

Open Source and the Architecture of AI

A Foo Camp for the new agentic economy

Co-hosted by O'Reilly Media, AI Disclosures Project, and Partners

October 23 – 25, 2026
Lighthaven, Berkeley, CA¹

You are invited to join a weekend “Foo” Camp to discuss the future of open source AI.

A handful of firms today are competing to dominate frontier AI and, in turn, to decide who may build on their systems, at what price, and under what terms. Open source and distributed AI systems provide an essential counterweight to this future. But open weights alone do not create a fully open market.

The challenge remains how independent developers, companies, universities, public institutions, and nations can use AI systems without becoming tenants of another proprietary layer of the stack. Making openness usable and competitive, rather than merely feasible, requires a range of ingredients, including accessible compute; transparent data and training methods; reusable model harnesses and evaluations; portable context and memory; standardized telemetry and execution traces; identity and permissions for users and agents; and loyal user agents that act for the user rather than for the platforms they encounter.

Over the weekend, we will consider what it would take to build this wider architecture of participation: not merely open models, but an open, federated, and composable AI production system.

Background

The Web was generative in large part because its core standards – HTTP for communication, HTML for publishing, and URLs for universal addressing – were open, interoperable, and implementable without permission from a platform owner. Linux, Apache, PostgreSQL, and the browser stack followed a related design principle: complex systems were divided into

¹ Want to suggest people who should attend? Contact: marsee@oreilly.com; Would you like to be on the convening committee? Contact: ilan@aidisclosures.org

components that different communities could understand, improve, and recombine through stable interfaces. Common interfaces allowed [specialized](#) producers to emerge, thereby expanding the market for everyone. Modularity reduced dependence on monolithic architectures, while open interfaces enabled third-party ecosystems in which external participants could both create and capture value. Tim O'Reilly has called this an [architecture of participation](#).

Open-weight AI models are approaching an important threshold. These models are increasingly competitive for many workloads and are being adopted by companies seeking greater independence from proprietary APIs. But their deployment may still depend on concentrated cloud infrastructure, proprietary fine-tuning services, closed evaluation systems, or applications that separately control the user's tools, memory, and accumulated context.

Mozilla's recent [State of Open Source AI](#) report finds that teams using open models reach production at lower rates than teams using closed models. The gap is not explained by model quality alone. Developers working with open models reported high infrastructure and compute costs, security and compliance concerns, ongoing maintenance, deployment complexity, and limited specialized support. Mozilla's assessment of the wider open stack identifies standardization and enterprise readiness as persistent weaknesses.

In other words, open-weight models are not enough to build a competitive open source ecosystem. Producers using open models need to be able to procure specialized products from across the stack to make their usage competitive and [task efficient](#) compared with [vertically integrated](#) options. This includes having access to competitive open harnesses that organize agency and execution efficiency; as a nominally open (weight) model inside a closed harness may offer little practical freedom to the developer, enterprise, or end user.

[Open protocols such as MCP](#) matter too. MCP can make tools and external data sources accessible across different AI applications, reducing the need to hard-wire each integration to a single platform. But model interchangeability, portable memory, identity, permissions, telemetry, and execution traces require additional standards and open infrastructure.

Under a more distributed or federated architecture, intelligence can emerge through combinations of specialized models, domain experts, local institutions, shared datasets, open tools, federated compute, and human communities. But distributed production requires mechanisms through which components can interoperate and contributions can be evaluated, trusted, and compensated.

This convening will bring together 200 builders, researchers, maintainers, funders, economists, policymakers, and institutional users to explore these and other questions:

From Open Models to Open Systems

- What must be open or interoperable for developers to use, modify, deploy, and switch between models, tools, and providers without rebuilding their systems?

- How can smaller models, knowledge distillation, and specialized intelligence make open systems competitive with frontier APIs?
- Where are open source alternatives most urgently needed across the AI stack?

Building a Competitive, Secure Open AI Market

- What technical, financial, and institutional barriers most constrain the deployment of open models?
- How can startups, maintainers, developers, and data contributors create and capture value without becoming dependent on new platform gatekeepers?
- What governance, provenance, and security practices can make open systems trustworthy without recreating centralized control?

Open Infrastructure for Agents

- What should be portable across AI applications, including: tools, memory, identity, permissions, payments, and execution history?
- How can agents move between models and services without losing context, while remaining inspectable and accountable?
- Which shared interfaces and infrastructure would allow open components to compete with tightly integrated proprietary systems?

The ultimate test of “openness” is whether independent actors can build, deploy, inspect, modify, switch, and meaningfully contribute – and whether the value created through their participation grows the market for everyone, rather than concentrating at a significant chokepoint. This Foo Camp aims to identify the concrete technical, economic, and institutional choices needed to make such an open ecosystem real.