

# Awake Venture Studio Launches ForgeOS, the First Open-Source Multiplatform Operating System for AI Agents

CHICAGO, IL, UNITED STATES, May 19, 2026 /EINPresswire.com/ -- [Awake Venture Studio](#) Launches ForgeOS, the First Open-Source Multiplatform Operating System for AI Agents



Built by Awake Venture Studio, ForgeOS solves the biggest problem in enterprise AI: controlling what agents do when they act autonomously — across Google, Anthropic, OpenAI, and any AI platform

“

You need to guarantee that Agent A cannot access Client B's data, that spending stays within budget, and that sensitive decisions get human review. No AI platform offered this. So we built it.”

*José Antonio Martínez Aguilar*

With a kernel architecture modeled on UNIX, ForgeOS is the first agentic harness — a governance layer that enforces budgets, permissions, and human oversight for AI agents on any framework

[Making Science](#), a global leader in digital acceleration powered by AI, today announced the open-source launch of ForgeOS, a new AI governance platform from its innovation hub, Awake Venture Studio. ForgeOS is the first operating system designed specifically for AI agents — giving companies the ability to control spending,

permissions, and oversight for AI agents across every major platform, without changing their existing code.

## The Problem: AI Agents Without Rules

As companies move from experimenting with AI to deploying autonomous agents that take real actions — sending emails, managing budgets, accessing customer data, modifying infrastructure — a critical gap has emerged. Every AI platform today gives you a way to build agents. None gives you a way to govern them.

What happens when a company deploys 200 AI agents across sales, marketing, finance, and operations — and one of them exceeds its budget by 10x? accesses data it was never supposed to see? Or takes an action that requires human approval?

ForgeOS answers this question.

### The Solution: A Harness for AI Agents

ForgeOS is not another AI tool or chatbot. It is the control layer that sits between the AI platform and the agent, managing what the agent is allowed to do — regardless of whether it runs on Google's technology, Anthropic's Claude, OpenAI's GPT, or any other AI system.

Think of it like the operating system on your computer. The applications run on top of it, but the operating system controls which files they can access, how much memory they can use, and when they need your permission. ForgeOS does the same for AI agents: it controls which tools they can use, how much they can spend, and when they need a human to approve their actions.

The key capabilities include:

- Spending controls — Set daily and per-task budgets for each agent. When the limit is reached, the agent stops automatically.
- Permission management — Define exactly which actions and tools each agent is allowed to take and use. An agent designed for research cannot send emails. An agent managing Client A cannot access Client B's data.
- Human-in-the-loop — Certain actions automatically require human approval before proceeding. Medical content needs a doctor's review. Financial decisions above a threshold need CFO sign-off.
- Complete audit trail — Every action, every decision, every approval is recorded. When regulators or clients ask "what did the AI do?", you have the answer.
- Crash recovery — If an agent fails mid-task, it resumes from where it left off instead of starting over.

### Works With Every Major AI Platform

ForgeOS is the only governance solution that works across all major AI providers simultaneously. Companies today use Google's Gemini for cost-efficient tasks, Anthropic's Claude for complex reasoning, and OpenAI's GPT for specific capabilities. ForgeOS applies the same governance rules to all of them through a single control panel.

This means a company can mix and match AI providers for different tasks while maintaining one set of rules for budgets, permissions, and oversight.

"The next frontier of AI is not just about creating smarter agents, but about creating systems that govern them safely," said José Antonio Martínez Aguilar, CEO and Founder of Making Science. "Every company deploying AI agents today faces the same problem we did — no spending controls, no permissions, no audit trail. ForgeOS solves this for every AI platform at once."

### Proven on Real Infrastructure

ForgeOS is already operational within Making Science in different internal processes and

workloads. In its first deployment, an AI agent governed by ForgeOS audited all of Making Science's cloud infrastructure across hundreds of projects in minutes.

Beyond this security auditor another three production agents have been built and deployed:

- A content operations pipeline that produces marketing content for multiple client brands, with automatic compliance review and human approval for regulated industries like healthcare and financial services
- A code security scanner that reviews every code change for security vulnerabilities before it reaches production
- A competitive intelligence analyst that researches market trends using multiple AI models, with budget controls to prevent overspending

These four agents are available also as open source.

"We built ForgeOS because we needed it ourselves," said Martínez Aguilar. "When you run AI agents for clients, you need to guarantee that Agent A cannot access Client B's data, that spending stays within budget, and that sensitive decisions get human review. No AI platform offered this. So we built it."

#### Open Source, Enterprise-Ready

ForgeOS is released as open source, reflecting Making Science's commitment to transparency and the European values of data sovereignty and auditability. Every line of code is readable and auditable — there are no hidden components.

"We believe AI governance should be transparent," said Martínez Aguilar. "Companies trusting AI agents with their operations and their clients' data deserve to see exactly how those agents are controlled. Open source is not just a licensing choice — it is a trust signal."

The platform is particularly relevant as the EU AI Act comes into effect, requiring companies to demonstrate oversight, auditability, and human control over AI systems. ForgeOS provides the technical infrastructure to meet these requirements.

ForgeOS ships with Mission Control, a real-time dashboard that gives operations teams complete visibility into their agent fleet — how many agents are running, what they're doing, how much they're spending, and which decisions need human approval.

#### Pricing and Availability

ForgeOS is available immediately on [GitHub](https://github.com/makingscience-awake/forgeos) (github.com/makingscience-awake/forgeos). The full technical vision is published as "The ForgeOS Manifesto" on GitHub.

Monica Gonzalez

monica.gonzalez@awakeventurestudio.com

Making Science

+1 877 698 9977

---

This press release can be viewed online at: <https://www.einpresswire.com/article/913677194>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.