



OpenSource

Volume: 14 | Issue: 02 | Pages: 100 | December 2025

THE COMPLETE MAGAZINE ON OPEN SOURCE

ForYou

An **EFY** GROUP Publication

Grafana: The Comprehensive Monitoring Platform

Cloud Cost Governance Using Kubecost, OpenCost And Infracost

Job Opportunities: What's Hot In The Cloud Space?

Optimising Cloud Services



Setting Up Snort To Secure Your Network

Open Source India 2025: A Pictorial Representation



Linkerd adds native support for the Model Context Protocol

Buoyant, the creator of Linkerd, the open source service mesh for cloud-native and Kubernetes environments, has announced upcoming support for the Model Context Protocol (MCP). This move positions Linkerd as the first open source service mesh to natively support MCP, offering a unified and secure foundation for both traditional microservices traffic and agentic AI-driven communication.

By adding MCP support, Linkerd will extend its lightweight security, observability and traffic management capabilities to agentic AI workflows. This includes metrics on resource and tool usage, latency and failure rates, along with fine-grained, identity-based access controls built on Linkerd's existing zero-trust framework. The aim is to provide a single mesh layer across all workloads. Enterprises are eager to embrace AI, but they cannot do so at the expense of their security posture and application reliability," said William Morgan, CEO of Buoyant. "Linkerd solves this problem by extending its proven capabilities to MCP traffic. We are giving organisations the tools to accelerate their usage with confidence."

Full MCP support will be available in 2026 across both the open source Linkerd distribution and Buoyant's enterprise offering.



linkerd

The move will make the company's service ecosystem publicly accessible to local garages, fleet operators, and independent mechanics, allowing them to repair and maintain Ola's electric two-wheelers without relying solely on its in-house network.



Under Phase 1, Ola Electric will sell genuine spare parts directly through its app and website. In Phase 2, the company plans to launch online diagnostic tools and technician certification programmes, empowering third-party repairers to service Ola vehicles independently.

"We have built our service ecosystem from first principles, using technology to make it fast, transparent, and efficient. With Hyperservice scale-up, we are opening up this capability to everyone," said Bhavish Aggarwal, chairman and CEO at Ola Electric.

"Every garage, fleet and customer can now access the same high-quality parts and systems that power Ola's own network. Genuine parts, full transparency, no middlemen. This is true service freedom, only possible in a D2C model."

Microsoft open sources Azure MCP Server v1.0

Microsoft has released the stable version 1.0.0 of Azure MCP Server, implementing the Model Context Protocol (MCP) to serve as an interface between AI agents and



Azure services. The move enables developers to query, manage, and automate Azure cloud resources using natural language or code. Significantly, Microsoft has open sourced the entire project, making both documentation and source code available on GitHub.

The launch underscores Microsoft's growing commitment to open source collaboration. The Azure MCP Server allows developers to customise, extend, and integrate MCP not only with Azure-based AI frameworks but also with third-party systems. The underlying Model Context Protocol, an open standard, unifies communication between large language models (LLMs) and backend systems, regardless of the agent frameworks used. This open source release invites the wider AI and DevOps community to build, test, and contribute integrations across multiple cloud ecosystems.

The server currently supports over 47 Azure services, including Azure AI Foundry, AI Search, Event Hubs, Service Bus, PostgreSQL, Kusto, Function Apps,

Storage, and Log Analytics. It features 170 structured command functions and offers three operational modes: namespace, full activation, or selective functions, designed to simplify onboarding and testing.

Developers can access a Docker image through Microsoft's container registry for CI/CD pipeline integration, alongside extensions for Visual Studio Code, Visual Studio, and IntelliJ. Security-critical operations rely on mandatory user confirmations, while .NET Ahead-of-Time (AOT) compilation enhances performance efficiency.

According to Microsoft, future updates will include closer integration with Azure tools and extended support for container workloads, thereby strengthening open source collaboration in AI-cloud interoperability.

Nvidia open sources its Aerial 6G software stack

At GTC DC 2025, Nvidia announced a major open source milestone by releasing its Aerial 6G software stack, marking a decisive shift towards programmable, open



AI factories. The move enables researchers, telecom startups, and hyperscalers to prototype full-stack 5G/6G networks and AI-RANs on desktop supercomputers such as DGX Spark. By open sourcing Aerial, Nvidia is democratising access to telecom innovation and accelerating the

transition from proprietary systems to programmable, composable architectures.

Nvidia chief executive Jensen Huang underscored a strong national message:

"The government is not standing still and it needs to keep up the pace to win the AI race." The company's open innovation model aligns with the goal to embed AI into America's infrastructure, manufacturing, and telecommunications, strengthening national competitiveness.

The open source announcement was supported by a wave of industrial deployments.

Pharma giant Eli Lilly unveiled the Blackwell SuperPOD, the world's largest AI factory in the sector, powered by over 1,000 Nvidia Blackwell Ultra GPUs and designed to accelerate drug discovery through federated learning and foundation models.

China's Kimi K2 Thinking model outperforms competitive closed and open models

China's Moonshot AI has unveiled Kimi K2 Thinking, an open-weights reasoning model that has outperformed leading closed systems from OpenAI and Anthropic, marking a major inflection point in the evolution of open AI. Released under a modified MIT licence, the model scored 44.9% on Humanity's Last Exam, surpassing GPT-5 and Claude Sonnet 4.5, and achieved 60.2% on BrowseComp, a key benchmark for agentic search and browsing. It further demonstrated its coding strength with 61.1% on SWE-Multilingual, 71.3% on SWE-bench Verified, and an impressive 83.1% on LiveCodeBench V6.

K2 Thinking excels in agentic tasks requiring long autonomous operations, capable of performing 200-300 sequential tool calls without human input. Built with a 256K context window and designed for test-time scaling, it expands both thinking tokens and tool-calling abilities. The model is currently accessible on kimi.com in chat mode, with a full agentic mode coming soon and API access already available.

For more news, visit www.opensourceforu.com



ClickHouse acquires LibreChat to accelerate agentic analytics

ClickHouse has acquired LibreChat, bringing founder Danny Avila and his team into the organisation and establishing the foundation for the open source Agentic Data Stack. The combined stack enables enterprises to deploy AI agents that analyse large datasets through natural language, act on insights, and integrate seamlessly with existing data workflows.

This move accelerates the industry shift towards agentic analytics, where intelligent agents replace conventional dashboard-driven interactions. These agents democratise access to data, reduce time-to-insight, and enable more intuitive human-AI collaboration. Organisations such as Shopify, Daimler, Fetch, and cBioPortal are leveraging the ClickHouse and LibreChat architecture to solve real-world problems.

In healthcare research, cBioPortal has used the combined stack to open new data exploration pathways. "It puts discovery at cancer researchers' fingertips," said Ino de Bruijn, manager bioinformatics software engineering at cBioPortal.

ClickHouse will continue to support and invest in LibreChat as an open source project, ensuring transparent deployment, multi-LLM flexibility, and enterprise-scale integration without vendor lock-in.

