Lyrion joins the Open Secure AI Alliance

Lyrion is a member of the Open Secure AI Alliance, formed to build and share open tools, models and techniques for defending software and agents. This is our statement on why we joined and what we intend to put into it.

FOUNDING MEMBERS · OPEN SECURE AI ALLIANCE

Open source software is one of the load bearing parts of the economy. It runs cloud computing, financial services, manufacturing, telecoms, government and most of what people think of as the internet, and it does so in the open, where anyone qualified can read it, test it and fix it. Security is one of the clearest beneficiaries of that arrangement.

The same question is now open in artificial intelligence: whether the systems that defend critical infrastructure sit inside a small number of closed products, or are built on models, harnesses and tools that any defender can study, adapt and run.

Why defenders need open systems

The world needs both closed and open models. For cybersecurity in particular, open models and open harnesses matter because they put defensive capability in more hands, make behaviour visible to the people responsible for it, and let an organisation run analysis on its own infrastructure without handing its data to anyone.

That last point is not theoretical. Defenders regularly need to inspect an incident at a depth and a speed that a closed tool, unable to tell an attacker from an investigator, will refuse. When a defender cannot inspect, adapt and run advanced AI on their own machines, their ability to respond is constrained at exactly the moment speed matters most.

The risks, stated plainly

Open models can be misused. Safeguards can be stripped, and capabilities can be repurposed. Those risks are real and they should be argued with honestly rather than waved away.

They are also not unique to open systems, and keeping weights closed does not stop a determined attacker from finding capable AI. The answer is to pair openness with strong safeguards, clear rules against malicious use, serious evaluation and fast remediation, rather than to leave defenders with less capability than the people they are defending against.

What the alliance is for

The Open Secure AI Alliance is a group of companies and foundations committing to develop and share open technologies, techniques and tools for safeguarding software and agents. It builds on work already underway in the open source community, including the Linux Foundation's Akrites initiative and OpenSSF, and it is oriented towards the unglamorous end of the problem: finding vulnerabilities, disclosing them, and remediating them using tools everyone can inspect.

What Lyrion brings to it

Lyrion's interest in the alliance is narrow and practical. Lyrion runs agents inside other companies' systems, with their data and their permissions. The standards for how an agent is scoped, what it is allowed to touch, how its actions are recorded and how a mistake is reversed should be worked out in public, by more than one vendor, and not left for each company to invent on its own.

So that is what we will contribute: the controls, the audit trail and the evaluation work that come out of running agents in production, published openly, for anyone in the alliance to take apart and improve.

Lyrion

Member, Open Secure AI Alliance
contact@lyrion.io · lyrion.io

This letter is published by Lyrion as a member of the alliance. The Open Secure AI Alliance was announced by NVIDIA on July 27, 2026; the announcement is at blogs.nvidia.com/blog/open-secure-ai-alliance/.