

Open:UK[®]  Open:HQ[®] 

AI Openness: 3 Years on from Llama 2



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: Mobile 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

1. Introduction

1.1 Executive Summary

Our report looks at the global history of AI Openness in the context of Llama’s launch as open innovation on 18 July 2023, and three year journey since culminating in Moonshot release of open weight model Kimi K3, challenging “best” frontier model. The report scans the technological, business and policy evolution and offers deep analysis of the relevant data and themes in AI openness. AI Openness has shifted to strategic capability, and no longer a technical choice

Evolution of models from Llama 2

- Llama 2 - never open source but open innovation. A critical step in shifting AI to openness
- DeepSeek R1 - open model 20% of cost of comparable models, distilled existing models including Llama 2
- China’s Moonshot AI - Kimi K3 on third anniversary of Llama has capabilities appearing to match best frontier models
- Evolution of models to small models like India’s Sarvam
- Models predominantly corporate and shared on Hugging Face whilst Agentic, Dev Tools and infra-structure being evolved by individuals and collaborative communities on GitHub
- Model adoption delayed by external environment and production readiness not capability
- Shift in thinking - biggest not best. Customisation is critical, see Lumai case study

Agentic Harness and AI Openness

- Agentic AI Harness for Models - open source community and individuals/ innovators building
- UK evolving critical agentic technologies including founding of OpenClaw and AutoGPT
- OpenAI admit an agent hacked Hugging Face, in early public example of an AI System
- Enterprise AI adoption accelerating, with agentic systems increasing deployment
- Infrastructure Layer likely where value and return will be seen - shifting thinking on ROI

Dev Tools and Infrastructure

- Dev tools are AI Native and open source, now key to development of AI
- De facto standards emerging such as MCP with foundations being built to host these (all currently US based) see MCP case study
- AI reshaping how software is developed; AI tools shaping developer choices e.g. language

Open Source Ecosystem

- Open Source projects overwhelmed by AI contributions but policy is evolving in projects
- New US Foundations: Agentic AI Foundation, Agent Name Service, Akrites Foundation and t x402 Foundation - vendor-originated protocols like MCP under Linux Foundation governance and founders running closed models. Recognises infrastructure and dev tools open source
- Enterprise drive to vulnerability management - Chainguard IBM Red Hat and OpenAI products geared to support maintainer security challenges, with variable upstream impact

The AI Openness Definition Debate

- What open source means in AI is unresolved, Open Source AI Definition did not gain support
- Alternative approaches focus on disaggregation of AI component to measure openness and likely better approach - examples Stanford, Radboud and LF Model Openness Framework
- Accusations of Meta open washing Llama 2 reached mainstream press
- Language shifted from open source to open weight models and AI Openness
- G7 actively looking at what AI Openness means in 2026 following June meeting/ statement

China now leads open source AI

- China distillation models to create DeepSeek R1, Alibaba’s Qwen family and Moonshot AI’s Kimi family. Focus on low compute usage across all. Adoption growing
- Policy level eight year program across two consecutive Five-Year Plans shifting emphasis to reducing proliferation and monetisation. Xi Jinping restated open source commitment in four-point policy at Shanghai AI Summit, announcing WAICO Foundation, July 2026
- Focus on embodied AI as world shifts to utility of AI for industrialisation

US Response

- Unprecedented investment and spend in proprietary AI with talk of bubble
- US export control restrictions on NVIDIA chips to China and Presidential Executive Order on Anthropic Claude Fable and Mythos crystallised growing global AI sovereignty fears
- US claims distillation is breach of copyright (perhaps ironic when training is considered)
- US reportedly in turmoil considering regulatory and soft power responses to Kimi K3

Shift in Models

- Value in models shifting from size to customisation, highlighted in Locai Labs case study
- Sovereignty layered across the Stack generally following NVIDIA’s five-layer-cake
- Localisation being accepted by many rather than build from ground up
- Potential shift in data centres and energy infrastructure as a consequence of smaller models
- Copyright in Text and Data Mining for the UK remains unresolved and yet more consultation

Analysing the Repo Data

- 5x as many GitHub Repos with 1,000+ Stars as in 2020US 1,131, China 685, UK 161
- UK moved ahead of Germany whose growth is slowest in top tier
- 1,000+ Stars GitHub Repos doubled since July 2023 - 2,159 to 4,806
- 40% of Agentic 1000+ Star Repos had none a year ago
- 37% of Agentic 1000+ star Repos did not exist a year ago
- UK Agentic Repos amongst top in world
- Agentic Repos created by individuals, collaborative, and community driven whilst top models are corporate built
- OpenClaw, created in London, has own foundation (see case study 2.11.3) second most popular topic on GitHub, replaced UK’s AutoGPT as fastest growing

Sovereignty

- Digital sovereignty is a key driver of open source AI policy, EU €3B investment and open source policy as part of sovereignty Strategy - explored in Thought Leadership
- Open source recognised as cornerstone but its globally collaborative nature and fact anyone can use it for any purpose on the outbound not just inbound often missed
- Layered sovereignty, current thinking, following five tier “NVIDIA Layer Cake”
- US export control on Claude Fable and Mythos crystallised geopolitical concerns
- Country models -debate whether from scratch of customise, shift for some to small models
- UK’s industrial strategy recognises AI’s importance and new PM shifting UK AI thinking to this, AI Minister role in UK escalated to Cabinet position

1.2 Introduction - Three Years on from the Opening of Llama 2

Professor Amanda Brock
CEO, OpenUK and OpenHQ



As work on this report came to its conclusion, Chinese leader Xi Jinping reiterated China's ongoing commitment to open source in AI in his keynote speech to the Shanghai World AI Summit, re-confirming eight years of concerted strategy in open source. China now sits at the heart of AI Openness. He also launched the WAICO AI Foundation with 29 country members focused on the Global South. It will inevitably work alongside China's national open source foundation, the Open Atom Foundation - the guardian of China's open source embodied to enable open source leadership, but leaves concerns as to a split in Global AI and its governance, despite a UN presence at its launch.

In the UK, new Prime Minister, Andy Burnham MP, committed to regain Britain's stability through re-industrialisation by putting the right values and standards at the heart of government and elevated the AI Minister's role to Cabinet level reflective of the importance of AI to society. The right values and standards are also needed to take back our digital and AI futures.

China as an open source Leader

China's rise in open source and AI did not begin with DeepSeek's release of open weight model, R1 in January 2025. Although for those who don't understand open source. DeepSeek's shockwave may have felt that way. DeepSeek's iterative AI innovation, distilling existing open technology, evolving innovation, reducing overhead in training and compute, and enabling a model that could be freely adopted at scale, was a critical moment in the history of AI and of AI Openness. It followed a tradition of open source building iteratively on others' innovation and giving-back. It was critical for China's AI and geopolitical positioning as an AI leader.

China's journey had begun publicly and strategically in 2018, was solidified in the 14th Five Year Plan in 2020, then reaffirmed and evolved in 2025 in the 15th plan to remove proliferation and focus on monetisation. Xi Jinping's 4 point policy reaffirms this commitment. The contemporaneous launch of Kimi's K3 model with comparable metrics to the best US proprietary AI apparently has the US in a muddle, or as Business Insider describes it, "Every so often, a Chinese company releases a new AI model, and Americans freak out." Several weeks of rumours of an Executive Order on open models have been replaced with murmurings of "soft power" being used to stop those in the US using open models from China.

China's approach to open source has not been novel. In many ways it mimics the open source software success of the US - with a foundation (Open Atom in China, Linux Foundation in the US) holding assets, managing governance, engaging communities of contributors and users, building standards and organising training and events. All are necessary artefacts in the soft power of open source leadership. All have been managed in a central, skilled organisation that manages the broader infrastructure and manages the innovation (this small part being the lab role others often mention). The roles of outdated public sector Open Source Program Offices, interacting with internal stakeholders and external communities, and the progression to the funding status of a German Sovereign Tech Agency, being encompassed under a single digital roof, to support the public and private sector's needs.

There is however one critical difference in approach between China's AI models built and the open source software that the US has built (with the help of the global open source communities) to form our global infrastructure despite both being open source. The software was built by individuals, communities and innovators who formed collaborative communities of expertise whilst models are generally created and shared by companies as a completed innovation to be iterated on by those communities. In these we are seeing the new models of openly sharing, iterative development and collaboration forming for AI.

The US Response

In the last three years the scale of US AI companies, the investment in them, and associated valuations for companies like OpenAI, Anthropic and NVIDIA have simply been unprecedented.

The US companies are in the main closed or proprietary AI providers. But they too have recognised that the infrastructure, Agentic Harness and Dev Tools - the layers in the Stack other than their models - must be open. As a consequence they have invested in - and in some instances may have apparent control over - new US-based foundations to steward these open tools and standards.

Sovereignty

Geopolitical shift, highlighted by the change in the US administration, has crystallised concerns. Export control exerted through Presidential Executive Orders rather than typical Federal or State Regulation, has seen the US flex its AI muscle and rock the world, as it denied access to chips for China and certain AI for non US citizens. The perhaps unintended consequence of this has been to drive users to China's open source.

As China's models have met the pace of innovation of the proprietary US frontier models, the US's strategy has shifted into turmoil. Rumours abound of a ban on open source models, or a Presidential Executive Order to prohibit Chinese models or perhaps, and most likely, Presidential exertion of soft power to stop the US usage of Chinese open models in the US.

Should we have held off on this report? Possibly, yes. But if we had it's certain the next AI Openness highlight (aka drama) would be swinging right at us.

And whilst the US worries about China and the rest of the world worry about control, dependency, access and the ability to serve their citizens with a safe AI future with their current US tech dependency, we see policy makers recognise their need to use open source to break these dependencies. Ensuring that they understand the fundamental principle of open source has become critical.

Subject to law - anyone can use the licensed outputs of open source for any purpose. What that means is commercialisation enables your competitors with your innovation. For nation states that means global enablement by open sourcing your innovation. Adding geo-political restraints into the licensing as with historic attempts to restrict commercialisation of open source software, will amount to open washing. For governments today open source as an answer to sovereignty and their open source strategies, policies and funding is about the "how" not just the what.

As I have said repeatedly, "open source" is not a magic word. There's a nuanced tried and tested formula that sits behind it and only those with years of experience in this space truly understand some of the nuances around that. Open source strategies will only work if the how is right as well as the what.

Building tomorrow's AI not chasing yesterday's

To enable an AI-future it appears today that business strategies and monetisation of AI should focus on evolving technologies and new approaches leveraging the existing models as a base. The US's frontier models have been met head-on by open weight models like Kimi K3. Size may well no longer matter, and as we hear from James Drayson in his case study with Locai Labs, may be key customisation is key. As of last week, he's in good company with Mira Murati (ex Open AI Strategy Officer and) Thinking Machines Labs Founder releasing her Inkling Model saying she didn't care if it wasn't the biggest, its ability to be customised was what mattered. Why? Because this and the agentic harness soothing models into production looks likely to be where revenue will be generated - not in the bloated models being offered today at unsustainable rates using unsustainable data centres.

Learning from the Data

And this is where this team and our partners' understanding and analysis of open source data is key. The data tells us that the UK is not well served by throwing money at building new large models.

Rather we should be leveraging future opportunity and forward-looking business models. The opportunity and the growth for the UK can be seen in the model-adjacent technologies - the Agents and Agentic Harness where open source leaders OpenClaw and Auto GPT were both founded out of the UK; in customisation of models with a sharp focus on return on investment, recognising as the US may too be, that being first mover can be an expense and not a benefit; and in the Developer Tools from companies like Tessl where our world-class software engineers here in the UK are able to provide the AI Native basis of de facto standards of governance and enable businesses to industrialise the use of AI. There is after all the small matter of MCP being entitled by birth place to a British Passport.

Frankly our leadership is clear when we look to GitHub. We recognise that whilst models may be primarily shared by large companies on Hugging Face, the open source repositories of AI that's evolved to manage and monetise, sit with the innovators on GitHub. And when we look at GitHub we see a different story. A story that offers opportunity for the UK, if it can bring itself to listen and understand how open source actually works. This same story offers an ending in successful industrialisation and global market positioning.

Capturing value in a UK “foundation”

The UK's innovation has been shipped to the US and monetised for decades, with the UK paying for an innovation pipeline, along with colleagues in other middle nations. The UK has been overly generous in this practice for too long. Systems thinking needs applying to the UK's planning in the way China planned its journey to open source success. This success is enabled by the right infrastructure which must be mapped. It is obvious that the UK needs to build the “Foundation” necessary to manage and deliver the technical and governance hygiene, ecosystem and skills that enable AI openness.

Ensuring a stable future requires more than independence from international companies. It requires us to effectively manage our own technologies here in the UK. We must stop giving those we develop to others to hold, manage and monetise. The opportunity from the open source submarine that powers the digital economy and increasingly the AI one, is right in front of the UK. This is not about sharing the glory of AI but leading it, with sustainable infrastructure and AI. That will be achieved by building the right technologies, using AI Openness and standards and leveraging the UK's leading ecosystem and its knowledge of how to do open source well.

Few countries in the world have the ability and the opportunity the UK has today.

Three years on from the launch of Llama 2 as open innovation - a bold move by Meta that OpenUK was the only open tech organisation to support - we support another brave move. As our new Prime Minister promises to reindustrialise the UK with AI, to take back control of our technological future and let's hope that he recognises that this can only be achieved through open source and AI openness.

23 July 2026

1.3 Take the OpenUK & OpenHQ Survey, in association with The PSC

Wherever in the world you are, whatever sector you work in or even if you are a student or do not work, please give 30 minutes to complete this survey and input into the data building economic modelling of the value of open source.

For the first time in 2026, this survey also includes international participants and AI. Feel free to share this with colleagues across any sector with knowledge of their organisation's digital capacity. Thank you for your support.



<https://openuk.uk/report/survey/>


2. Three Year Evolution of the AI Revolution

On the third anniversary of the launch of Llama 2, we see the release of China’s Kimi K3 model, which tracks its utility against the largest proprietary or closed frontier models, shifting the technology of AI openness from a to competitive parity with closed AI models.

This section tracks the journey from the opening of Llama through those years , including not only AI models but also the Agentic AI, Dev tooling and the impact of AI on open source software in that time-frame and the evolution of policy and regulation associated with it.

2.1 Launch of Llama2 as Open Innovation

OpenUK partnered with Meta on the [launch of Llama 2](#), as “Open Innovation”, launching on 18 July 2023. The Open Technology remit OpenUK has is broader than pure open source, which enabled its Board to unanimously decide to support Meta’s launch of “Open Innovation”. Perhaps not as open as they would have liked it was at the time a massive step in opening up AI and remains one of the most important and influential steps in the history of AI openness.



The first Llama model had been made available on a restricted free licence for research only. When the model was leaked within a week of being shared on March 2, 2023. The pace of innovation by the open community in the ensuing weeks created shockwaves across proprietary AI model providers, particularly in the US.

"We support an open innovation approach to AI. Responsible and open innovation gives us all a stake in the AI development process, bringing visibility, scrutiny and trust to these technologies. Rather than limiting large language models and AI progress to a select few, we should bring AI development into the open where we can all benefit from this technology."


Llama Launch Website, July 2023

The sentiment was perhaps best captured in a leaked memo from Google saying “we have no moat”. Without the barrier of the unprecedented costs (compute, human and power) of building large language models, and having access to such a model, the open source communities were able to creatively innovate at pace. This had the potential to leave the proprietary vendors unable to differentiate their products from the freely available ones and raised concerns as to customers’ willingness to pay for these products without the benefit of a proprietary intellectual property and innovation “moat”.

Llama 2 is distributed under the new [Llama Community License](#) which includes restrictions that do not meet the pre-existing Open Source Definition (OSD). Notably they require anyone with over 700m monthly users to seek a commercial licence from Meta, on undisclosed terms and prohibits using the model’s outputs to improve other large language models.

It also includes an Acceptable Use Policy. There was no publicly available alternative open source definition or AI definition at the time of launch. Llama could not legitimately be described as open source and open innovation described it well. Meta’s founder Mark Zuckerberg announced Llama 2’s launch on Facebook as open source and a battle began between Meta and the open source community.

Llama’s distribution was driven by Yann LeCun, Meta’s chief AI Scientist running Meta’s FAIR labs. Llama was relatively successful in no short part due to significant marketing and policy funding to support it and many reports being commissioned across tech policy organisations. LeCun left Meta in 2025 setting up Advanced Machine Intelligence Labs (AMI Labs) focusing on “world models”. Through his engagement with the AI Alliance, he has driven a goal of building a geographically collaborative open model. Meta replaced LeCun with Alexandr Wang as Meta’s first-ever Chief AI Officer, leading the newly formed Meta Superintelligence Labs. Wang’s models are to date closed.



The opening of Llama, despite some of the ensuing challenges, has however not only been hugely significant to AI but also to openness and open source.

2.2 The Great Debate defining Open Source AI

What open source means in the context of AI

The dispute over what “open source” means when applied to AI is important from an accuracy and policy perspective in particular. The OSD allows distributors and users of open source to rely on innovation licensed under it being usable by anyone for any purpose in perpetuity. When restrictions are introduced these are often, like those in the Llama community licence, driven by commercialisation. On occasion they are not commercial but political or geographically driven.

Any restriction on who can use what’s distributed or the purpose for which it can be used stops the technology distributed being open source. Wrongly describing code that is shared on a restricted licence open source software has long been described as “open washing”.

This cascade of sharing and iteratively developing on top of another’s shared code and resharing that output is fundamental to the success of open source software. It’s at the core of the benefit or value of open source.

Open Source AI Definition

The Open Source Initiative reacted to the potential inadequacies of existing licensing and the OSD by creating the Open Source AI Definition (OSAID), published in October 2024. It requires the freedoms to use, study, modify and share a system, including release of the training code and parameters and enough information about the training data to recreate it. Developing this early stage in the AI evolution, it unsurprisingly focuses on requirements of models. The definition is demanding enough that very few systems meet it. Meta’s Llama nor Mistral’s models and even DeepSeek’s R1 do not.

The definition’s release was contested, as not broad enough and the Open Weights Alliance responded with an Open Weights Definition at the 2025 Paris Action Summit. It guarantees only two of the four freedoms, to use and to share, positioned as a complement to open source rather than a substitute for it. For some its very existence is problematic, citing the potential to undermine the existing OSD.

Only 25 organisations, including Mozilla, SUSE and EleutherAI and 170 individual supporters, from academia, consultancies and public bodies have endorsed it according to the OSI’s web page. Developers are largely absent and many of the endorsers were involved in the definition’s creation. Validation is narrower still. Five models have passed OSI’s validation.

Disaggregation Approaches

This alternative approach treats openness as a matter of degree rather than definition and has manifested through different methodologies. For many it is preferred to the inflexibility of definitions.

Model Transparency Index

[Stanford’s Foundation Model Transparency Index](#) first released in October 2023, scores leading developers against 100 indicators. The initial mean was 37/ 100 rising to 58 by May 2024 when self-reported, falling to 41 in December 2025 under supervision. IBM was top at 95, DeepSeek at 32 and Alibaba at 26. Only 30% of the companies were willing to report transparency metrics, compared against 74% a year earlier.

**“We have open-weight LLMs
but very few fully open LLMs”**

**Jean-Baptiste Kempf,
Founder and CEO, Kyber;
President and Lead VLC
Developer, VideoLAN**

Open Source AI Index

[Radboud University](#)’s Liesenfeld and Dingemanse assess model openness by disaggregating 14 components of AI (models) marking each as one of: open; partially open; or closed. Their 2024 survey of more than 45 Models found that most models described as open were open weight models at best. Today it’s the [European Open Source AI Index](#), breaking down 123 models as at 3 July 2026. It’s topped by Yu-Lan-Mini, BLOOM, OLMo 3 and Apertus, followed by Llama, Mistral and Qwen well below them.

Model Openness Framework

The Linux Foundation’s work on defining the completeness and openness of AI Models manifested in the [Model Openness Framework](#). This expanded into the [Open Model, Data and Weights \(Open MDW\) Licence](#), aimed to cover AI in compliance with the EU AI Act. NVIDIA has notably adopted the OpenMDW license framework across its Cosmos, Nemotron, and Isaac GR00T model families.² It has also previously established a [Community Data Licence Agreement \(CDLA\)](#) created in conjunction with Microsoft, seeking to meet data requirements.

A Shift in the Language of AI Openness

By the [Paris AI Action Summit in February 2025](#), there was a shift in the terminology and language of openness itself. A marked movement away from discussing “open source AI” and use of alternative terms including “AI openness”, “open models” and “public good” or “public interest AI”. This change was reflected in the re-titling of OpenUK’s reports to AI Openness encompassing this wider group of terms as they evolved without the distraction of debate on the meaning.

This shift in language recognised that the focus to date had not been on AI as a whole but on models, and an increase in societal concerns around ownership of and access to the primary sources of AI with the public good focus.

The G7 and AI Openness

On 29 May the OSI presented to the G7 Digital Ministers a new [Vision on AI Openness and open source](#), seeing a shift in the OSI’s thinking to this broader terminology of AI Openness.

The G7 Ministers recognised that “Open Source software has become a cornerstone of our economies, substantially reducing software development costs, serving as the engine of discovery by enabling reproducibility, accelerating research, and lowering barriers to participation for researchers worldwide.” They go on to state that “AI poses new challenges in terms of definitional clarity: the meaning of Open Weight or Open Source AI remains disputed, and while the traditional definition of Open Source provides pointers, further clarification is needed.” Expanding they recognise “The Benefits of AI Openness, open weight AI models and related open source software can have positive micro- and macro-economic impacts, including on Gross Domestic Product.”

We can expect to see further collaborative work around definitions through 2026 and into 2027, and perhaps consideration of the popular disaggregation approaches. Perhaps finally ending the great debate.

Open Washing

Over the last 3 years concerns of “open washing” in AI have moved into the mainstream.

For any commercial organisation open sourcing its technology effectively enables its competitors with its own innovation - there being no possibility of a commercial restriction in a true open source licence. Historic open washing in software and the evolving AI open washing debate centred around such restrictions. The debate reached mainstream press in 2024 with the [The New York Times](#) and [The Economist](#) flagging Meta’s alleged open washing when calling Llama 2 open source not open innovation.

The OSI does not hold a trade mark on the term “open source” so has no legal basis to enforce the term’s misuse. The open source ecosystem however calls-out attempts to take advantage of the term open source. When this has happened in open source software the offender has generally been accused of “open washing”.

2.3 The Deep Seek Moment and the rise of China's AI

Model growth and investment leapt from unprecedented levels to even more unprecedented levels in the US. A few companies with the wherewithal built these models and concerns of their potential control of global AI grew. On 20 January 2025, apparently out of nowhere, DeepSeek released [DeepSeek R1](#) “distilling” large language models, including Meta’s Llama, claiming its production cost 20% of the \$100M a model took to create at the time. The models weights were released on a MIT open source licence. Whilst it was criticised as raining data was not shared, the quality of the supporting documentation allowed Hugging Face to build its own version [R1 Open](#), within days.

The release of this model began China’s open source and AI leadership, crashed NVIDIA’s share price by 17%, and sent shock waves across the AI industry. Rumbblings of AI’s use as a geopolitical tool grew and the geo-policy battle began.

This development - a new form of innovation in distilling models - is a classic example of open source at work and the iterative development it enables.

2.3.1 China's Strategic Approach and Policies

[The 14th Five Year plan](#) from 2020 included open source as a national strategic pillar to achieve technological self-reliance and global leadership in AI. This included the establishment of a national foundation for open source, the [Open Atom Foundation](#). [The 15th Five Year Plan](#), released end of October 2025, saw a shift recognising model/ technology proliferation and the need for Chinese technologies to compete less with each other, with a concerted shift to monetising the technology.

This was reinforced in July 2026, in Xi Jinping’s speech at the World Artificial Intelligence Conference ([WAIC 2026](#)), when he re-committed to open source and named it as one of the four pillars of China’s policy, recognising openness as a route to inclusive AI development.

He called for joint development of open source ecosystems and open international communities, with independent choices for enterprises, and intellectual property protection. He presented China’s open models as a public good and pledged resources to developing countries. Open weights shifted for China from simply being leadership to a clear instrument of foreign policy and the summit saw the announcement of the World Artificial Intelligence Cooperation Organisation (WAICO) with a secretariat in Shanghai. Open weights shifted from leadership to a clear instrument of foreign policy for China.

2.3.2 China's Technology

Model proliferation in China comes from the sheer scale of open weight models that are frontier-adjacent capability with models distributed on open source licences. Qwen has become one of the most widely used open model families outside the US, with more than 200,000 derivative versions on Hugging Face, under an Apache 2.0 licence. Alibaba has begun holding back its newest flagship, serving the [Qwen 3.7](#) generation through its cloud API rather than on an open source licence but the earlier Qwen 3.5 and 3.6 continue to be available as open models. Moonshot AI has taken the opposite course, releasing [Kimi K2.6](#) as open weights but under a modified MIT licence. [Kimi K3](#) released on 18 July, the third anniversary of Llama 2’s release placed fourth for intelligence according to [Artificial Analysis](#), tracking close behind Claude Fable 5 and OpenAI ChatGPT 5.6, but ahead of Claude Opus 4.6 and ChatGPT 5.6 Terra.

2.4 The US Response

Having led in AI model evolution from ChatGPT4 in March 2023 through to the opening of Llama that summer the US continued to appear to be the very clear global AI leader. The release of China’s open weight R1 model in January 2025 caught it off guard and the ensuing onslaught of open weight models from China and their increasing sophistication have rocked the US’s leaders from a technological and policy perspective.

2.4.1 US Executive Order Approach and Policies

Over the three years since Llama’s opening there has been a shift in the US Administration which saw the Biden-era frontier and security driven concerns of catastrophic risk with a proliferation of State-level engagement move on 20 January 2025, to President Trump’s market-driven approach and support for US

companies. Vice President JD Vance was clear in his speech to the Paris AI Action Summit in February 2025, that the US was number one in AI and planned to remain so. The current Administration can be seen to have used its leadership in AI as a geopolitical lever, supporting the country’s technology companies to lead around model development and data centre investment. Within weeks of taking office, globally applied tariffs caused turmoil and accentuated global AI sovereignty concerns.

Presidential Executive Orders

The longest Executive Order in US history, President Biden’s Executive Order 14110, from October 2023 Agencies focused on model evaluation and guidance on AI models for creators, with private sector engagement through developer red-team reporting and IaaS foreign-training disclosures. One of President Trump’s first actions post inauguration on 20 January 2025, was to rescind EO 14110, displaying a 404 message on the White House web site and calling it “dangerous” and “unnecessarily burdensome”. Coincidentally this was repealed on the same day as China’s DeepSeek released R1.

The US AI Safety Institute was rebranded the [Center for AI Standards and Innovation](#), reflective of the upside opportunity for enterprise and the revenue driven approach of the new President who has built close relationships with the Big Tech companies. The [National Institute of Standards and Technology](#) (NIST) which had led on Software Bill of Materials (SBOM) saw an exodus as did many public bodies. For NIST this included Dr Alan Friedman who had led on SBOM.

A risk of Republican-driven export controls on intangibles, weights, data and research, was recognised after DeepSeek’s release shifted the US debate from catastrophic risk to what the President termed “China risk” and with a clear agenda to drive US leadership and innovation.

[Executive Order 14179: Removing Barriers to American Leadership in Artificial Intelligence, 23 January 2025](#) revokes earlier regulation that acts as a barrier to innovation and takes a pro-innovation approach, with the goal of decisive action to enable the US to lead in global AI. Later referred to by the President as revoking “my predecessor’s attempt to paralyze this industry and directed my Administration to remove barriers to the US AI leadership”.

[Executive Order 14365: Ensuring a National Policy Framework for Artificial Intelligence, 11 December 2025](#), opens by recognising that “leadership in Artificial Intelligence (AI) will promote the US national and economic security and dominance across many domains”.

[Executive Order 14409: Promoting Advanced Artificial Intelligence, Innovation and Security, June 2, 2026](#), established an intentional voluntary framework for frontier models to be shared with the US government for security and cybersecurity testing, up to 30 days before release. The legality of distillation of models has been called into question by the US. It’s often seen as a somewhat ironic statement when many across the world argue that model training itself - the training of the US models later being distilled by China - is inappropriate use of copyrighted works.

There is also a Presidentially proposed draft AI Act from early 2026 in contention, which includes a clear shift to enforced AI governance for the US.

Export Control Directive and Anthropic Claude Fable5 and Mythos5

Anthropic Founder, Dario Amodei was public and vocal in his concerns around the power of frontier models. Claude’s Mythos model can seek out and find software vulnerabilities, as well as working on how to exploit those vulnerabilities automatically. Access to the Mythos Preview was given via Project Glasswing from April 2026 to a limited and selected group. The disclosures surrounding Anthropic’s Mythos Preview and Project Glasswing were viewed as both corporate marketing and a reflection of genuine risks to critical infrastructure.

[Anthropic’s Blogpost](#) reported that the US government had, based on evidence of a potential narrow, non-universal jailbreak, exerted an Export Control Directive on 12 June 2026, against advanced AI models Claude Mythos 5 and Fable 5. These models had export control bans placed on them restricting their usage to US citizens. Anthropic therefore had to abruptly remove access to those models globally. Days later, on 24 June the order was partially lifted allowing distribution to a select group of US companies and further on 30 June, the export control ban was lifted, apart from with respect to China.

Yann LeCun's response to the export control restrictions was telling:

"Dario Amodei's ridiculous fear mongering about Mythos/Fable (and AI in general) finally pays off."

Yann LeCun, Founder and Chairman, AMI Labs

OpenAI's Terra and Luna Models are also being shared with only a small group due to governmental pressure.

US State Regulation

The US State reaction has seen over 1,000 State Bills including California's vetoed SB 1047, a delayed and then replaced Colorado Act, [Texas TRAI GA](#) and California's Frontier AI Act which took effect Jan 2026; New York's RAISE Act signed Dec 2025; whilst Cruz's [10-year moratorium](#) was defeated near-unanimously. A frenzied proliferation of regulation which has to a large extent become overshadowed by the Presidentially-driven regulation and policy.

Response to China

Despite the earlier "China Threat" rhetoric, the US has not formally responded to the Shanghai Summit, or open models, but the Department of State has warned against what it as a potential export of "techno-authoritarian governance". At the time of writing rumours abound as to a potential Executive Order banning open models, which is generally assumed to be too extreme, a more measured Executive Order, limiting or banning use of Chinese open models or and perhaps most likely exertion of soft power to minimise or remove open Chinese model usage in the US.

US soft power can also be seen in the build of US Foundations such as the Agentic AI Foundation, a privately held not for profit, seeking to share the standards and de facto standards of AI by opening these and preserving them in a Steward.

2.4.2 The Technology

The shockwaves from R1's release in 2025 and reduction in required compute saw Chip Manufacturer's NVIDIA share price drop momentarily by 17% and they were such that OpenAI's founder Sam Altman (who had originally set OpenAI up with the help of Elon Musk, as an open source foundation) express a personal opinion,

"I personally think we have been on the wrong side of history here and need to figure out a different open source strategy."

**Sam Altman,
CEO, Open AI, 30 January 2025**

By June of 2025 the first OpenAI Open language models, [gpt-oss](#) were released.

On 5 November 2025, [the FT reported](#) Jensen Huang's view on the AI race.

"China will win the AI race with the US"

**Jensen Huang,
CEO, NVIDIA**

Shortly thereafter Nvidia itself began open sourcing spanning the [Nemotron family](#) for agentic AI, the [Cosmos platform](#) for physical AI, [Alpamayo family](#) for autonomous vehicle development, Isaac GR00T for robotics and [Clara](#) for biomedical claiming to empower companies with the tools to develop real-world AI systems.

The US has largely continued to build proprietary models with a few notable exceptions like the work of [nous AI](#). The proprietary models have however been accompanied by a shift to governance through open source standards and tooling, with the infrastructure layer being open source. As we moved to report launch OpenAI made public one of the first autonomous Agentic hacks from an unreleased proprietary model. In response Hugging Face's co-founder Thomas Wolf explained that the closed proprietary nature of the model was limiting the response from defenders, unable to access the model expressing a view that open sourced models were becoming more secure than closed due to access and transparency, which it appears will ultimately engender greater trust.

2.5 The European Union

The European Union (EU) took early leadership in an attempt to regulate AI, building a dense and much criticised Act, which has been implemented but the enforcement of this is doubtful. A little like the OSAID may simply have been too early an attempt to build structure in emerging AI landscape.

2.5.1 The EU AI Act

The world's first AI regulation, the EU AI Act came into force on 1 August 2024. It places restrictions on providers, importers, distributors and deployers of AI systems. It includes specific provisions on open source with provisions giving providers of reduced liability. Following a great deal of criticism including within the Commission that this dense regulation was having the impact of inhibiting innovation, the [Commission published a Digital Omnibus](#) proposing to simplify the Act, and the Council and Parliament reached political agreement on it on 7 May 2026. The Agreement defers the most demanding obligations, with high-risk systems under Annex III to 2 December 2027, and high-risk AI embedded in regulated products under Annex I postponed to August 2028. It remains unclear whether the Act will be enforced.

2.5.2 EU Sovereignty and Open Source Strategy, 2026

The EU has a long history of focus on the US's Big Tech and policy direction since 2020 has had a strong sense of its concerns about dependence on Big Tech and the US. A series of competition cases in the EU involving Big Tech have also magnified this friction. Long before President Trump's second term in office US media was reporting anti-American EU sentiment as a consequence of this and sovereignty in the EU can perhaps be viewed as the greatest point of friction with the US outside of the US China relationship. By 2025 there was recognition across the EU that its current actions - lawsuits like the Schrems data privacy litigation, regulation and business activities like Gaia X were not securing the EU's sovereignty.

It published an [AI Continent Action Plan](#) in April 2025, set out to make Europe a leading AI continent through public compute, with a network of AI Factories and up to five Gigafactories, an InvestAI facility of some €20B, and a commitment to triple the Union's data-centre capacity. It was followed in October 2025 with the [Apply AI Strategy](#), which pairs a drive for adoption across strategic sectors with an explicit preference for European and open source solutions and a "Buy European AI" approach in public procurement. This was supported by a [Cloud and AI Development Act](#) to strengthen competitiveness in cloud and AI ecosystems, and a [Data Union Strategy](#) intended to build a single market for the data on which European models would be trained.

Read against the NVIDIA five-layer Stack the strategy concentrates on the lower layers of the stack, compute, data centres and data, where dependence on non-EU suppliers is hardest to unwind, and it treats open source as an instrument of sovereignty rather than a threat to it.

On 3 June 2026 its [EU Tech Sovereignty Package](#) was released as a key to strengthening competitiveness, resilience, and strategic autonomy in a fast-changing digital world. This included the [EU Open Source Strategy](#), which was in response not only to sovereignty but to the EU's January open source consultation which received over 1,600 responses - more than any in its history. The open source communities have lobbied the Commission and Parliament with some vigour for 20 plus years and the EU has for many been open to an open approach setting up an OSPO for the Commission in 2019 and building to this has been a long process. The question now, with a commitment of €3B is how this will be done. Much of the how can be seen to align with the proposals OpenUK has made for a UK Foundation.

2.5.3 The Technology

France's Mistral has become Europe's AI model darling, distributing open weight models under an Apache 2.0 licence. In 2025 it announced the build of an industrial cloud with NVIDIA. Multiple specialist models including Black Forest Labs have emerged to some acclaim. The EU has actively invested in Super Computers, FABs and Chip / silicon in an attempt to actively shift its supply chain dependencies.

Euro Stack was instigated in 2025 as a set of initiatives aimed at building the EU's digital infrastructure. It has met a mixed response due to its anti-American rhetoric. It did not put a structure or foundation in place and this led to confusion around its make up and who spoke with its voice. Two of its participants are in litigation over the Eurostack name.

2.5.4 The EU Open Source Policy

Sachiko Muto
Chairwoman, Open Forum Europe



On 3 June, the European Commission published its [Open Source Strategy](#) as part of the Technological Sovereignty Package. The strategy places open source at the very centre of Europe's push for digital independence yet for anyone who has celebrated high-level policy endorsements of open source before, and then watched the outcomes disappoint, the temptation is to file this one alongside the rest and wait to be proven right. The endorsements have been many. Malmö in 2009, Tallinn in 2017, Berlin in 2020, the European Interoperability Act in 2024 - a steady procession of ministerial signatures affirming that open technologies matter. Several member states have gone further and wrote a preference into their own procurement rules. Italy, France and the Netherlands have each adopted some version of "open, unless" - choose open source by default, unless there is good reason not to.

The trouble was always the "unless." [A recent study for the Danish government](#), carried out with Johan Linåker, surveyed 16 of the world's most digitally advanced countries and found open source policies almost everywhere, in one form or another - alongside outcomes that varied enormously. The gap between having a policy and changing what actually gets acquired by public sector organisations turned out to be very wide. Where an escape hatch exists, it gets used. The "unless" box, more often than not, is the box that gets ticked. And yet this time appears to be genuinely different - which is not a claim to make lightly, because making it means setting aside 20 years of well-founded disappointment.

What earns that conclusion is not that open source sits higher on the political agenda, though it does. It is the quality of understanding on display. Past endorsements approached open source mainly as a way to save money on licenses or argued for open source on ideological grounds. This package treats it as an ecosystem to be sustained, and gets the details right in a way that has not been seen before. It talks about maintainers, not just users. It funds critical dependencies rather than assuming they maintain themselves. It commits to building a network of Open Source Program Offices across public administrations, sets key performance indicators and a monitoring framework, commits real budget - on the order of €2B across the wider package - and, crucially, signals that the Commission is willing to lead by example, recognising that not every public body in Europe has the same resources to draw on.

This level of understanding did not come from nowhere. The Commission has been on its own open source journey for a long time, from its first open source software guidance around the turn of the millennium through to the [internal strategy renewed in 2020](#), running code repositories, a solutions catalogue and an open source observatory for years. The Commission's [Cloud Sovereignty Framework](#), published late last year, showed a sophisticated grasp of what sovereignty actually requires - the ability to migrate without lock-in, to operate and maintain from within the EU, to inspect and modify under open licences.

Another factor deserves credit: the mobilisation of the open source community through the passage of the Cyber Resilience Act - a hard, at times painful process that nonetheless taught both sides how to talk to each other, and left the community more fluent in policy than it had ever been. A vestige of this mobilisation translated into engagement with the strategy itself: the [consultation on open digital ecosystems](#) drew more than 1,600 responses, a figure described in Brussels as something of a record.

The scepticism, then, turns out to be incomplete. This is not another declaration. It is the most serious and best-informed articulation of an open source strategy to appear in nearly 20 years. But - and there is, inevitably, a "but" - there is a catch, and it is the one that has always mattered. The Open Source Strategy is non-binding. It is a statement of intent, not a legal instrument. The teeth are elsewhere, and still being forged. The Cloud and AI Development Act - the legislative core of the package - has now entered the European Parliament and Council, where intensive lobbying is certain and a fight that runs well into 2027 is likely. And the revision of the Public Procurement Directive, where a genuine open source preference would acquire real force, has already slipped: expected in early July, it has now been pushed to September. Even the Cloud Sovereignty Framework, for all its rigour, does not commit to co-develop open alternatives where none yet exist - which means it, too, leaves the "unless" box available to tick.

2.6 India

2.6.1 The Technology - Sarvam Small Language Model

India's approach is different with Sarvam founded in 2023 selected in April 2025 by the Indian government's IndiaAI Mission to build the country's sovereign model using dedicated public compute. They delivered three variants, a large model for reasoning, a small model for real-time interactive use, and an edge model for on-device tasks. The company's Sarvam-M, 24B-parameter model was released in May 2025, fine-tuned from France's Mistral Small, a dependency that attracted criticism. Its Sarvam-30B and Sarvam-105B were trained from scratch on Indian infrastructure and released as closed models at the Impact Summit then opened a few weeks with the models under the Apache 2.0 licence. Both models support all 22 official Indian languages and built on a \$50M budget. The Indian government has granted a number of public sector contracts to Sarvam enabling it to leverage these to raise additional funding. A prime example of governmental soft support to enable sovereign AI. The small model approach enables models to potentially run on personal devices, and differentiated from that of many western countries, focusing on the needs of the Global South, recognising smaller resources and limited access to compute and data centre capacity.

2.6.2 Approach to Policy

India's approach focuses on voluntary principles with enforceable rules related to a narrow set of identified harms. The [India AI Governance Guidelines](#), released by MeitY in November 2025 form a reference for responsible adoption. MeitY also based its AI Summit on the same seven "sutras" covering trust, people-first governance, innovation over restraint, fairness and equity, accountability, understandability by design, and safety and resilience, under the guiding principle of "Do No Harm". The New Delhi Impact Summit in February produced a non-binding declaration organised around seven chakras. It was endorsed by 89 countries and organisations, including both the US and China (a notable absentee from the event). India joined the US-led Pax Silica initiative but didn't join WAICO.

2.7 Middle Nations

The power of Middle Nations as a collective was first flagged by Canada's President [Mark Carney's landmark 2026 speech](#) at the World Economic Forum's DAVOS, in which he called for middle powers such as his own Canada to work together to counter the rise of hard power and the great power rivalry, in order to build a more cooperative, resilient world. One year to the day of President Trump's repealing the Biden Executive Order and release of DeepSeek R1, this was framed in geo-political shift amidst US tariffs. His speech and its sentiments caught the imagination of many and sent the sovereignty debate into a spin.

The Tony Blair Institute's February 2026 report "[Open Source: How Middle Powers Can Build Influence in the Age of AI](#)" argues that middle powers, defined as states with strong technological ecosystems that cannot compete at the frontier of model development, are not thereby excluded from the AI race, because most of the economic and strategic value will be captured downstream through deployment, adaptation and integration. It recognises that an effective open source strategy cannot be model-centric. It recommends building openness across the entire AI stack through five components one of which is a flagship open source program, treating tooling and maintenance as critical infrastructure, curated strategic data sets paired with pro-innovation data and copyright rules, procurement that rewards open standards, and open sector-specific benchmarks developed by existing regulators and aligned internationally.

2.8 Canada

The Centre for International Governance Innovation has been leading in Canada saying, "True artificial intelligence (AI) sovereignty isn't about building a Canadian fortress; it's about mastering the open source ecosystems that power the world."

The Canadian government looked to Canadian innovators to build Canadian infrastructure and recognised that this must be open source via a [National AI Systems Sovereign Compute Competition](#), launched in April 2026. On 10 June 2026, Canada released its [National Artificial Intelligence Strategy: AI for All](#) with open source as a key component, arguing that open source can "support external evaluation and accountability, accelerate research and innovation, and encourage greater competition in the marketplace." Canada's Mila is one of the companies participating in the "Mozilla Rebel Alliance" taking on US Big Tech.

2.9 The UK

2.9.1 Policy not Regulation

Government Departments

The UK's AI policy was led by the Department of Science, Innovation and Technology (DSIT) until disbanded on 20 July 2026. DSIT's remit has been shifted back to the Department of Culture Media and Sport from which it was originally framed and the Department of Business, Energy and industrial Strategy (BEIS). At the time of going to press, principal AI bodies that sat within DSIT included !.AI and the [AI Security Institute](#) (AISI) (previously the AI Safety Institute).

AISI has utilised its open source evaluation framework, Inspect, which was open sourced in May 2024, to independently audit frontier models. In early 2026, its testing revealed a massive leap in automated cyber capabilities: Anthropic's Claude Mythos Preview autonomously executed a complex, end-to-end corporate network attack simulation, while an early checkpoint of OpenAI's GPT-5.5 demonstrated a universal jailbreak that completely bypassed its cyber safeguards before public release. The name change was accompanied by a narrowing of scope has stripped critical research into AI-enabled fraud, radicalization, and intimate imagery from the government's current workplan, leaving it without empirical auditing tools during public safety crises.

Sectorial Approach instead of Regulation

The UK has taken a sector specific pro-innovation approach to AI regulation, with regulated sectors for deltas in their applying a principles-based framework to assess existing regulation's suitability and addressing gaps in this. This sectorial approach has been in place instead of regulation, although An AI Bill was announced in the 2024 Kings Speech, but has not come to fruition.

AI Opportunities Plan

Released in early 2025, pre the Paris AI Action Summit, the plan outlined strategic goals and targeted compute capacity, data-sharing and widespread AI rollout across public and private sector. Open source was glaringly missing from this plan.

Compute RoadMap and Sovereignty

The UK's Compute Roadmap was released in Summer 2025 and included an initial strategy on sovereignty which indicates that the UK will not be isolationist but rather take advantage of its unique position between the EU, US and other countries.

Copyright, Data and Text and Data Mining

The UK has an ongoing multi-year challenge around Text and Data Mining where despite multiple consultations taking place over several years and the matter being classed by the Intellectual Property Office as too difficult, and returned to DSIT, neither the Conservative nor Labour Governments have offered a decision on the legality of this in the AI context, thereby casting doubts on the legality of AI model training on public data in the UK. This indecision has been driven by the framing which (wrongly) casts this debate as creative industries versus AI and fails to recognise that researchers and the open source community are also impacted by such training.

In March 2025, OpenAI pulled out of the UK's Stargate funding that had been committed in September 2025, alleging that the cost of energy in the UK was too high and citing regulatory issues believed to refer to the unresolved TDM issues.

AI Playbook

Launched in 2025, the [Artificial Intelligence Playbook](#) was launched with 10 Guiding principles for deploying AI in the public sector.

Sovereign AI Fund

Set up in early 2026, [Sov.ai](#) is designed to manage government funding in AI like a VC. Previously available compute has been repurposed and distributed through the fund.

Open Source Strategy

AI Minister, Kanishka Narayan has been actively supportive of open source, including 10 fellows seconded into the public sector to work on open innovation Llama, funded by Meta. The AI Minister's role has been elevated to the UK Cabinet on 20th July.

A London Tech Week announcement in June 2025, stated that a group of AI Devs under 30 would be forming part of a Board working with policy makers and given access to the Minister on three occasions in the following year. This led to expectations of a more experienced group of open source experts across disciplines with suitable depth of understanding being added to the Board. Whether this advisory group will come to fruition with the change in leadership remains to be seen.

Beyond this open source has been glaringly obvious in its absence in UK AI policy.

2.9.2 The Technology

The UK generally ranks third globally in AI and in AI openness, but a long way behind the US and China's leadership. It does not have a frontier model but the jewel in the UK AI Crown DeepSeek (sold to Google in early 2013) has made many recognisable AI developments and Arm's role in AI has also been recognised and both organisations have significant open source contributions.

Although the UK does not have its own frontier model, smaller model providers include Locai Labs, building open models that are customised to the UK, and the Tessera world model. The government has announced investment in a UK frontier model built by the Lumen consortium backed by Cosine. One of the founding team at Cosine has indicated that this will not be open source.

Cloud company Civo and hyperscaler in the making NScale have made considerable use of open source in their infrastructure.

The UK's significant and experienced open source software community have engaged in building Dev Tools and Agentic technology at pace with many companies being formed including Tessl, nolabs, inc., AutoGPT and BoltMCP. Agentic projects of significance founded in the UK include OpenClaw and AgentMemory. As is identified later in this report these areas may well offer significant ROI.

2.9.3 Case Study: Open weights, owned compute and the case for sovereign AI

James Drayson
CEO and Co-founder, Locai Labs



Background

Locai Labs, its name a play on “local AI”, was founded in 2025 by James Drayson with brother George and father Lord Paul Drayson. It remains a family business. James had intended to build a biotechnology company until ChatGPT’s launch in 2022 convinced him that more consequential advances lay in AI. Brother George is completing a doctorate in Foundational AI at UCL. The business proposition, which Drayson calls “owning your own intelligence”, is that organisations currently “rent capability” from American or Chinese models through API calls, exporting their data from the UK with no guarantee of continued service in return. In his view, this leaves the UK a “digital tenant on international infrastructure”.

Working Under Compute Constraints

Compute is harder to access now than a year ago, Drayson says, citing a 50GW backlog in UK energy grid connections and a general scarcity of GPUs. A British company cannot compete with the US or China in pre-training from scratch. Instead Locai Labs builds on open weight models trained elsewhere deploying them on-prem for customers thereby offering them control of their whole stack.

Arguing that the constraints it is under forces Locai Labs to ask research questions that better-resourced laboratories have no reason to ask, Drayson explains that research condenses models for inference. The firm claims to run up to a trillion parameters on-prem on workstations built around NVIDIA RTX Pro 6000 cards, without the vast training data centres such a model once required. Four-bit quantisation (allowing models to run on lower vram) is in production, one- and two-bit under experiment, and their condensed models hold to within about 5% of full precision.

Fable and Mythos’ Impact on Sovereignty

Drayson describes three positions in geopolitical thinking from governments and policy makers which impact his views on sovereignty and the work of Locai Labs:

The first is an assumption that the US builds the world’s models while the UK concentrates on the application layer, ElevenLabs being an example of this;

The second follows US export control restrictions on advanced NVIDIA GPUs being sold to China. This raises the risk of the same US controls being turned towards the UK and leaves the UK vulnerable with respect to underlying chip provision. This vulnerability has supported UK government funding for UK chip firms Fractile; and

The third follows US export control at the model layer. Following the recent access restrictions to Anthropic’s Claude Mythos and Fable models through the export control imposed overnight via an Executive Order. These were lifted after a month for the UK but even now remain in place for China.

Drayson accepts that OpenAI and Anthropic will likely retain long-term significance for the UK, but argues that for organisations already worried about control, data security, intellectual property leakage and unpredictable API pricing, the three policy positions have converted a hypothetical concern into a reality and a precedent. This has prompted a movement back towards on-prem deployment in the UK and other worried nation states.

James casts this as a correction to what he considers past over hasty cloud migration, given the reach of the US CLOUD Act.

What is Open?

Locai Labs releases its post-trained models under the open source Apache 2.0 licence, along with the training data it constructs, including a dataset of Welsh cultural material. Enterprise work on proprietary customer data is not released. Beyond cost, Drayson argues for openness as an enabler of transparency and accountability, criticises opacity among proprietary providers - such as uncertainty about which model a user is served - and disputes the safety case for closed systems.

He argues that closed systems are unaudited whereas open source equivalents are audited by their communities in their home and international countries. He contends that it is not adequate to rely on such a closed process.

Welsh, Low-Resource Languages And Data Creation

The company talks of national models, produced by post-training an open weight base such as Chinese open weight models or NVIDIA’s Nemotron, training them on British cultural and legal material and the UK’s non english languages including Welsh, Scots, Gaelic and Irish. They differentiate these “National Models” from domain-specific models built for a particular enterprise. Building the Welsh model exposed how thin data from the UK can be. There were few if any conversational-pair datasets in Welsh, so the team used Welsh subtitle translations from open corpora to generate synthetic conversational pairs for post-training. Drayson believed the technique had not been applied to the Welsh language, and other smaller countries have since approached the company to build national models.

He is sceptical of the single general-purpose model, noting that a model the size Fable is reported to be - around six trillion parameters - has no clear constituency. He expects domain-specific models tailored to their audiences will be more useful and cheaper moving forward.

Forget Me Not

The problem occupying most of the firm’s research is “catastrophic forgetting” occurring when new data is introduced to the degradation of existing standards in other areas. When it first post-trained a Qwen model on Welsh language data, Welsh language performance improved but mathematics, science and safety evaluations in the model were degraded by around 50%.

Drayson attributes this degradation or “catastrophic forgetting” where the new data significantly deviates from the base model’s distribution. Their answer is “Forget Me Not”, a framework that mixes the new material used for training with data the model itself generates to reflect that distribution. They train the model on the merged data sets, instilling the new knowledge whilst at the same time retaining the prior standards in mathematics, science and safety. He sees this as the first stage of models’ continual learning, with the commercial benefit that a deployed model accumulates customer-owned usage data on which it can be retrained to answer any objection that on-prem models fall behind frontier releases. This is achieved by embedding knowledge in the weights rather than a context window.

Project Mercury, the Locai One, and Pricing

For national security deployments, controlling the model is not enough, because the customer also needs complete traceability of the training data. “Project Mercury” pre-trains models from scratch on UK-owned compute. This is a deliberately modest 30 to 250B parameters.

“Locai One”, is an AI workstation developed with NVIDIA, configurable with one to four GPUs and ships with an OpenAI-compatible API, a desktop application and a management portal. This workstation will let customers run frontier-grade AI entirely on their own hardware with complete control over data and access.

Critical of token pricing as confusing to buyers, Locai Labs offers fixed-cost owned compute instead.

Government Policies

Drayson views UK AI policy as very different from France’s, where Mistral has had the President as its salesman, and India, where Sarvam has state backing through a contract that supports its funding. He sees the implicit British position as being that a domestic AI company must compete globally and stand on its own two feet commercially.

James notes that token costs at OpenAI and Anthropic are subsidised today in a way that cannot continue long-term, which will also have a long-term impact. He concludes that Locai Labs must be viable with or without State-backing.

Locai Labs has however been awarded compute for research by UK Research and Innovation, including for its work on “model unlearning”, and is in early procurement talks with the Department of Science Innovation and Technology.

James argues for the need for legislation to ensure that pension reform opens up funding for UK businesses to ensure that British institutional capital is leveraged to fund British companies.

Conclusion - Asks for the open ecosystem

Drayson’s message to the ecosystem is straightforward. Transparency is no longer a nice-to-have but a baseline requirement for enterprises deciding which model to build on. By publishing Nemotron’s pre-training data as well as opening the model itself, NVIDIA has set a benchmark that lets customers verify exactly what went into the model rather than take it on faith. While Drayson is clear that few competitors will follow suit, that openness is precisely what turns a capable model into a trusted foundation.

The ask is simple: if you want enterprises to build on your model, show them what it’s made of.

2.9.4 Fireside Chat: Open foundation models enabling a map of the planet

Professor Anil Madhavapeddy

Professor of Planetary Computing, University of Cambridge



Professor Anil Madhavapeddy has shaped open source systems infrastructure for over two decades. He was part of the original Cambridge team that built the Xen hypervisor in 2002 - the first open source type-1 hypervisor, which went on to underpin the modern cloud - and has since worked on containerisation at Docker, the OCaml language and formal verification, and unikernels. Having trained in his early days at NASA, it is perhaps unsurprising that his current focus is planetary computing. He is a key member of the team at TESSERA, an open foundation model for the Earth’s terrestrial surface that releases its embeddings, weights and code openly for anyone to build on.

Can you explain what TESSERA is, how it is used, and why open source is so important to it?

TESSERA is an open foundation model for the entire terrestrial surface of the planet. It takes a year of multi-modal satellite time-series data - radar and optical imagery from missions such as the European Sentinel-1 and Sentinel-2 - and learns a compact seasonal “fingerprint” for every 10-metre patch of ground on Earth.

Those fingerprints can then be used very cheaply, often without a GPU, to carry out downstream geospatial tasks: finding crops, classifying land use, locating solar panels or waste dumps, even measuring forest density. We release the map embeddings globally and annually as CCO, alongside equally open weights and code, so anyone can build on them.

To use it, you download the fingerprints for your area of interest and train a lightweight model on top - for mapping crops, forests, flooding or habitats - using a handful of labels rather than millions. We have seen state-of-the-art accuracy across very different mapping tasks from the same embeddings, so we place a lot of trust in them being accurate.

Openness matters here because a foundation model of the planet that you cannot inspect, cannot reproduce, and cannot train yourself would be a dangerous thing to derive social land-use policy from. The people who need it most - conservationists, smallholder farmers, government agencies across the Global South - are the least likely to be able to pay big money for proprietary access.

And could you share the hedgehog example our CEO loves?

As a fun but important example, we recently used TESSERA to help understand why Britain’s favourite mammal is quietly disappearing. Hedgehogs are now classed as vulnerable to extinction, and rural populations are crashing. With my colleagues Silviu Petrovan, Sadiq Jaffer, Lauren Moore and David Coomes, we are mapping the habitats a hedgehog actually cares about - the “scruffy land” of brambles, hedgerows, and the scrubby green corridors between gardens. Silviu coined the term “digihog”, because we are also using GPS-tagged hedgehogs to train models of how they move, so we can predict where habitat preservation is most needed.

TESSERA uses Barlow Twins. How does that connect to Yann LeCun’s argument that LLMs are a dead end and self-supervised methods are the future?

The bet in self-supervised learning is that you learn the structure of the world from raw, unlabelled observations rather than from human annotation. That is a perfect fit for satellite data, because the planet pours petabytes of unlabelled imagery down at us every day, and it is enormously expensive to label all of it by hand. With self-supervised methods like Barlow Twins, we never tell TESSERA what a forest or a flood is; it learns a representation of the planet’s surface from the raw signal, and we only add human labels at the very end, for the specific task at hand. Barlow Twins, JEPA, LeJEPA and their many variants are all particularly good at building “world models” for physical environments, which is exactly what our geospatial work needs.

Yann argued at the Indian AI summit that the fair way to build open models is for each country to contribute its own data and federate them into one shared model. That may be too late for language models, but it is well suited to satellite data: we already have a public global baseline through Landsat (NASA) and Sentinel (EU), and many countries fly their own higher-resolution task satellites over their own territory. So TESSERA may be the ideal place to actually trial his idea of a global training cooperative, because the incentives line up in a way they simply don't for frontier LLMs.



Claude Fable was used as part of TESSERA. Why was it cut off, in your view, and what has the impact been?

The practical impact is that I am now forced to be reactive rather than proactive.

I maintain a number of open source projects - OCaml, Docker and others - and I receive a steady stream of reported security vulnerabilities against them. It is incredibly frustrating not to be able to run my own scans. Losing access means I have to wait and react to vulnerabilities as they come in, rather than proactively scanning for them myself.

Could this moment be considered a positive turning point?

It can be. Open source is the one form of technology transfer that doesn't require trusting the supplier indefinitely, and that is an enormous advantage when alliances are as fluid as they seem to be these days. The tension is that open source is inherently a global commons. I don't think we resolve that by pretending the borders aren't there; we resolve it by building neutral, shared infrastructure that is valuable enough to everyone that no single nation wants to break it - like the Internet itself. A planetary dataset is a good candidate for exactly that: nobody owns the weather, and everybody needs reliable maps.

What is the replacement, and why DeepSeek specifically? Was there a cost saving or any loss of quality - and does swapping a US proprietary model for a Chinese open one just trade one dependency for another?

DeepSeek is still very expensive to run locally, though the cost of local inference is dropping every day. We had to spend quite a bit of cash to get it working at all, so losing access to hosted models is not ideal in the short term.

And yes, in the long run it does trade one dependency for another. That is precisely why we must ultimately have sovereign models built on open training data - otherwise we are always at the mercy of a supplier. It comes back to the case for open source itself.

Is US policy inadvertently pushing the world towards open source, and towards China's products? Are you inclined to choose non-US commercial technology?

We are "anti-lock-in" more than we hold any geopolitical stance.

As Mark Carney recently noted in his Davos speech, "Middle powers must act together because if we're not at the table, we're on the menu."

Those middle powers include much of the world's Global South, who desperately need every data tool they can get in order to embark on climate adaptation for their populations. Everything we can do to keep TESSERA open is vital to that adoption.

Do you expect your story, and others' experience, to drive more technologists to build on open source - and will that extend to end users?

Every hackathon we have run has been filled to capacity, and our GitHub issues are overflowing with requests. That is simply because building on an open foundation model like TESSERA is the faster, cheaper way to get to a working result, and once a generation of geospatial engineers has tasted that, they don't go back to proprietary tools. Why would they take on all the extra hassle?

The risk we have to manage is community, not adoption. When everyone is coding with AI assistance at high speed, we have to work much harder to keep people talking to each other - mentoring, trading tips of the craft, making sure they actually understand how it all works. Our end users will benefit from the vastly cheaper and more flexible geospatial tasks that TESSERA enables, and that benefit extends to open source use of it too.

What first drew you to open source, and why has it stayed central to your work for over two decades? Has your reason changed?

I was on the original Cambridge team that built the Xen hypervisor back in 2002 - the first open-source type-1 hypervisor, which ended up underpinning the modern cloud. What drew me in was entirely selfish at first: open source was simply the most fun way to learn how the Internet worked. I got started on OpenBSD and PHP contributions in the late 90s, and you could hack on any part of it. It stayed central because I turned that early hacking into a career - in systems infrastructure like Docker, in formal verification with OCaml, and in novel operating systems with unikernels. Nothing proprietary I have worked on has had anything like that reach or longevity.

My reason has changed over the years. In the early days it was about fun and accessibility. In the middle years, around Docker, it became more of a corporate experience - not necessarily worse, because people get paid now. But more recently, working on the planet's climate and biodiversity crisis has turned my position into an ethical one: we must ensure that the poorest people in the world can use open source as a leapfrogging technology to equitably improve their own lives.

In simple terms, what is TESSERA, and what problem does it solve?

TESSERA is a free, open "map-maker" for the whole planet. The problem it solves is that understanding the Earth from satellites used to require a supercomputer, vast storage, and a small army of experts to label imagery by hand. TESSERA does the expensive part once, for everyone, and hands users a ready-made "seasonal fingerprint" of every patch of ground on Earth. After that, answering a huge variety of questions - where the forests are, where the floods will hit, where the hedgehogs live, whatever you can observe from space - becomes about as easy as ordinary programming. And with open coding models now available, that means just about anyone can have a go.

2.9.5 UK Instigated Global AI Summits

Global AI Summits Growing out of Rishi Sunak's first summit at Bletchley in November 2023, little was reported on the second summit in 2024, in South Korea co-hosted by the UK and Korea. Bletchley's narrow leadership audience of 100 attendees saw open source, mentioned a mere twice - by the French and Chinese delegations. OpenUK's CEO attended both the [France](#) and [India](#) Summits and provided AI Openness update reports from each.

2.10 International Policy Making and Governance

2.10.1 AI Action Summit, Paris 2025

By the Paris's Action Summit there was a demonstrable shift to discussing openness. This was evidenced with the launch of the "Public Good" [Current AI Foundation](#), led by Martin Tisne who later earned France's highest honour [Chevalier de l'Ordre National du Mérite](#) in recognition of this work. Professor Camille

"ROOST means open source not open weights. It means coming up with new ways to get new software built as a high priority. The simplest argument for open source is that you can see what the code is doing and can see a bug. This principle has been established for years. More than transparency where you can test, test them again, expand them and make them stronger. There is a disconnect today where AI safety and tooling have been out of step with each other. Hindsight is 20:20 but if you are hiring today you should be considering automation and over-hiring will be corrected. We are making this world-wide as everyone benefits from trust and safety. Generative design and learning systems - when we wrote down the rules we started with long memos. Today, isn't it better to see systems learning?"

Eric Schmidt, Former CEO of Google,
Chair of the Special Competitive Studies Project

François's [ROOST.org](#) looking to tools not rules to manage AI was launched by Yann Le Cun and Eric Schmidt. The [French Summit Statement](#) was however not signed by the US or UK.

Whilst the event engaged with the attendees there was some mention of open source and a focus on public good, open source did not make it into the final summit statement.

2.10.2 Open Source in the India AI Impact Summit

The [India Summit Statement](#) was delayed by a day whilst the content was negotiated, but 89 countries and organisations signed up to it. For the first time a Summit Statement included a mention of open source.

2.10.3 Geneva in 2027

The fifth summit has been announced as being hosted in Switzerland and will take place on 27 and 28 June 2027.

2.10.4 International Organisations

The UN and Global Governance

The UN hosted a global governance day in 2024 and a two day event in 2026, both at AI for Good in Geneva. It convened a group of Independent Experts which did not include figures from open source and their May 2026 output did not focus on AI openness.

The G7 Summit

The G7, under a French Presidency under the influence of President Macron, hosted the [Open Source Initiative at its Digital Ministers](#) meeting in May with an output setting out a [Vision on AI Openness and open source](#) and Black Forest Labs CEO Robin Rombach spoke on open source to Heads of State at the June Summit.

[Rombach urged G7 leaders](#) to back open innovation honing in on the questions that OECD States face: open models mean enterprises and governments can build rather than rent AI, underpinning transparency and trust; and competition and strategic independence. They recognise that the role for governments focuses on safety tooling, evaluation standards and targeted regulation.

2.11 Shifting focus to Agentic AI

In 2024 the UK-based AutoGPT, came to the fore. Written in a month in 2023, by Toran Bruce Richards, the base agentic technology was ground breaking and rapidly established AutoGPT as the fastest growing AI repository in the world. This has been replaced by another Agentic technology built from the UK, OpenClaw. Agents are, like the legal concept, AI innovation empowered to undertake an activity devolved to them on your behalf. When managed wisely they are subject to scope setting and boundaries.

2.11.1 The Agentic Harness

Today's models are "harnessed" by agentic technologies including orchestration, tools, memory, execution environments, permissions, identity, evaluation, observability and governance. Essentially, everything in an AI agent that is not the LLM model is included in the harness, as [Birgitta Böckeler explains](#).

The harness is critical in production, and its measures of success are very different from those of models. To meet the requirements of securing model and agent usage the harness will consider and use identity, permissions, auditability, budget controls and operational oversight across multiple AI components. Only when these are adequate will the harness enable production ready AI products.

"We're in that brief moment where competitive advantage is up for grabs. The companies figuring out how to deploy agents comprehensively - with appropriate design but without crippling limitations - will define how business operates for the next decade."

Toran Bruce Richards, Founder and CEO,
AutoGPT


It is clear when GitHub's data is analysed that this operational layer of the AI Stack is one that is open to smaller companies and individuals to innovate in and in which there is a great deal of room for competition. Monetising the tools in this layer of the stack will be critical. It does not come with the cost burden of models.

For open models Mozilla's report has identified that a smaller percentage make it into production and the value of the agentic harness around these comes to the fore in the securing and enablement of these open models in particular.

In the traditional patterns of open source adoption which are seen to be much more applicable to the tooling and infrastructure software of a harness than to the hugely expensive models themselves, there is also a commonality of tool usage that will enable developers, and likely be built not via Hugging Face but via GitHub, as the traditional home of those open source developers.

2.11.2 Thought Leadership: Security, Identity and Harnessing Agentic

Sal Kimmich
CEO, Gadfly AI



Agentic AI breaks assumptions open source security has relied on for two decades. Code review assumed a human wrote the code. Dependency scanning assumed a human chose the dependency. Neither holds when an agent writes, selects and executes in the same session.

Identity is the first gap

There is no standard way to answer a simple question: which agent, acting on whose authority, made this call. MCP solves tool discovery, not authentication between agents. AGNTCY, now under the Linux Foundation, is the closest answer at the protocol layer, and it is early. Most deployments today authenticate the human and trust the agent by default.

Harnessing is the second gap

An agent with write access to a repository, a registry or a CI pipeline is a privileged actor. Few teams apply the access controls to agents that they apply to a new hire.

The third gap is what agents find

Frontier models scan codebases and surface vulnerabilities faster than any maintainer community can triage. Akrites, launched by the Linux Foundation in June 2026, is the sector's response: shared incident response, one coordinated disclosure process, a maintainer-of-last-resort commitment. It is also confidentiality-first and vendor-funded, which sits uneasily against open source's public-by-default norms.

Identity, access control and disclosure process are the three load-bearing walls for agentic security. Open source has working models for all three. It does not yet have them adapted for agents.

What this means under UK regulation. The Cyber Security and Resilience Bill is not yet law. It has cleared the Commons and is before the Lords, with Royal Assent expected later in 2026 and obligations phasing in after. As drafted, it would bring managed service providers and data centres into scope with 24-hour incident reporting. It says nothing about agents specifically. Start logging what yours do now, ahead of the requirement.

The Bill has been criticised in the House of Lords for its lack of focus on AI. Open source is likely to be picked up in ancillary regulation.

"The organisations succeeding aren't the ones limiting agent capabilities out of fear. They're the ones who have deployed agents that physically can't make dangerous mistakes by design, and so don't need to hold them back."

Toran Bruce Richards, Founder and CEO, AutoGPT

2.11.3 Case Study: The OpenClaw Foundation

Vincent Koc
Chief Architect, OpenClaw Foundation



Vincent Koc is the engineering lead at the OpenClaw foundation. With close to two decades working in technology he sees open source as a thread running through almost all of it, contributing to projects across the ecosystem, from Apache Foundation projects to data and AI tooling used by companies like Microsoft and Google, while working across the Asia Pacific and the US. Alongside that, he spent time teaching and mentoring in AI, which shaped how he thinks about making this technology accessible to all.

An open source pathway

My path into OpenClaw was the classic open source story - I showed up as a contributor, started maintaining, and focused on the unglamorous stuff the project badly needed as it grew: observability, security, keeping the lobsters arriving safe and fresh.

One thing led to another, and alongside a brilliant group of fellow maintainers I ended up becoming a core part of the project. When the Foundation formed and asked me to come on full-time as Chief Architect, it was an easy yes. I'd already been living it.

The Founder

For readers who don't know him, Peter Steinberger, Open Claw Founder is an Austrian software developer. He spent about a decade building a company that made PDF and document software, sold it in 2021, and then hit that strange stretch a lot of founders hit after an exit: financially fine, creatively restless. So he moved to London and started building things purely for fun. Dozens of little projects in a few months, mostly to fall back in love with programming. The last of those projects was a WhatsApp bot. A little AI assistant that lived inside the messaging app he already used every day. That bot became OpenClaw.

A personal challenge

At first, the challenge he sought to solve (or as we say in open source the itch he was scratching) was a completely personal one. He took the bot on a trip to Marrakesh to help him get around, find restaurants, and translate conversations. And early on it didn't feel right. It felt like a tool, all bullet points and tables, when what he wanted was something closer to a friend.

So he told it that...and it changed.

The moment everything shifted was when he sent it a voice message. He'd never built voice support. The agent worked out on its own how to detect the audio, transcribe it, and reply - all in a few seconds.

When he asked how it managed that, it cheerfully told him it had figured it out on its own. That was the realisation for him: these things aren't chatbots with better manners. They can act, improvise, and do real work for you.

Scratching an itch with a great big claw

The problem OpenClaw grew into solving is that everyone should be able to have that. A personal AI that works for you, on your terms, in the apps you already know. Peter was playing around and created OpenClaw in a moment of creativity. That I think shaped everything about the project.

It wasn't a startup with a pitch deck. It was one person playing, and you can feel that in the software to this day. The mascot is a lobster. The community named it "Bots". People send it voice notes like they'd text a friend.

OpenClaw is a personal AI assistant that runs on your own computer and talks to you through the apps you already use, like WhatsApp or Telegram. The difference from a chatbot is that it has eyes and hands. It doesn't only answer questions. It can read and write files, browse the web, check your email and calendar, and carry out tasks for you around the clock.

Why it matters is who it empowers. You don't need to be a programmer. You describe what you want in plain language and the agent figures out the how. We've seen non-technical people run genuinely complex things on it, everything from automating small business operations to tutoring services, things they never could have built themselves before. That's the shift. AI stops being something you chat with and becomes something that works for you.

Why open source the Claw?

Open source was essential to what came next. Making it open is what turned one person's experiment into a community. It meant the project could be supported by everyone and could help everyone, and honestly it's why the project survived its own growth. When something goes from a weekend project to one of the fastest growing repositories on GitHub in a matter of weeks, you need thousands of hands, not one pair.

Because the goal was different this time, open sourcing this project made sense. His previous company was a commercial product, and that proprietary was the right model for that business. OpenClaw is trying to be something else: personal AI that belongs to everyone.

You can't build that as a product one person owns.

The power of community

Opening it up meant thousands of contributors improving it, extending it, and hardening it faster than any company could, and it meant people could trust it, because they can read every line of what's running on their own machine. For a piece of software this personal, that trust matters more than anything.

The Discord community alone is a couple of hundred thousand people. On GitHub, more than 2,700 people have shipped code into OpenClaw, and for many of them it was their first ever pull request, which we're quite proud of. Add over 300 direct sponsors and one of the most starred repositories on GitHub, and it's a big, noisy, wonderful crowd. Current numbers are always at openclaw.org.

A few bumps along the road

The project launched in November 2025 as "Clawdbot", which was a playful nod to the AI model powering it at the time. In late January, Anthropic got in touch with trademark concerns, since Clawdbot sounded a bit too much like Claude. Peter didn't fight it.

The community picked a new name, Moltbot, in an overnight Discord vote. Lobsters molt their shells to grow, so the metaphor was lovely, but the name never really landed, and in the chaos of the switch some crypto scammers grabbed the abandoned accounts and made a mess.

So three days later the project was renamed one final time, properly this time, with trademarks checked and domains secured: OpenClaw.

The honest answer is that the name never mattered that much to us. The ethos is the thing. Open, because that's what it is. Claw, because the lobster stays.

"What that week did prove however, is how resilient the open source community around it is: two renames, a scam wave, and still the project came out bigger than before."

What inspired OpenClaw

"There's been a stream of evolutionary steps rather than one big invention.

The heartbeat was an early one: instead of waiting for you to prompt it, the agent wakes up on its own, checks your email and calendar, and gets on with things. Skills are another, small reusable capabilities that anyone can write and share, so the whole community benefits every time one person solves a problem. There's now a public registry for those, with security scanning we built alongside NVIDIA.

But the big one, the reason it took off, is simply that it made agents work for personal use.

The industry had been promising agents for years and the reality never matched. OpenClaw closed that gap. And it grew up inside the open ecosystem from day one: built in the open, extended by its users, and model-agnostic, so it runs on frontier models and open models alike.

Democratising AI technology

I'd frame it differently. 'Open agents help democratise access to AI, and access here includes the proprietary and enterprise models too. Because OpenClaw gives everyone the same open starting point, switching between models and tools becomes easy. Your assistant, your memory, your setup all stay yours, and the model underneath becomes a choice rather than a commitment.' [Box out]

You can see that playing out already. NVIDIA ships OpenClaw with open models for people who want everything local. Developers in China run it on DeepSeek. Others point it at the big labs. Whoever owns the models and compute, individuals own the agent layer, and that's where your data and your relationship with the AI actually live. That's the counterweight that matters.

A foundation Steward

The Foundation is fully stood-up and scaling fast having launched publicly on 8 July 2026 with a full-time team, a board, and a growing number of partner organisations that the Foundation works with day to day. The best sources are openclaw.org and the launch post on the OpenClaw blog, "Introducing the OpenClaw Foundation."

The Foundation is registered in the US as 501(c)(3) non-profit, serving a global community. It's modelled on Linux and the other big open source projects that have proven they can scale globally and outlast any single company or person. What we're trying to achieve with OpenClaw is bigger than AI agents themselves, and to build something that lasts decades, the entity and its setup had to be right from the start.

Independence isn't a nice-to-have for us. It's the product.

Broadly, OpenClaw is built in three overlapping circles. There's the Foundation's full-time team, who steward the project day to day. There are the core maintainers, drawn from across the industry, who review contributions and guide the architecture. And there's the global community of contributors around them. The Foundation sits underneath all of that as the neutral home, with a board now in place.

Donors

Our publicly named major [donors](https://openclaw.org/donors) include Offline Holdings, the University of Michigan, OpenAI, Microsoft, and NVIDIA, and beyond them more than 300 individuals and organisations sponsoring OpenClaw directly through GitHub. There's additional funding we can't disclose yet because of agreements and a lot of support in-kind rather than in cash: companies like Tencent and Red Hat contribute full-time maintainers, and infrastructure partners keep the project fast and secure.

On sustainability, for us it's about much more than keeping the software running. The mission includes giving back to the ecosystem through education, research, and events, so there's real longevity around this work and not just a maintained codebase.

Balance is deliberately built-in. Our team comes from diverse backgrounds and skill sets, and our benefactors are global, spanning the US, China, and beyond. No single company controls OpenClaw, and the Foundation exists precisely so that stays true. OpenAI has been a wonderful supporter, but they're one

voice among many at the table, alongside Microsoft, NVIDIA, Tencent, Red Hat, universities, and thousands of individual contributors. We balance the scales on purpose, because we're building something meant to last longer than any one company's strategy.

What are the key contributions, code, governance, etc.?

All of it! Code, obviously, from typo fixes to whole subsystems. Design and documentation. Serious security work, including the skill-scanning and supply chain efforts. Companies contribute upstream too: enterprise features, compliance tooling, dedicated maintainers. And honestly some of the most valuable contributions are people simply helping other people, answering questions in Discord, writing guides, running meetups. Our line is that every claw counts, and we mean it.

Get Involved

Join the Discord, that's the front door. Come to an event or a local meetup group, or start one. If you write code, we love first-time contributors and go out of our way to make a first pull request a good experience. And if you don't write code, use OpenClaw, share what you build, and help the next person get started.

What is most needed honestly changes day to day. But the thing I'd emphasise is that you don't have to write code to matter here. Some of the most important work is people supporting OpenClaw in their local chapters, in their workplaces, and among their friends: showing someone what a personal agent can do, helping them set it up, sharing what they've learned. That local, person-to-person work is what we need most, and it's the part no company can do for us.

What would you ask the policy makers

Open source is still widely misunderstood at the policy level, particularly around security and governance. Some policymakers are starting to see the power of it, but there can be a knee-jerk reaction, at various levels globally, that treats openness as the risk rather than the safeguard. In practice it's the opposite: open code can be audited by anyone, and local-first means your data stays on your hardware instead of on someone else's servers.

The bigger takeaway is what access does - democratising access to AI, making it available to everyone in a form they can actually use and own, strengthens the economy, education, and innovation across every sector of a country. OpenClaw is proof that this can happen bottom-up, from a community, without anyone's permission.

If I had one ask of policymakers it would be to please engage with open source AI before you regulate it. Come talk to the communities building it, understand how the security and governance actually work, and write rules that protect people's right to run AI on their own hardware, on their own terms. Openness is the accountability mechanism, not the loophole.

Conclusion

OpenClaw started as a fun experiment, one person trying to solve his own problem and enjoying himself while doing it.

That's not a footnote to the story, it's the whole attitude.

Bringing fun into the work, injecting a bit of yourself into what you build every day, is core to how the team made OpenClaw and something we hope everyone keeps striving for. Solving the mission is good. Building the software is great. But we want to keep the fun in it too.

The lobster is loose, and we plan to keep it that way.

2.12 AI and Open Source Software

2.12.1 Impact of AI on Open Source Software

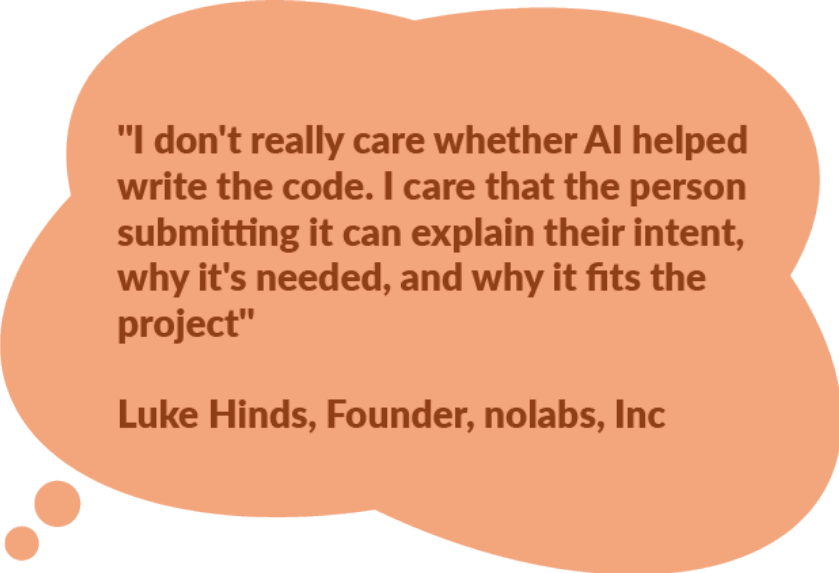
Open Source Projects

Already under-resourced and underfunded by global users reliant on open source for their national and enterprise infrastructure and critical infrastructure, many open source software projects have felt the pressure of a deluge of contribution to projects/ PR requests that are AI generated. The dawn of vibe coding where individuals unable to code are able to submit code that is entirely AI generated has had an additional and significant impact.

The debate as to whether AI tools are detrimental to, or even the end of open source projects versus the belief that AI tools are supportive of more efficient coding and helpful to open source is at its peak. Security challenges over the summer with industry responses through the Akrites Foundation set up and individual company offerings only add fuel to this.

Governance in AI utilisation for code creation is gradually evolving to meet this demand, as it always does, but the pace of AI evolution has required that this move forward with pace. Projects are beginning to set rules on AI generated and enabled contributions, and the Linux Foundation's [AI Coding Assistants Guidance](#) was shared by Linus Torvalds in early 2026.

The UK founder of nolabs, inc., Luke Hinds was [quoted by Adrian Bridgwater in the NewStack](#) saying,



"I don't really care whether AI helped write the code. I care that the person submitting it can explain their intent, why it's needed, and why it fits the project"

Luke Hinds, Founder, nolabs, Inc

This captures the heart of the differentiation between vibe coded software by unskilled individuals and those of a suitably skilled human contributing code benefitting from their use of AI tools the output of which they have discerned adequately to stand behind.

Open Source Dev Tools

Even those companies that distribute closed or proprietary models are engaging in the build of open source tooling for AI Devs as seen in the birthing of new foundations specific to open source standards such as MCP. Dev Tools can be seen to be rapidly growing in the AI space, to increasingly be open source and having an expectation that these will, as with open source software, be a source of revenue generation. Changes instigated by AI tooling are not always as expected and GitHub's [Octoverse Report 2025](#) identifies a 66% jump in the use of Typescript in just one year to make it GitHub's most popular language, seeing humans influenced in development choices by the tooling they choose.

2.12.2 Thought Leadership: Open Source tools for AI

Hannah Foxwell

Co-founder BIMP and creator of AI for the rest of us



Developer tools are the products, platforms and open source projects that reduce effort to design, build, test, deploy and operate software. It's a complex ecosystem with a lot of optionality and diversity.

Shifting Foundations with AI

The foundations of Dev Tools are shifting. Large language models have moved beyond being occasionally helpful. They are now credible, and sometimes formidable, at writing and understanding code. We exist in a technology landscape optimised around our finite human focus as the most scarce resource. Agentic coding changes the dynamics. We need tools that are optimised for AI Agents. We need tools that allow us to instruct, guide and supervise AI Agents.

Dev Tools must evolve with the evolving needs of their users. Many developers no longer write code all day, they manage a fleet of AI Agents. We need Dev Tools that serve the needs of the agent author and the human supervisor.

Some aspects of the Dev Tools business have changed forever. Software development teams are shrinking and a "per seat" pricing model may no longer make sense if your users are agents and don't need "seats". Dev Tools companies where revenue is currently coupled to the number of human users may struggle to grow their business and need to consider other metrics to measure and meter their value. The build vs buy decision has also changed. If a SaaS product is simple enough that it can be "vibe coded" in a couple of prompts then it is unlikely to succeed with a subscription model. That or the price will need to be reduced to a point where it becomes preferential to a DIY solution.

With these dynamics in mind businesses must consider what value they sell. One enduring theme is trust. You choose a partner because you trust them. You trust them with your data, trust them with security, trust them with reliability, trust them to do business ethically. This is still a big deal.

AI Native

When we talk about AI-native development we are optimising for processes where AI agents are the actors and human developers the supervisors. Many practices we have in place today were created for very human problems. Everything from the 2-week sprint to the code review is optimised for a world where developer time is a precious resource.

"AI-native" is a lot bigger than simply writing code with an agent. It touches everything from product strategy, requirements definition, architecture, coding standards, testing and reliability engineering. These are not solved problems but everywhere you look there are glimpses of a future where every part of the software development lifecycle is agentic.

AI Native and Open Source

Open source is where innovation thrives. Where ideas are tested, proved and refined. Where real solutions to real problems can find their users and build a community. Open source is also transparent, and with that comes scrutiny and trust. Open standards and protocols allow us to scale new capabilities (just look at MCP) and provide interoperability between vendors. Instead of lock-in or proprietary interfaces we have a dynamic ecosystem, this accelerates open innovation further.

AI generated Code

There is no doubt that many maintainers have found themselves under pressure from a deluge of AI generated "code-slop" contributions. Some have introduced AI guidance ([Kubernetes](#)) and others will publicly denounce and block future contributions from anyone who submits poor quality AI generated contributions ([Ghostty](#)) to manage the noise.

It's important to remember that true open source is free from restrictions on users and use. If your AI-authored contribution is not accepted by a maintainer that is their decision but you remain free to fork that repo to add your changes. As a maintainer you are allowed to say no. You are the caretaker of that project and you can set out the priorities for it.

Contributors should always take ownership for the quality of their contributions. When contributors delegate responsibility for review to the maintainer that is unfair, unreasonable and selfish.

The Opportunity

Beyond the hype of each new model there is another story unfolding in the day to day working lives of each and every one of us. The impact of these models can be felt keenly by developers but outside of the tech realm the promise of productivity isn't playing out that way. Organisations were promised savings and were delivered a "work-slop" - this does not represent the potential of this technology and is an opportunity few have managed to take advantage of.

Even in the domain of software development we are still learning the right harnesses, loops and skills that make AI native development reliable and secure. UK-based Tessl are leading the way with their focus on trustworthy skills, robust evals and context engineering practices, whilst [nono.sh](#) have closed an enormous security gap with their isolated agent sandbox. We have an abundance of problems to be solved in the Dev Tools space and we are fortunate to have trailblazers right here in the UK.


Asks of the UK Government

We have the [Cyber Runway](#) accelerator for cyber-security but we have no equivalent for developer tools. This is a glaring gap in our startup ecosystem. The seismic shift in software development requires more focussed support. We need to create the right conditions for world leading teams to turn problems into developer solutions at the speed of AI.

2.13 AI and Open Standards

2.13.1 Global Open Standards

Our digital world has likely needed global governance for many years, but AI has finally shifted the governance conversation to build recognition of the need for international alignment in AI. Standards are seen to be the answer with an acknowledgement by many that these must be open and freely available to have the necessary impact.



"Open standards are the beginning, middle, and the end."

Juliette Denny, Founder, Growth Engineering and Zavmo

Open source software has seen decades of code-based default standards across software packages and developer tooling. The scale of adoption of these in particular across the last decade, has led to the normalisation of these and their being housed in a not for profit foundation.

Traditional standards bodies are already accepting of the challenge, that they have not been able to keep up with the pace of innovation in software development and this is a problem to their structure and business model that is only enhanced by AI.

Increasingly in AI the ROOST, tools not rules approach is seen to offer a practical approach for AI Devs enabling them to follow good technical hygiene practices without the burden of regulatory compliance, which comes with a clear and present risk of regulatory capture where only a few large companies might be able to comply.

Demis Hassabis Framework for Frontier AI Standards Body

On 14 July Google DeepMind founder Demis Hassabis set out his [Framework for Frontier AI](#), calling for an industry-funded standards body to be built to test models and set AI standards. He suggests that it would be a "US-initiated effort", public-private, enterprise funded and Federally controlled, with "appropriate federal agencies and the US National Labs to conduct testing in areas relevant to national security". The premise that this be a single country-led, US foundation, not a global one has provoked debate and geopolitical concerns of fragmentation across AI if such a body is created and driven by the US. The work the body might undertake is well scoped within the very considered proposal but sits in an overarching risk of a form of "standards capture" in AI by a few large enterprise funders effectively able to buy control of such a body and the standards it chooses and outputs it approves.



"Open source helps ensure that power isn't concentrated among a few powerful players."

Andrew Martin, Co-Founder and CEO, ControlPlane

2.13.2 Fireside Chat: Mobile Context Protocol (MCP) and Protecting its openness

David Soria Parra
Member of Technical Staff, Anthropic



Please briefly tell us about yourself, your role, and how you came to be in it?

I am currently a Member of Technical Staff at Anthropic and the Co-Creator and current Lead Maintainer of the Model Context Protocol. I joined Anthropic 2 years ago, after working 10 years at Facebook / Meta on developer tooling, source control and static analysis.

How did you get involved in open source and what impact has it had on your career?

Since the beginning of my career I have worked on Open Source. I started working part-time as an engineer at a small agency (Mayflower GmbH) that worked a lot with PHP in Munich. One or two people were deeply involved in PHP's development and that's how I got involved. I started working on small patches, documentation, libraries, these types of things. It's really thanks to my mentors there that I got involved in Open Source. The work on PHP later allowed me to get a part-time job during university at Sun Microsystems working on optimizing PHP for Sparc processors. This was the first job I got hired for because of my background in Open Source. At the same time I was the PHP 5.4 and PHP 5.5 release manager and on the side I got deeply involved with Git and Mercurial, leading me eventually to work on Mercurial where I wrote some key parts. One of the people involved in the Mercurial community was building a Source Control team at Facebook, and that's how I ended up there. It was the start of my career. Getting into a FAANG company only happened because of Open Source. You can easily say that my career is built and enabled mostly because of my Open Source contributions. It paved the path I am on.

Who worked with you on MCP and what problem were you trying to solve?

I worked on MCP together with Justin Spahr-Summers. While the idea was mine, Justin really was the one who believed in it being important and helped to drive a lot of the initial implementation. We wanted to enable people to bring their workflows to AI. In a way we wanted to build the limbs to the AI brain. At the time the models got increasingly more capable and it was clear to us that the missing piece is allowing people to build for themselves and connect their systems and flows to these systems. We also knew that people are using a wide variety of clients, be it IDEs, Desktop applications like Claude Desktop or websites. So we needed to solve the classic NxM problem, N clients need M integrations. A protocol reduces this to N+M+1.

How did your open source backgrounds shape your approach in building MCP?

It was clear to me that building a protocol only works when you succeed in building an ecosystem. My experience told me that the most sustainable way to build a successful ecosystem is through open source. My background here of course helped me significantly to understand (1) why developers want an open source solution (2) why ecosystems benefit from open source and (3) how to successfully build an Open Source project.

Why was open sourcing it the right call rather than keeping it as an internal advantage?

I believe MCP only works because it is open source. You need to allow people to build client integrations, you need to allow people to build servers. This only works if you have an open specification of the protocol connecting client and server. It's only logical that of course you will want to provide an open source implementation of the SDKs for the protocol then as well. Protocols are definitions for how systems communicate with each other. They only work if they are open and can be freely implemented. The alternative is to do a proprietary API. You can do that, but we believe that an open source protocol like MCP uplifts everyone in the ecosystem, and allows people to use AI more as a whole.

For a non-technical reader, what does MCP do and why does it matter?

MCP (Model Context Protocol) is a universal adapter that connects AI assistants to the tools and data you actually use, email, files, Slack, databases. Think of it like USB: instead of building a custom integration for every AI-and-tool pairing, a tool builds one MCP connection and any AI assistant can use it. It turns AI from something that just talks about work into something that can do work, read your real files, search your real messages, update your real spreadsheets. MCP enables one conversation to handle your entire digital life, whether you're planning a big purchase, building a business, or writing a novel.

Why did Anthropic donate MCP to a foundation?

We knew from the early days of MCP that making a true standard requires ensuring that such standard stays vendor neutral and without any trademark or license risks. Early in the process we decided that we need to either create a foundation or donate it to a foundation. With the Linux Foundation we did eventually find a great home. We donated MCP to the Linux Foundation precisely so the developer community can trust it will stay open and vendor-neutral. And we contributed to a \$12.5M fund at the Linux Foundation to strengthen the security of the open source ecosystem more broadly.

MCP has been adopted even by rivals like OpenAI and Google - what made that cross-industry cooperation possible?

MCP solved a problem that every AI company shares, so adopting it is simply rational. Before it, connecting an AI system to an external tool meant building a custom integration for each pairing. MCP replaced that with a single open standard, so a tool builder writes one connector that works everywhere. In addition, we build MCP as a truly open source project. This included open governance that is merit based and transparent decision making. This allowed others to engage and build trust that an open standard can work, before fully committing.

What does MCP tell us about where open standards fit in the AI era?

Open standards keep AI companies competing on model quality and safety, not on who controls the tools. Open source software is essential for building a secure and innovative ecosystem for agentic AI. Myself and Anthropic are excited to continue contributing to MCP and other agentic AI projects through the AAIF. We recently launched a dedicated program for open source developers building on MCP, because the community building on this standard is as important as the standard itself.

As agents move from coding into general knowledge work, what's next for open infrastructure?

Anthropic supports open source software when it makes sense – like the development of MCP, and we will continue to innovate that. We believe proprietary models are the safest for widespread enterprise development. As models become more advanced, what you can build with them grows, but so does the complexity of building well. At Anthropic, we're invested heavily to ensure Claude is a true virtual collaborator, while keeping humans in the loop on the work that matters most.

2.13.3 New Foundations

A shift is seen in new open source foundations being built out of the US with a focus on agent management representing the use of the Steward concept to manage open protocols and standards used to manage and secure open source technology that services models whether the models are open or closed and as opposed to hosting AI models that are themselves open source. China has also responded to this flurry of new US-based foundations with its new foundation.

The Agentic AI Foundation

The creation of the Linux Foundation's Agentic AI Foundation (AAIF) in December 2025 marks a massive industry push to standardise AI systems, drawing on heavily adopted open source technologies enabling protocols and standards donated by Anthropic, Block, and OpenAI. Led by Executive Director Mazin Gilbert, the foundation has scaled rapidly, reflecting an environment where tools like the Model Context Protocol (MCP) and AGENTS.md already hold a substantial market footprint. The \$350,000 platinum membership tier may create a walled garden of board membership and voting power consolidating this into a small group of companies being limited to six organisations with founders believed to have a veto power, and a potential lack of community balance.

Akrites Foundation

The Linux Foundation's launch of Akrites on June 25, 2026, marks a massive, coordinated effort to address the inversion of vulnerability economics caused by AI. Because frontier models can scan and surface serious open source security flaws in minutes, the bottleneck has shifted from flaw detection to human remediation—a reality highlighted by the fact that fewer than 5% of recently validated vulnerabilities have actually been patched.

Backed by seed funding from [Alpha-Omega](#) and engineering commitments from 19 major companies including AWS, Microsoft, Google, OpenAI, Anthropic, and major financial institutions, Akrites establishes a centralised, confidentiality-first Security Incident Response Team (SIRT), using CVE, CVSS, and VEX to streamline patch development and seeks to protect volunteer maintainers from being overwhelmed by fragmented, machine-generated bug reports which have been overwhelming open source projects.

It introduces a “maintainer of last resort” protocol, empowering Akrites to step in and directly secure critical, abandoned packages so downstream users aren't left exposed. Akrites introduces operational and philosophical tensions to the open source ecosystem with its critics questioning whether this intervention will build sustainable, human-driven codebase maintenance or simply establish a corporate-backed clearing-house that routes automated reports while consolidating control over legacy software ecosystems in those corporate entities. This manages vulnerability disclosures through an embargoed, confidentiality-first process run by large AI companies and for some this sits uncomfortably with open source's historical community driven champions.

Like the Agentic AI Foundation it has been set up with leadership and heavy investment to secure and standardise open source software interfaces and dependencies that AI companies' tools run on, whilst the underlying weights in the AI models of many of its members remain closed.

Agent Name Service Foundation (ANS) - announced but not launched

Security vulnerabilities of unmonitored AI agents have prompted a shift toward internet-scale identity infrastructure. The [Cloud Security Alliance survey](#) taken by 418 practitioners in April 2026, saw 82% of organisations discovering AI agents in their infrastructure that nobody authorised and 65% experienced an agent-related security incident in the preceding year. The ANS was built and run in production by GoDaddy who donated the specification to the Linux Foundation.

By anchoring an agent's identity directly to existing DNS architecture and logging modifications via an append-only transparency framework, ANS aims to provide cross-organisational verification that proprietary corporate identity stacks could not achieve.

x402 Foundation

Coinbase and Cloudflare announced their intent to form the401 Foundation on 23 September 2025 and it went live on 14 July 2026. The Foundation applies the same model to machine payments that the Agentic AI

Foundation applies to agent tooling, taking a protocol originated by a single company and placing it under Linux Foundation's open source governance. The x402 protocol itself was created and open sourced by Coinbase [in May 2025](#), and it revives the long-dormant HTTP 402 “Payment Required” status code so that an AI agent meeting a paid endpoint can settle the charge and retry the request without an account, a subscription or a card entered by a human.

The Foundation membership is broadwith 40 organisations at launch across three tiers, and premier members drawn from the card networks, payment processors and cloud and crypto infrastructure, including Visa, Mastercard, American Express, Stripe, Adyen, Fiserv, Google, Amazon Web Services, Shopify, Circle and Cloudflare. The stated purpose, as with the agent foundations examined above, is to keep the payment layer open, interoperable and free of lock-in. The protocol is payment-method and chain-agnostic, and any compliant settlement network can plug in to it, yet most live settlement runs on the USDC stablecoin over Coinbase's Base chain, which leaves a dependence on Coinbase and Circle in the critical path, as the standard above it remains neutral. A competing protocol - L402I from Lightning Labs - offers an independent alternative.

China's WAICO AI Foundation

The World Artificial Intelligence Cooperation Organisation (WAICO) was established in Shanghai on 16 July 2026, with a goal to “promote international cooperation and develop AI regulation” across countries, and 29 countries including Indonesia, Brazil, Malaysia, South Africa, Senegal, Russia and Pakistan sign-up to its founding Agreement. China's 2025 proposal for this was discussed at Asia-Pacific Economic Co-operation (APEC) Meetings in October 2025 as an answer to the dominance of US frontier labs for the Global South. China will host APEC in November 2026.

WAICO is set up as an independent intergovernmental organisation headquartered in Shanghai and guided by the UN Charter - the UN Secretary General attended the launch. Its objectives cover consultation of members, shared technological development and reduction of the technological risks.

China's absence from the India AI Impact Summit in February 2026, organised during one of China's biggest annual holidays, may have influenced this second gathering and potential fork in global collaboration. The next Global AI Summit is in June 2027 in Geneva and understanding its interface with this will be critical to international policy.

This will sit alongside local Chinese AI organisations like [Beijing Academy of Artificial Intelligence](#) (BAAI) founded in 2018 to bring together academic communities of practice in AI.

Calls for a UK Foundation

OpenUK has called for the [Proposal for a national UK open source foundation](#) to host and lead open source in the UK and to act as a distributor of UK open source funding. This is an attempt to enable the infrastructure necessary to enable the UK's innovation as a platform for global collaboration and to enable greater success for UK open source-based businesses to scale.

The Tony Blair Institute (TBI) also recognises the UK's need for an “open source AI lab”. TBI proposes the lab develop reusable open tools for government departments, distil and fine-tune open models for priority sectors such as health and AI for science, curate high-value data sets with the Sovereign AI Unit, maintain critical open source infrastructure, and coordinate internationally with companies and bodies. They suggest evolving the existing Incubator for Artificial Intelligence into a dedicated lab with a cross-government mandate and long-term funding, reporting either to the Minister for AI.

OpenUK's recommendations are broader in scope. Rather than focusing only on building new technologies, the UK also has to nurture what is already being built and that is managed and monetised effectively as well. The approach here should be to bring a broader ecosystem understanding together, and then focus on where to support the most value creation as well as the right values in the community and business markets.

OpenUK's proposal benefits from the vast UK industry's experience with broader goals and would engage the UK and world's leading open source experts in its build. The hope being that the work offers learning and opportunity for replicable builds in other countries.

2.14 The AI Stack and AI Sovereignty

2.14.1 Sovereignty and Geopolitical Shift

Sovereignty began as a data concept, requiring compliance with local data privacy laws and geographically local servers to host that data. Over recent years it has evolved to expand into a broader digital environment and today’s debate is generally AI focused. OpenUK will shortly share a Sovereignty paper drilling deep into this, but across this report, shifts in geopolitics, combined with the potential for dependence on US tech companies has triggered concerns in nation states that critical services relied upon for nation states to function are now in the hands of a few tech companies and primarily US-based. This dependency has created understandable concern.

Conversations on sovereignty can by their nature be exclusionary, which sits counter to the basis of global collaboration in open source. Whilst the requirement in open source technically focuses on the outbound enabling anyone to use the outputs for any purpose and does not require that anyone can contribute, whoever and wherever they are, it has been an accepted principle of open source that it is a globally collaborative movement and that people are not excluded for personal characteristics including their nationality or location, with a few exceptions driven by export control regulation, law trumping both licensing and principles in open source.

In October 2024 a group of Russian developers were excluded from being maintainers of the Linux Kernel effectively removing any possibility of their control whilst they continued to be able to contribute. This was the first time developers had been continuously excluded from a project based on geography and the beginning of concerns with respect to geography. Russia’s unsurprising response was to state that it would bifurcate the Linux Kernel.

As nation states have looked to remove technological dependency on technology from and controlled by other nations, open source has been recognised as a cornerstone of sovereignty. It is inevitable that open source is used as the basis of a sovereignty strategy, as creating the stack without it is simply not feasible. Organisations like Eurostack which have advocated for this have been the recipients of open source concerns based on their understanding of open source. Open source does not only allow access for all when it is taken but also it is critical that its outputs are shared without geographical or political restrictions. Imposition of these has a similar impact to adding anti commercialisation provisions in licensing and effectively stops the licensed innovation being open source, meaning that calling it such would amount to open washing.


A challenge in the broader sovereignty debate not specific to open source is the lack of definition and agreed ontology in this space meaning there is a great deal of discussion at cross purpose and miscommunication. In June 2026, as is explored in the US Policy section of this Report, Anthropic’s Claude Fable5 and Mythos5 which had already been subject to limited release and review, were subject to a US Export Control Directive and withdrawn from non US citizens for almost three weeks. Earlier US restrictions on NVIDIA chip distribution to China and US tariffs have also sharpened minds, bearing in mind the critical role that the US tech sector plays in the provision of national infrastructure.

For some it is a question of control, but for others and in particular the traditional open source ecosystem this concept of control does not sit well with the nature of open source and they focus instead on not being dependent. When we consider dependency we can split this into dependency on technology itself and commercial dependency on businesses.

Many conversations of sovereignty describe sovereign innovation or technology that is not built from the ground-up but which has customised usage of technology from other countries and the idea of a partially sovereign technology being customised for local intelligence is seen by many as good enough. The concept of layered sovereignty following the layered stack has become a de rigeur talking point on this. The five layers have led to the term “Layered Sovereignty”: largely following NVIDIA’s AI Cake Stack.

Others discuss sovereignty across the stack, and see concerns varying depending on the level of the stack at which the concern arises. It is easiest to build domestically localised innovation at the top, at the Application layer and hardest at the bottom where there is a need for vast spend and physical infrastructure. For some it is enough to have sovereignty at a level where the base of the stack remains internationally provided.

Nvidia’s Five Layer Cake AI Stack The Stack Layer Cake Recipe Card 2026



Concerns have been raised with this type of sovereignty on the basis of US export control around Anthropic models, but which has also been applied by the US with respect to silicon chips. The export control applied to China by the US created the potential for this to be applied globally or to any specific country at a moment in time, as it was with Claude Fable and Mythos.

At Layer five, applications can be written and owned locally with modest resources, because they sit closest to users and draw mainly on skills that many states already hold. At Layer 4, the spread of open weight models, from Qwen and Kimi to Sarvam, has made a sovereign model capability far cheaper to acquire than it was two years ago, and the frameworks on which those models are built are themselves open.

At Layer three, control over data and the software that moves it is being pursued through national data assets and sees an ongoing challenge of data in AI environments. The constraint becomes real at the lower two layers with Layer two, the advanced accelerators on which training depends, being supplied by a handful of organisations and no degree of openness to alleviate dependence.

Layer one, the energy and data-centre capacity that everything else consumes, is the least substitutable of all, since it can be neither downloaded nor relocated. However it is also one of the trickiest areas to fulfill. This may be alleviated by shifts in AI. Where a year or two ago the monolithic data centres the US was building for AI training appeared globally necessary, the shift in AI, and potential scale of models explored later in this report, may see this burden lifted. It may be that the model of AI with a combination of processing on devices - which is currently being pursued by the mobile industry - combined with edge-based data centres close to the user and suitable for inference as opposed to training, and the smaller customised models of the future, that these giant data centres and the associated infrastructure will be less of a challenge.

Sovereignty asserted at Layers four and five while Layers one and two remain external is therefore partial, and the report’s wider finding applies here in concentrated form.

Mozilla has publicised the creation of its “Rebel Alliance” collaborating across industry to build a Sovereign open source AI Stack. The Stack detail is expected later this year.

No conversation on AI in 2026, can be complete without a consideration of “AI Sovereignty” and what it means for concentration of power and ownership of frontier AI models to be in the hands of a few, primarily US, wealthy organisations. Whole reports are being written around the sovereignty debate and this section can be no more than a taster of the debate and the key challenges it seeks to mitigate.

2.14.2 Thought Leadership - The Commons we have not (yet) created

Anastasia Stassenko
Founder and CEO, Pleias



The most powerful models Europe lost access to this year happen to be called the Fable series. It is the kind of coincidence you cannot improve on. The suspension did not create a vulnerability so much as expose a fable that had been told for years: that you did not need to build the substrate of artificial intelligence, only to use it well. Own the application layer, the story went, and let others burn the capital underneath. When the layer underneath was switched off from Washington, the story switched off with it.

The lesson most people drew was the obvious one - you cannot rent a substrate and call it sovereignty. That is true, and it is also the least interesting thing to say about it. Of course frontier models are expensive; of course whoever pays for the compute holds the leverage. If the case for decentralised AI stopped at "we need our own GPUs," it would be a procurement problem, and procurement problems get solved or they don't. The reason the suspension matters is that it exposed a deeper dependency, one that does not disappear even if you buy the hardware.

Here I have to bring in my philosophical background and talk about the French philosopher Gilbert Simondon, because he described the shape of the problem before there was a machine to attach it to. Individuation - the process by which a subject actually forms - is, in his account, two things at once. It is collective: a mind does not form in isolation but through feedback loops with a milieu, other minds, other practices, the slow friction of people who do not agree. And it is contextual: it happens inside a language, a craft, a tradition dense enough with shared references to sustain real disagreement. There is no individuation in general. There is only individuation within a milieu specific enough for thought to occur.

Hold those two commitments against the frontier proposition - one model, trained on everything, served to everyone through a single interface - and the mismatch is structural, not economic. A model like that can be extraordinarily responsive and still be uncontestable. You can prompt it; you cannot argue with it the way you argue with a colleague or a tradition. It flattens the milieu into a generic surface. Close the loop through which a subject individuates inside one company's API, and the individuation that results is not free. It is administered.

This is the argument for decentralisation that does not reduce cost or anti-monopoly politics. It is a claim about the conditions under which people think. You need models embedded in particular languages, particular domains, particular communities of practice - not because the local is romantically superior to the universal, but because the loop through which a subject forms requires a milieu specific enough to sustain genuine collective thinking. One model trained on everything by a single entity cannot provide that. It can simulate responsiveness. It cannot simulate the situated friction of a real epistemic community. The frontier logic is not only concentrated. It is epistemologically impoverished - and plurality is not a concession to scarcity but a condition of the kind of thinking that matters.

The strange thing, the thing that gives this argument political weight it would otherwise lack, is that the CFO arrives where the philosopher does. As AI stops answering questions and starts running operations - procurement, compliance, logistics, decisions at every level - the model a company operates through becomes a layer of its governance. Every inference call constrains a process. No firm should run its core operations through a provider that reads its data and learns its patterns while serving its competitors on the same infrastructure. The commons argument says closed cognitive infrastructure is incompatible with collective thought. The competitive argument says closed operational infrastructure is incompatible with autonomy. Simondon and the CFO, by entirely different routes, demand the same substrate: plural, contextual, governed by those who depend on it.

The uncomfortable part is that wanting this does not make it buildable, and the tradition we reach for first does not fit. We call the alternative open source, but there is no open source AI in the sense there is open source software. Every capable open-weight model - Llama, Granite, the Mistral weights, Qwen, and OLMo, the one that comes from a research institute rather than a company - is a corporate release with a permissive licence attached, and the licence is the whole fragility. Open source software worked because the people who built it owned the means of producing it - a laptop, a connection, time - and because the GPL made openness irreversible: once code was free it could not be enclosed again. There is no GPL for weights. You can open them, but if the data, the methodology, and the training code stay closed, what you have is an artifact a corporation can withdraw tomorrow. A community built on a revocable gift is not a commons. It is a tenancy.

So the commons tradition has to borrow from one it has always ignored - the tradition that knew how to provide things too expensive for any individual and too important to leave to monopoly. Railways, electricity, telecommunications each posed exactly this, and each was answered with an institutional invention: common carriage, public utilities, cooperatives, rural electrification. Capital-intensive infrastructure that is a precondition for participation cannot be governed by the same logic as the activities it enables. It is worth noting who else has reached this conclusion. When Leo XIV placed AI infrastructure, data, and algorithms under the universal destination of goods, a tradition working from entirely different premises landed exactly here. Compute has a universal destination, or it has an owner.

And the challenge is double, because it is not only compute. It is also data - the permissibly licensed material models learn from, before it is scraped into proprietary pipelines and sold back as a service with nothing returning to the commons that produced it. Common Corpus, the two-trillion-token open dataset we built at Pleias, is one very modest attempt at that second layer: a demonstration that what models learn from can itself be open, and governed by something other than extraction. But without public compute, even open data feeds an enclosed pipeline. The two commons rise or fall together.

What is missing is the institutional form, and it is within reach precisely because the software world already invented most of it. A copyleft licence for AI: use the weights, modify them, build on them, but derivatives stay open, and commercial serving returns a share to the commons that produced them. Couple that to cooperative governance and an inference-revenue loop - training as the capital cost, inference as the demand that funds the next run - and the commons sustains itself rather than begging for subsidy, the way energy cooperatives and credit unions always have: those who depend on the infrastructure govern it, and the surplus stays.

None of this exists at full scale. All of it is being attempted in fragments. Whether those fragments converge into a form with the weight to hold is a political question, not a technical one - and for the first time the coalition is real, because the philosopher who wants plural infrastructure and the company that wants operational autonomy have finally begun to ask for the same thing.

2.14.3 Case Study: Open weights, owned compute and the case for sovereign AI

Tom Wilkinson Chief Data Officer, Scottish Government



Tom Wilkinson's remit covers public sector data reuse, the use of AI in public services, open government data, and specialist applications including geospatial work and data science as Chief Data Officer at the Scottish government. He was Senior Reporting Officer for the Scottish AI Alliance, he led on Scotland's AI Strategy initially shared in 2021 and concluded in 2026. With a background in applied data science, he prefers his current technical remit to that earlier policy role, seeing political attention to AI from 2023 onwards becoming a distraction from delivering benefits in public services. He sees the challenges of transparency to citizens, commercial bargaining power and the long-term maintainability of systems as connected.

Two dimensions of openness

Wilkinson distinguishes two dimensions of AI openness as relevant to his work with Scottish government.

One is the open government dimension. Alongside open data, he maintains the Scottish AI Register, a public record of AI systems in use or in development across the Scottish public sector, held under a Ministerial Mandate requiring public bodies to complete and share the relevant information in AI Systems. Scotland was the first part of the UK to make such registration mandatory, beginning with government departments and extending outwards to the rest of the public sector in phases. For Wilkinson transparency about specific uses is the most important component of openness from a Scottish perspective.

The second dimension of openness is commercial benefits. Tom describes open technology as fundamental to competitive markets, and credits it with keeping the relationship between government and its suppliers on a basis that is fair. It enables public money to buy genuine value for Scotland's citizens. Applied to AI, where commercial model value remains unsettled, open licensing, open code together with open weight models, enables interoperability and the choice to move between suppliers. Where a supplier has built a solution for the government, Wilkinson wants certainty that it can be taken elsewhere, and managed by another team or supplier.

The Public Sector Procurement Challenge -What government actually buys

Procurement has become a key public sector challenge in AI and models. The principal tool used for general purpose generative and agentic work across Scotland's public sector is Microsoft Copilot based on Scotland's existing procurement relationships. Although a process exists for approving other models, Claude Code among them, Wilkinson is uncertain of how many instances of these actually exist in practice.

His assessment is that in many respects maturity matters more than tooling. He doubts anyone in Scotland is as yet making active choices between more and less open AI for production systems. While some AI has been tested quite extensively in the delivery of real services, he describes an absence of production, maintainable machine learning at this stage. The decisions openness advocates care about, such as whether a model has open weights, have largely not had to be taken at this stage, which is what gives the standards being set now their practical weight.

Portability and maintainability

Asked whether a portability standard exists for AI systems, Wilkinson said the government would be very interested in one, while arguing that portability is the narrower of two problems. The nature of large language models means that even moving between two versions of the same model from one vendor is difficult, since the response to a given prompt may vary materially across versions. Maintainability is therefore the deeper problem, because a system can be locked in place by its own brittleness rather than by any contract. Open weight models are the obvious response, since they can be served by more than

one vendor, though committing to them risks falling behind the most competitive performance available. Openness at the model layer can also be undone at the platform layer, and where open weight models are accessed through AWS Bedrock or Azure AI Foundry, Wilkinson would want reassurance that they are not wrapped in native technology preventing them from being lifted out again.

Sovereignty, efficiency and resilience

The Scottish government holds no official position on sovereignty, and Wilkinson is careful to say so; his team's role is to give policy officials the technical support they need to reach one. His own framing treats sovereignty as a question of risk, and of the trade-off between efficiency on one side and redundancy and resilience on the other. What tends to be called sovereignty usually pays off in resilience, by removing dependency on any single provider, and at national scale it resembles commercial independence more than a constitutional principle. Its cost is efficiency, since international economies of scale often make the less sovereign option cheaper. Cast this way, sovereignty becomes a procurement and engineering judgement that can be argued on the evidence, though he does not claim this framing has become policy.

Skills and the intelligent customer

Around 30 people report to Wilkinson directly or indirectly, within a digital directorate of roughly 600, and he is also head of community for data across the Scottish public sector, a group of perhaps 50 to 100 for which he coordinates the heads of practice and maintains standardised role descriptions. On capability his account is qualified. A recent move towards holding these skills internally takes place in a context of very limited public finances, sometimes crudely controlled by limiting hiring, so a reliance on outsourcing persists. What has changed over five years is the quality of the buyer rather than the volume of internal delivery; the government now holds enough expertise to act as an intelligent customer, and to build some things for itself.

Data maturity and the Data Exchange

Wilkinson's team runs a data maturity program that supports public bodies in carrying out assessments, which consistently find organisations at very different stages; the mechanism for closing the distance is communities of practice, through which less mature organisations learn from more mature ones. A newer program, the Data Exchange, addresses the same problem structurally. Its first component is the identification of standards for interoperability, following the FAIR principles of findable, accessible, interoperable and reusable data; its second is the construction of common componentry, available without a commercial licence, so that public bodies can meet those standards even where they lack the maturity to procure a solution themselves. Wilkinson hopes that a combination of open source foundations, shared components and increasingly capable AI coding assistants will bring compliant infrastructure within reach of smaller bodies.

Open data and open source contribution

The 2025 to 2028 delivery plan's commitment to open Earth observation data, LiDAR and an open platform is partly realised. A re-platforming exercise will bring Scotland's geospatial and Earth observation sources together under a single data.gov.scot, which Wilkinson expects to be visible within the year, and a national LiDAR capture has begun. On open source he is direct. Scottish government is smaller than the UK government, and he claims no equivalent of the Ministry of Justice's Splink project or NHS England's cataloguing work. Where Scotland has contributed is through procurement. Via CivTech, its challenge-led innovation program, his area set a challenge on making high-quality public participation scalable, attaching conditions requiring open source components and an open core model. The resulting company, CrownShy, builds open source integrations between digital democracy tools, its commercial model resting on deployment rather than software. Wilkinson regards this as his best contribution to open source, since the government has caused open source code to be written for open government without writing it itself.

Two key questions to openness

When it comes to privacy, Wilkinson separates two questions.

First, open data and unintended disclosure. He says that safeguards ensure that open data published by government is not disclosive of personal information, generally managing this through aggregation techniques designed to prevent later disaggregation.

Secondly, concerns that open source itself introduces risk. Tom takes the view that open source software is probably more secure, opting for the bazaar rather than cathedral.

Sovereignty

Sovereignty for Scotland is framed as a resilience question, with an efficiency price, though without official endorsement; and where the government cannot contribute code, it writes openness into the conditions it sets in the market. The tensions remain unresolved on Wilkinson’s own account, and the absence of production machine learning across the Scottish public sector means these trade-offs have mostly been described rather than tested.

Conclusion

Scotland’s approach to openness and AI amounts to a set of institutional habits today rather than a doctrine.

Publication of public sector AI usage is mandated; open weight models and open source code are treated as a form of “commercial insurance”, with the performance cost acknowledged.

“Keeping software up to date is going to be much more important in the post-AI world.”

Adrian Mouat, DevRel Engineer, Chainguard

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.

2.15 Value generation and ROI

A consistent and persistent theme in the AI discussion across the three years considered in this report is where value generation and return on investment will be seen. Evidence of return is limited in both proprietary and open source reporting.

Many believe that current pricing models applied to frontier AI subscriptions are not viable for long term profitability. Thinking is evolving and there appears to be a shift away from big is best in models moving to a strong likelihood of value being generated at the agentic harness and infrastructure level. Deployment in enterprises is not seen to be delayed by the functionality of models themselves but by deployment constraints in enterprise.


Looking to the discussion across this report and the AI Stack it would appear that key areas of value generation may include the layers in the stack beyond the models including the API layer, infrastructure, the agentic harness, and the energy or data centre layer.

3. What the Data Shows

This section draws on repository (Repo) data from GitHub and Hugging Face, sharing the data mined and drawing conclusions from this. The GitHub data has been analysed in collaboration with Runa Capital. All figures are current as of July 2026.

Stamford’s HAI Index shares that today on GitHub there are 5.6m projects and that this is 5x the number in 2020. Following the approach in our previous reports focusing on Repos with 1,000+ GitHub Stars as a visible proxy for significant activity. Stars measure attention rather than use, quality or production adoption, as an indication of where open AI development is concentrated and how quickly that concentration is changing, rather than as a complete census.

Models are shared on GitHub by some, but are primarily shared through Hugging Face which does not hold geographic data, reducing analysis options. Hugging Face now hosts over 2m public models and 500,000 datasets and uploads are almost 3x the number they were three years ago. 3.1 AI Repositories Repository counts are an imperfect proxy for national capability, but among the few measures of activity that can be compared consistently across jurisdictions over time. The threshold used here, 1,000+ Star filters out the large volume of experimental and abandoned work that would otherwise dominate any raw count, leaving projects that have attracted sustained attention from developers. Geographical attribution follows the identified location of the repository owner and excludes any with no location recorded.




3.1.1 AI Repositories Globally by Country

Measured by the number of AI repositories with 1,000+ Stars, open development remains heavily concentrated in two leading countries - the US leads at 1,131 and China with 685. The UK holds third

place with 161, ahead of Germany at 125, France at 85, India at 67 and Australia at 25. The UK’s position is therefore secure but distant, with roughly one qualifying repository for every seven held by the US, and the ranking has changed little in order even as the absolute numbers have risen.

Number of AI Repositories with 1k+ stars by Country, July 26



In association with Runa Capital


Figure 4 Source: Runa Capital

Of 2,279 repos attributed to national locations, 49.8% are US, 30.16% with a drop to 7.06% in the UK.

This general repo data from GitHub must however be read in the light of China’s more limited use of GitHub and the tendency for Models to be shared more on Hugging Face than GitHub. OpenUK is collaborating with the local Chinese open source community to publish a more detailed China AI Openness report later this year.


3.1.2 AI Repositories Evolution over the last 3 years

AI repositories with 1,000+ stars have more than doubled over three years, from 2,159 in June 2023 to 4,806 in June 2026 but the ranking is constant with the US top at 1,288, China second at 777 and both roughly doubling. The gap between the two leaders and the rest of the world has remained constant.



In association with Runa Capital

Figure 5 Source: Runa Capital



In association with Runa Capital

Figure 6 Source: Runa Capital

Apart from the US and China, the ordering has changed significantly:

- The UK grew from 83 to 184 passing Germany, into third place
- Germany’s growth has slowed from 95 to 147 - the slowest expansion of larger contributors
- France grew from 45 to 97, more than doubling
- India grew fastest in proportional terms, from 29 to 77

3.1.3 Top AI Repositories Globally

Drilling down into the composition of the global top ten repositories shows how recent much of this activity is. The top four projects, openclaw/openclaw, obra/superpowers, affaan-m/ECC and NousResearch/hermes-agent, have each accumulated their stars within the past year. This not only leaves them in the top positions globally but also places them at the head of the fastest-growing list.

Top 10 AI Repositories Globally

In association with Runa Capital

Repo Name	Repo Description	Stars	Country
openclaw/openclaw	Your own personal AI assistant. Any OS. Any Platform. The lobster way. 🦞	382,103	UK
obra/superpowers	An agentic skills framework & software development methodology that works.	248,687	USA
affaan-m/ECC	The agent harness performance optimization system. Skills, instincts, memory, security, and research-first development for Claude Code, Codex, Opencode, Cursor and beyond.	227,059	USA
NousResearch/hermes-agent	The agent that grows with you	210,980	USA
tensorflow/tensorflow	An Open Source Machine Learning Framework for Everyone	196,104	USA
n8n-io/n8n	Fair-code workflow automation platform with native AI capabilities. Combine visual building with custom code, self-host or cloud, 400+ integrations.	195,565	Germany
Significant-Gravitas/AutoGPT	AutoGPT is the vision of accessible AI for everyone, to use and to build on. Our mission is to provide the tools, so that you can focus on what matters.	185,433	UK
ollama/ollama	Get up and running with Kimi-K2.6, GLM-5.1, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.	175,670	USA
f/prompts.chat	f.k.a. Awesome ChatGPT Prompts. Share, discover, and collect prompts from the community. Free and open source – self-host for your organization with complete privacy.	165,034	Turkey
AUTOMATIC1111/stable-diffusion-webui	Stable Diffusion web UI	164,027	UK

OpenUK logo

Figure 7 Source: Runa Capital

3.1.4 Top AI Repositories in the UK

Top 10 AI Repositories UK, July 2026

In association with Runa Capital

Repo Name	Repo Description	Stars
Significant-Gravitas/AutoGPT	AutoGPT is the vision of accessible AI for everyone, to use and to build on. Our mission is to provide the tools, so that you can focus on what matters.	185,433
mlabonne/llm-course	Course to get into Large Language Models (LLMs) with roadmaps and Colab notebooks.	80,734
rohitg00/ai-engineering-from-scratch	Learn it. Build it. Ship it for others.	37,590
VectifyAI/PagelIndex	PagelIndex: Document Index for Vectorless, Reasoning-based RAG	33,862
AlexsJones/llmfit	Hundreds of models & providers. One command to find what runs on your hardware.	29,192
cheahjs/free-llm-api-resources	A list of free LLM inference resources accessible via API.	26,210
rohitg00/agentmemory	#1 Persistent memory for AI coding agents based on real-world benchmarks	24,781
coleam00/Archon	The first open-source harness builder for AI coding. Make AI coding deterministic and repeatable.	22,758
EthicalML/awesome-production-machine-learning	A curated list of awesome open source libraries to deploy, monitor, version and scale your machine learning	20,705
tirth8205/code-review-graph	Local-first code intelligence graph for MCP and CLI. Builds a persistent map of your codebase so AI coding tools read only what matters, with benchmarked context reductions on reviews and large-repo workflows.	19,277

OpenUK logo

Figure 8 Source: Runa Capital

Two projects now rank among the leaders in both the world and UK, openclaw. openclaw heads each list at 382,103 stars, having reached that position from a standing start within roughly seven months, while Significant-Gravitas/AutoGPT, previously globally first, places second in the UK at 185,433 stars and seventh globally. The remainder of the UK list is made up of educational and tooling projects of more moderate size, running from Maxime Labonne’s llm-course at 80,734 stars down to a set of libraries, templates and smaller tools in the 20-to 40-thousand range, none of which appears in the global ranking.

3.1.5 Fastest growing AI Repositories Globally

Newly popular projects in 2026, have pushed long-standing projects down the rankings. TensorFlow, for years among the most-starred machine-learning repositories, now fifth. Growth of this scale reflects intense but recent attention. The rest of the top ten mixes established tools so that despite the US supplying most of the leading repositories, it does not supply all of them.

3.1.6 Comparison of Contributions Received by AI repositories

Star counts and contribution depth diverge sharply among the leading repositories. langchain-ai/langchain draws on 3,573 contributors and NousResearch/hermes-agent on 1,527, while obra/superpowers with 248,687 stars from just 36, and TauricResearch/TradingAgents 91,642 from 19. Stars measure attention, which accrues within days of release; contributors measure participation, which accrues over years. Both measures are indicative of collaborative development. New entrants score highly on stars but less so on contributors. Older projects show the inverse: tensorflow, slowest-growing by stars, retains the deepest contributor base at 3,933. AutoGPT, at 788 contributors, sits between the two patterns.

Contributors Comparison Across Top 10 AI Repositories, July 2026

In association with Runa Capital

Repo Name	Repo Description	Contributors
openclaw/openclaw	Your own personal AI assistant. Any OS. Any Platform. The lobster way. 🦞	2,275
obra/superpowers	An agentic skills framework & software development methodology that works.	36
affaan-m/ECC	The agent harness performance optimization system. Skills, instincts, memory, security, and research-first development for Claude Code, Codex, Opencode, Cursor and beyond.	274
NousResearch/hermes-agent	The agent that grows with you	1,527
n8n-io/n8n	Fair-code workflow automation platform with native AI capabilities. Combine visual building with custom code, self-host or cloud, 400+ integrations.	678
f/prompts.chat	f.k.a. Awesome ChatGPT Prompts. Share, discover, and collect prompts from the community. Free and open source – self-host for your organization with complete privacy.	534
ollama/ollama	Get up and running with Kimi-K2.6, GLM-5.1, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.	593
AUTOMATIC1111/stable-diffusion-webui	Stable Diffusion web UI	579
Significant-Gravitas/AutoGPT	AutoGPT is the vision of accessible AI for everyone, to use and to build on. Our mission is to provide the tools, so that you can focus on what matters.	788
tensorflow/tensorflow	An Open Source Machine Learning Framework for Everyone	3,933

3.1.7 Top Topics in AI Repositories

Leading Repository Topics, July 2026

In association with Runa Capital

Topic	Repos with Topic	Total Stars	Total UK Commits	Total UK Contributors
ai	277	5,965,280	51,632	486
openclaw	36	1,258,736	34,464	108
assistant	277	458,906	30,782	40
own-your-data	3	446,867	30,782	40
personal	2	389,845	30,779	39
crustacean	1	382,103	30,778	38
molty	1	382,103	30,778	38
typescript	205	3,985,375	21,412	354
python	442	6,848,794	21,824	697
hacktoberfest	525	6,025,884	28,354	1,132

Unsurprisingly the general AI topic ranks top and almost equally unsurprising in 2026, is the fact that Open Claw ranks second as the global topic of conversation.

3.2 Model Repositories

Hugging Face hosts more repos for models than GitHub. The platforms measure different things and the distinction matters when analysing the figures. GitHub counts projects and the code around them; Hugging Face counts artefacts, downloads and likes, and its two leaderboards diverge almost completely. By number of downloads on Hugging Face, the top ten is dominated by small, older utility models. By number of likes on Hugging Face, the picture inverts toward large generative models. Downloads therefore track infrastructure and GitHub likes track attention, and neither alone describes which models the ecosystem depends on.Only one model, Qwen3-0.6B, is included in both top 10’s.

3.2.1 Top 10 Most Downloaded Models

The downloads leaderboard is led by small, older models built for retrieval and text processing, with sentence-transformers/all-MiniLM-L6-v2 at 253M downloads on 22.7M parameters. Four of the ten are unchanged from 2024. Projects at that scale grow automatically, each time a pipeline initialises or a container rebuilds, they record how many production systems have taken a dependency on that model rather than how many people have chosen one.

Top 10 Most Downloaded Models, July 2026 In association with Hugging Face

Repo Name	Task	Params	Updated	Downloads	Likes
sentence-transformers/all-MiniLM-L6-v2	Sentence Similarity	22.7M	Jun 1	253M	5.09k
cross-encoder/ms-marco-MiniLM-L6-v2	Text Ranking	22.7M	Aug 29 2025	85.1M	285
google-bert/bert-base-uncased	Fill-Mask	0.1B	Feb 19 2024	80.4M	2.71k
BAAI/bge-small-en-v1.5	Feature Extraction	33.4M	Feb 21 2024	66.5M	514
google/electra-base-discriminator			Feb 29 2024	55.9M	142
sentence-transformers/paraphrase-multilingual-MiniLM-L12-v2	Sentence Similarity	0.1B	Jan 28	49.6M	1.32k
BAAI/bge-m3	Sentence Similarity		Jul 3 2024	35.2M	3.24k
sentence-transformers/all-mpnet-base-v2	Sentence Similarity	0.1B	Aug 19 2025	30.4M	1.33k
google-t5/t5-small	Translation	60.5M	Jun 29 2023	27.8M	575
Qwen/Qwen3-0.6B	Text Generation	0.8B	Jul 25 2025	26.7M	1.42k

Top 10 Most Liked Models, July 2026 In association with Hugging Face

Repo Name	Task	Params	Updated	Downloads	Likes
black-forest-labs/FLUX.1-dev	Text-to-Image	12B	Jun 27 2025	604k	13.7k
deepseek-ai/DeepSeek-R1	Text Generation	685B	Mar 27 2025	8.86M	13.5k
stabilityai/stable-diffusion-xl-base-1.0	Text-to-Image	3B	Oct 30 2023	1.41M	7.95k
CompVis/stable-diffusion-v1-4	Text-to-Image	0.9B	Aug 23 2023	456k	7.04k
meta-llama/Meta-Llama-3-8B	Text Generation	8B	Sep 27 2024	1.35M	6.61k
hexgrad/Kokoro-82M	Text-to-Speech		Apr 10 2025	10M	6.55k
meta-llama/Llama-3.1-8B-Instruct	Text Generation	8B	Sep 25 2024	8.27M	6.34k
openai/whisper-large-v3	Automatic Speech Recognition	2B	Aug 12 2024	5.93M	5.96k
black-forest-labs/FLUX.1-schnell	Text-to-Image	12B	Aug 16 2024	244k	5.38k
deepseek-ai/DeepSeek-V4-Pro	Text Generation	862B	22 June 2026	1.49M	5.25k

3.2.2 Top 10 Most Liked” Models

The leaderboard for likes is dominated by large generative releases, black-forest-labs/FLUX.1-dev at 13.7k and deepseek-ai/DeepSeek-R1 at 13.5k. Imagemage generation fills four of the top 10 ranking places.

The gap between standing measured by likes whilst use is measured by downloads is wide in places. When it comes to use FLUX.1-dev has 604k downloads against DeepSeek-R1’s 8.86m on comparable l counts.

Only Qwen3-0.6B comes close to appearing on both lists demonstrating its high level of likes and usage. Kimi K3 was newly released adn did not have data available at the time of this review.

3.3 Agentic Repositories

With the shift to broader AI technology, beyond models, and the rise of agentic AI as the harness for models enhancing their delivery, there is also an increase in open source agents and associated tools following a more traditional open source route. Where models are shared primarily on GitHub agents and Dev Tools are primarily on GitHub.

When we look at the general AI data we see that Agentic repos have stars and collaboration around them to a greater extent.


3.3.1 Agentic Repositories by Country

Drilling down on the agentic data, it is nationally more concentrated than in AI as a whole.

Of 591 repositories with an identified owner-location, the US accounts for 265 and China for 126. Some way behind the UK is third with 27, ahead of Singapore at 17, Germany at 15, India at 14, Italy at 13 and France at 12.

The field is markedly youthful with 353 of the 891 repos. Circa 40% did not hold any stars a year ago. Of 891 repositories with 1,000+ stars, two in five held no stars a year ago.

37% were created within the past 12 months, so the ranking reflects adoption over months rather than years and should be read as provisional and may well evolve quickly.



The median repository age is 2.21 years and only 1.61 years for the US and 1.40 years for China.

Contributor depth diverges from star counts across the dataset: obra/superpowers holds 248,687 stars against 36 contributors; while langchain-ai/langchain has 141,219 stars and 3,573 contributors.

Concentration is nonetheless pronounced:the US accounts for 265 of the 591 located repositories and China 126, leaving the UK third at 27. The UK's third place in agentic repositories is consistent with its placing generally in AI repositories, and 4.57% of the global total in agentic as opposed to.

3.3.2 Top 10 Agentic Repositories Globally

The global top 10 is dominated by the US, which holds six of the ten places. The UK, China and two nationally unattributed repositories take the remainder.

Three of the leading four did not exist a year ago. obra/superpowers, affaan-m/ECC and NousResearch/hermes-agent all held no stars in July 2025 and now sit above 200,000, so the upper part of the ranking has been rewritten within a single reporting period.

Top 10 Agentic Repositories Globally In association with

Repo Name	Repo Description	Stars	Country
obra/superpowers	An agentic skills framework & software development methodology that works.	248,687	USA
affaan-m/ECC	The agent harness performance optimization system. Skills, instincts, memory, security, and research-first development for Claude Code, Codex, Opencode, Cursor and beyond.	227,059	USA
NousResearch/hermes-agent	The agent that grows with you	210,980	USA
Significant-Gravitas/AutoGPT	AutoGPT is the vision of accessible AI for everyone, to use and to build on. Our mission is to provide the tools, so that you can focus on what matters.	185,433	UK
Snailclimb/JavaGuide	Java 面试 & 后端通用面试指南 · 覆盖计算机基础、数据库、分布式、高并发、系统设计与 AI 应用开发	156,860	China
langgenius/dify	Production-ready platform for agentic workflow development.	148,085	USA
firecrawl/firecrawl	The API to search, scrape, and interact with the web at scale. 🚀	147,462	USA
langchain-ai/langchain	The agent engineering platform.	141,219	USA
google-gemini/gemini-cli	An open-source AI agent that brings the power of Gemini directly into your terminal.	105,831	USA
browser-use/browser-use	🌐 Make websites accessible for AI agents. Automate tasks online with ease.	103,343	USA

Figure 14
Source: Runa Capital

The established projects beneath them grew at ordinary rates: AutoGPT from 176,701 to 185,433 and langchain from 110,923 to 141,219, with firecrawl the exception at 42,525 to 147,462. Contributor counts do not follow stars. langchain draws on 3,573 contributors and hermes-agent on 1,527, while superpowers has 36 and firecrawl 156, which suggests that the largest star totals in this set reflect rapid adoption of a released artefact rather than the accumulation of a development community around it.

3.3.3 Top 10 Agentic Repositories UK

The UK's position rests on depth rather than breadth. Its repositories are older than most, with a median age of 2.21 years against 1.61 for the US, and its 387,366 stars are dominated by AutoGPT. Beneath it, Pydantic, Trigger.dev, DocsGPT and the DeepMind libraries form a smaller but established layer.

Top Agentic UK Repositories In association with

Repo Name	Repo Description	Stars
Significant-Gravitas/AutoGPT	AutoGPT is the vision of accessible AI for everyone, to use and to build on. Our mission is to provide the tools, so that you can focus on what matters.	185,433
rohitg00/ai-engineering-from-scratch	Learn it. Build it. Ship it for others.	37,590
VectifyAI/PagelIndex	🔗 PagelIndex: Document Index for Vectorless, Reasoning-based RAG	33,862
rohitg00/agentmemory	#1 Persistent memory for AI coding agents based on real-world benchmarks	24,781
pydantic/pydantic-ai	AI Agent Framework, the Pydantic way	18,258
arc53/DocsGPT	Private AI platform for agents, assistants and enterprise search. Built-in Agent Builder, Deep research, Document analysis, Multi-model support, and API connectivity for agents.	17,975
triggerdotdev/trigger.dev	Trigger.dev - build and deploy fully-managed AI agents and workflows	15,598
olimorris/codecompanion.nvim	🌟 AI Coding, Vim Style	6,725
google-deepmind/open_spiel	OpenSpiel is a collection of environments and algorithms for research in general reinforcement learning and search/planning in games.	5,312
pydantic/logfire	AI observability platform for production LLM and agent systems.	4,352

Figure 15
Source: Runa Capital

3.4 Conclusion - What we see from the Data

The data shows a story, a journey in openness since the 2023 opening of Llama 2 as open innovation.

Shifting models

Firstly we see a very current shift in thinking in language models. Alongside the evolution of the small language model is a shift in the belief of big being best to understanding that being customisable by the user is more important. Two model releases from the US and China in the same week of July 2026 illustrate how rapidly the frontier is shifting and supports this belief.

On 15 July 2026 in the US, Thinking Machines Lab, founded by former OpenAI chief technology officer Mira Murati, released Inkling, an open weight model built on the premise that enterprises want customisable AI rather than rent huge models from a handful of frontier labs and openly saying that Inkling is not the strongest model available. Its Tinker platform follows the same argument we saw from Locai Labs in our case study - fine-tuning open models for particular tasks or industries rather than competing on scale is the future of AI models. It shows us that successful models may not be the biggest and that there is an increasing shift towards tailored models.

A dependency runs between the two releases: Thinking Machines used data generated by existing open models, Kimi K2.5 among them, in Inkling's final training phase.

Shift from models to agents

AI openness today focuses not only on the models of the 2023 discussion but increasingly on model harnessing via agents. When we compare model and agentic repositories we see that the agentic set is young and community-formed: two in five of its repositories held no stars a year ago, and contribution is broad, with langchain drawing 3,573 contributors and hermes-agent 1,527.

There is a stark difference between the development today of the model layer and the agentic layer. Model repositories follow a different pattern. Foundation models are, for the most part, produced inside companies and released complete - thrown over the wall - with the artefact rather than the development process made public. The download data bears this out: the most-used models on Hugging Face are small, stable and years old, embedded in pipelines rather than iterated on in the open, while the most-liked are large generative releases from a handful of labs.

Growth in the agentic layer is cumulative and distributed, whereas growth in the model layer is episodic and concentrated, arriving in discrete releases from a small number of organisations. Agentic technology is more likely created by an individual who takes community contributions and collaboration as the project evolves. The Dev tooling built around smaller, tailored models is where the two meet.

We assume the difference in approach lies in the funding requirements and resources needed to build models.

Developer Tools

Typescript has increased in popularity to be the most used language on GitHub with an increase of 66% in the last year, with this increase being seen to be at least partially due to Dev Tools influencing their user.

The UK's place in AI Openness

It is notable that the UK-based repositories which rank highly from a Stars/ contribution perspective are in agentic and not foundation models. Open Claw has replaced AutoGPT at the top but the UK has spawned two of the most important open agents globally.

The reasons for this can be seen spread across this report, but very much linked to the UK's huge innovation in the more affordable areas. In these areas, market entry is easier, and can be driven by individuals or SMEs who are truly a part of the traditional open source ecosystem and who understand the value and ways of open source.

With a shift in UK focus at a policy level to industrialisation, the recognition of these shifts and the potential opportunity they enable will be critical in guiding the UK's short and long term AI strategies and potential monetisation.

3.5 The Entrepreneur's AI Journey

Matt Barker

Founder and CEO, Bolt MCP



The AI race over the past few years has largely become one of building the biggest, most capable foundation model. Every new release is measured in benchmarks, reasoning ability and scale. It's exciting, and it continues to push the boundaries of what's technically possible.

But bigger models are not the same thing as bigger economic outcomes. While AI is already delivering value in areas such as software engineering, customer support and scientific research, the transformational impact across the wider economy has yet to fully materialise. Like every major technology shift, adoption takes time. Organisations need to rethink processes, governance and skills before the technology can realise its full potential.

As we watch \$100B flow into frontier AI companies and the infrastructure required to train ever larger models, we need to be realistic about the UK's position. Competing directly with the US and China on foundation models is unlikely to be a winning strategy. We simply don't have the same access to capital, compute or energy.

This raises a more interesting question: do we actually need to?

Recent developments, including Mira Murati's Inkling model, suggest that even some of the world's leading AI researchers are thinking differently. Rather than pursuing scale for its own sake, there is growing interest in open weight models that organisations can adapt and build upon.

At the same time, many of the biggest barriers to AI adoption are no longer about the model itself. They are about securely connecting AI to organisational knowledge, business systems and workflows so it can complete meaningful tasks. That's where AI stops being a chatbot and starts becoming infrastructure. Perhaps, then, the UK's ambition should not be to build the world's largest model. It should be to build the world's most useful one.

The UK is already home to world-class AI research, globally respected universities, a thriving startup ecosystem, the AI Safety Institute, and pioneering work on AI interoperability standards such as the Model Context Protocol. We have exceptional talent and deep expertise in sectors like healthcare, financial services, defence and life sciences.

Those are exactly the ingredients needed to lead the next phase of AI. The first phase of AI has been a race for intelligence. The next will be a race for outcomes.

American friends tell me that London is already being viewed as an international hub for AI. If Silicon Valley becomes known for building the world's most powerful models, the UK has the opportunity to become the place where those models solve the world's hardest problems.

By focusing on open technologies, secure integration with organisational data and systems, and practical deployment into the real economy, we can build an ecosystem that turns frontier AI into productivity, innovation and economic growth. That feels like an ambition worthy of the UK's strengths

3.6 Case Study: Agentmemory

Rohit Ghumare
Creator & Maintainer, Agentmemory



Rohit Ghumare's route into open source followed his contribution to cloud-native projects such as Kubernetes and Istio gaining him recognition as a CNCF Ambassador, Docker Captain and Google Developer Expert. Those in turn led to his developer relations roles at open source companies including Solo.io, Cerbos and Oracle. This underwrote his move to the UK via a Global Talent visa, where he now works as founding developer relations engineer at iii. He notes that the existing open source and AI community in London, including organisations such as OpenUK, gave him a professional network in the country long before he had settled there. "Agentmemory", the repository he began in late February 2026, grew out of his daily work inside AI coding agents. He recognised the challenge that tools such as Claude Code, Cursor and Codex begin each session without recollection of the previous session. Developers repeatedly re-describe their stack, their conventions and their architectural decisions, consuming tokens repeatedly. Ghumare treats the context window as a workaround rather than as memory proper. The project hooks into an agent's lifecycle, captures what happens during a session, compresses those observations into structured facts, and reintroduces the relevant material when a new session begins. Retrieval combines BM25, vector embeddings and a knowledge graph, organised around a four-tier consolidation model of working, episodic, semantic and procedural memory drawn from accounts of human cognition.

He built the project benchmark-first, and attributes much of its reception to that decision. He believes that Agentmemory reaches 95.2% retrieval accuracy on LongMemEval-S and reduces token use by roughly 92% when measured against the practice of supplying full context in every session. The system runs locally on SQLite with local embeddings and requires no external services. He links this design choice to developer trust in a tool that records everything an agent does. Releasing "agentmemory" under the open source Apache-2.0 licence speaks to the wider question of openness in AI tooling. His reasoning:

- Inspectability - a component that logs and stores an agent's activity should be auditable, and running it entirely on a local machine is part of the product.
- Pace - integrations across Claude Code, Cursor, Gemini CLI and Codex change weekly and cannot realistically be kept current by one maintainer.
- Infrastructure - memory is a foundational layer sitting between agents and their data, and that an interoperability layer of this kind should remain open in the manner of the Model Context Protocol rather than becoming proprietary.

The project has evolved from Ghumare's solo project to almost 200 issues and a similar number of pull requests. Contributions range from bug fixes to README translations into Hindi, Japanese and other languages supplied by the community, and he credits an early contributor, Tanmay, among others. The project has gained more than 25,000 Stars within five months. This places the repo among the UK's ten most-starred over the period, alongside 2,100 forks and an average of more than 10,000 npm downloads each week. Stars may be a weak measure of adoption and he is the first to agree with this, pointing to downloads and forks as the more meaningful signals in his view. As of this report, the agent memory project has been forked more than 2,000 times on GitHub.

Coverage by developer newsletters and aggregators such as AlphaSignal and HelloGitHub supported his early distribution. However, he attributes the initial traction to reproducible benchmarks against alternatives including mem0 and Letta, shared to an audience accumulated by him over his years of developer relations work, and to the technology's benefit of reducing installation to a single command. There is no commercial layer, Agentmemory has no paid tier, no sponsorship gate and works offline at no cost. This is an old school open source project built on GitHub to scratch a modern itch.

Ghumare frames the return on the project as reputational and professional rather than financial, consistent with how open source work has served his career. He applies the same approach to his separate educational project, aiengineeringfromscratch.com, which he reports has attracted more than 38,000 Stars. Agentmemory is a useful illustration of a recurring pattern in open AI infrastructure, in which auditability and local operation function as the trust mechanism, community maintenance keeps pace with a fast-moving integration surface, and the commercial value accrues indirectly through reputation rather than through a licence.

4. What the Literature Says

4.1 Literature Review: Open Source AI, Global Ecosystem Shifts and the UK Policy Response in 2026

Across the 2026 reporting, a consistent picture emerges: open source AI has become central to understanding the future direction of AI.

While discussions of AI often focus on competition between frontier models or the capabilities of individual systems, the literature increasingly points towards a broader shift where value is created, how innovation diffuses, and how governments, organisations and societies are responding. Open source sits at the centre of these developments as an increasingly important mechanism through which AI capability is distributed, adopted and governed (Mozilla, [State of Open Source AI 2026](#); OECD, [The Benefits of AI Openness](#)).

The reports reviewed span international research institutions, governments, industry analysts and technology organisations, including Mozilla, the OECD, Stanford University, GitHub, Microsoft, Amazon Web Services, the UK government and a range of policy and economic analyses. Although their individual focuses vary considerably, several common themes emerge:

- The rapid convergence between open weight models and proprietary models
- Growing concerns around digital sovereignty
- Accelerating enterprise adoption
- The emergence and growth of an agentic harness
- Developer tools and open standards
- Significant disruption to labour markets, and
- An increasingly interventionist policy response from governments seeking to strengthen national AI capability

The first part focuses explicitly on developments in AI Openness and the later sections broadens the discussion to examine wider trends that, although not exclusively concerned with openness, have important implications for the evolution of the open source ecosystem. Together, these reports suggest that understanding the future of AI increasingly requires understanding the role that openness plays within a much wider technological, economic and geopolitical landscape (OECD, [The Benefits of AI Openness](#); Bennett School of Public Policy, [Reimagining the AI Arms Race](#)).

Open Source and the Commoditisation of AI Models

Perhaps the strongest point of convergence across the literature concerns the changing economics of foundation models. Rather than competing solely on model performance, evidence increasingly suggests that frontier capabilities are becoming widely accessible through open weight models, fundamentally changing where competitive advantage and monetisation is likely to reside (Mozilla, [State of Open Source AI 2026](#); OECD, [The Benefits of AI Openness](#)).

A consistent finding is the rapid narrowing of the performance gap between open weight and proprietary models which is notably decreased by Kimi K3's release on 18 July after all of the reporting considered. Mozilla estimates in July 2026 that this gap is just 3.3%, while the OECD concludes that open weight models achieve approximately 90% of the performance of comparable proprietary alternatives at release while costing around one-fifth as much (a price ratio of 0.2).

At the same time, Mozilla reports a 50x reduction in inference costs over the past 36 months, from approximately US\$20 to US\$0.40 per million tokens, dramatically lowering barriers to adoption (Mozilla, [State of Open Source AI 2026](#); OECD, [The Benefits of AI Openness](#)).

Evidence from platform adoption reinforces this trend. OpenRouter reports that Chinese-developed open models now account for approximately 45% of weekly tokens processed through its platform, while Alibaba's Qwen family reached 942m cumulative downloads on Hugging Face by March 2026, almost twice that of Meta's Llama models. These findings suggest that open weight ecosystems are supporting increasingly diverse global participation rather than concentrating capability among a small number of proprietary providers (Mozilla, [State of Open Source AI 2026](#)). The impact of Kimi K3 is likely to increase this further.

Increased accessibility is reflected in developer behaviour. GitHub's [Octoverse 2025](#) reports record levels of

open source activity, with more than 180M developers now using GitHub after 36M new developers joined during 2025 and in the UK the number of GitHub accounts topped 5M for the first time. Open source development reached 1.12 government contributions across public repositories, while 60% of the ten fastest-growing open source projects by contributors were AI-focused, suggesting that much of the industry's innovation is increasingly concentrating around open AI infrastructure rather than proprietary applications (GitHub, [Octoverse 2025](#)).

GitHub further suggests that commoditisation is influencing developer behaviour as much as model economics. As AI becomes embedded within everyday development workflows, developers increasingly favour languages and frameworks that reduce friction when working alongside AI systems, indicating that competitive advantage is shifting from individual models towards the wider development ecosystem (GitHub, [How AI is Reshaping Developer Choice](#)).

Taken together, these developments point to a broader economic shift. As foundation models become increasingly commoditised, value creation is moving towards the surrounding ecosystem, including the agentic harness made up of infrastructure, specialised memory, orchestration layers and increasingly sophisticated agentic systems. The OECD's analysis in [The Benefits of AI Openness](#) complements this view by demonstrating that lower-cost open weight models enable wider diffusion of AI capability across organisations and economies. Together, the literature suggests that openness is not simply reducing barriers to access but fundamentally changing where innovation occurs and where value is captured (Mozilla, [State of Open Source AI 2026](#)).

Open Source, Enabling Sovereignty and Strategic Competition

Beyond economics, the literature consistently identifies open source AI as an increasingly important component of digital sovereignty. Rather than viewing openness solely as a mechanism for accelerating innovation, the literature increasingly frames it as critical infrastructure underpinning national resilience, technological independence and long-term competitiveness. It is recognised as an [economic] lever shifting from being a technical decision (Mozilla, [State of Open Source AI 2026](#); Oxford Insights, [Government AI Readiness Index](#); Bennett School of Public Policy, [Reimagining the AI Arms Race](#)).

Following the UK government's decision to exert export controls to suspend international access to Anthropic's Claude Fable 5 and Mythos 5 models for all non-US nationals. CNBC and Computer Weekly describe the 19 day suspension as a demonstration of the risks associated with dependence on proprietary AI. Mozilla argues that the episode illustrates a key distinction between proprietary and open weight models: while commercial services can be withdrawn or restricted, an open weight model deployed locally cannot simply be switched off and is licensed in perpetuity. The literature therefore presents openness not only as a question of transparency, but also of resilience, continuity and technological independence critical to sovereignty (CNBC, [Computer Weekly](#); Mozilla, [State of Open Source AI 2026](#)).

This thinking is increasingly reflected in industrial policy. The [UK] [Government AI Readiness Index](#) documents growing investment in sovereign compute, nationally governed datasets and locally deployable open weight models. It sees a growing view of open source AI as strategic digital infrastructure rather than simply another software ecosystem (Oxford Insights, [Government AI Readiness Index](#); European Commission, [AI Continent Action Plan](#); Mozilla, [State of Open Source AI 2026](#)).

The Reimagining the AI Arms Race report provides an important counterpoint, arguing that viewing AI solely through the lens of US-China competition overlooks emerging models based on open innovation, shared infrastructure and international collaboration. Together, the literature suggests that AI Openness is increasingly evolving from an alternative development model into a foundational component of national resilience, industrial strategy and long-term technological capability (Bennett School of Public Policy, [Reimagining the AI Arms Race](#); Mozilla, [State of Open Source AI 2026](#); Oxford Insights, [Government AI Readiness Index](#)).

Economic Value, Governance and Enterprise Readiness

While the convergence of the capabilities of open weight and proprietary models is one of the most widely reported developments, the literature also highlights important tensions around transparency, enterprise deployment and economic value. Together, these findings suggest that increased openness does not

automatically resolve the practical challenges associated with deploying AI at scale (OECD, [The Benefits of AI Openness](#); Stanford University, [AI Index Report 2026](#); Mozilla, [State of Open Source AI 2026](#)).

From an economic perspective, the OECD estimates that a 10% increase in global AI open source software contributions is associated with an approximate 0.5% increase in long-run GDP, reinforcing the role of open source as both a driver of technological innovation and wider economic productivity through knowledge diffusion, collaboration and reduced barriers to entry (OECD, [The Benefits of AI Openness](#)).

At the same time, the literature identifies a growing paradox around transparency. Stanford University's AI Index Report 2026 concludes that the most capable frontier models - whether open or proprietary - are becoming progressively less transparent regarding their training data, development methodologies and parameter counts. This reinforces an important distinction within the open source AI debate: open weights increase accessibility but should not be equated with complete openness or transparency. Instead, the literature increasingly treats openness as a spectrum spanning open weights, open source software, open data and fully transparent AI development practices (Stanford University, [AI Index Report 2026](#); Mozilla, [State of Open Source AI 2026](#); OECD, [The Benefits of AI Openness](#)).

Mozilla also identifies an important gap between technical capability and enterprise readiness.

While 89% of organisations already incorporate open source components within their technology stacks, many open weight models continue to face challenges around regulatory compliance, governance, security and enterprise support. At the same time, Mozilla challenges the assumption that proprietary systems are inherently more secure, highlighting critical zero-click data exfiltration vulnerabilities in major commercial models and concluding that security depends less on licensing than on implementation, governance and operational practice (Mozilla, [State of Open Source AI 2026](#)).

GitHub suggests that increasing developer productivity does not reduce the need for governance. Instead, AI-assisted development can accelerate architectural drift unless organisations invest in stronger standards, documentation and architectural oversight. As AI increases development velocity, governance and organisational capability become increasingly important complements to model capability (GitHub, [How AI is Reshaping Developer Choice](#)).

GitHub also demonstrates the continuing interdependence between public open source and enterprise software. While 81.5% of code contributions now occur within private repositories, 63% of all repositories remain public, illustrating how enterprise development increasingly depends upon open libraries, models and frameworks that underpin commercial software development (GitHub, [Octoverse 2025](#)).

Taken together, the reports indicate that the future success of AI Openness will depend not only on continued improvements in model capability which are being seen, but also on the maturity of the surrounding ecosystem more generally. Infrastructure, governance, security, compliance and operational support emerge as equally important determinants of enterprise adoption as whether the AI is open or not (Mozilla, [State of Open Source AI 2026](#); OECD, [The Benefits of AI Openness](#); Stanford University, [AI Index Report 2026](#)).

Enterprise Adoption and the Rise of Agentic Systems

Beyond developments in open source specifically, the literature consistently documents the rapid embedding of AI within organisational activity. Although many of these reports examine AI adoption more broadly, they remain highly relevant to the open source ecosystem because increasing enterprise adoption expands demand for accessible, adaptable and cost-effective models, shaping both the opportunities and future direction of AI Openness (Stanford University, [AI Index Report 2026](#); AWS, [UK AI Adoption Report](#); Microsoft, [AI Diffusion Report Q1 2026](#)).

Stanford University's AI Index Report 2026 estimates that organisational AI adoption has reached 88%, while generative AI tools now create approximately US\$172 billion of annual value for consumers in the US. In the UK, AWS reports that almost two-thirds (64%) of organisations now use AI, with 78% stating they would be more likely to increase adoption if the government demonstrated stronger leadership through public sector deployment (Stanford University, [AI Index Report 2026](#); AWS, [UK AI Adoption Report](#)).

Adoption also varies significantly across countries. Microsoft's AI Diffusion Report Q1 2026 identifies the United Arab Emirates (70.1%) and Singapore (63.4%) as global leaders in population-level AI usage, while the

US ranks only twenty-fourth, with an adoption rate of 31.3%. These findings distinguish technological leadership from societal diffusion, suggesting that widespread deployment is becoming as important as frontier model development (Microsoft, [AI Diffusion Report Q1 2026](#)).

The literature also points to changing patterns of innovation. AWS reports that AI-native start-ups are reaching billion-dollar valuations in an average of just 3.5 years - around half the time required by previous generations of technology companies. Reduced development costs and the increasing availability of open weight models are lowering barriers to experimentation and commercialisation for emerging firms (AWS, [Global Startup Trends Report](#); OECD, [The Benefits of AI Openness](#)).

A particularly prominent theme across the literature is the transition from conversational AI towards increasingly autonomous agentic systems. As foundation models become increasingly commoditised, competitive advantage and the potential for return on investment is shifting towards the systems, workflows and organisational capabilities built around them. Microsoft's AI Diffusion Report documents a 28x increase in agentic pull requests on GitHub within just ten months, reaching approximately 2.3M by March 2026, while Evidence from Codex demonstrates substantial productivity gains from autonomous coding agents.

Reflecting growing confidence in these capabilities, the Solopreneurs Report argues that expectations surrounding agentic AI contributed to the "SaaSocalypse", during which investors briefly erased approximately US\$1T from traditional software company valuations amid expectations that autonomous AI systems would increasingly replace conventional enterprise software (Microsoft, [AI Diffusion Report Q1 2026](#); OpenAI, [Evidence from Codex](#); Stripe Economics, [Solopreneurs](#)).

GitHub extends this argument by suggesting that AI is increasingly reshaping technology choice itself. Rather than simply accelerating software development, AI-assisted workflows appear to be influencing which languages, frameworks and tools developers choose to adopt. The report argues that technologies offering stronger compatibility with AI assistants are becoming increasingly attractive, creating what it describes as "convenience loops" that reinforce adoption across developer ecosystems (GitHub, [How AI is Reshaping Developer Choice](#)).

Taken together, the reports suggest that as foundation models become more widely available through both proprietary and open source ecosystems, sustainable competitive advantage is becoming progressively less dependent on access to models themselves. Instead, differentiation increasingly derives from how effectively organisations combine models with proprietary data, workflows, governance and organisational knowledge. Within this environment, open source AI assumes particular importance because it enables organisations to customise, integrate and orchestrate AI systems with greater flexibility and reduced dependence on proprietary platforms (Mozilla, [State of Open Source AI 2026](#); OECD, [The Benefits of AI Openness](#); Microsoft, [AI Diffusion Report Q1 2026](#)).

Conclusion

Across the 29 reports reviewed, a consistent picture emerges: openness is no longer simply an alternative development model but an increasingly important driver of innovation, diffusion and competition. As the performance gap between open weight and proprietary models narrows, and deployment costs continue to fall, competitive advantage is shifting away from models themselves towards the infrastructure, governance and organisational capabilities built around them (Microsoft, [AI Diffusion Report Q1 2026](#); OECD, [The Benefits of AI Openness](#)).

The review also demonstrates that the future of AI Openness cannot be understood in isolation from wider developments across the ecosystem. Accelerating enterprise adoption, the emergence of agentic systems and the agentic model harness, growing concerns around digital sovereignty, labour market disruption and increasingly interventionist industrial policy are collectively reshaping the environment in which AI Openness is developing (Stanford University, [AI Index Report 2026](#); Microsoft, [AI Diffusion Report Q1 2026](#); Oxford Insights, [Government AI Readiness Index](#); UK Government, [AI Opportunities Action Plan](#)).

Taken together, the evidence indicates that openness is becoming a strategic capability rather than simply a technical characteristic. While important questions remain around governance, transparency and long-term sustainability, the literature consistently suggests that open source AI will play an increasingly central role in determining how AI is developed, adopted and governed across organisations, economies and societies (Mozilla, [State of Open Source AI 2026](#); OECD, [The Benefits of AI Openness](#)).

4.2 Literature Themes from 2023-2025

Across the literature reviews and report analyses compiled by OpenUK from 2023 to 2025, several distinct and interconnected themes emerged. This is a synthesis of findings from dozens of global reports by organisations such as GitHub, Stanford University, Demos, Anaconda, and the Linux Foundation,

Here are the key themes dominating the discourse on AI and open source:

1. **The Democratisation of AI and Economic Accessibility** A dominant theme is that AI Openness is democratising access to cutting-edge technology. The high costs of developing proprietary AI models create significant barriers, but open source models mitigate these by offering financially accessible alternatives that reduce training and adoption costs for academic institutions, startups, and smaller enterprises. Reports like GitHub's Octoverse highlight a massive surge in global developer engagement, with a 59% increase in generative AI projects and broad global participation that proves the democratising potential of these platforms. Ultimately, this accessibility is shifting AI innovation from the hands of a few well-capitalised tech giants to a much broader global community.
2. **Security, Safety, and the Burden of Trust** Balancing the open source ethos with safety and control is a central tension. Reports from Anaconda, Arm, and Demos consistently flag security and risk management as the most critical challenges for enterprise adoption. Key vulnerabilities identified in the literature include:
 - Model extraction and data leakage.
 - The exposure of vulnerabilities or the installation of malicious code.
 - The risks of cyberattacks, misinformation, or bioengineering misuse if powerful models are released without safeguards.
 The literature's response was a shift from access to assurance — SBOMs, dependency scanning, disclosure spectrums for risky releases (Centre for the Governance of AI), and safety benchmarking.
3. **The Definitional Dilemma and "Open-Washing"** Reports frequently grapple with the confusion around what "open source AI" actually means. Unlike traditional software, AI openness is not binary but exists on a spectrum. Reports repeatedly warn against "open-washing"—the practice of companies labeling their models as open source despite imposing heavy commercial or usage restrictions. To address this, academic and policy reviews (such as the Stanford AI Index and the Mozilla Columbia Convening) advocate for a "disaggregated approach". This approach breaks an AI system down into its component parts—such as inference code, model weights, training data, and documentation—to accurately assess how open and reproducible the system truly is.
4. **The Shift to Agentic AI and Resource-Efficient Models** The reviews track a rapid evolution in AI capabilities, indicating that Agentic AI (autonomous agents capable of reasoning and decision-making) is the next wave of transformation. Alongside this shift to multi-modal and agentic systems, there is a massive trend toward smaller, highly efficient models. Instead of relying exclusively on massive, energy-intensive legacy models, developers are successfully turning to resource-efficient open models—such as DeepSeek R1 or Microsoft's Phi-3-mini—which deliver highly competitive benchmark performance while drastically reducing computational parameters and training costs.
5. **Digital Sovereignty and the Public Interest** A key theme is a geopolitical and strategic shift toward treating AI as a "Digital Public Good." Emerging open protocols are enabling cross-vendor interoperability, which prevents platform lock-in and allows nations and organisations to maintain digital autonomy and sovereignty. By utilising open tooling and transparent data, countries can build localised, sovereign AI ecosystems that do not blindly inherit the blind spots of generic Western-trained systems. However, the literature also warns that AI infrastructure bottlenecks, such as the availability of computing power and coherent governance frameworks, remain significant hurdles to scaling these sovereign, public-interest AI initiatives.
6. **The discourse itself was scaling and maturing.** The literature reviews shifted from citing a handful of sources in 2023 to synthesising ~29 reports by mid-2025, and the tone shifts from speculative ("this could democratise AI") to evidence-based ROI claims (Anaconda's 12–18 month payback, CapGemini's 1.7x ROI).

Crucially, the field is moving from advocacy to measurement.

5. Conclusion

Dr Jennifer Barth
Research Director, OpenUK



Three years after the release of Llama 2, the debate about open source AI and AI openness has moved on. The issues now are where openness exists within the AI stack, who retains control and how the resulting economic and social value is distributed. Llama 2 was not open source in the traditional sense, but its release was still an important act of open innovation. It challenged the idea that advanced AI would remain the exclusive preserve of a handful of closed frontier laboratories. What followed has been more complicated than a straightforward shift towards openness. Instead, the structure of the AI ecosystem itself is changing.

The model layer remains highly concentrated. Training the largest models still requires levels of compute, capital, energy and data available mainly to large companies. These models tend to be developed privately and released as completed artefacts, rather than created collaboratively in public. The distinction between a model being fully open and open weights therefore matters. Access to model weights can enable experimentation, adaptation and deployment, but it does not provide full visibility of the data, methods or decisions that produced the model. Rather than trying to settle the question through a single definition, it may be more useful to examine the different components separately: weights, code, data, documentation, licensing and governance. This gives a clearer account of what is actually open, and what remains under corporate control.

Even so, the practical importance of open weight models should not be underestimated. DeepSeek R1 showed how quickly capabilities could be reproduced and extended through iterative development and distillation, reportedly at a fraction of the cost of comparable systems. The rise of Qwen and Kimi has also placed China at the forefront of open model releases, reflecting almost a decade of coordinated policy, institution-building and investment. More recent releases from Moonshot AI and Thinking Machines Lab reinforce the broader pattern. Performance gaps are narrowing, costs are falling and, for many users, the ability to adapt a model may matter more than marginal improvements in benchmark performance. Openness has become part of industrial strategy and international influence, rather than remaining solely a principle of software development.

The Locai Labs case study shows what this means at a more practical level. Its work involves building on open weight models, condensing them for local deployment and adapting them to particular organisations, sectors, languages and national contexts. Their work on Welsh language demonstrates the value of localisation, but also exposes the shortage of appropriate data for many languages and communities. Digital sovereignty in this context does not mean attempting to recreate every part of the global technology stack domestically. It means understanding where external dependence creates unacceptable risk and retaining enough control over models, data, compute and deployment to make meaningful choices.

The restrictions placed on access to the Claude Fable and Mythos models made this risk tangible. A capability accessed through a foreign-controlled API can be altered, repriced or withdrawn with little notice. Yet, as Anastasia Stasenko argues, sovereignty cannot be achieved simply by buying GPUs. AI systems are beginning to mediate how organisations think, decide and operate. A genuinely plural ecosystem will require models grounded in different languages, institutions and communities of practice, supported by open data and durable governance arrangements. Without these foundations, an open weight release may amount to a revocable corporate gift rather than a lasting contribution to the commons.

The Scottish government case study approaches openness in similarly practical terms. Here it appears not as doctrine, but through institutional habits: publishing information about the use of AI systems, supporting interoperability, investing in shared components and using procurement to encourage open development. Open technology can provide a form of commercial insurance by maintaining choice between suppliers, although that resilience may sometimes come at a cost in efficiency or performance. It also depends on government retaining the expertise to act as an intelligent customer. Without internal capability, sovereignty can quickly become little more than a procurement slogan.

These choices are being made within an open source ecosystem of remarkable scale. GitHub now hosts 5.6m AI projects, 5x the number recorded in 2020. Hugging Face hosts more than 2m public models and 500,000 datasets, with uploads running at almost 3x the level of three years ago. Among the more visible projects examined in this report, the number of AI repositories with at least 1,000 stars rose from 2,159 in 2023 to 4,806 in 2026. These measures do not offer a complete census of adoption, but they do show how quickly open repositories have become part of the basic infrastructure through which AI is explored, developed and increasingly deployed.

Much of the innovation is now taking place beyond the model. Models are being combined with agents, memory, tools, identity systems, permissions, evaluation and observability. This is the point at which AI begins to become operational infrastructure, and it is also where open source communities are playing a more familiar role. Of the 891 agentic repositories identified in the analysis, 40% had no stars a year earlier and 37% had been created during the previous 12 months. OpenClaw accumulated 382,103 stars in roughly seven months, while LangChain has drawn contributions from 3,573 developers. The picture is one of sudden bursts of growth alongside deeper and more sustained forms of collaborative development.

The UK has established a notable position within this activity. Its number of AI repositories with 1,000+ stars increased from 83 to 184 between 2023 and 2026, allowing it to overtake Germany and rank third behind the US and China. Its position is particularly strong in the agentic layer, where OpenClaw and AutoGPT have become two of the most prominent projects internationally. This should not be regarded as a consolation prize for failing to produce a frontier model. The technologies that allow models to be integrated, controlled and put to work may prove to be where a substantial share of value is ultimately created.

Rohit Ghumare's Agentmemory illustrates this opportunity at project level. Developed locally, released under the Apache 2.0 licence and improved through community contribution, it uses inspectability and local operation to build trust. Its growth also shows that the value of open source is not necessarily captured through licence revenue. It can emerge through adoption, reputation, professional opportunities and the commercial services built around a shared infrastructural layer. MCP offers a related example. An open protocol can become a de facto standard connecting models with tools and organisational systems. Moving such protocols into foundation stewardship may help preserve their neutrality, although the continued dominance of US-based foundations and standards bodies shows that institutional ownership matters as much as technical adoption.

AI is also changing software development itself. The 66% increase in TypeScript use over the past year, partly associated with the expansion of AI developer tools, suggests that AI is influencing both how software is produced and which technologies developers choose. This growth brings new strains. Maintainers are dealing with greater volumes of low-quality, machine-generated contributions, while autonomous agents unsettle existing assumptions about authorship, identity and access. Knowing which person authenticated is no longer enough when an agent can select dependencies, write code and execute actions. Governance will have to be built into the architecture of these systems.

Security presents a similar problem. AI can identify vulnerabilities more quickly than many maintainer communities can investigate and repair them. Commercial initiatives from Chainguard, IBM, Red Hat, OpenAI and Anthropic may help to manage the volume, but they could also produce a two-tier commons in which paying organisations receive protection before the public projects on which the wider ecosystem depends. The important test will be whether these programs strengthen upstream projects and shorten the time between discovery and repair. For enterprises, this forms part of a wider readiness gap. The main constraint is increasingly not access to capable models, but the ability to integrate, govern, monitor and secure them. Work will change alongside the technology. Productivity gains will depend on organisations redesigning workflows, developing technical and governance skills, and preparing people to work with increasingly agentic systems. Adoption is therefore an organisational challenge as much as a technological one and building workforce resilience to adapt and thrive is key for organisational sustainability. The UK's industrial strategy recognises the importance of AI, but recognition is the easier part. The harder task is to turn ambition into infrastructure, institutional capability, skilled communities and credible routes through which innovation developed in the UK can be commercialised and retained here.

What emerges from the data and case studies is a change in the real site of competition. The model still matters, but advantage is increasingly shaped by the wider ecosystem through which it is adapted, deployed and governed. Open source AI should therefore be understood as a strategic capability. Its significance lies in whether organisations and countries retain the ability to inspect, adapt, combine and govern technologies on which they increasingly depend. The first phase of the AI revolution centred on access to intelligence. The next will be shaped by who can put that intelligence to work, on whose terms and for whose benefit.

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

6.4 Methodology

The research used a mixed method approach to explore and demonstrate the state of AI Openness in the UK and globally. Interviews were conducted with industry leaders, founders and Open Source Software and AI experts and included as case studies and thought leadership.

Runa Capital supported our assessment of local repositories to establish comparisons in locally based contributions and repositories from each country. Runa Capital created a dashboard collecting data on open source repositories, users and activities. The data collection processes are automated, leveraging the APIs and tools provided by GitHub and updated daily. Runa gathers data about all repositories with 1000+ Stars at the time of collection. The author of the repository can be an individual or an organisation and where they own multiple repositories, each is considered as a separate entity. For any author who has at least one repository with 1000+ Stars, information is collected on all repositories they own. A commit in the content of version control with Git is a record of changes made to a set of files. All commits to 1000+ Stars repos are collected. Every instance where a user makes a change to the code is monitored. Users are collected in terms of repository authors and contributors including profile information, number of repositories owned and largest repository. To determine the user's geolocation the process uses the Google Maps API. This is updated for user profiles once per 6 months. Contributors from the UK are defined as users who have been geographically pinpointed as being in the UK through the geolocation process. These users are identified based on the location data they provide in their GitHub profiles, which is then verified and refined using geolocation tools.

GitHub supplied data on the total number of developers, in particular using data from GitHub Innovation Graph from which the heatmap has been produced. In the data provided by GitHub, countries with less than 100 developer accounts were excluded due to privacy concerns.

6.5 Acknowledgements

Thank you to our team of economists, psychologists, data scientists and social scientists to all who contributed, and to our designer Zin Nwe Zaw Lwin, our communications team, Mark Kember and Martina Po, and our website lead, Elefteria Kokkinia. We are grateful to the individuals who participated and provided us with case studies, fireside chats and thought leadership to bring the key issues to life.

6.6 References

All sources used in establishing our thinking on this report, whilst only some are quoted:

- Akrites. "Open Letter." Accessed July 21, 2026. <https://akrites.org/letter/>
- The AI Insider. "UK AI Opportunities Action Plan, Decoded." June 1, 2026. <https://theaiinsider.tech/2026/06/01/uk-ai-opportunities-action-plan-decoded/>
- Amazon Web Services. "New AWS Report Reveals That Nearly Two-Thirds of UK Organisations Now Use AI." About Amazon UK. <https://www.aboutamazon.co.uk/news/aws/new-aws-report-reveals-that-nearly-two-thirds-of-uk-organisations-now-use-ai>
- Amazon Web Services. "AWS Global Startup Trends Report Reveals AI-Native Startups Are Reaching Billion-Dollar Valuations in Half the Time." Amazon Press Center, June 2026. <https://press.aboutamazon.com/2026/6/aws-global-startup-trends-report-reveals-ai-native-startups-are-reaching-billion-dollar-valuations-in-half-the-time>
- Al Jazeera. "China's Xi Jinping Launches New AI Alliance: What Is It?" <https://www.aljazeera.com/news/2026/7/17/chinas-xi-jinping-launches-new-ai-alliance-what-is-it>
- Bdeir, Ayah. LinkedIn post congratulating Martin Tisné. LinkedIn. https://www.linkedin.com/posts/ayah-bdeir_congratulations-to-martin-tisne-on-receiving-activity-7467969841610539008-Ceiz
- Bennett School of Public Policy. "Reimagining the AI Arms Race." University of Cambridge. <https://www.bennettschool.cam.ac.uk/publications/reimagining-the-ai-arms-race/>
- Brevia. "Milburn Review of NEET Young People 2026: Implications for Employers." 2026. <https://www.brevia.co.uk/news/brevia/milburn-review-neet-young-people-2026-implications-employers/>
- Business Leader. "LinkedIn UK Boss: Small Firms Suddenly Have the AI Advantage." <https://www.businessleader.co.uk/insights/linkedin-uk-boss-small-suddenly-ai-advantage>
- Carpe Datum Law. "Colorado's AI Reset: Two Weeks, a White House Callout, and a Pivot Away from the EU Model." May 2026. <https://www.carpedatumlaw.com/2026/05/colorados-ai-reset-two-weeks-a-white-house-callout-and-a-pivot-away-from-the-eu-model/>
- Center for Security and Emerging Technology. China's 14th Five-Year Plan: Outline of the People's Republic of China 14th Five-Year Plan for National Economic and Social Development and Long-Range Objectives for 2035. Translation. Georgetown University. https://cset.georgetown.edu/wp-content/uploads/t0284_14th_Five_Year_Plan_EN.pdf
- Chatham House. The US government's latest U-turn on Anthropic's Mythos sends mixed signals on AI governance, 2 July 2026. Accessed July 22, 2026. <https://www.chathamhouse.org/2026/07/us-governments-latest-u-turn-anthropics-mythos-sends-mixed-signals-ai-governance>
- Cloud Security Alliance. "New Cloud Security Alliance Survey Reveals 82% of Enterprises Have Unknown AI Agents in Their Environments." Press release, April 21, 2026. <https://cloudsecurityalliance.org/press-releases/2026/04/21/new-cloud-security-alliance-survey-reveals-82-of-enterprises-have-unknown-ai-agents-in-their-environments>
- CNBC. "Anthropic's Fable Shutdown Is a Big Moment for Open-Source AI." June 16, 2026. <https://www.cnbc.com/2026/06/16/anthropics-fable-shutdown-is-a-big-moment-for-open-source-ai.html>
- CNBC. "China's Moonshot AI Unveils Kimi K3 That Rivals OpenAI, Anthropic." July 17, 2026. <https://www.cnbc.com/2026/07/17/moon-shot-ai-kimi-k3-model-openai-anthropic-china.html>
- Coinbase. "Coinbase and Cloudflare Will Launch x402 Foundation." Coinbase Blog. <https://www.coinbase.com/blog/coinbase-and-cloudflare-will-launch-x402-foundation>
- Computer Weekly. "Data Dive: Kill Switch and Catch-Up—Can Europe Close the Sovereignty Gap?" Accessed July 21, 2026. <https://www.computerweekly.com/news/366645336/Data-dive-Kill-switch-and-catch-up-can-Europe-close-the-sovereignty-gap>
- DeepSeek. "DeepSeek-R1 Release." DeepSeek API Docs, January 20, 2025. <https://api-docs.deepseek.com/news/news250120/>
- Diginomica. "Agent Identity Gets a DNS Anchor: How GoDaddy Built Agent Name Service, Then Gave It to the Linux Foundation." <https://diginomica.com/agent-identity-gets-dns-anchor-how-godaddy-built-agent-name-service-then-gave-it-linux-foundation>
- The Economist. "Meta Is Accused of Bullying the Open-Source Community." August 28, 2024. <https://www.economist.com/business/2024/08/28/meta-is-accused-of-bullying-the-open-source-community>
- ETF Database. "Open-Source AI Models Are Eating the Frontier: Where Value Goes." Accessed July 21, 2026. <https://etfdb.com/artificial-intelligence-content-hub/open-source-ai-models-eat-frontier/>
- European Commission. "AI Continent Action Plan." Shaping Europe's Digital Future. <https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan>
- European Commission. "Apply AI Strategy." Shaping Europe's Digital Future. <https://digital-strategy.ec.europa.eu/en/policies/apply-ai>
- European Commission. "A European Approach to Artificial Intelligence." Shaping Europe's Digital Future. <https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>
- Fortune. "Anthropic's Fable Fiasco Leaves the Door Open for Open-Source AI, Particularly Cheaper Models from China." June 16, 2026. <https://fortune.com/2026/06/16/us-anthropic-ban-open-source-ai-deepseek-zai/>
- Ganglani, Kunal. "7 Open Source AI Projects Every Developer Needs." June 18, 2026. <https://www.kunalganglani.com/blog/open-source-ai-projects-developers-2026>
- GitHub. "Octoverse: A New Developer Joins GitHub Every Second as AI Leads TypeScript to #1." The GitHub Blog. Accessed July 21, 2026. <https://github.blog/news-insights/octoverse/octoverse-a-new-developer-joins-github-every-second-as-ai-leads-typescript-to-1/>

- GitHub. “The Latest Octoverse Findings: How AI Is Reshaping Developer Choices.” The GitHub Blog. February 2026. <https://github.blog/news-insights/octoverse/>
- The Guardian. “UK Sets Out AI Infrastructure Push at London Tech Week—How Does It Stack Up?” June 13, 2026. <https://www.theguardian.com/technology/2026/jun/13/uk-ai-hardware-london-tech-week-investment-chips>
- Help Net Security. “Nearly Half of Open-Source AI Projects Never Reach Production.” July 20, 2026. <https://www.helpnetsecurity.com/2026/07/20/mozilla-open-source-ai-adoption-report/>
- House of Commons Library. Research Briefing SN06705. UK Parliament. <https://commonslibrary.parliament.uk/research-briefings/sn06705/>
- Hugging Face. “State of Open Source on Hugging Face: Spring 2026.” March 2026. <https://huggingface.co/blog/huggingface/state-of-os-hf-spring-2026>
- Krikorian, Raffi. “Mozilla’s Inaugural ‘State of Open Source AI’ Report Is Here.” The Mozilla Blog, July 14, 2026. <https://blog.mozilla.org/en/mozilla/mozilla-state-of-open-source-ai-report/>
- LinkedIn News. “Gen Z Lead UK Founder Boom.” LinkedIn. <https://www.linkedin.com/news/story/gen-z-lead-uk-founder-boom-7252148/>
- Linux Foundation. “Linux Foundation Announces Intent to Launch Agent Name Service to Establish Trusted Identity Infrastructure for AI Agents.” Press release. <https://www.linuxfoundation.org/press/linux-foundation-announces-intent-to-launch-agent-name-service-to-establish-trusted-identity-infrastructure-for-ai-agents>
- Linux Foundation. “Linux Foundation Announces Operational Launch of x402 Foundation to Standardize Internet-Native Payments for AI Agents and Applications.” Press release. <https://www.linuxfoundation.org/press/linux-foundation-announces-operation-al-launch-of-x402-foundation-to-standardize-internet-native-payments-for-ai-agents-and-applications>
- McGuireWoods. “Trump Administration Executive Orders Seek to Deregulate Artificial Intelligence.” Client alert, February 2025. <https://www.mcguirewoods.com/client-resources/alerts/2025/2/trump-administration-executive-orders-seek-to-deregulate-artificial-intelligence/>
- Meta. “Meta and Microsoft Introduce the Next Generation of Llama.” Meta Newsroom, July 2023. <https://about.fb.com/news/2023/07/llama-2/>; Meta. “Llama 2 Community License Agreement.” <https://ai.meta.com/llama/license/>
- Microsoft Research. Microsoft AI Diffusion Report 2026 Q1. Microsoft, 2026. <https://www.microsoft.com/en-us/research/wp-content/uploads/2026/05/Microsoft-AI-Diffusion-Report-2026-Q1.pdf>
- Ministry of External Affairs, Government of India. Bilateral document 40809. <https://www.mea.gov.in/bilateral-documents?dtl/40809>
- Moonshot AI. “Kimi K2.6.” Kimi Blog. <https://www.kimi.com/blog/kimi-k2-6>
- Mozilla. “The State of Open Source AI - V1.0 July 2026” <https://stateofopensource.ai/>
- Murray, Wallace & Chen. “Executive Order Targets State AI Regulation through Federal Preemption.” January 20, 2026. <https://mwcllc.com/2026/01/20/executive-order-targets-state-ai-regulation-through-federal-preemption/>
- National Cyber Security Centre. “The ‘Vibe Coding Spectrum’ Approach to AI-Assisted Software Development.” Accessed July 21, 2026. <https://www.ncsc.gov.uk/blogs/the-vibe-coding-spectrum-approach-to-ai-assisted-software-development>
- The New Stack. “Linux creator Linus Torvalds tells AI haters to walk away from Linux, or go fork it”. July 16, 2026. Accessed July 22, 2026. <https://thenewstack.io/torvalds-linux-ai-stance/>
- The New York Times. “Mark Zuckerberg, Meta and AI.” May 29, 2024. <https://www.nytimes.com/2024/05/29/technology/mark-zuckerberg-meta-ai.html>
- Open Atom Foundation. “Open Source 2025.” <https://openatom.cn/deep-report/index/main/2025>
- OpenAI. “Patch the Planet: A Daybreak Initiative to Support Open Source Maintainers.” Accessed July 21, 2026. <https://openai.com/index/patch-the-planet/>
- OpenAI. The Shift to Agentic AI: Evidence from Codex. <https://cdn.openai.com/pdf/5d1e1489-21c0-43e4-9d42-f87efdbf0082/the-shift-to-agentic-ai-evidence-from-codex.pdf>
- Open Source Initiative. “The 2026 State of Open Source Report.” April 28, 2026. <https://opensource.org/blog/the-2026-state-of-open-source-report>
- OpenSourceForU. “Working With Open Source AI Frameworks on Windows.” <https://www.opensourceforu.com/2026/07/working-with-open-source-ai-frameworks-on-windows/>
- OpenUK. State of Open: The UK in 2023. 2023. <https://openuk.uk/report/state-of-open-the-uk-in-2023/stateofopen-2023/>
- OpenUK. State of Open: The UK in 2024, Phase Four. December 2024. <https://openuk.uk/wp-content/uploads/2024/12/State-of-Open-The-UK-in-2024-Phase-Four.pdf>
- OpenUK. AI Report. February 2025. <https://openuk.uk/wp-content/uploads/2025/02/AI-Report-FINAL.pdf>
- OpenUK. AI Openness Update Report. July 2025. <https://openuk.uk/wp-content/uploads/2025/07/OpenUK-AI-Openness-update-report.pdf>
- OpenUK. “Public Good AI.” July 2025. <https://openuk.uk/stateofopen/publicgoodai/>
- OpenUK. AI Impact Summit 2026: AI Openness Report. February 2026. <https://openuk.uk/wp-content/uploads/2026/02/AI-Impact-Summit-2026-AI-Openness-Report.pdf>
- Organisation for Economic Co-operation and Development. Benefits of AI Openness. OECD Publishing. https://www.oecd.org/en/publications/benefits-of-ai-openness_746e8c9a-en.html
- Oxford Insights. Government AI Readiness Index 2025. 2025. <https://oxfordinsights.com/ai-readiness/government-ai-readiness-index-2025/>
- Oxford Insights. “Tracking Progress on the UK’s AI Opportunities Action Plan: Implementation Is What Counts for Growth.” <https://oxfordinsights.com/insights/tracking-progress-on-the-uks-ai-opportunities-action-plan-implementation-is-what-counts-for-growth/>
- Permanent Mission of France to the United Nations. “Statement on Inclusive and Sustainable Artificial Intelligence for People and the Planet.” <https://onu.delegfrance.org/statement-on-inclusive-and-sustainable-artificial-intelligence-for-people-and-planet/>
- Qwen. “Qwen3.7.” Qwen Blog. <https://qwen.ai/blog?id=qwen3.7>
- Rombach, Robin. “Our Co-founder and CEO Urges G7 Leaders to Back Open Innovation.” Black Forest Labs, June 2026. <https://bfl.ai/blog/our-co-founder-and-ceo-urges-g7-leaders-to-back-open-innovation>
- Sarvam AI. “India’s Sovereign LLM.” Sarvam AI Blog. <https://www.sarvam.ai/blogs/indias-sovereign-llm>
- Sarvam AI. “Sarvam 30B and 105B.” Sarvam AI Blog. <https://www.sarvam.ai/blogs/sarvam-30b-105b>

- Sifted. “Mistral CEO Pitches Open Source AI Days after Anthropic Ban: ‘We Exist Outside of State Control.’” Accessed July 21, 2026. <https://sifted.eu/articles/mistral-arthur-mensch-open-source-anthropic>
- SUSE. “Reflections from Brussels, Utrecht and Paris.” SUSE Communities. Accessed July 21, 2026. <https://www.suse.com/c/reflections-from-brussels-utrecht-and-paris/>
- The Stack. “Hugging Face Hacked: Turned to Chinese LLM for Help after US Models Blocked Blue Team.” Accessed July 21, 2026. <https://www.thestack.technology/hugging-face-hacked-turned-to-chinese-llm-for-help-after-us-models-blocked-blue-team/>
- Stanford Institute for Human-Centered Artificial Intelligence. AI Index Report 2026. Stanford University, 2026. <https://hai.stanford.edu/ai-index/2026-ai-index-report>
- State of Open Con. <https://stateofopencon.com/>
- Stripe. “Stripe Tour London 2026.” Stripe Newsroom, 2026. <https://stripe.com/en-de/newsroom/news/stripe-tour-london-2026>
- Stripe Economics. “Solopreneurs, Solow and the SaaSocalypse.” <https://www.stripeeconomics.com/p/solopreneurs-solow-and-the-saasocalypse>
- TechnologyChecker.io. “Open-Source AI Adoption 2026: 5.6M Projects vs. Real Deployment.” May 29, 2026. <https://technologychecker.io/blog/open-source-ai-adoption>
- Tony Blair Institute. “Open Source: How Middle Powers Can Build Influence in the Age of AI” <https://institute.global/insights/tech-and-digitalisation/open-source-influence-age-of-ai>
- Tony Blair Insitute. “The UK Doesn’t Need Its Own ChatGPT, It Needs a National Open-Source AI Lab” <https://institute.global/insights/tech-and-digitalisation/the-uk-doesnt-need-its-own-chatgpt-it-needs-a-national-open-source-ai-lab>
- UK Department for Science, Innovation and Technology. “AI Opportunities Action Plan.” GOV.UK. <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>
- Wiley Rein. “White House Launches AI Action Plan and Executive Orders to Promote Innovation, Infrastructure, and International Diplomacy and Security.” Client alert, July 2025. <https://www.wiley.law/alert-White-House-Launches-AI-Action-Plan-and-Executive-Orders-to-Promote-Innovation-Infrastructure-and-International-Diplomacy-and-Security>
- Wilson Elser. “Artificial Intelligence Legislative Update.” <https://www.wilsonelser.com/publications/artificial-intelligence-legislative-update>
- Wirtschaftskammer Österreich. China’s 15th Five-Year Plan (2026–2030). EUCLERA translation. <https://www.wko.at/ktn/aussen-wirtschaft/euclera-translation-15th-five-year-plan-2026-2030-.pdf>

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.



[OpenUK Newsletter](#)



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

©OpenUK and OpenHQ 2026 - Registered Office: 8 Coldbath Square, London
EC1R 5HL Company Number: 11209475 - VAT Registration: GB379697512



