



AI Openness: 3 Years on from Llama 2

Security and Vendor response

Karan Saini
Research Manager, OpenUK



Index

1. Introduction	4
1.1 Executive Summary	4
1.2 Introduction - Three Years on from the Opening of Llama 2 Professor Amanda Brock, CEO, OpenUK and OpenHQ	6
1.3 Take the OpenUK & OpenHQ Survey, in association with	8
2. Three Year Evolution of the AI Revolution	9
2.1 Launch of Llama2 as Open Innovation	9
2.2 The Great Debate defining Open Source AI	11
2.3 The Deep Seek Moment and the rise of China’s AI	13
2.3.1 China’s Strategic Approach and Policies	13
2.3.2 China’s Technology	13
2.4 The US Response	13
2.4.1 US Executive Order Approach and Policies	13
2.4.2 The Technology	16
2.5 The European Union	17
2.5.1 The EU AI Act	17
2.5.2 EU Sovereignty and Open Source Strategy, 2026	17
2.5.3 The Technology	17
2.5.4 The EU Open Source Policy, Sachiko Muto, Chairwoman, Open Forum Europe	18
2.6 India	19
2.6.1 The Technology - Sarvam Small Language Model	19
2.6.2 Approach to Policy	19
2.7 Middle Nations	19
2.8 Canada	19
2.9 The UK	20
2.9.1 Policy not Regulation	20
2.9.2 The Technology	21
2.9.3 Case Study: Open weights, owned compute and the case for sovereign AI, James Drayson, CEO and Co-founder, Locai Labs	22
2.9.4 Fireside Chat: Open foundation models enabling a map of the planet, Professor Anil Madhavapeddy, Professor of Planetary Computing, University of Cambridge	25
2.9.5 UK Instigated Global AI Summits	27
2.10 International Policy Making and Governance	28
2.10.1 AI Action Summit, Paris 2025	28
2.10.2 Open Source in the India AI Impact Summit	28
2.10.3 Geneva in 2027	28
2.10.4 International Organisations	28
2.11 Shifting focus to Agentic AI	29
2.11.1 The Agentic Harness	29
2.11.2 Thought Leadership: Security, Identity and Harnessing Agentic, Sal Kimmich, CEO, Gadfly AI	30
2.11.3 Case Study: The OpenClaw Foundation, Vincent Koc, Chief Architect, OpenClaw Foundation	31
2.12 AI and Open Source Software	35
2.12.1 Impact of AI on Open Source Software	35

Index

2.12.2 Thought Leadership: Open Source tools for AI, Hannah Foxwell, Co-founder BIMP and creator of AI for the rest of us	36
2.13 AI and Open Standards	38
2.13.1 Global Open Standards	38
2.13.2 Fireside Chat: Model Context Protocol (MCP) and Protecting its openness, David Soria Parra, Member of Technical Staff, Anthropic	39
2.13.3 New Foundations	41
2.14 The AI Stack and AI Sovereignty	43
2.14.1 Sovereignty and Geopolitical Shift	43
2.14.2 Thought Leadership: The Commons we have not (yet) created, Anastasia Stasenko, Founder and CEO, Pleias	45
2.14.3 Case Study: Open weights, owned compute and the case for sovereign AI, Tom Wilkinson, Chief Data Officer, Scottish Government	47
2.14.4 Thought Leadership: Security and Vendor response, Karan Saini, Research Manager, OpenUK	50
2.15 Value generation and ROI	52
3. What the Data Shows	52
3.1.1 AI Repositories Globally by Country	52
3.1.2 AI Repositories Evolution over the last 3 years	53
3.1.4 Top AI Repositories in the UK	54
3.1.5 Fastest growing AI Repositories Globally	55
3.1.6 Comparison of Contributions Received by AI repositories	55
3.1.7 Top Topics in AI Repositories	55
3.2 Model Repositories	56
3.2.1 Top 10 Most Downloaded Models	56
3.2.2 Top 10 Most Liked” Models	56
3.3 Agentic Repositories	57
3.3.1 Agentic Repositories by Country	57
3.3.2 Top 10 Agentic Repositories Globally	58
3.3.3 Top 10 Agentic Repositories UK	58
3.4 Conclusion - What we see from the Data	59
3.5 The Entrepreneur’s AI Journey, Matt Barker, Founder and CEO, Bolt MCP	60
3.6 Case Study: Agentmemory, Rohit Ghumare, Creator & Maintainer, Agentmemory	61
4. What the Literature Says	62
4.1 Literature Review: Open Source AI, Global Ecosystem Shifts and the UK Policy Response in 2026	62
4.2 Literature Themes from 2023-2025	66
5. Conclusion	67
6. Formalities	69
6.1 Authors of this Report	69
6.2 Creators of this Report	69
6.3 Contributors	70
6.4 Methodology	71
6.5 Acknowledgements	72
6.6 References	72
6.7 Sponsors	74

2.14.4 Thought Leadership: Security and Vendor response

Karan Saini
Research Manager, OpenUK



A common architecture has emerged across the companies now selling responses to AI-discovered vulnerabilities. Findings produced by frontier models are pooled into a private clearing house, fixes are developed and validated under embargo, and remediated packages reach paying subscribers or vetted coalition members ahead of public disclosure. The justification is that a published patch can be reverse-engineered into a working exploit within hours, so prompt public disclosure now favours the attacker. In effect, the security of open source software will now increasingly be mediated by commercial subscriptions and membership schemes, with access to models capable of finding the flaws itself being rationed by verification regimes.

Chainguard: Athena

Chainguard launched Athena on 15 June 2026 with members including BNY, Cisco, Cloudflare, Corridor, DepthFirst, Docker, JPMorganChase, Kyndryl, LTIMindtree and PwC as a coalition for the orchestrated defence of open source software. Members submit pre-disclosure findings, many of them generated through Anthropic's Project Glasswing and OpenAI's Daybreak, through an encrypted portal. Submitters choose the sharing scope and embargo timeline. Athena deduplicates and enriches the findings, traces when a flaw was introduced and whether it is already fixed at head, and publishes the resulting metadata as an OSV feed. Patched builds and hardened private forks reach members through Chainguard Libraries before public disclosure. The commercial mechanism sits alongside the coalition and fixes generated are available to Chainguard customers as a product before it is available to everyone else as a public release. At launch Chainguard reported more than 20,000 findings processed and over 2,000 patches shipped across 500 projects; three weeks later it reported over 40,000 vulnerabilities processed, of which it says 42% were critical or high severity and 86% network reachable, and added Akamai, JFrog, Morgan Stanley and Qualys as members.

IBM and Red Hat: Lightwell

IBM and Red Hat announced Project Lightwell on 28 May 2026 as a \$5B commitment backed by more than 20,000 engineers to address a specific operational problem - the standard remedy for a vulnerable dependency is to upgrade to a fixed upstream release. Large organisations running older or heavily customised versions in production frequently cannot or do not wish to do this. Lightwell backports fixes to the versions customers already run, shipping them as digitally signed, certified artefacts with SBOMs delivered into existing build pipelines. Lightwell Network is generally available on an annual subscription whilst Lightwell Clearinghouse Premier is restricted at launch to financial services, with expansion to government, healthcare and telecommunications planned in later phases. Partners include Bank of America, BNY, Citi, Goldman Sachs, JPMorganChase, Mastercard, Morgan Stanley, Royal Bank of Canada, State Street, Visa and Wells Fargo. Red Hat's upstream-always commitment, sees fixes developed for Lightwell submitted back to the originating project for review and acceptance, which is intended to prevent the catalogue from becoming a permanent private fork of the public ecosystem.

OpenAI: Daybreak

OpenAI's launched Daybreak, launched on 11 May 2026 making Standard GPT-5.5 generally available; GPT-5.5 with Trusted Access for Cyber is shared only with verified defenders for authorised defensive work; and the more permissive GPT-5.5-Cyber is reserved for approved workflows such as red teaming and exploit validation. Approval is discretionary, turning on identity verification, intended use and organisational capability, and may not be extended to third-party or customer-facing traffic.

OpenAI presents this limitation of distributions as an alternative to withholding capability.

The openly directed component, Patch the Planet, funds researchers and equips them with Codex Security to carry findings through to merged fixes, with more than 30 projects committed including cURL, Go, Python and Sigstore, and an early five-day sprint producing dozens of merged fixes. It is the only part that transfers capability to maintainers rather than to their downstream consumers, and it is small beside the commercial tiers.

Anthropic

Anthropic has taken the opposite approach by keeping Claude Mythos Preview restricted under Project Glasswing on safety and national security grounds.

The Volume Claim

The vendors' case rests on vast figures for vulnerabilities that are largely self-reported and not directly comparable. Chainguard's June 2026 quarterly report records 18,016 vulnerability instances across 886 distinct CVEs in its customer base, with high-severity findings 63.1% of instances, a rise of 13% on the preceding quarter. Anthropic estimated that its tooling surfaced 6,202 high- or critical-severity issues across more than 1000 open source projects. IBM's Lightwell cites the Mythos Preview model as identifying 3,900 high- or critical-severity vulnerabilities in open source software.

GitHub published 1,560 reviewed advisories in May 2026 and processed more than 6,000 advisory decisions a month between March and May, while still reporting that incoming private reports, repository advisories and CVE requests exceeded its review capacity.

Impact on Openness

The three commercial responses narrow access in order to widen protection through limiting access to the models gated by verification, in OpenAI's case through Trusted Access and in Anthropic's through outright restriction. Then access to fixes is gated by subscription or membership, through Lightwell's tiers and Chainguard Libraries. Both are on the same reasoning, that publication in the current conditions arms attackers faster than it protects defenders.

That reasoning is not obviously wrong. Its cost is that the remediation of a public commons is now being performed first for the organisations that can pay for it, with upstream contribution as a consequence rather than the primary channel. The upstream-always commitments are attempts to prevent that outcome, and it is too early to say whether they are sufficient. What is measurable, is the interval between a member receiving a hardened build and the corresponding fix landing in the public upstream repository, together with the proportion of processed findings that reach an upstream merge at all. Neither figure is currently published by any vendor.

Conclusion

Security work on open source is moving from a public activity to one intermediated by commercial clearing houses sharing first with paying members, rationing either the models that find flaws or the fixes that resolve them, each paired with a commitment to return value upstream on the basis that machine-generated findings now exceed what maintainer communities can absorb. They suggest that pooling under embargo is the only workable answer. Its corollary is that the commons becomes dependent on the arrangement, and that those able to pay are served first on software everyone else relies on. Whether this becomes accepted practice requires two unreported measures: the interval between a member receiving a hardened build and the fix reaching the public repository, and the share of processed findings that reach an upstream merge at all.

At the time of publication, the technology press shared that Microsoft is preparing to launch Project Perception, to compete with these products in securing open source.

6. Formalities

6.1 Authors of this Report

Professor Amanda Brock, CEO, OpenUK

OpenUK CEO, Amanda's built one of open source's most recognised and impactful organisations. Executive Producer of State of Open Con (2023- 2025), Amanda's a globally sought-after keynote speaker. A lawyer with 25 years' experience, 5 as GC of Canonical, she's been instrumental in shaping open source's legal frameworks, as she was internet law during the early 2000's. Regularly contributing to tech press, she edited 'Open Source: Law, Policy and Practice', (2022).

Recognition: Computer Weekly 50 Most Influential Women in UK Tech (2023, 2024); Computing IT Leaders 100 (2023, 2024); Lifetime Achievement Award WIPL (2022); Women Who Will Changemaker (2023); INvolve Heroes (2022, 2023); Novi Awards (2024) and Ambassador, Open Charge Alliance. Advisory Appointments: UK Cabinet Office Open Standards Board; UKRI Digital Research Infrastructure; UKRI Exascale; KDE; commercial boards – Mimoto, Scarf, FerretDB and Space Aye; and is Fellow Open Forum Academy; Distinguished Fellow Rust Foundation; and European Representative, OIN.

Dr. Jennifer Barth, Head of Research, Symmetry Research at FSP and Chief Research Officer, OpenUK

Jenn holds a DPhil from the University of Oxford and is Head of Research at Symmetry Research at FSP, where she brings a rare blend of ethnographic depth, commercial acuity, and ethical clarity to the AI conversation. Jenn's work sits at the intersection of technology, organisational culture, and socioeconomic change. She is known for translating complex technological shifts into human-centred insights that resonate deeply with executive audiences. Her measured, critically informed perspective — grounded in qualitative research and lived experience — enables leaders to move from abstract concern to practical, principled action. Alongside her role at FSP, Jenn serves as Chief Research Officer at OpenUK.

Karan Saini, Research Manager, OpenHQ

Karan Saini is a New Delhi-based security researcher, public interest technologist, and open source contributor. He was most recently a Senior Information Controls Fellow with the Open Technology Fund, hosted by the Internet Governance Project at Georgia Tech, where he researched the scale of DNS-based web censorship in India. He regularly writes for leading publications and has contributed to award-winning media projects.

6.2 Creators of this Report

OpenUK

[OpenUK](#) is the unique open tech industry organisation for the business of open technology in the UK. It spans the opens – software, hardware, data, standards and AI and is the convening point for the UK's business, academic and contributing communities across open tech. Our work supports the UK's journey to become "The State of Open". Our organisation is run with the support of our volunteer community and their leadership in the tradition of open source manner delivering on three pillars: community, legal and policy and learning. OpenUK's Community is recognised through our world-leading recognition program including the Open UK Awards (the Oscars of Open Source) now in their 7th year, New Year's Honours Lists and Ambassador Scheme.

OpenUK undertakes research and reporting both on its own account through its "State of Open Reports" and on a commissioned basis for third parties. Case studies, Thought Leadership, Surveys and desk-based research are included in our reporting which pushes the envelope and leads the way. Our Research and Reporting Show and Tell events coalesce the global open source research communities digitally to regularly update and share research practices and topics. .

The State of Open Con has become one of the world's leading open source conferences since its inception by OpenUK in 2023. In 2027 we expect to host 1000 people across 8 tracks and plenary sessions, with at least 50 partners in our delegate experience space and over 200 speakers in London. For 2026 we are on the road, with smaller events across the UK, meeting our audience. Contact <mailto:admin@openuk.uk>

Symmetry Research

Symmetry Research, an FSP company, looks beyond the surface and behind the curtain of the fundamental innovations and trends shaping our society, markets, culture, and values. We are academics and researchers looking at the intersections of emerging technology and socioeconomic impact, producing independent research for thought leadership and business solutions. We take a research-led approach to helping organisations transform through and with technology. Symmetry Research's mission is to share and grow knowledge about the interaction of technology and everyday lives. We want to understand the past, present, and future of human interaction with emerging technologies and socioeconomic changes—from behaviour to context, nature to nurture, origin to experiences—we do independent research for high-stakes business decisions.

6.3 Contributors

Anastasia Stasenko

Anastasia Stasenko is the co-founder and CEO of Pleias, a Franco-German startup developing full open source small reasoning models trained exclusively on permissively licensed data. A senior associate lecturer in data analysis and digital strategy at Sorbonne-Nouvelle University, she coordinated the creation of Common Corpus, the largest fully open dataset for training LLMs, with over two trillion tokens.

Professor Anil Madhavapeddy

Anil Madhavapeddy is the Professor of Planetary Computing at the Department of Computer Science and Technology in the University of Cambridge.

David Soria Parra

David Soria Parra is a Member of Technical Staff at Anthropic based in London. He co-created the Model Context Protocol (MCP). He leads its technical direction, governance, and ecosystem development. Before joining Anthropic in 2024, David spent nearly a decade at Meta leading engineering organizations, and served as an Entrepreneur in Residence at Sutter Hill Ventures. He is a long-time open source contributor with a track record spanning PHP, Git, and Mercurial.

Hannah Foxwell

With over a decade of experience in engineering leadership, Hannah Foxwell has always advocated for the human aspects of technology transformation and evolution. Hannah is relentlessly curious about the tools, technologies, processes, and practices that make life better for the people involved in software development. This curiosity has naturally evolved alongside technological changes—from Agile to DevOps, DevSecOps, Cloud Native, and Platform Engineering. In 2026 Hannah co-founded BIMP, bringing together her knowledge of cybersecurity, platform engineering, and AI to deliver software supply chain security at the speed of AI. As the creator of "AI for the rest of us" Hannah leads a community of over 2,000 enthusiasts in the UK, providing free monthly events and an annual conference aimed at demystifying the technology and improving AI fluency.

James Drayson

James Drayson is the CEO and co-founder of Locai Labs, a UK company based in London that develops on-prem sovereign AI so enterprises can own their intelligence. James holds a Bachelor of Science in Biology from the University of California, Los Angeles, and a Master of Science in Innovation, Entrepreneurship and Management from Imperial College London. He is Forbes featured and a UKRI Young Innovator Award winner.

Matthew Barker

Matt Barker is co-founder and CEO of BoltMCP, focused on helping enterprises securely connect AI to their data and control how it's used. Previously, he was CEO and co-founder of Jetstack, where he helped build one of the most widely adopted projects in the cloud native ecosystem, cert-manager. Under his leadership, Jetstack bootstrapped to a team of 50, and was acquired by Venafi in 2020.

Rohit Ghumare

Creator & Maintainer, Agentmemory, CreateMVP, and AI Engineering from Scratch AI-native course. Open source Leader and Cloud-native engineer. He writes and talks about AI agents, harnesses, AI Inference, and the infrastructure underneath, and is currently building iii.dev. Contributor to major open source global projects.

Sachiko Muto

Sachiko Muto is the Chair of OpenForum Europe, a senior researcher at RISE Research Institutes of Sweden, and a Board member of the Drupal Association. She originally joined OFE in 2007 and served for several years as Director with responsibility for government relations and then as CEO. Sachiko has degrees in Political Science from the University of Toronto and the London School of Economics; she received her doctorate in standardisation policy from TU Delft.

Sal Kimmich

Sal's work is centered on advancing secure computing practices and pioneering privacy-enhancing technologies to safeguard sensitive data. At the Confidential Computing Consortium (CCC) within the Linux Foundation, they play a pivotal role in the development and implementation of Open Source projects that establish Trusted Execution Environments (TEEs). These environments are designed to provide secure areas within hardware that protect data in use, addressing critical security challenges across various industries.

Vincent Koc

Vincent Koc is Chief Architect at the OpenClaw Foundation, where he leads engineering for one of the fastest-growing open source projects in AI. An engineer and technologist with nearly two decades of experience across the Asia Pacific and the US, he has contributed to major open source projects across the data and AI ecosystem and teaches applied AI at universities including MIT. He joined OpenClaw as a community maintainer focused on security and observability before moving to the Foundation full-time.

Quoted Individuals


Andrew Martin, Co-Founder and CEO, ControlPlane; Adrian Mouat, DevRel Engineer, Chainguard; Chris Yiu, Head of Policy Northern Europe, Meta; Eric Schmidt, Former CEO of Google, Chair of the Special Competitive Studies Project; Jean-Baptiste Kempf, President, VideoLAN and Kyber; Jensen Huang, CEO, NVIDIA; Juliette Denny, Founder, Growth Engineering and Zavmo; Luke Hinds, Founder, nolabs, Inc; Sam Altman, Founder and CEO, OpenAI; Toran Bruce Richards, Founder and CEO, AutoGPT; Yann LeCun, Founder and Chairman, AMI Labs

6.7 Sponsors


OpenUK is grateful to its many sponsors and supporters without whose support its work including its report would not have been possible in particular, its general sponsors Arm, GitHub, Google, Linaro, Microsoft, Red Hat.



[Download the full report here](#)



openuk.uk/report/aiopenness2026



[OpenUK Newsletter](#)



[Our Conferences - State of Open Con](#)

