



OLF NETWORKING

2024 Year In Review

The Collaboration
Hub Takes Flight

Table of Contents

- Exec Summary & Introduction..... 03
- The Collaboration Hub Takes Flight..... 05
- Message from the LFN Board..... 11
 - LFN Board Members 12
 - LFN Members 13
- Message from the General Manager..... 14
- Message from the CTO 16
- Message from the Technical Advisory Council..... 17
- LFX Insights..... 18
- 5G Super Blue Print: Centerpiece of Open Integration..... 20
- LFN Project Updates 22
 - Anuket..... 22
 - CNTi 24
 - FD.io..... 25
 - L3AF 27
 - Nephio 29
 - OpenDaylight 30
 - ONAP..... 31
 - Parglider 34
- Industry Impact..... 35
- Conclusion 38

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.”

The Linux Foundation Networking (LF Networking, LFN) is the largest set of open source networking projects in the world formed by a broad industry coalition with the goal of fostering a commercial-ready networking ecosystem that embraces open, emerging and evolving technologies.

With nearly seven years of development and solid project-market fit, LF Networking software and initiatives continue to lay the groundwork for digital transformation. This includes foundational support for network infrastructure and services across service providers, cloud providers, enterprises, vendors, and system integrators, facilitating rapid interoperability, deployment, and adoption. Key tenets of 2024 progress include:

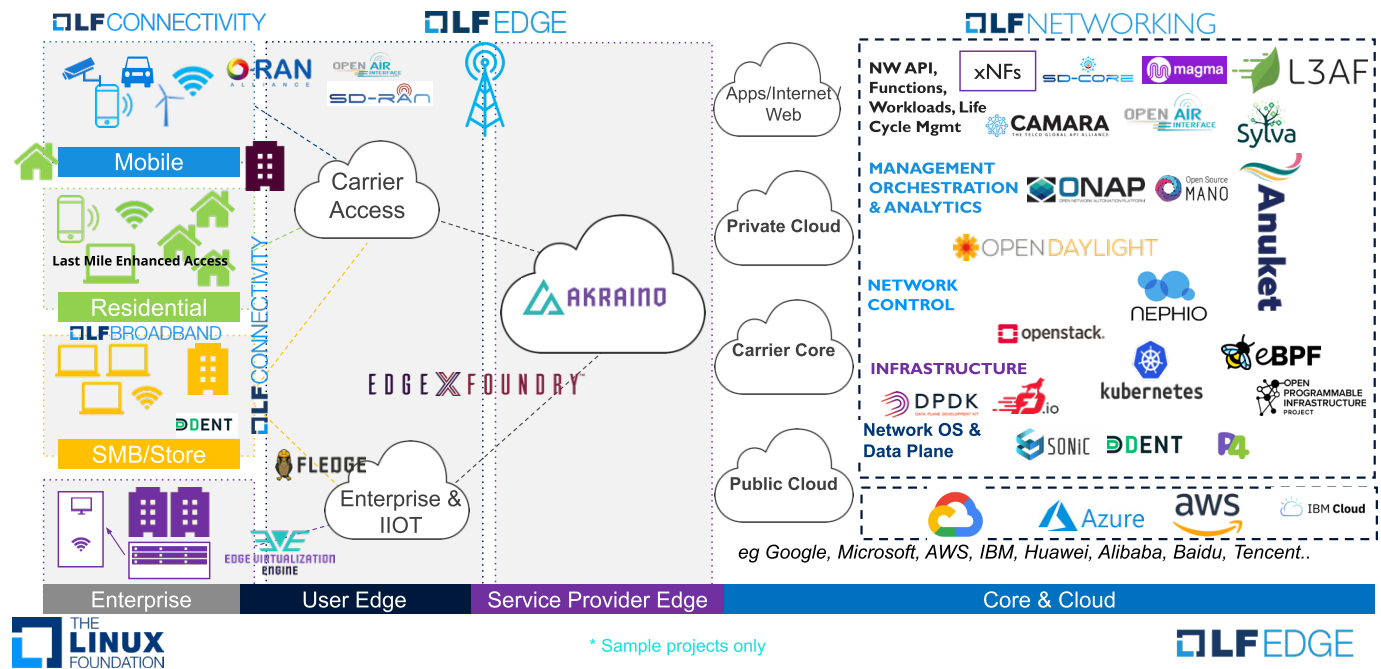
- ▶ The rise of domain-specific AI for networking
- ▶ Evolution of the Cloud Native Telco Initiative, which combined CNCF’s cloud native networking program with existing LFN work, with a certification program now in Beta.
- ▶ Nephio advanced cloud-native network automation in 2024 with expanded use cases.

Introduction

► Real-World Deployments were on full display at this year's **ONE Summit**, with diverse organizations like Walmart, TELUS, and the United States government showcasing practical deployments of open source solutions that address real-world challenges in interoperability and network modernization

We'll share more details throughout the report, and as we look towards 2025, we invite you to dive deeper into open-source networking by exploring our projects and becoming an active participant. We look forward to collaborating with you on advancing open innovation!

Open Source Networking, Edge and Access



The LFN Collaboration Hub Takes Flight

LF Networking has established itself as the central collaboration hub for the open networking industry by bringing together a diverse range of stakeholders to drive innovation and create open, interoperable solutions. With projects like ONAP (Open Networking Automation Project), which provides a comprehensive framework for automating network services, and L3AF, which provides complete lifecycle management of eBPF programs (enabling dynamic composition and orchestration of these programs across network infrastructures), LFN empowers industry leaders, developers, and enterprises to work together seamlessly. Through other initiatives such as Anuket, focused on integrating and testing cloud-native and telecom infrastructure, LFN accelerates the adoption of open standards and architectures. LFN's neutral environment for collaboration also helps to unify industry efforts toward next-generation networking, making it a pivotal organization for advancing 5G and Next-G, edge computing, and cloud-native network functions across the global Telecom, Cloud and Enterprise networking landscape.

LF Networking's longevity, reach, and diverse set of project communities and initiatives have created a hub of open industry collaboration that has taken off in spades over the past year. Read below for examples of how the community has pushed boundaries over the past year as it continues to move the industry towards greater digital transformation.

Taking Flight: Open AI, Open RAN and Beyond

LF Networking initiates key partnerships that promote progress across the open networking domain, with a focus this year on domain-specific artificial intelligence (AI). A recent example is this year's [memorandum of understanding with the Open Air Interface \(OAI\)](#), a significant step toward integrating AI into telecommunications. The partnership emphasizes AI's potential impact on open network automation and management, particularly in areas like Open RAN.

The LFN Collaboration Hub Takes Flight

In addition, the community's ties with other open source projects like Sylva, through initiatives like Anuket and CNTi, as well as CAMARA and its collaboration with GSMA's Open Gateway, showcases commitment to creating standards and frameworks that support open ecosystems for a more connected future. And this is just the beginning.

Network Automation & the Journey through the Cloud

Network automation has vastly transformed telco and cloud environments, setting the stage for additional advancements in enterprise automation as the third major market. Paraglider, a new flagship initiative in enterprise cloud automation and a new project under the umbrella in 2024, is a testament to the community's efforts in expanding cloud native automation beyond telecommunications to enterprises. A good example of enterprise adoption is Walmart and its use of L3AF within its technology stack of retail and e-commerce platform technology stack. L3AF helps manage needed flexibility and observability to manage Walmart's complex, large-scale networks; it enables the ability to perform seamless upgrades that enhance its network operations without disruption.

This year, the Cloud Native Telecom Initiative (CNTi), developed in partnership with the Cloud Native Computing Foundation's (CNCF) Telecom User Group (TUG), exemplifies the organization's holistic approach to cloud native solutions. This initiative, which has now become an official LFN project, promotes unified approaches to cloud native networking and automation. The group has launched a Certification program (currently in beta) to ensure that products and solutions meet the stringent requirements to operate effectively in a cloud-native environment.

Further pushing cloud native boundaries, the Nephio project showcases a commitment to building interoperable and automated cloud-native networking environments for diverse industry applications. This was demonstrated this year

The LFN Collaboration Hub Takes Flight

through the creation of the Open RAN orchestration use case jointly with the Open Air Interface (OAI) community. The two communities worked in harmony to build a cloud-native Radio Access Network that can be replicated and extended by any researcher or developer innovating in the 5G/6G domain.

The Rise of Domain-Specific AI & the Future of Connectivity

Following the initial excitement surrounding large language models (LLMs) and generic AI technologies, 2024 brought into the era of domain-specific AI, which is becoming a priority for operators and cloud service providers due to challenges in data sharing needed to train models. These models are poised to support top use cases that both drive new revenue streams and significantly reduce operational expenses.

Numerous opportunities have been identified and are being actively implemented by telcos, such as using generative AI to draft postmortems after network incidents in minutes instead of hours. Many companies have also developed advanced systems that consolidate their data into a unified platform, enabling integration, normalization, efficient querying, and seamless access to insights.

Some examples of how LFN is pursuing opportunities across the future of connectivity starts with Anuket's Thoth, a collaborative project for developing data sets and AI models tailored to networking. Efforts here align with our larger vision of leveraging AI to improve network performance, efficiency, and scalability. Check out a series of blog posts that go into this in more detail ([here](#) and [here](#)).

We're also actively engaged with enterprises and government entities, with a focus on the future of connectivity and innovation in telecommunications. A highpoint this year was the 6G Innovation Day workshop held as part of Open Networking and Edge Summit, which brought together 100+ leaders from across U.S government agencies and commercial partners working to advance Open RAN and AI. For a summary of those discussions and some key takeaways, check out the [Event Brief](#), which covers

The LFN Collaboration Hub Takes Flight

open 5G deployments, requirements for 6G design and implementations, public digital infrastructure, with real

world, practical examples of putting the regulatory and best practices examples into production.

Challenges & Opportunities

This past year, we've continued to address evolving challenges with the broad open source networking ecosystem as adoption accelerates. Some key areas include:

- ▶ **Production-Readiness and Customization:** Projects such as ONAP and Nephio are pivotal for network automation and cloud-native operations, serving as essential reference implementations. Contributing organizations get the opportunity to add the necessary functionality and create products that are production-grade.
- ▶ **Ensuring Interoperability:** Achieving seamless interoperability across diverse network environments requires us to ensure our projects are robust and adaptable to different telecom infrastructures, a key tenet of the LFN mission.
- ▶ **Balancing Innovation with Vendor Lock-In Risks:** While open innovation is essential, custom vendor distributions and reliance on a limited set of contributors can introduce risks of vendor lock-in, diminishing openness. Our community fosters an open, community-driven ecosystem that minimizes these risks and promotes true vendor-neutrality.
- ▶ **Vulnerability to Shifting Licensing Terms:** Sudden changes in licensing terms can disrupt operations, so CSPs and others need assurances of project continuity and freedom from restrictive licenses. LFN's lasting tenure of transparent governance, use of stable open-source licenses, Contributor License Agreements (CLAs), and adherence to a neutral, collaborative framework under the Linux Foundation help appease these concerns.
- ▶ **Security:** LFN projects have been in the forefront of adopting secure software development standards for the past several years by adopting many of the

The LFN Collaboration Hub Takes Flight

tools and best practices of the [OpenSSF](#) initiative. This year the trend continued with the establishment of the security and quality workstream within the Technical Advisory Council (TAC). This workstream is taking security and quality to the next level by setting quantifiable goals that all LFN projects will work to meet or exceed.

► **Next-G Transition:** As it becomes clear that the transition towards the next generation of wireless communications

will be a gradual one, LFN projects provide an excellent reference implementation for researchers and developers who are working on Next-G innovation.

Individual projects like Nephio and ONAP are simplifying deployment of wireless networks, and the 5G Super Blueprint is providing a reference architecture for integrating the necessary components of both the current generation of wireless as well as the next one.

What's Next?


As the LFN networking community looks ahead, several key developments are on the horizon. AI use cases in networking will mature, with a shift in focus from generalized artificial intelligence to more specific, ROI-driven applications such as employee assistants. Predictive AI will play a significant role in designing and maintaining networks by leveraging its capability to process streaming time series data. Data availability will become crucial for success, opening opportunities for the community to expand into data sharing and innovation platforms.

The definition of next-generation wireless will commence in earnest, utilizing open source software as a bedrock for innovation. This transition will be more gradual, with open source projects, particularly in the Radio Access domain, acting as pivotal enablers for both research and commercial products. Contributions to open source networking projects will continue, with a trend towards consolidation. Projects dealing with adjacent technologies will collaborate more closely, sometimes merging code bases to offer stronger, more comprehensive solutions. Companies will streamline their contributions, gravitating towards larger, well-established projects.

The LFN Collaboration Hub Takes Flight

AI workloads will drive the need for true multi-cloud connectivity and portability. Given the costs and resource constraints associated with AI training and inference, end users will seek optimal storage and execution environments for their data, guided by considerations of data privacy, energy consumption, and costs. This will create a device-edge-cloud continuum for running AI and data workloads, with edge and networking open source projects leading in providing necessary frameworks and connectivity solutions.

Network service providers will leverage open source technologies to develop new revenue streams, with easy-to-consume APIs accelerating the adoption of new services. Open source projects will facilitate the monetization of service provider assets, and edge infrastructure for AI, managed by open source technology, will emerge as a novel offering.



Message from the LFN Board



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

The major trend I see shaping the industry today is the rapid evolution of AI. This transformative technology offers fantastic opportunities but comes with incredibly fast lifecycles—something new emerges almost every week!

Key areas like generative AI, synthetic data, and foundation models are becoming more mature. At the same time, newer innovations like Agentic AI, decision intelligence, and AI simulation are continuously entering the scene.

AI in the Context of Networking

Traditionally, networks have long lifecycles and are built on established standards. The challenge now is how to effectively leverage these fast-evolving AI technologies within the realm of networking.

It's clear that waiting for formal, du jour standards is no longer feasible. Enter LFN, which has successfully created de-facto standards in the past and is well-positioned to do so again.

Just as with the transition to virtual network functions (VNF) and cloud-native management, the investment needed to apply AI to networks is too significant for any single company to undertake alone. This highlights the value of collaboration within the networking industry. By working together, we can benefit our companies, our community, and society as a whole.

What sets LFN apart is its commitment to open, global collaboration. With its deep networking expertise, hands-on approach, and agility in development, LFN occupies a unique and valuable position in the networking ecosystem.

As the year comes to a close, here are some aspirations for the future:

- That AI can take on tedious tasks, giving us more time for meaningful and creative pursuits in both work and personal life.
- For LFN projects to harness the power of AI, becoming even more productive and effective.
- That we overcome the challenges of data donation, enabling the community to advance AI for networks using real-world data.
- That LFN continues to deliver on its vision and grow as a hub for collaboration.

Finally, a heartfelt thank you to every contributor in the LFN community and the LFN staff for their dedication and hard work in 2024.

Here's to an exciting and impactful 2025!

LFN Board Members



Ankur Jain
Google



Christian Olrog
Ericsson



Chuanyu Chen
Huawei



Douglas Knisely
Qualcomm



Frank Brockners
Cisco



Frederic Desjarlais
Walmart Global
Technology



Sreekanth Sasidharan
Infosys



Hanen Garcia
Red Hat



Hari Vasudev
Walmart Global
Tech



Jérôme Birot
Telco Cloud



Jonne Soininen
Nokia



Junlan Feng
China Mobile



Kandan Kathirvel
Google Cloud



Manish Mangal
Tech Mahindra



Olaf Renner
Nokia



Praveen Atreya
Verizon Global Network
And Technology



Sebastian Scheele
Kubermatic



Sebastian Zechlin
Deutsche Telekom



Sun Qiong
China Telecom



Walter Kozlowski
Anuket Community



Ryuji Wakikawa
Keio University

LFN Members

Platinum Members



Gold



Silver



Associate



Message from the General Manager



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

As we wrap another year of collaboration, I'd like to take a moment to reflect on the tremendous progress we've seen from our community over the past year and also take a brief look at what's ahead.

Our goals this year for the Collaboration Hub included a focus on empowering our ecosystem to embrace digital transformation, specifically around domain-specific AI and cloud native automation. The community really stepped up in both domains, notably with work on Anuket's Thoth project and work to increasingly incorporate AI and ML to automate network operations, predictive analytics, and anomaly detection. The addition of the CNTi initiative into the LFN project fold, and the group's work on cloud native certification, best practices and a test catalog is another great example of this collaboration.

Other notable takeaways from 2024 include growth in strategic collaborations and innovative initiatives. A notable development was the inclusion of Qualcomm Technologies as a Platinum member in May, bringing expertise in intelligent communication and computing to the open source ecosystem. This partnership aims to empower the open ecosystem to scale digital transformation across Open RAN, AI, 5G/6G, and cloud-native networking. Additionally, we began a [formal partnership with Open Air Interface \(OAI\) to integrate](#) Network Automation for Cloud Native Applications in the RAN and Core Network.

We also hosted the Open Networking & Edge (ONE) Summit in April, returning to the event's Silicon Valley roots. The event featured a robust, cross-industry ecosystem gathering, focused on AI implementation across telecom, cloud, enterprise, and edge, as well as the future of open innovation and securing and automating open end-to-end networks. A key tenet of the event was the OpenGovCon: 6G Innovation day, which brought together 100+ leaders from across U.S government agencies and commercial partners, all working to advance Open RAN and AI for 6G readiness, in collaboration with open source communities. The sub-event produced a report on [The State of Open RAN Among U.S. Government Entities](#). Keynote speakers from industry leaders such as Google, Walmart,

Verizon, and ZEVEDA shared insights on the latest innovations through an open source lens.

Further strengthening collaborative efforts and strategic alliances, we welcomed SoftBank Corp. as a Platinum member in November. SoftBank's deep experience in network technologies and innovation is expected to advance open-source network innovation, particularly in areas like 5G, AI, edge computing, and advanced connectivity solutions. This strategic addition underscores LFN's commitment to fostering innovation and open-source collaboration across the telecommunications and technology sectors.

As we look to 2025 and beyond, we face both an increasingly complex geopolitical and economic landscape and an immense potential for new technologies to transform how we live our lives. What we know for certain is that open source demonstrates remarkable resiliency in uncertain times by fostering collaborative innovation, adaptability, and a shared commitment to solving global challenges through transparent and accessible technology.



Message from the CTO



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

There are two dominant trends driving innovation in the networking industry - the transition to Cloud Native technology and the aspiration for AI-Native networks. In 2024 those were the two overarching stories for most of the technical strategic work carried out by our open source communities.

On the Cloud Native front, we solidified the work of the Cloud Native Telecom initiative (CNTi) and made it an official project under LF Networking. With an elected community leadership, the project now focuses on providing best practices and test tools for ensuring modern network functions are compliant with the Cloud Native concepts. This is highly valuable for network operators that want to make sure the network functions supplied by their vendors are ready to leverage the Cloud infrastructure with the scalability and resilience built into it. It is also crucial for the CNF vendors and integrators to have common tools that can be used across the industry to ensure conformance to the Cloud Native best practices.

There was, and still is to some extent, a great degree of hype around AI, especially GenAI, for networking. The communities under LF Networking tried this year to figure out the tangible uses for AI in networking. Last year we spent our time understanding what use cases may benefit from applying AI and which layers of the technology stack need development. This year, through the work of the LFN AI Taskforce we identified such use cases and technologies and shared our findings in an updated LF Networking AI [whitepaper](#) earlier in 2024. For the rest of the year, we focused on identifying the main roadblocks that the open source community can help move out of the way and zeroed-in on data processing. We continued the work on the Thoth data anonymization project, so that more network related data can be shared while still protecting personally identifiable information (PII). Another topic we invested in this year was identifying efficient methods and platforms for sharing network data among AI researchers. These initiatives are expected to accelerate innovation around networking AI.

As 2025 is rapidly approaching, there is still a lot to be done in the domain of enabling efficient data sharing, but we believe that these efforts will bear fruit next year, allowing the industry to unleash the power of AI towards creating true autonomous and energy efficient networks. AI is not only affecting how we build and operate networks, but is also changing how workloads are using the network. We expect that in 2025 we will need to redesign some of our networks to best serve these unique types of workloads.

Message from the Technical Advisory Council



BY OLAF RENNER

*Open Source Program Office,
Nokia – Chair, LFN Technical
Advisory Council*

In 2024 the Technical Advisory Council improved its project onboarding pipeline and successfully approved the induction of new projects.

We introduced a project candidate stage that allows projects looking to integrate under the LFN umbrella to get early onboarding support to help graduation to a Sandbox stage. The first project that completed the move from a candidate to the new LFN sandbox project was Paraglider. Paraglider is a control plane for cloud networking resources that simplifies the network configuration across different cloud providers through a unified controller.

Another project that was successfully inducted as an LFN sandbox project was the Cloud Native Telecom Initiative (CNTi). The origins of CNTi have been CNCF CNF Certification and CNF working group as well as the CNF Testbed. These resources now moved under LFN and reformed as CNTi project to develop and maintain best practices, test catalogue and self-certification for cloud native networking functions. CNTi will work in close collaboration with other projects like Nephio (Kubernetes based infrastructure and CNF automation) and Anuket (infrastructure and conformance specifications).

Additionally the LFN 5G Super Blueprint which originated as a demonstrator for use and integration of open source projects in 5G networks extended the use case library and is now also governed as an LFN incubation project.

With new projects coming into LFN the TAC started to review common quality and security goals that should be applied to all LFN projects. The importance of supply chain security in open source and implementation of proper project security policies resulted in the documentation of best practices already in 2023.

In 2024 the work started to define specific security goals and measurements to verify if goals are reached and additional project quality goals have been added. As the list became quite extensive, goals have been prioritized to start implementation in 2025.

With the established LFN AI task force the work on network domain specific AI did advance in LFN projects with Intent Driven Orchestration for Autonomous Networks, Telco Data Anonymization and Dataset sharing platform use cases. To further support the AI work and advise the TAC, a dedicated AI subject matter expert position has been opened.

LFX Insights

Open source, like any other industry vertical, tracks its performance and growth through the use of a set of objective performance indicators. The Linux Foundation Networking 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.

Stable Communities

Over time, the LFN project community has ebbed and flowed as is the natural course of development. Some of the more long-standing projects (such as OpenDaylight and FD.io) have held steady with contributions over the past few years (with peaks at milestone releases as is the case with FD.io):

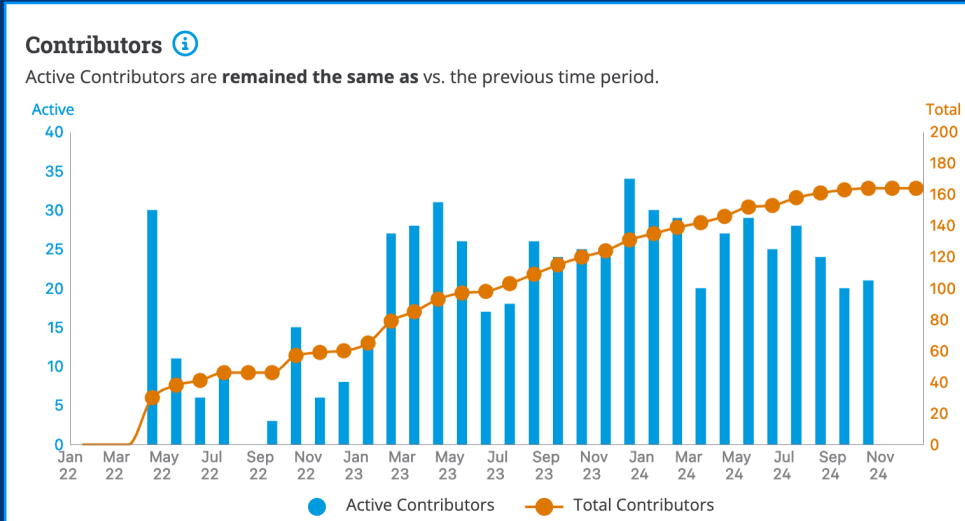
Currently, LFN hosts:

Has over **8200** contributors | Across **327** organizations

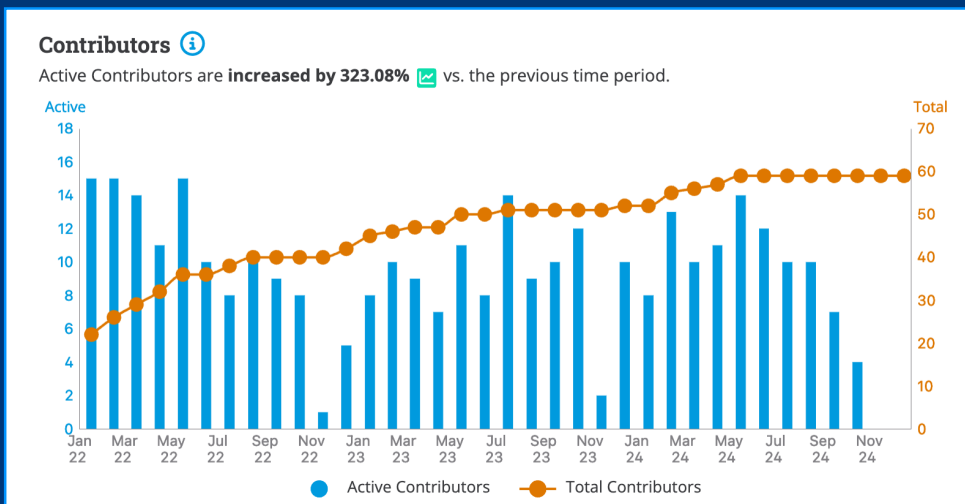


FD.io contributions, Jan. 2022 - Dec. 2024

Meanwhile, newer projects (such as Nephio and L3AF) have more recent peaks in activity, particularly centered around release launches:



Nephio contributions, Jan. 2022 - Dec. 2024



L3AF contributors, Jan. 2022 - Dec. 2024

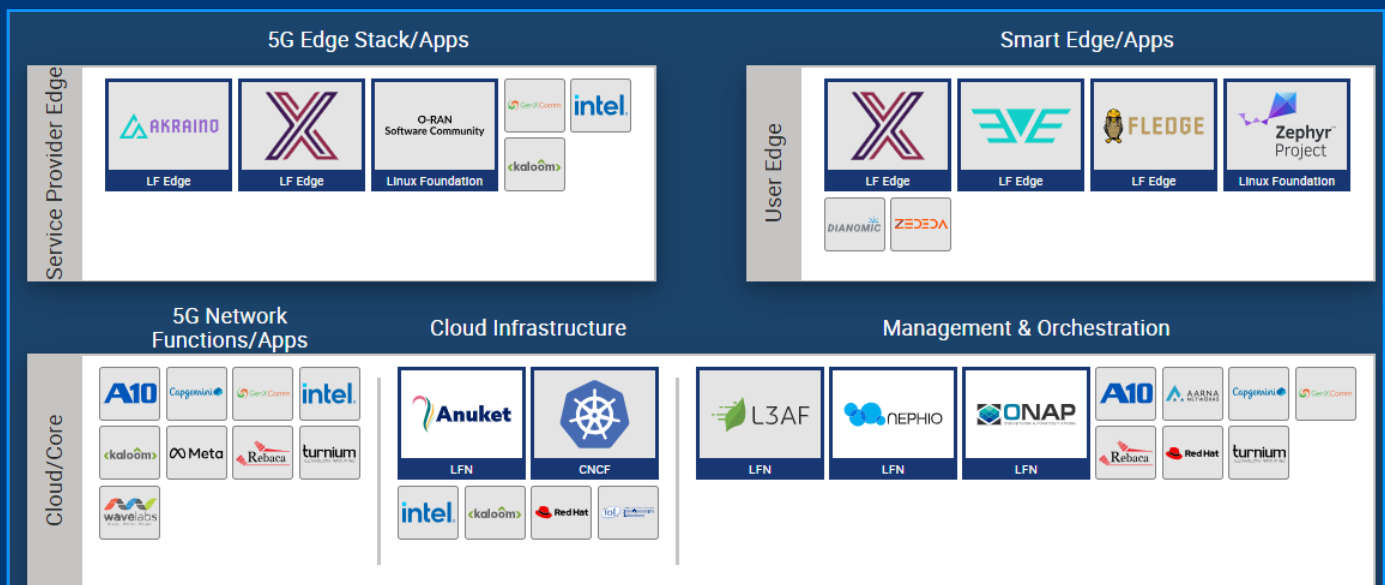
Check back in with [LFX Insights](#) often, as improvements and new functionalities are underway.

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.

5G Super Blueprints Landscape



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, the LFN telecom cloud infrastructure architecture modeling and testing project continues its mission to meet operator requirements for normalization, software and test development environments. The project completed its Pieman release in the first half of 2024 and plans to release its Quinnipiac release in December. Due to continuing significant resource constraints over the past years, only two specification sub-projects (RM and RA2) and two implementation and test projects are active (Thoth and Functest). In the last months of 2024 the project has issues even to reach quorum in its TSC meetings. Anuket projects and work streams focused on strengthening container-based open infrastructure specifications and implementations.

Some of other active projects include:

- **Thoth** – The team is working on making data sets available for AI model and application development, by focusing on novel data anonymization techniques
- **Functest, functest-kubernetes** — A suite of functional tests for OpenStack and Kubernetes deployments, now includes support for CNTi Testsuite.

The project continued its vital work and launched the **Pieman Release on May 24, 2024**. Anuket is being adopted as the reference infrastructure for other open source projects such as Sylva as well as part of the GSMA permanent document Reference Model as NG.126,

Reference Architecture 2 (Kubernetes based) as NG.139, and Reference Architecture 1 (OpenStack based) as NG.133.

Reference Model (RM) highlights include:

- Updated contents in Hardware Infrastructure Manager section, aligning and referencing the Redfish standard. Prior to Pieman, Hardware Infrastructure Manager existed in RM as a concept only. Now, it can be implemented, e.g., using ODIM as a basis.
- Added load balancer into RM technology agnostic definition, aligned to RA2 PR #3415. It is based on a functional description, not on a specific implementation as typical load balancer technologies are defined.
- Started a technology agnostic guidance section to assist in planning more efficient open source based cloud infrastructure energy consumption. The RM focus is on the energy efficiency of the workload/ infrastructure/management interactions. Energy savings are becoming a major issue for the telecom industry, with the massive use of data, energy efficiency and sustainability in the Hybrid, Multi-Cloud environment needs a consolidated approach by the ecosystem players.
- Added section on automated TLS certificate lifecycle management for workloads (from the infrastructure perspective).

Reference Architecture 2 (RA2) The RA2 project focuses on building a Kubernetes based architecture that has now been consolidated into the most comprehensive industry wide set of specifications for Kubernetes Telco Cloud – acceptance as a GSMA standard is confirming this status.

- Added Express Data Path (AF_XDP) network acceleration specs, enabling faster and more efficient data plane traffic processing. This provides an agreed network acceleration implementation, which benefits all users with containerised data plane functions
- Update to Kubernetes 1.29, including API specs and new functionality. Since PaaS services are essential operational add-ons to Kubernetes, which implement containerised network function monitoring, logging, traffic steering, etc.
- Add cert-manager for TLS certificate lifecycle management. A number of network functions require TLS certificates, this plugin allows management of their lifecycle automatically
- Add Load Balancer specs, for ingress traffic distribution across microservices. Load balancers

steer traffic across the pods that compose containerised network functions. RA2 is specifying what is required of a load balancer to integrate with a Kubernetes based Telco cloud.

What Worked

Anuket projects and work streams continued their efforts to strengthen container-based open infrastructure specifications and implementations. We were able to take advantage of the LFN Intern program and had a successful collaboration with an intern from India who was able to help the project with several document migration and clean up efforts.

Next Steps and Decisions

- Need to drive more support from within LFN and larger Telecom community
- Need to continue the review of the health of sub-projects

BY BETH COHEN, TSC Co-chair, and Software Defined Networking Product Technologist at Verizon and **GERGELY CSATARI**, TSC-Co-chair, and Senior Specialist, Open Source at Nokia





The mission of the Cloud Native Telecom Initiative (CNTi) is to accelerate the adoption of cloud-native technologies and practices in telecom networking. Achieving widespread adoption requires overcoming various ecosystem challenges and roadblocks, such as those highlighted by Communications Service Providers (CSPs) in the NGMN Cloud Native Manifesto. To address these, the CNTi collaborates closely with other open-source networking projects to establish best practices for cloud-native networking, improve an open-source Cloud Native Network Function (CNF) testing framework, and develop a vendor-neutral CNF conformance program. Through these efforts, the CNTi aims to empower operators to confidently adopt cloud-native patterns, unlocking the benefits of performance, scalability, and security while enhancing predictability across diverse cloud-native infrastructures.

Summary

The Cloud Native Telco Initiative (CNTi), under the umbrella of LF Networking, achieved remarkable progress in 2024, focusing on accelerating the adoption of cloud-native workload best practices within the telecommunications industry.

Key highlights include:

- Onboarded as an official LF Networking Sandbox project
- Successfully transferred agreed CNCF assets to CNTi providing a solid foundation for best practices, cloud native workload testing, and a certification program

- Achieved monthly releases of the CNTi test suite bringing a wide range of new features and improvements including:
 - ✓ Documentation simplifications and enhancements
 - ✓ Test case run time optimizations
 - ✓ Enhancement of console output to be more informative during execution
 - ✓ Addressed multiple test issues and provided fixes
 - ✓ Conducted thorough internal code refactoring
 - ✓ Improved the stability of our CI/CD pipeline
 - ✓ Hardening of code and verification workflows
 - ✓ Added new functionality in support of collaboration efforts with the project Sylva Validation program
- CNTi has launched their first beta Cloud-Native Network Functions (CNF) Certification Program, to ensure greater alignment and standardization across telecom applications.
- CNTi participated in multiple open-source programs speaking at ONE Summit NA and participating in the Open Source Summit in Vienna

In 2024, the Cloud Native Telco Initiative solidified its role as cornerstone for modernizing telecom networks. Its advancements in cloud-native technologies, standardization efforts, and ecosystem collaboration continue to drive innovation, enabling telcos to deliver more agile, cost effective, and high-performance networks.

LFN Project Updates



The **FD.io** release train maintained its momentum in 2024, delivering another year of outstanding performance with timely VPP releases and comprehensive CSIT release reports. The FD.io Technical Steering Committee (TSC) played a pivotal role in driving the community's success by ensuring balanced representation through nominated membership, implementing cost-saving measures, and facilitating the acquisition and integration of cutting-edge hardware into the FD.io performance and CI lab (which will continue into 2025). FD.io's core projects, VPP and CSIT, continued to set new benchmarks in data plane performance and testing excellence.

Some key highlights from 2024 include:

- **Regular Release Cadence:** FD.io maintained its commitment to timely updates, delivering multiple releases throughout the year. Notably, the Vector Packet Processing (VPP)

Project releases:

- VPP 24.02 release on February 28 introduced 262 commits since the previous release, including 22 new features and 123 fixes.
- VPP 24.06 release on June 28 introduced 284 commits since the previous release, including 14 new features and 124 fixes.
- VPP 24.10 release on October 30 introduced 241 commits since the previous release, including 19 new features and 86 fixes.

- **Enhanced Testing and Performance:** The Continuous System Integration and Testing (CSIT) team executed comprehensive testing plans to ensure robust performance and functionality, including periodic trending reports and CSIT Release reports which include comprehensive test coverage as well as performance comparison with the previous VPP release:

- CSIT 24.02 Release report available on March 13.
- CSIT 24.06 Release report available on July 10.
- CSIT 24.10 Release report available on November 13.

- **Community Engagement and Outreach:** FD.io actively participated in industry events, including presentations at LFN Developer and Testing Forum, KubeCon, and EnvoyCon, to showcase its advancements and foster collaboration with other open-source communities.

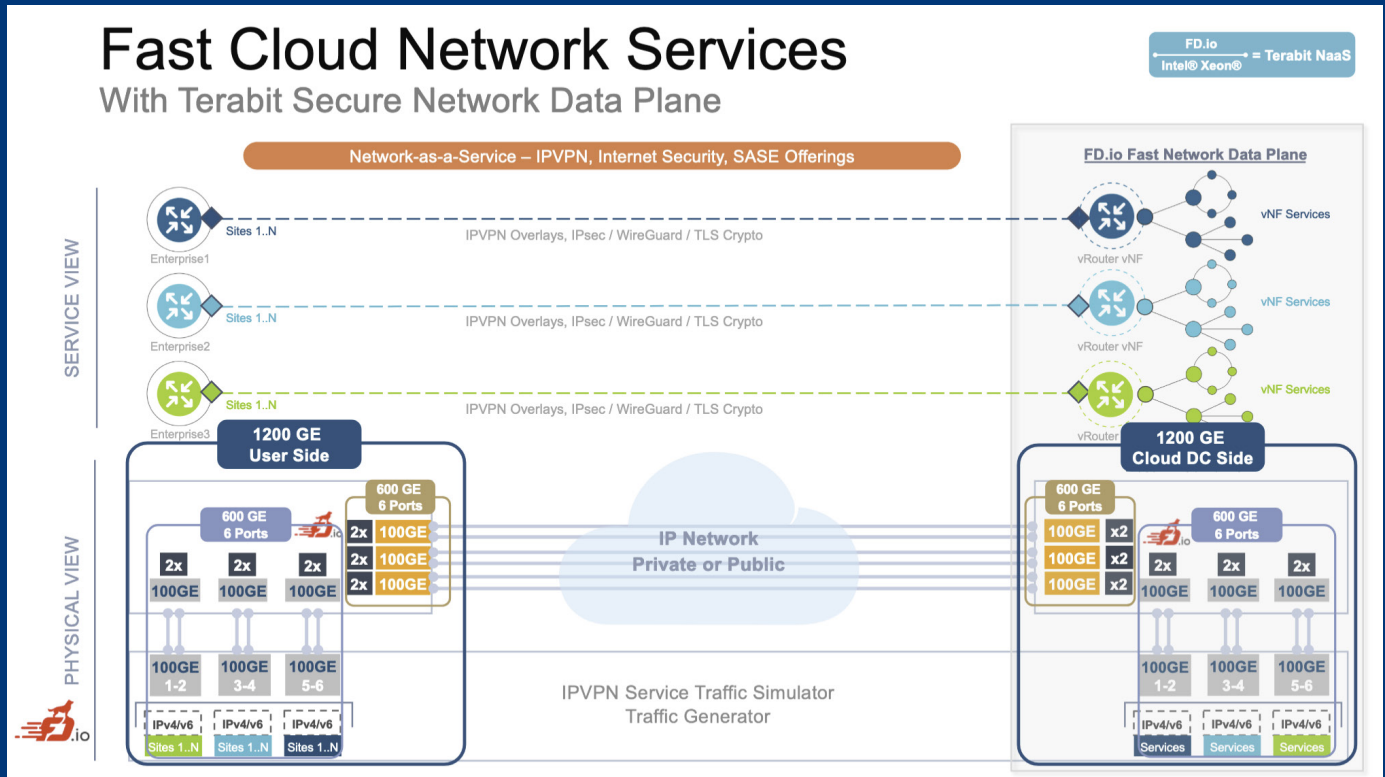
- **Strategic Goals and Initiatives:** The Technical Steering Committee outlined objectives for 2025, focusing on expanding testing coverage to the latest generation chips, improving outbound marketing, and enhancing FD.io's visibility.

These efforts underscore FD.io's dedication to advancing high-performance, flexible, and scalable software-defined infrastructures.

I am grateful for all the contributions to the FD.io Community and would like to thank each and every one who contributed to FD.io not just in 2024, but all along the way. Special thanks to Tom Jones, FreeBSD committer and developer specializing in network performance, for his

excellence in porting VPP to FreeBSD. I am confident that FD.io will continue to attract and retain the contributions of those with a passion for excellence and continue to keep things moving throughout the coming years, starting with 2025.

BY DAVE WALLACE | *Chair, FD.io Technical Steering Committee*



LFN Project Updates



In 2024, **L3AF** made significant strides in enhancing eBPF (extended Berkeley Packet Filter) program management and networking capabilities, in addition to being promoted to an “Incubation” project within LF Networking, representing a growth in maturity.

L3AF in Action

The L3AF platform provides full lifecycle management for running ebpf programs at several hook points to defend against DDoS attacks, it offers deep visibility into network infrastructure, which are usually hidden outside of the Linux kernel, and enables complex functions to be performed directly in the traffic flow within the technology stack of Walmart’s retail and e-commerce platform.

A notable milestone was the release of L3AF R2 in January, which introduced several key improvements, followed by the availability of L3AF 2.1 in the fourth quarter.

Collective updates over both releases include:

- **Native eBPF Program Loading:** L3AF now supports native loading and management of eBPF programs, eliminating the previous necessity for user-space programs to handle initial loading. This streamlines the deployment process and enhances efficiency.
- **Enhanced Chaining and Monitoring:** The update improved eBPF program chaining across multiple network interfaces and introduced monitoring capabilities for eBPF program versions, providing greater flexibility and control over network configurations.

- **Open-Source Traffic Mirroring:** An open-source traffic mirroring program was added, offering network administrators a powerful tool for monitoring and analyzing network traffic.
- **L3AF 2.1 introduces Graceful Restart functionality,** allowing seamless upgrades of the L3AF control plane without impacting any running eBPF programs in the data plane.
- **Container Support for l3afd:** L3AF can now run within a container, equipped to operate in cloud native environments, improving orchestration and scalability for users leveraging Kubernetes and other cloud-native platforms.
- **BPF CO-RE in the eBPF Package Repository:** L3AF 2.1 now supports BPF CO-RE, enabling portable BPF applications that run across different Linux kernels without modifications.
- **Support for KProbes and Tracepoints:** Enhanced observability with support for kprobes and tracepoints, providing deeper kernel-level insights for better eBPF troubleshooting.
- **Dynamically add programs to new interfaces:** L3AF 2.1 allows dynamic program attachment to new interfaces, benefiting complex network environments like multi-VM hypervisors.
- **Alternative traffic management options:** L3AF now supports attaching eBPF programs to HTB qdisc hooks for refined traffic shaping and resource allocation.

- **Enhanced Logging and Storage for Easier Debugging:** Improved logging with local filesystem storage simplifies debugging, accelerates issue resolution, and supports integration with centralized log analysis systems for better observability.
- **L3AF Goes to Dockerhub:** L3AF is now available on Docker Hub, making deployment easier within

containerized environments for streamlined DevOps integration.

These advancements have solidified L3AF's role in simplifying and enhancing networking performance and security. The project's commitment to open-source collaboration continues to drive innovation in the networking landscape.

L3AF/eBPF programs at Walmart

Observability

- Kernel level metrics visibility
- Improved MTTD & MTTR

Network Functions

- Prioritization for Payment in Stores
- Traffic Mirroring & Load Balancing



Security

- Traffic Control for DOS/DDOS mitigation
- Traffic Flow Exporter

Business Impact

- Accelerated site speed
- Eliminates traffic bottlenecks
- Value addition & ROI



In 2024, **Nephio** continued to expand its role as a leading open source initiative for multi-cloud and network automation, achieving milestones across technology, community growth, and adoption. Launched in 2022, Nephio now boasts contributions from over 70 organizations with a social media reach of over 80K and 2,000+ event attendees at eight in-person events.

In 2024, Nephio advanced network automation with new capabilities, expanded use cases, and broader industry adoption:

- **Use Cases & Verticals:** Nephio enabled flexible RAN, Core, and Transport configurations, creating a unified, vendor-neutral network management approach applicable across industries.
- **Evolving Technologies:** Enhanced Kubernetes-native Resource Management (KRM), SDK extensibility, and AI-driven automation boosted operational efficiency and adaptability.
- **Synergies with Other Projects:** Collaborations with O-RAN, ETSI, TMForum, Project Sylva, CAMARA, and LFN AI & Data fostered a robust ecosystem supporting open source network automation.
- **Community Growth & Adoption:** Positive market feedback and lighthouse deployments accelerated Nephio's adoption, establishing it as a leader in cloud-native automation.
- **Expanding Beyond Telecom:** Nephio community is working on addressing automation needs across retail, enterprise, finance, government, and other sectors, enabling large-scale edge deployments and operational improvements.

Roadmap

- Looking ahead, Nephio R4, R5 will focus on production readiness with enhanced multi-cloud support, strengthened security, lifecycle management for network functions, O-RAN integrations, support for various GitOps technologies, NB API integration, user experience upgrades.
- With the progress of GenAI efforts within the Nephio there is a growing interest in the community to solve many tedious, manual automation templates generation and configuration management through GenAI. The community will publish white papers soon. There are also evolving opportunities to create synergies with LFN broader initiatives to support GenAI vision for nephio solving data and analytics problems for improving customer experience across both telcos and enterprises.

Nephio's 2024 accomplishments mark major progress in network infrastructure automation, addressing both telecom and cross-industry needs. With R4 on the horizon and advancements in O-RAN, Kubernetes-based automation, and GenAI integration, Nephio is positioned to lead cloud-native automation for years to come. Its community will continue driving adoption and innovation, solidifying Nephio as a key platform for automation across industries.





OpenDaylight celebrated its eleventh anniversary as a project this year, and has had a fascinating journey as one of the most long lived open source SDNs.

Highlights this year include:

- › ODL issued its 21st release, "2024.09 Scandium," which includes:
 - Java 21
 - Lightweight Netty-based RESTCONF Northbound
 - OpenAPI is now highly scalable
 - Preparatory work for better K8s integration
- › Community is focused on 22nd release, "2025.03 Titanium"

- Migration from Akka to Apache Pekko ready
- RFC8639 Subscribed Notifications being worked on
- Datastore persistence & scalability improvements planned

Looking ahead, the community is focused on bringing:

- BGP/LS container images still planned
- Developments continue on optical network control to extend support for most widely adopted standards
- Additional support from downstream consumers continues to be of critical importance



LFN Project Updates



The **ONAP**, as a collection of Network Automation functions, offers advanced and individual network automation capabilities. It focuses on component selection, AI-assisted network orchestration, natural language processing, communication and infrastructure security, modernization, and adherence to industry standards.

ONAP is no longer a platform; instead, it provides various individual network automation functions and security reference configurations by leveraging Service Mesh, Ingress and Gateway, OIDC IAM. It supports the development of individual ONAP functions and component deployment through CI/CD under the ONAP Streamlining evolution.

ONAP will evolve to become more intent-based and declarative, integrate greater use of Generative AI, and support a secure software supply chain.

Key Developments in 2024:

- **Security Enhancements:** ONAP projects have addressed critical security concerns by converting ports to HTTPS, removing hard-coded passwords, enabling Kubernetes pods to operate with non-root privileges, and mitigating Common Vulnerabilities and Exposures (CVEs). These measures have significantly bolstered the platform's security. Additionally, by leveraging industry-standard/de facto security security protocol and mechanisms such as Istio Service Mesh and Ingress Gateway, ONAP ensures secure inter- and intra-component communications.
- **Platform Modernization:** Components such as the Common Controller Software Development Kit

(CCSDK), Configuration Persistence Service (CPS), Usecase User Interface (UII), Portal-NG and Policy Framework were upgraded to Java 17. Additionally, various software versions updates ensure that ONAP leverages the latest software development frameworks.

- **ONAP Streamlining Evolution:** This initiative makes ONAP components modular and independent through interface abstraction, loose coupling and CI/CD. As a result, ONAP has evolved into a collection of individual network orchestration functions, allowing the industry to pick and choose specific components and enabling flexible and dynamic function adoption.
- **Intent-based Declarative and GenAI solutions:** Supports generative AI solutions powered by large language models (LLMs), and includes data service enhancements (domain-specific datasets) of Intent-driven networks.
- **Industry Standard-Based Network Interface Upgrade:** CCSDK/SDNC now supports an RFC8040-compliant network interface.
- **OpenSSF Gold Standard Achievement:** The CPS and Policy Framework projects have achieved the Open Source Security Foundation (OpenSSF) Gold Badging standard, demonstrating ONAP's commitment to high-quality, secure, and reliable open-source software development.

Recent Releases:

Montreal Release: The Montreal release introduced several enhancements:

- **Streamlining-Based Release Plan:** The ONAP release process has been simplified and made more agile by including only Start, RC, and Sign Off stages, without fixed milestones. PTLs and communities can incorporate

LFN Project Updates

multiple agile deployment cycles as needed within the release.

- **Modeling Improvements:** Automated tools for YANG modules have been introduced to enhance development efficiency and improve version control.
- **Operations Manager (OOM) Updates:** Added support for Gateway-API in Ingress templates and integrated with the MariaDB operator, ensuring ONAP MariaDB instances are created with the latest version by default.
- **Service Design and Creation (SDC) Enhancements:** Added capabilities for service updates directly in YAML, comprehensive support for defining interface operations, and adaptability in defining CSAR structure to meet individual model requirements.
- **For other key updates, see <https://lf-onap.atlassian.net/wiki/spaces/DW/pages/16518372/Montreal+Release+Key+Updates>**

NewDelhi Release: The most recent ONAP release, NewDelhi, introduced several enhancements:

- **Key Updates:** see <https://wiki.onap.org/display/DW/New+Delhi+Release+Key+Updates>
- **Leveraging Industry De Facto Technologies:** Replaced ONAP-specific component communication mechanisms, such as MSB and DMaaP with Strimzi/Kafka.
- **Security enhancements:** Improved security using Istio Ingress, Istio Service Mesh and Keycloak-based authentication and authorization. Added support for Gateway-API in Ingress templates and integrated with the MariaDB operator, ensuring ONAP MariaDB instances are created with the latest version by default.
- **GenAI Solutions:** Supports generative AI (GenAI) solutions powered by LLMs, and Data Service (domain-specific tuning) enhancements for intent-driven networks.
- **Infrastructure enhancements:** Updated Kubernetes, Istio, Keycloak, Strimzi-Kafka, K8ssandra-operator and Mariadb-operator.

Oslo Release: The Oslo release introduced several enhancements:

- **Key Updates:** see <https://lf-onap.atlassian.net/wiki/spaces/DW/pages/16552136/Oslo+Release+Key+Updates>
- **OpenSSF Gold badging status:** Both CPS and Policy Framework have achieved the highest OpenSSF Gold badging status!
- **Lightweight ONAP:** Developed use cases and PoC for compositing selected individual ONAP components to create higher-level solutions, such as intent-driven network automation.
- **ONAP Streamlining Evolution:** Continued the component individual build and deployment of components using continuous deployment. For example, Deutsche Telekom (DT) completed an Argo-CD based ONAP component deployments as a proof of concept, with this feature set to be officially available in the Paris release.
- **Leverage Industry De Facto Technologies:** Continued replacing ONAP-specific infrastructure and communication mechanisms with industry-standard / de facto solutions, such as Strimzi/Kafka, Realm client authorization, Oauth2-proxy.
- **Infrastructure enhancements:** Continued updating Kubernetes, Istio, Keycloak, Strimzi-Kafka, K8ssandra-operator and Mariadb-operator
- **Production-quality ONAP Deployment Readiness:** Enhanced OOM scripts and other components to achieve production-quality deployment readiness, as demonstrated in Deutsche Telekom's TNAP environment.

Looking Ahead to 2025:

ONAP plans to continue its trajectory of innovation and improvement in 2025, focusing on:

- **Cloud-Native Transformation:** Building upon previous releases, ONAP aims to further enhance its cloud-native capabilities, facilitating seamless integration and deployment within modern cloud environments.

LFN Project Updates

- **Modularity and Flexibility:** The project intends to increase the modularity of its platform, allowing users to select and deploy specific components tailored to their unique use cases, thereby improving flexibility and scalability.
- **Standardization and Compliance:** ONAP is committed to aligning more closely with industry standards, including those set by 3GPP and O-RAN, to ensure interoperability and compliance across diverse network environments.

To achieve these capabilities and ensure compliance, ONAP TSC has set the following goals for 2025:

- Ensure ONAP core components are focused and operate independently, from build to runtime

- Make ONAP core components more autonomous and ready for use by both ONAP, LF and other external users
- Incorporate more GenAI capabilities and use cases into the ONAP components, and promote the adoption of open-source LLM models and frameworks aligned with LF AI & Data and GenAI Commons
- Foster inter-community collaboration with other LF communities, such as O-RAN and Nephio
- Ensure the security of ONAP components and operations
- Define a secure LFN CI/CD pipeline by leveraging OpenSSF-associated reference tools

These strategic initiatives position ONAP to remain at the forefront of network automation solutions, delivering enhanced value to its users and stakeholders in the evolving technological landscape.

BY PAWEL PAWLAK | *Chair, ONAP TSC & SECCOM, and Product Manager, Incognito Software Systems*



LFN Project Updates



In April 2024, **Paraglider**, an open-source initiative aimed at simplifying network setup and management within and across multiple cloud environments, joined the LFN project roster as the first “Candidate” project. Collaborating with industry leaders such as Aarna Networks, Broadcom, Google Cloud, IBM, Intel, Microsoft, Uber, and UC Berkeley, Paraglider offers a unified, high-level API that aligns with users’ connectivity intentions, reducing the need to navigate complex low-level networking components. This approach enables users to focus on their desired network outcomes rather than the intricacies of network configuration.

By abstracting the complexities of cloud networking, Paraglider facilitates seamless operations across

various cloud providers, allowing users to select the most suitable services for their workloads without the burden of managing intricate inter-cloud connectivity. The project’s open-source nature encourages community participation, inviting developers and organizations to contribute to its development, explore its capabilities, and extend its functionality to additional cloud and on-premise infrastructures.

During 2024, the project graduated from a “Candidate” to “Sandbox” maturity level, reflecting the evolution of the community and establishments of roles and processes for building and releasing the code. At the close of the year, the project is getting ready to release a production-ready solution for cloud tenant networking.



Linux Foundation Launches **Paraglider** to Reduce Complexity for Developers and Network Operators Within and Across Clouds



Industry Impact

By the Numbers



9

PRESS RELEASES



800

PRESS CLIPS/
MENTIONS



14

MEDIA/ANALYST
BRIEFINGS



2.3B

POTENTIAL
AGGREGATE REACH



350K

LINKEDIN
IMPRESSIONS

Industry Event Participation:

- FutureNetWorld (Paris)
- IOWN
- Open Networking & Edge Summit
- GTI Summit
- Open Telco LATAM Summit
- Telco Cloud & Edge Day
- IEEE Open RAN Summit
- OAI 10th Anniversary
- Open Source Summit EU
- Telco Cloud Native Summit (TelecomTV)
- OCP Global Summit
- ORAN Alliance F2F meeting (Montreal)



TOP QUOTES

“Walmart is... heavily involved in L3AF, an open source orchestration platform and lifecycle manager for eBPF programs. At Walmart, L3AF and eBPF are implemented for a range of network functions, including observability, network functions (such as prioritizing payment in stores) and security...But even an organization with the scale of a Walmart — or Amazon, Google or another hyperscaler — can’t develop all the technology it needs internally... Open source makes it possible for companies to collaborate on their IT needs.”

— MITCH WAGNER, FIERCE NETWORK




“Joshiyura pointed to a need for the development of more domain-specific artificial intelligence (AI) architectures. This means the need to scale current large language model (LLM) architectures to work in more diverse telecom environments.”

— DAN MEYER, SDXCENTRAL

“[NextG networks will] be open and integrated across domains, connecting space-based, air, ground and undersea networks, and connecting networks of sensors, compute and communications.”

— MITCH WAGNER, FIERCE NETWORKS



Top Headlines



Walmart scales stupendous global network on open source [Read Now ▶](#)



LF Networking's Thoth enables telcos to leverage domain-specific AI [Read Now ▶](#)



Where does the Linux Foundation see telecom open network gaps? [Read Now ▶](#)



Domain-Specific AI for Networking [Read Now ▶](#)



Open Source in Telecom Networking [Read Now ▶](#)



The US is losing the 5G race; but NextG is just getting started [Read Now ▶](#)



SONiC, Open RAN, Gen AI in Networking: Insights from LF ONE Summit 2024 [Read Now ▶](#)



Nephio R2 pushes open-source cloud-native network automation [Read Now ▶](#)

Social Media

2024 brought an increased focus on LinkedIn with growing audience and engagement rate.

350K

LinkedIn impressions

23%

Growth in LinkedIn Engagement Rate

63%

growth in LinkedIn Followers

54K

YouTube views

251

Videos uploaded

31%

Growth in channel subscribers

Most viewed content focuses on:

- OpenGovCon: 6G Innovation Day
- Intelligent Networking, AI and Machine Learning for Telecommunication Operators
- Nephio R2 & R3 Releases
- Nephio & GenAI
- 5G Super Blueprints
- Sister project releases and project updates including CAMARA, SONiC and more.

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.

Join LF Networking to drive innovation, reduce costs, and drive new revenue opportunities by actively supporting our projects and contributing to the future of open source networking. Learn more about [LF Networking membership](#).





OLF NETWORKING

Thank you all for a successful 2024
and here's to embracing the limitless
skies in 2025!