to support newer kernel baselines such as "Noble" (kernel 6.8+), with a 2.2 added support to probe which increases observability use cases this release targeted around Q1 of 2026. In parallel, the project has identified coexistence with Cilium CNI in Kubernetes environments as an important future milestone, aligning L3AF's ePPF-Program-as-a- Service model with widely adopted cloud-native networking stacks.

Together, these advances position L3AF as a practical, production-hardened control plane for eBPF—one that simplifies how operators and developers deploy, chain, observe, and evolve kernel-resident programs across diverse cloud and edge environments.

NEPHIO

Nephio solidified its position in 2025 as a strong LF Networking project for cloud native, intent-driven network automation, delivering two major releases—Release 4 (R4) and Release 5 (R5)—that moved the project squarely into production-ready territory for O-RAN, multi-cloud, and GitOps-driven deployments. R4, launched in early 2025, was the result of collaboration from 250+ contributors across 45 organizations, underscoring the breadth of the ecosystem rallying around Nephio. It introduced key advancements such as (pre-standard) O-RAN O2 IMS and SMO FOCOM interface support for O-Cloud lifecycle management, native FluxCD GitOps automation, and major platform, UI, and security improvements aimed at large-scale, carrier-grade deployments.

Building on this foundation, Nephio Release 5 (R5), announced mid-year, further strengthened scalability, flexibility, and operability. R5 added support for multiple reconciliation agent types (ArgoCD and FluxCD), delivered significant robustness and scalability enhancements to the Porch platform, introduced bare metal cluster support, and expanded observability parameters to pave the way for future assurance and closed-loop automation use cases. The release also advanced O-RAN and transport orchestration scenarios, continuing Nephio's focus on end-to-end RAN, core, and transport automation across multi-vendor, multi-cloud environments.

Nephio's vision for GenAI-assisted automation also progressed from concept to concrete roadmap in 2025. Building on the "Nephio & Gen AI: Empowering Efficiency in 5G/O-RAN and Beyond" work and an accompanying white paper, the community continued exploring how generative AI can simplify intent capture, template generation, and configuration workflows—enabling AI-authored configurations, smarter troubleshooting, and data-driven optimization over time.

Looking ahead, Nephio's roadmap calls for deeper O-RAN integrations, additional security and performance hardening, and tighter collaboration with projects such as Sylva and CAMARA, along with broader cross-industry adoption (including managed implementations like Google's Telecom Network Automation).

Together, R4, R5, and the GenAI and ecosystem workstreams position Nephio as a cornerstone of AI-ready, cloud-native automation for telecom and beyond.

OPEN DAYLIGHT

OpenDaylight marked another milestone year in 2025 as one of the longest-running and most widely deployed open source SDN controllers, balancing platform modernization with a strong focus on stability and usability. The project delivered the Scandium SR2 (and later SR3) service releases, which were primarily maintenance-focused, tightening overall stability through targeted bug fixes and dependency upgrades—including bumping odlparent to 14.0.7, along with fixes in core components such as YANG Tools (e.g., issues YANGTOOLS-1655 and YANGTOOLS-1651). These updates reinforced OpenDaylight's move to a modern toolchain, with Java 21 now required for both runtime and development and Maven 3.9.5+ as the baseline build environment, helping downstream integrators align on a consistent, future-ready platform.

In parallel, the community delivered on its roadmap to the 22nd major release, "2025.03 Titanium." Titanium formalized OpenDaylight's next-generation platform baseline: standardizing on Java 21 end-to-end, upgrading core building blocks (odlparent, YANG Tools, MD-SAL, controller, NETCONF, AAA, and others), and providing a documented upgrade path from Scandium via the Titanium platform upgrade guide

The release also completed the long-planned migration from Akka to Apache Pekko in clustering, simplifying future maintenance while preserving high availability for multi-node controller deployments.

By October, Titanium-SR1 and Scandium-SR3 were both available as supported releases, giving users a choice between the latest feature line and a hardened Scandium stream, each with aligned documentation and downloads.

Beyond code, OpenDaylight invested in making SDN experimentation and learning more accessible. A new ODL tutorial authored by Pedro David Arjona ("Testing Software Defined Networks with OpenDaylight and Mininet") was highlighted to the community, offering a practical, step-by-step guide to building an SDN laboratory using the Mininet network emulator and the OpenDaylight controller. The tutorial walks readers through installing OpenDaylight, wiring it to Mininet, and designing and testing their own SDN experiments—lowering the barrier for students, researchers, and operators to explore SDN concepts hands-on.

Collectively, the Scandium service releases, Titanium platform upgrade, and new educational material underline OpenDaylight's continued evolution as a stable, modern, and well-documented SDN control platform within the LF Networking ecosystem.

ONAP

BY BYUNG-WOO JUN | *ONAP TSC Chair and Principal Engineer, Ericsson Software Technology*

ONAP's second major evolution accelerated in 2025, as the project firmly completed its transition from a monolithic platform to a modular collection of network automation functions. The year was anchored by two major releases: Oslo (15.0.0) on January 16, 2025 and Paris (16.0.0) with marketing release sign-off on July 10, 2025 (with progress on the forthcoming Quebec release well underway!). Both releases advanced the ONAP Streamlining Evolution with stronger security, modernized runtimes, and deeper support for intent-based, GenAI-driven automation, while reinforcing ONAP's position as reusable network automation components rather than a single, tightly coupled platform.

The Oslo release delivered tangible progress on the roadmap laid out in the 2024 LF Networking annual report. It hardened ONAP's security posture by defaulting component ports to HTTPS, removing hard-coded credentials, enabling non-root Kubernetes pods, and systematically addressing CVEs, all underpinned by Istio service mesh and ingress gateways for secure intra- and inter-component traffic. Key building blocks such as SO, SDC, DCAE, Multi-Cloud, A&AI, CDS, CCSDK/SDNC, CPS, UUI, Portal-NG, and the Policy Framework were upgraded and refreshed to current ecosystem versions, aligning ONAP with modern cloud-native stacks. At the same time, Oslo expanded support for intent-based, declarative automation and GenAI, including domain-specific data services and Model-as-a-Service (MAAS) capabilities in UUI, while CPS and the Policy Framework achieved OpenSSF Gold badging, marking ONAP as a leader in secure, best-practice open source network automation.

The Paris release built on this foundation to operationalize ONAP's modular future. Paris continued the Streamlining Evolution by enabling independent build–to–runtime operation of ONAP components, supported by GitOps pipelines using Argo CD for faster, automated deployment across multi-cluster and multi-cloud environments. Packages and intents can now be onboarded via Git, triggering repository-driven automation flows that align with ONAP's intent-driven architecture. Paris also introduced Policy–OPA-PDP integration for cloud-native, standards-aligned policy decisioning, expanded MAAS in UUI for GenAI to manage AI artifacts, and delivered coordinated package upgrades to reduce vulnerabilities. Importantly, Paris marked the transition to a more flexible release cadence, after which individual ONAP projects manage their own agile release cycles while still reporting into ONAP governance for architectural consistency and interoperability.

From a governance and organizational standpoint, 2025 was also the year that "semi-standalone ONAP projects" became the norm rather than a concept. New organizational and architectural models clarified how PTLs, the TSC, ARCCOM, SECCOM, OOM, and Documentation teams coordinate: common services (security framework, Istio/Keycloak, Argo CD, CI/CD, documentation) remain centralized, while individual projects own their features, lifecycles, and versioning, and expose their functions directly to consumers. This modular structure allows operators and vendors to select only the ONAP components they need (e.g., "lightweight ONAP" deployments, CCVPN and 5G slicing use cases, or AI-assisted intent-driven automation), while still benefiting from shared security, architecture, and quality guidance.

LFN Project Updates

The community remained healthy and active throughout 2025. According to LFX Insights, ONAP maintained over 110 active contributors throughout 2025, with a 55% quarterly contributor retention rate and a diverse contributor base where eight individuals account for just over half of contributions, which are indicators of a stable yet well-distributed ecosystem.

Together, the Oslo and Paris releases, the move to semi-standalone projects, and the strong contributor base position ONAP to remain a reference set of intent-driven, GenAI-ready network automation functions for operators, cloud providers, and enterprises as they build AI-native, standards-aligned 5G and next-generation networks.

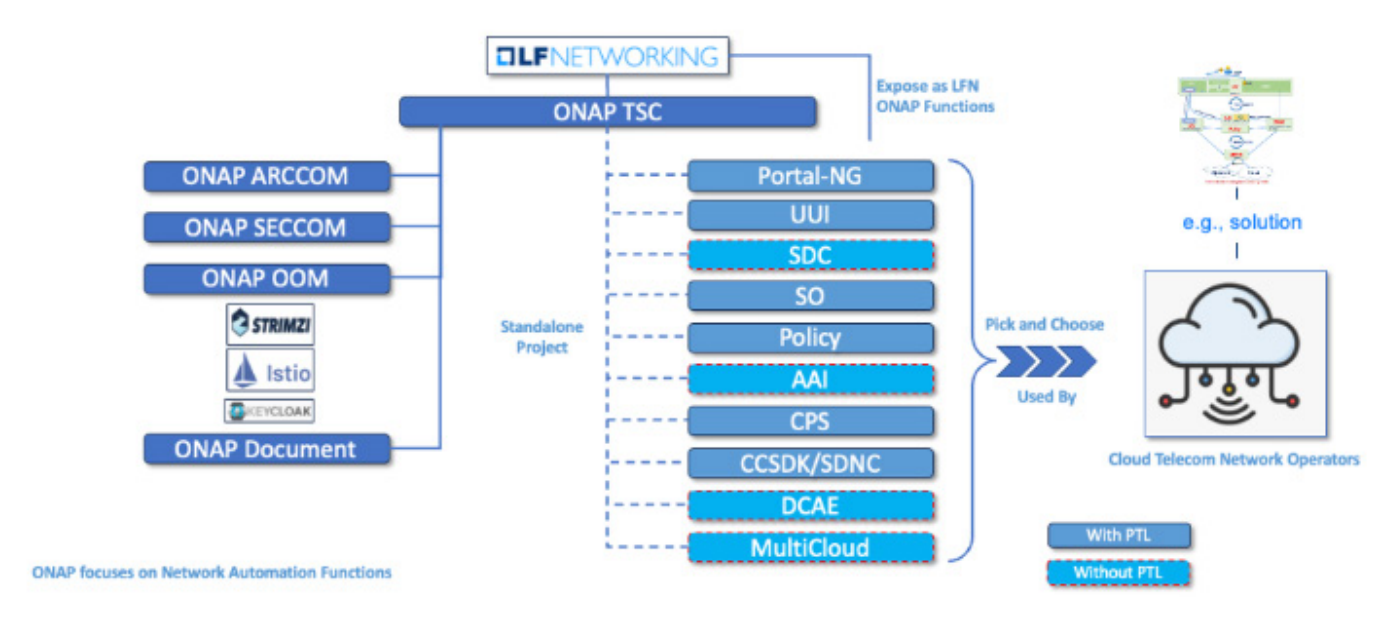

LFN Project Updates

Paraglider's first full year as an LF Networking project in 2025 was all about turning its vision—a unified, intent-centric control plane that simplifies single cloud and multi-cloud networking by abstracting low-level and cloud-specific building blocks—into a concrete, usable open source implementation.

The key milestone for the year was the Paraglider 0.1.0 release, announced in May 2025 as the first official (though not yet production-ready) version since open sourcing under LFN. Building on the 2024 launch of Paraglider as LFN's first "Candidate" project, the 0.1.0 release delivered core multi-cloud capabilities including Kubernetes cluster and private endpoint creation, support for more connection combinations across clouds, and the ability to attach existing compute resources into Paraglider-managed topologies. It also incorporated bug fixes from early community feedback and introduced feature flags so users can selectively enable experimental features as the control plane evolves. Together, these capabilities expand the set of scenarios in which Paraglider helps developers and operators express their desired connectivity at a high level and let Paraglider orchestrate the underlying cloud-specific resources.

Throughout 2025, the community focused on deepening functionality and preparing for broader adoption. Work progressed to improve feature parity across cloud plugins (including support for private endpoints and existing resources in each provider), to strengthen networking resource compliance validation and remediation, and to explore an AWS plugin, reflecting strong community feedback to add support for this major cloud. In parallel, contributors kicked off research into AI integration and "reverse-engineering" intent from existing network deployments—mapping real-world cloud configurations back into Paraglider's model to simplify onboarding and future automation. Looking ahead, priorities include expanding engagement with potential users and contributors through conferences and user groups, adding further cloud plugins such as Oracle, and formalizing project representation in LFN's technical governance. These efforts set the stage for Paraglider's progression from prototype to production-ready multi-cloud tenant networking in the coming cycles.

Industry Impact

By the Numbers

9
PRESS RELEASES

2.4K
PRESS CLIPS/ MENTIONS

15
MEDIA/ANALYST BRIEFINGS

845M
POTENTIAL AGGREGATE REACH

107K
LINKEDIN IMPRESSIONS

A Sampling of Industry Event Participation:

- Mobile World Congress Barcelona
- KubeCon + CloudNativeCon EU
- Cloud Native Telco Day
- Open Source Summit Europe
- DTW Ignite
- AI_Dev
- AI Agent Internet Architecture and Technology Seminar
- Open Source Summit India
- FutureNetWorld
- Open Networking & Edge Summit
- Open Telco LATAM Summit
- Telco Cloud Native Summit (TelecomTV)
- ORAN Alliance Workshop
- ORAN Alliance and OSFG Workshop

Industry Impact

TOP PRESS QUOTES

"The open-source project from LF Networking allows organizations to leverage their current AI/ML platform investments while gaining networking-specific capabilities that general-purpose platforms cannot provide effectively."

— NETWORK WORLD

"By aligning with LFN and other Linux Foundation projects, ATIS says 'we are combining the strengths of open standards and open source to accelerate deployment of secure, sustainable and intelligent next-generation networks."

— RCR WIRELESS NEWS

"LF Networking (LFN)... has struck several alliances to advance development in AI-native networks such as 6G and open radio access network (open RAN)."

— SDXCENTRAL

"LF Networking... emphasizes the convergence of open source, artificial intelligence, and cloud-native technologies as the foundation for next-generation networking infrastructure."

— NETWORK WORLD

"Enabling the adoption of cloud-native applications and artificial intelligence, open source is proving to be an important driver of innovation in enterprise networks."

— CXOTODAY.COM

Industry Impact

Top Headlines

NETWORKWORLD

Linux Foundation launches Essedum 1.0 to simplify AI integration in network operations Read Now ▸

sdxcentral

Open RAN, AI-RAN boosted by Linux Foundation partnerships Read Now ▸

CXOtoday.com
IT Perspective for Decision Makers

Transforming Network Operations with Open Source Read Now ▸

TELECOM TV

What Telcos Want from Open Source Read Now ▸

TELECOM TV

Reducing the Risks of AI in the Network Read Now ▸

sdxcentral

Linux Foundation extends AI model leadership with new frameworks Read Now ▸

NETWORKWORLD

Linux Foundation shares New AI Projects, Milestone Releases Read Now ▸

RCR Wireless News
NEWS. INSIGHT. IMPACT.

Test and Measurements: ATIS, Linux Foundation seal 5G/6G collaboration Read Now ▸

Social Media

LF Networking achieved a strong year of social visibility, engagement, and video-driven reach across LinkedIn and YouTube.

+107K
Linkedin impressions

127.4K
YouTube impressions

8%
Growth in LinkedIn Engagement Rate

24
Videos uploaded

1.3K
Watch hours

39%
growth in LinkedIn Followers

Most viewed content focused on:

- Open source leadership and ecosystem direction
- AI-driven transformation across telecom and networking
- Cloud, edge, and data-centric future of service providers
- Strategic industry outlooks and thought leadership from experts

Follow and Join Linux Foundation Networking

To learn more about open source networking projects under the LF Networking umbrella, please visit the LF Networking website.

Follow us on LinkedIn (LF Networking LinkedIn and ONE Summit by LFN LinkedIn) and X/Twitter so you don't miss out on what's happening in the world of open source networking.

Bookmark and subscribe to the LF Networking YouTube channel to watch all recorded sessions from LF Networking events.


LFNETWORKING
Thank you all for a successful 2025
and here's to embracing the
limitless skies in 2026!