Hidden in Plain Sight

How Governments Can Unlock the Next Wave of Open Source AI Innovation

April 2026

STARTUP
C*ALITION

Contents

Introduction

Governments across Europe are racing to define their place in the global AI economy and work out what AI sovereignty means to them. In doing so, they are in danger of ignoring a key tool of strategic leverage for both concerns - open source. Open source - systems developed in ways that enable users to access, modify, tailor and build on and share them freely - can give greater degrees of control and transparency than proprietary alternatives.

Open source software already underpins much of the modern digital economy. Yet government engagement with the ecosystem that produces this infrastructure - and the startup ecosystem that makes use of it - can be shallow. Open source doesn't look like other tech projects that governments are used to engaging with. This is not because the case for openness is weak, but because its ecosystem can be hard to engage with, and has been systemically misunderstood. As a result, governments have too long seen developer communities and side projects, when instead they're looking at the early stages of novel strategic infrastructure. With the rise of open source AI, this misperception and sidelining is a growing strategic problem. A new generation of AI companies is emerging and the Government needs to recognise them. This report explores why this is an ecosystem that nations who want to seize the opportunity AI offers cannot afford to ignore. Openness in AI can unlock the next startup wave.

Arguably open source AI's public tipping point came in January 2025, when Chinese AI laboratory DeepSeek released a model that matched or exceeded the performance of America's leading proprietary systems - and at a fraction of the cost. So disruptive was its release that it wiped billions from US AI stocks. DeepSeek had such impact in part because it was 'open' when compared to closed, proprietary, AI. Unlike a fully open sourced AI system - DeepSeek didn't disclose the datasets or training code used to create its models - but it did release its model weights and made its models free to download, use and modify. Within days, DeepSeek's models had been downloaded millions of times, and used by developers worldwide. Over 1000 models based on DeepSeek-V3 and R1 are now available on platforms such as HuggingFace. DeepSeek did not merely demonstrate Chinese technical capability; it announced a new phase of technological competition in which open source has become a shaper of geopolitical tensions.

Britain's political leadership is increasingly discussing the merits of open source AI as a strategy for Britain to pursue, with AI Minister, Kanishka Narayan, declaring in a speech in February 2026, that he will "reaffirm today what we have felt deeply in government: Britain will be the home of global open source AI talent" and that the Government, in institutions like the UK AI Security Institute (AISI), were building on open source infrastructure (this in reference the AISI platform, *Inspect* - an open source framework for LLM evaluations).

The strategic case for open source AI grows more urgent when we examine what Europe stands to lose if it fails to act, while nations with explicit pro-open-source strategies gain ground. Europe must pursue a distinctive path to the great AI powers, one that delivers strategic sovereignty capabilities.

Why Write This Report?

Despite growing public investments in AI across Europe and renewed political interest in tech sovereignty and innovation, policymakers lack a compelling, human story about the role open source plays in powering startups, driving competition, and lowering market barriers.

This report, supported by Mozilla, gives a new perspective on the role of open source AI models in providing open innovation - a foundational ingredient in Europe and the US's innovation economies - on par with access to capital. It aims to give policymakers and officials an understanding that goes beyond the abstract and explores what is at stake if we do not treat open innovation responsibly. While policymakers are beginning to talk more about the merits of open source AI and open innovation, talk is not enough, and policy change will be required.

Indeed, despite growing political interest in open source AI, government engagement with the open source ecosystem remains shallow. This is not because the case for open source is weak — it is overwhelming, but it is likely because the people building in open source do not look, sound, or operate like the technology companies the government is accustomed to dealing with. As many reading this report will know, open source founders are more likely to describe themselves as maintainers than CEOs, and their projects often resemble community experiments before they resemble businesses. The government has, as a result, placed them in a box marked "niche developer community" and moved on.

This report argues that this is a strategic error. Beneath the surface of what looks like 'tinkerer culture', commercially viable companies are emerging, and some are already supplying European governments with critical infrastructure. The challenge is not whether open source AI matters. It is that the government has systematically misread the ecosystem that produces it, and in doing so, has failed to support the transition from open source project to open source business.

By interviewing startup founders, venture capitalists, and ecosystem stakeholders, the research presented here aims to highlight real-world evidence of how open source AI strategies are shaping product development, market entry, and long-term viability in AI ventures. The following is a result of these conversations.

Why Does Open Source AI Even Matter?

The Threat of Dependency

Open source AI encompasses systems that are created in ways that allow anyone to access, use, modify and build on them. This openness gives developers and organisations a degree of control and transparency not found anywhere else.

Because of this, open source AI models offer a route out of reliance and vendor lock-in. Many defenders of the merits of open source AI look to the current dominance of big tech AI firms. For example, Paul Drayson, the chairman of UK startup Locai Labs, described a dominance of a select few firms and their proprietary models and software:

> "We've seen how society has been changed by social media. It's going to get changed by AI. We'd better have the ability to have an influence on how AI changes our society. And we won't do that if all we can do is use other people's models."

It is this ability to have control over AI models which leads developers to open source. As Drayson explains:

> "The model, the biases in the model, the knowledge base, determines the characteristics of the intelligence that you're using. And this new intelligence is going to change society. We'd better make sure that we have the ability to influence how it's developed."

For many who advocate for the widespread adoption and use of open source AI models, the stakes extend to the fundamental question of who shapes how AI mediates services, information access, and institutional decision-making in the AI age.

Sovereign Capabilities

For many advocates of open source AI, Governments taking the investment and advocacy of open source seriously is critical for building sovereign capabilities, which can contribute to international relevance in the stack. Enterprise open source potentially represents not just a procurement option, but an enabler of strategic sovereignty, economic growth, security resilience, and global influence, all whilst reducing lock-in risk and controlling costs more effectively than traditional proprietary alternatives. Indeed, many challenger scaleups such as Mistral, have leveraged the competitive edge that open source offers.

The competitive threat is not merely about individual products but about systematic market capture across the entire technology stack. This ambition is increasingly unconstrained. The result is a race to capture every layer of the AI stack, from chips and cloud infrastructure through foundation models to application layers. As Gavin Owens from Tonk Labs observes, "the frontier model companies will eat all the margins" of companies building on top of proprietary foundations.

To understand how open source AI models can shape AI sovereignty today, it is critical to understand their place in the AI stack. As shown here, open source models themselves currently only make up a small portion of the wider stack required to run and deliver generative AI. Chips and compute infrastructure, for example, remain highly concentrated proprietary technology. As a result, as the Tony Blair Institute has pointed out, open source AI models are not a sovereignty “silver bullet”. However, open source can be a highly effective hedging strategy to ensure nations can secure reliable access to AI on terms that preserve agency and choice while preventing the loss of control over where value is created. In the future, the more open the whole stack is, the fewer rents will have to be paid to access and develop AI capabilities.

Even the most successful AI nations know the power of open source. America has long centred an open source AI strategy while China is increasingly using open source AI as a sophisticated geopolitical strategy - and means of commercial competition. By releasing capable open models, Chinese companies and institutions embed their approaches into global AI development, and developers worldwide build on these foundations.

The China Threat

China's approach to open source AI represents a potentially sophisticated geopolitical strategy, not necessarily just commercial competition. Open source functions here as a form of geopolitical soft power. By releasing capable open models, Chinese companies and institutions embed their approaches into global AI development, and developers worldwide build on these foundations.

There was much disagreement among many of the interviewees we spoke to over the extent to which the proliferation of open source AI models from China was a threat or concern to Europe. At the very least, it should be noted that huge American big tech companies, such as Airbnb, are building on Chinese open source models to deliver parts of their product or service. John Spindler, a venture capitalist at TwinPath Ventures, explained:

> "The Chinese view is that they may not be able to keep up with closed models, but they can sully the water by keeping releasing lots of open source models, making the economics of building out on the back of LLMs difficult if you've already given a free version."

Lucie-Aimée from Hugging Face captures the dynamic of open source AI models changing geopolitical tensions, saying in her interview that:

> "Opening up models can really put countries on the map, e.g. DeepSeek and China. China was just not as present on the map for AI until they decided to make a model publicly available and get people to really reuse it."

Yet this same dynamic exposes the limits of the current open source paradigm. As Raffi Krikorian, CTO of Mozilla, points out, Chinese models like Qwen — which he describes as "probably the best performing open source model on the planet" — are only open-weight. Their training data, methodology, and checkpoints remain opaque. Research has already shown that Chinese and Western models can make significantly different decisions on sensitive questions, and without visibility into training data, those divergences cannot be fully understood or corrected.

Chapter 1: What is Open Source AI, and How Can the Government Properly Understand it?

As Jonny Williams from Red Hat, one of the world's leading open source companies, observed: "Open source in government is often narrowly understood as publicly releasing code. However, this represents only the first tier of open source maturity." Indeed, it is an issue that policymakers' understanding of what actually is open source AI, and how it is different from open source code, or proprietary software and AI models, is critical to ensuring it can be successfully adopted and utilised.

Open Source

Think of open source software not as just free stuff, but as roads. As one founder interviewed for this report put it: roads generate little revenue on paper. However, entire economies are built on top of them. The same foundational value is seen in open source software. When you use your smartphone, browse the web, stream video, or pay with a contactless card, you are almost certainly using open source software at multiple layers of the stack. Linux, a free to use, open source operating system, powers the vast majority of web servers globally.

The Story of Open Source

The story of open source begins with California, specifically, with the coding counterculture of the 1960s and 1970s. What started as a radical idea about sharing code among academics became, over five decades, the invisible foundation of the digital economy.

> "If you trace back open source, it comes from academia," explained Sam Hields from OpenOcean. "That's really the roots, that's why the most common license is called the MIT license, literally from MIT."

What began as a philosophical stance rooted in the belief of all developers being able to view, edit and play with the code that ran computers and the internet, evolved into something with profound economic consequences.

Today, as Gavin Owens, CEO of Tonk Labs, put it bluntly: "If you took away all the open source software, literally nothing would run. The Internet would just disappear. Your computer wouldn't boot up." A 2024 Harvard Business School study attempted to quantify this value by calculating what it would cost to recreate open source software from scratch using paid developers. Their estimate: $8.8 trillion in societal value, not the value of "free stuff," but the value of shared infrastructure that eliminates wasteful duplication while enabling innovation.

However, Western governments have largely treated open source as something that happens in the private sector, not something for the government to take a stance on, whether to critique or endorse.

Meanwhile, China has integrated it into its national strategy. DeepSeek was not an accident but a statement of intent: China can compete at the frontier, and China will compete in the open. The irony is acute. Open source emerged from American universities and Californian counterculture. Its values, transparency, collaboration, and distributed innovation align with Western political traditions. Yet Western governments have neglected it, while an authoritarian state has grasped its strategic potential.

Open source AI is a development from the open source tradition, with much disagreement and debate about how it is different from open source, and Michael Webb has written an essay specifically on this topic that you can read here.

Open source has become ubiquitous in the modern economy for many reasons:

- Firstly, it provides cost efficient methods of building infrastructure, where developers need not start from scratch.
- Secondly, when code is open for all developers to see, it can be inspected by all, and flaws in the code can be noticed and shared, improving the security of the open source software used.
- Thirdly, it deters against vendor lock-in, enabling organisations to retain control over their technology stack.
- Finally, open source software leads to the creation of bustling communities of developers around the software, specifically, this may look like building programs on top of it.

Open Source AI

Open source AI is different from simple open source code, as it consists of different components (the complexity of what does consist of ‘open source AI’ means it is highly debated within the community). Open source AI models may consist of: model weights, model architectures, the training data, the training methodology, or the inference code. An open source AI model may contain some of these components, or all. For everything to be

open makes an open source model truly open, combined with it being freely available for anyone to use, modify, and distribute. Andreas Liesenfeld and Mark Dingemanse have a page on GitHub that keeps track of whether models are truly open or not, here.

An open weights model is when models release trained parameters but may keep the training data, methods and architectural details proprietary – this is often the most common form of open sourcing an AI model – open weights.

For the purposes of this report, we recognise that OSAI has a strict definition but colloquially for a lot of people when they refer to open source ai they are referring to open weights, so when we discuss open source AI, we're primarily talking about models where weights are available, can be customised, and can be deployed independently. Whether they meet every OSAI criterion is less important than whether they enable the kind of innovation, competition, and sovereignty that closed systems cannot.

There was constant contestation, even from the interviews we conducted for this report. For example, Matt Squire from Fuzzy Labs, an AI consultancy working with UK public sector clients, applied a stricter standard:

> "Even the likes of Mistral won't tell you everything they've used to train their model. There are very few truly open source AI models that actually share everything, absolutely everything. For a model to be really open source, it should be feasible for somebody to reproduce that model from scratch if they wanted to."

Not everyone in the ecosystem views this strict standard as the most productive framing. Raffi Krikorian, CTO of Mozilla, describes holding two positions simultaneously: a philosophical stance that true openness means everything, "not just all the tools, but also the datasets, also the evaluation sets, also the training mechanisms, all the checkpoints", and a pragmatic recognition that insisting on this standard right now risks missing the bigger picture. "Open weights might be a good middle ground," he argues, "because we have to worry about the openness of the entire system." In his view, focusing exclusively on model openness makes it "easy to suck up all the oxygen" when the real bottlenecks lie elsewhere in the stack — developer tooling, deployment infrastructure, and critically, the hardware layer, where "there is only one infrastructure provider for the entire world." For policymakers, this reframing matters: a strategy focused solely on open source models, without addressing the proprietary chokepoints in compute and tooling that surround them, risks solving only part of the problem.

It is this dynamic of playing with and changing open source models and weights that accelerates the speed of development in AI. This is what will supercharge the AI race, and is already doing so. This is demonstrated in the significant development of Western startups

using Chinese open source AI models to build on top of, with speed, in order to compete in the digital sector.

This paradigm shift means that nations and companies no longer need unlimited resources or have to necessarily acquire proprietary alternatives to compete at the frontier. Indeed, whether the proprietary models can even compete with these open source AI models is a question in itself.

Contributions and Reusing as Critical to Open Source

The UK's Government Digital Service does publish code on GitHub, and it does attract developers and a community to build on it. But fundamentally, these repositories exist as archives, and not living projects which genuinely contribute to the innovation seen within typical open source community development. This is an approach of simply, "write code in the open and hope someone picks it up". This is a fundamentally passive approach that misses the real opportunity which open source offers. It is based on the assumption that transparency alone generates value, that simply making code public will somehow attract a community of volunteers eager to improve government software in their spare time. It almost never works that way. As Matt Squire at Fuzzy Labs argued,

> "I think government understanding of open source is poor, but that's simply a lack of knowledge. The big vendors are spending all their marketing budgets and sales budgets on talking to the right people and convincing them to buy their technology. So it's not obvious that there are other ways of doing this."

This gap in understanding is not unique to the UK. Raffi Krikorian, whose role at Mozilla involves regular engagement with governments internationally, offered a blunt assessment, where, at the India AI Summit, he observed ministers across multiple countries arriving at an interest in open source AI but without a clear framework for acting on it — "bumbling their way into it," as he put it. "There would be a great service if someone could credibly go to all of them and be like, here is the state of play. But no one has done that in a credible way yet." The problem, in other words, is not that governments are hostile to open source AI. It is that they lack a clear, credible account of what it is, what it enables, and where to direct resources.

Indeed, a key concept which is left completely unaddressed in much existing Government policy strategy is that of reuse and development through open source. True open source is not about building something once and publishing it. It's about building on what others have built, contributing improvements back, and enabling others to build on your work in turn. This is "recommoning", the continuous process of taking from, and contributing to, shared

resources. Currently, UK government projects rarely participate in this cycle. They don't systematically build on existing open source tools, and when they do build something, it's rarely designed for others to easily adopt and adapt. This is what enabled open source AI to really accelerate the race.

Many interviewees discussed what they viewed as the Government's current approach to open source AI, which did not meet the moment, or embrace the potential of OSAI. One interviewee expressed frustration with the British government over the fact that it did not have a positive story on the potential of OSAI, he said:

> "The UK Government has done a pretty bad job globally speaking in like, celebrating the role of open source… there's been so much hammering about safety, existential risk, catastrophic risk, and a lot of that has implicitly cast a cloud over open source in Europe". Concerns about accelerating the AI race through the release of open source models have seemed to have inhibited the potential of the Government to actively embrace and support the release of good open source technology.

Indeed, Ben Brooks, affiliate with the Berkman Klein Center, agreed, saying:,

> "If any government is serious about sovereign AI, that doesn't just mean training the models in a particular jurisdiction. It also means establishing the capacity and the culture to take open technology, experiment with it, build on it, and figure out how to make it useful across the economy. That is a broad integration and deployment challenge, not just a narrow development challenge."

There is not a clear strategy emanating from the British government to address where they can capitalise on and support the proliferation of open source AI. Indeed, in particular in the case of the strategy of pursuing sovereign AI, it is notable that there is not much discussion about the potential of open source AI models to support this aim.

Chapter 2: Open Source AI Firms and Startups

Why should organisations build on open source AI?

Strategic Flexibility

Open source AI offers a much higher degree of flexibility for founders to build on top of. As CEO of Tonk Labs, Gavin Owens outlined in an interview, proprietary AI models have captured the foundation layer and have become the incumbents in the AI sphere. As a result, challenging these companies by building on proprietary models is nearly impossible. Recognising this reality, Owens suggests that the exception, or challenge to this, will come from the proliferation of open source AI, as a flurry of innovation will come from the use and reuse of these models, as opposed to the potential stasis due to the dominance of the existing proprietary AI incumbents.

For countries that cannot match American or Chinese levels of capital, open source AI offers something more than a cost-saving measure, it offers a collaborative development model that turns a collective weakness into a structural advantage. Raffi Krikorian, CTO of Mozilla, draws on the precedent of Linux: "In the United States, you could afford to have a Microsoft that builds its own proprietary system that everyone can use, but the rest of the world, you probably can't. And so Linux actually turned out to be a really good answer."

He envisions a similar model for European AI: shared open source infrastructure that individual nations fine-tune to their own needs. "You can imagine a world where the Brits build on top of it, but do their own thing. The Norwegians build on top of it, but do their own thing. But core infrastructure could be shared." Crucially, this model preserves data sovereignty, "the Brits don't have to give up their data. You could train your own model, but use this code base to go do it", while avoiding the duplication of effort that would come from each European nation attempting to build AI capabilities independently.

The alternative, Krikorian argues, is strategically untenable: "Every single European state trying to do everything soup to nuts requires both too much capital and too much talent." For European policymakers weighing the case for open source AI, this is perhaps the most consequential framing: not open source as ideology, but open source as the only viable architecture for collective competition.

Market Positioning and Credible Neutrality

Open source also serves as a powerful market-entry strategy, particularly for startups competing against established players. Lucie-Aimée from Hugging Face described how open-sourcing models gets companies "on the map" in developer communities.

> "There was this kind of mouth-to-mouth word propaganda for the people that really mattered," she notes. "If you want to develop AI models, you want AI developers to know you. The easiest way to do that is by making sure they have access to your models."

This strategy builds upon the development of open source software companies. For decades, conventional wisdom held that open source companies faced inherent disadvantages in raising capital and achieving exits. The logic seemed straightforward: if your core technology is freely available, investors reasoned, how can you build a defensible business? This assumption has shaped countless funding decisions and policy frameworks. It is also comprehensively wrong.

A landmark study published by Serena VC in collaboration with the Linux Foundation analysed 25 years of investment data across thousands of open source software companies. The findings challenge nearly every assumption about open source as a business strategy. Open source companies don't just compete with their closed-source counterparts; they also outperform them.

The data reveals that open source startups raise funding faster, command higher valuations, close larger funding rounds, and graduate to subsequent funding stages at superior rates. Perhaps most strikingly, 12% of open source companies achieve liquidity through mergers, acquisitions, or public offerings, a figure that significantly exceeds closed-source benchmarks. When these exits occur, they command higher valuations than comparable proprietary companies. The study found no timeline penalty for reaching these exits; the persistent myth that open source companies take longer to mature finds no support in the evidence.

Indeed, in one of our interviews, we spoke to Matthieu Lavergne, a partner at Serena VC who has become a significant advocate for open source AI:

> "We found out just the opposite of what everyone was saying... When you're an open source founder, the financing journey is actually better than the closed source software company in terms of speed, valuation, round size, graduation rates, and every KPIs of the fundraising funnel"

These findings align with broader economic research on the value of open source to the global economy.

Open source AI is already disrupting the economics of building - economist Frank Nagle has found open source AI models have become, on average, 6x cheaper to use than equivalent closed models, and they are narrowing the performance gap with each release.

Technical Advantages and Innovation

The technical benefits of open source AI extend well beyond cost savings. Red Hat, the world's largest open source company, operates as a platform provider that helps organisations fine-tune, quantise, and add retrieval-augmented generation (RAG) capabilities to open source models like Llama and Granite. Jonny Williams, Red Hat's Public Sector AI and Platform Lead, points to an Australian startup building a "sovereign AI factory" using Red Hat's open source foundation. This approach gives the company flexibility to deploy across edge computing and cloud environments with complete consistency, something much harder to achieve with proprietary, cloud-locked systems.

Tonk Labs views this flexibility of open source as central to their long-term strategy. Gavin Owens sees LLMs not as specialised tools for chat but as general computational power. "Having access to open source LLMs allows for better flexibility and ability to innovate," he argues. His technical bet is that within one to two years, smaller models running locally or in-browser will reach parity with datacenter models for many coding tasks. This vision of AI capabilities that don't require constant communication with proprietary servers makes sense in an open source context.

The ability to experiment rapidly without navigating licensing negotiations accelerates innovation cycles. Startups can test approaches, iterate on architectures, and validate assumptions without the friction of enterprise procurement processes. For founders, this velocity often makes the difference between discovering product-market fit and running out of runway. Indeed, Ben Brooks, affiliate with the Berkman Klein Center, explained the rationale for this experimentation:

> "Developers and researchers can inspect open models before deployment, modify them for better performance in particular tasks, fine-tune and run them securely, and do it all without spending hundreds of millions to train their own models from scratch. We have to remember these are early days, early innings, and no one can fully anticipate how these models will be used to solve real-world problems. Enabling people to experiment, not just as users but as developers, is critical to realising the potential of these technologies."

Indeed, beyond open source AI models, there have been new innovations in open source AI tools. For example, the open source agentic AI tool, Open Claw, developed by Peter Steinberger has become popular in 2026 (although now Steinberger has joined OpenAI). Indeed, the mass adoption of various forms of open source AI tools proves that they can compete with proprietary alternatives.

Collaboration, Infrastructure and Talent Attraction

Hugging Face aims to solve a genuine infrastructure problem in AI development. "Previously, people sent model training code via email between researchers," Lucie-Aimée recalls. "Now we enable effective collaboration between large groups, individuals, and companies, reducing duplicated work through model reuse and building on existing work." The platform hosts models, datasets, and evaluation frameworks that teams can build upon rather than recreating from scratch. Hugging Face has been a significant enabler of the acceleration in open source AI development.

The regulatory environment can either support or hinder these dynamics. An interviewee noted that EU frameworks around data regulation and interoperability could significantly help open source startups. GDPR requirements that force platforms to enable data export make it easier for companies to bring users' data into their ecosystem. "The regulatory framework in the EU could be very helpful for the proliferation of open source startups," he observes, though he notes this potential remains largely unrealised in current policy. What these interventions are, will be explored towards the end of the report.

Case Study: New British Startups Building Open Source AI Models

Perhaps no UK startup better illustrates the strategic possibilities of building on open source than Locai Labs. Founded in London, the company has developed a proprietary post-training methodology that allows it to enhance open source models without the catastrophic forgetting that typically occurs when models are fine-tuned.

As their AI Officer explained:

> "We wanted to compete with the best of the best, but we don't have the resources of America and China, where they can build from scratch in these silos,". "So we wanted to build on open source, and that's also something we really believe. It's a very positive way of developing science and research."

The result is a model that, according to Locai Labs' benchmarks, beats GPT-5 on Arena Hard conversational benchmarks while being approximately one-fifth of the size of leading frontier models. Crucially, Locai Labs had to develop techniques for "de-censoring" Chinese models, which undergo mandatory censorship that makes them unsuitable for Western markets. "To make our model neutrally aligned and open and honest, we went through part of our training process in de-censoring these models," George explained.

Their business model reflects the dual opportunity in open source AI: a consumer product competing with ChatGPT on independence and privacy grounds, and an enterprise offering that builds domain-specific models for healthcare, education, and science sectors. "We see this as a really interesting opportunity for both commercial enterprises that want to have a model for them," says Paul Drayson. "What we're offering is a model which is independent of big tech."

Chapter 3: Challenges in Building in OSAI

The performance data and founder strategies paint an optimistic picture of open source as a business model. But the reality founders and investors navigate is considerably more complex. Open source creates genuine economic value and competitive advantages, yet converting those advantages into sustainable revenue remains one of the most difficult challenges in the startup ecosystem. Understanding what works, what fails, and why European companies struggle more than their American counterparts is essential for designing effective policy interventions.

Models That Work

The most successful commercialisation approach in recent years has been managed cloud services with usage-based billing. Customers don't pay for the software itself but for the convenience, reliability, and scale of a platform, often with consumption-based pricing tied to data processing volumes. This creates strong net revenue retention as customer usage compounds over time.

The open core approach, maintaining an open source community edition while reserving certain features for the commercial versions, has also produced successful companies. Enterprise features like clustering, advanced security, collaboration tools, and single sign-on often sit behind commercial licenses while the core product remains free. This balances community adoption with revenue generation, though it requires careful judgment about where to draw the line.

Strategic Open Source from Big Tech

A further complication is that the largest contributors to open source AI are not startups but big technology companies pursuing strategic objectives. Sam Hields from Open Ocean explained that companies like Meta and Microsoft produce many of the most popular open source AI libraries. He noted that "When they open source something, competitors can't run away with it, and they reduce internal maintenance because the community maintains what they release."

As a result, Sam argued that successful open source companies must target "extreme technical complexity" and "sophisticated engineers who genuinely need help." These successful open source companies may, as a result, sit deep in the data infrastructure stack, far from end users, requiring specialised expertise to deploy at scale. This stiff and complex

competition may mean that only companies that can establish themselves as essential maintainers of complex infrastructure achieve pricing power.

The Tinkerer Problem: Why Governments Can't See What's in Front of Them

Will Bennett from Seedcamp identified what is perhaps the central tension in European open source AI: "A lot of European open source is embedded in what we call 'tinkerer land.' It doesn't get fully commercialised. It becomes a bit of a project or an artefact. Open source can be a synonym for 'that's their hobby' in Europe, whereas in the US, open source is a synonym for GTM [Go-to Market] strategy."

As a result, there is a cultural gap between the open source AI community and the Government, where the government does not necessarily understand the value that it can capture from the open source AI community. Open source culture has prioritised community reuse and collective development, often above strict go-to-market strategies or scaling ambitions. However, this historical reality need not determine the future.

From this, the Government sees a landscape of developer communities, side projects, and technical curiosities where it should see the early stages of critical infrastructure. The open source developers building encrypted communications protocols, AI evaluation frameworks, or sovereign data tools are doing work that directly serves the government's stated objectives around strategic sovereignty for countries and reducing dependency on foreign technology providers. But because they do not arrive with a pitch deck and a sales team, they remain invisible to the officials designing AI strategy.

Yet the assumption that open source in Europe is all a hobby and no business is itself increasingly outdated. A growing number of projects are making the transition from community experiment to commercial product — and some are already serving government clients. The Matrix protocol, for instance, began as an open source project for decentralised, encrypted communications. On first encounter, it looks like a classic tinkerer endeavour: a foundation-governed protocol maintained by a passionate community. But, in reality, Matrix is now used by the French and German governments, among others, for secure internal communications, as it offers an independent, fully inspectable alternative to proprietary platforms. Further, Linux and MySQL have also proved themselves to be internationally competitive European open source systems. This is the exact kind of model which the British Government should have in mind when considering what the height of their ambition should be for using and building on open source AI technologies.

Therefore, the challenge must be for policymakers to recognise that the open source ecosystem produces value through a different mechanism than the proprietary technology sector, and that the government's current engagement model — which waits for companies to look and behave like conventional enterprise vendors before taking them seriously — systematically excludes innovative open source AI opportunities.

The Deployment Gap

Distinct from the perception problem is an operational one. Even organisations that recognise the value of open source AI frequently abandon it during implementation. Raffi Krikorian of Mozilla, in our conversation, estimated that approximately 90 per cent of Fortune 500 companies have experimented with open source AI, but only around 15 per cent have sustained their adoption — a gap he attributes entirely to deployment difficulty. "They've all just decided it's too hard to do, even though the math clearly shows in the long run it's cheaper."

If large, well-resourced corporations struggle with deployment, the implications for government are stark. "If that's a Fortune 500 company saying that, governments have that on a crazier scale," Krikorian observed. The problem is not that open source AI models are insufficiently capable. It is that the tooling, documentation, and operational support required to move from experimentation to production remain immature. This has direct policy implications. Investment in developer experience, such as tooling, workforce training, deployment infrastructure, may yield higher returns than investment in model development itself. A model that cannot be deployed is a model that creates no value, regardless of its benchmark performance.

The European Funding Gap

The commercialisation challenge is amplified in Europe by structural weaknesses in the funding ecosystem. Multiple founders and investors report that European venture capital is strong for early-stage funding, and for rounds up to €30-50 million, but late-stage growth capital requires American investors. Matthieu from Serena VC noted that European pension funds deploy less than 1% of assets to venture capital, compared to significantly higher allocations in the US. This creates a predictable pattern: promising European open source companies that need patient capital for the "bake time" required to build communities and convert users to customers often end up either relocating to the US or dying before reaching sustainable scale.

The cultural differences between US and European investors compound this problem. Gavin Owens from Tonk experienced this firsthand: "US VCs have higher risk tolerance and invest more in 'vibes.' EU VCs demand more proof and are less imaginative." Multiple founders echoed this assessment. John Spindler from Twinpath described two main investor types: "'vibe' investors who pattern match with minimal due diligence, and thesis builders who want technical and commercial due diligence on signals." He notes that pre-Series A, there's insufficient data for rigorous modelling, so investment relies on validating a pathway from free to paid conversion. European investors, he suggests, struggle with this ambiguity more than their American counterparts.

The Commercialisation-Innovation Trade-off

Open source creates significant economic value; the performance data proves this. It enables startups to compete against incumbents, accelerates innovation cycles, and builds collaborative infrastructure that benefits entire sectors. Yet converting these advantages into sustainable businesses can remain genuinely difficult, particularly in Europe, where late-stage capital is scarce and investor appetite for patient, community-building approaches is limited.

Lucie-Aimée from Hugging Face suggests a policy solution:

> "Instead of high-risk moves from research into industry or making your own startup, governments could facilitate this easily by providing infrastructure with the requirement that models be open-sourced. We provide resources, and what we want back is not money but that you provide the model so the next startup doesn't have to start from scratch again."

This approach acknowledges the commercialisation challenge while addressing it through public investment rather than expecting every open source project to become a venture-scale business. Not every valuable contribution to the innovation ecosystem needs to be a unicorn. Some open source infrastructure functions best as a public good with the economic returns captured diffusely across the companies that build upon it rather than concentrated in a single firm's revenue line.

The question for European policymakers is whether they will design policies that work with the economic reality of open source, supporting both the community infrastructure that creates diffuse value and the commercial players that capture concentrated returns, or whether they will continue optimising for proprietary models that generate clearer, more immediate revenue at the cost of the collaborative advantages that open source uniquely enables.

Chapter 4: Policy Recommendations

The Regulatory Opportunity

Having left the EU, the UK is not bound by the AI Act's provisions on general-purpose AI models. The Act has proved itself to not provide meaningful regulation to encourage innovation, but rather as a stifling and jarring blanket approach without clear direction. Indeed, George Drayson explained the EU AI Act problem: "The policy on general-purpose AI models asks you to provide the data your model was trained on and how many FLOPs you used. If I want to build on top of a model like Qwen, I don't know what data it was trained on. So where does it put you? It pushes you towards having to build the model yourself from scratch."

This opportunity gap, therefore, creates a distinctive opportunity to design regulatory frameworks that accommodate and encourage building on open source foundations, rather than inadvertently penalising the most resource-efficient approaches to AI development.

Ben Brooks, affiliate at the Berkman Klein Center, suggested a positive approach:

> "We should not only see open source as an exception to the norm, but try to emphasise open source development, encourage it, and legislate in a way that promotes the open release of useful technologies. From a compliance perspective, it shouldn't be harder to release an open tool or model than to keep it behind a paywall."

Policy Changes

1. Build on Open Source for Government Infrastructure

Every major AI system, whether labelled "open" or "closed," relies on open source components throughout its stack. Yet government policy treats open source and open source AI as an afterthought rather than recognising this reality. The AI Opportunities Action Plan and Sovereign AI Unit have articulated ambitions for AI leadership and AI sovereignty, but have not translated these into policies which recognise the utility of open source and support its role in building digital architectures.

The Government should recognise open source software and open source AI models as preferred foundations for government digital infrastructure, and create a binding

policy framework that makes this operational. Government departments should adopt an 'open source first' policy for AI components, analogous to the GDS open source guidance for software, with a 'comply or explain' requirement when proprietary alternatives, or when proprietary implementations of open source are used. Specific measures should include:

(a) Update the Technology Code of Practice to cover AI explicitly. This should include specific guidance on the procurement and deployment of AI models. This guidance should establish a presumption in favour of open source AI models (as defined below) for government use cases, with a "comply or explain" requirement when proprietary AI models or proprietary implementations of open source models are selected instead. The point - to create a clear moment that forces officials to engage critically with the reflex of choosing a closed-model path.

(b) Publish a definition of open source AI for government use. The absence of a clear, government-endorsed definition of "open source AI" creates confusion in procurement and policy. For the purposes of government policy, open source AI should be defined as models where weights are publicly available under permissive terms that allow modification and redistribution, with sufficient documentation to enable meaningful reuse. This definition should be published as formal guidance and referenced in procurement frameworks. It should be informed by, but not necessarily identical to, the Open Source Initiative's Open Source AI Definition (OSAID), and should be reviewed annually to keep pace with the technology

As one interviewee pointed out, there is an international divergence between the UK, EU, USA, and China in their approach to open source. While China has embraced open source, and Congress in the USA has increasingly been making more noise about the potential of open source AI, it has taken the EU and the UK longer to publicly come out and support open source AI. This ought to change, and Governments should firstly be steadfast in their support of open source AI and the unique potential of autonomy it can give.

2. Public Money – Public Code

Publicly funded AI projects should be required to produce open outputs — models, datasets, and tools — that remain available to the UK research and commercial ecosystem regardless of what happens to the recipient organisation.

Publicly funded AI outputs should be delivered according to adapted FAIR principles (Findability, Accessibility, Interoperability, Reuse), originally developed for research data but requiring extension for AI-specific assets – we welcome the UKRI's recent AI strategy, committing themselves to this principle too. Further, we suggest that:

- Models should be released with sufficient documentation to enable meaningful reuse: architecture details, training methodology, evaluation benchmarks, and known limitations. The Open Source Initiative's Open Source AI Definition (OSAID) provides

a useful baseline for what constitutes genuinely open release versus superficial "open-washing."
- Datasets should be released with clear provenance documentation, licensing terms, and data cards describing composition, collection methodology, and known biases.
- Code should be released under OSI-approved licences via accessible repositories.
- Amend standard funding conditions for UKRI AI grants and Innovate UK AI programmes to include open release requirements. This does not require primary legislation — it can be implemented through updated grant terms and conditions.

Therefore, publicly funded AI projects should generate public assets: open models, open datasets, and open tools that subsequent researchers and companies can build upon.

(a) Amend standard UKRI and Innovate UK grant conditions to include open release requirements. This does not require primary legislation. UKRI already sets terms and conditions for its grants and has precedent for open access requirements in research publishing. The same principle should be extended to AI outputs. Specifically:

- Models funded with public money should be released with: architecture details, training methodology, evaluation benchmarks, and documented limitations. The OSAID provides a useful baseline for what constitutes genuine open release versus superficial "open-washing." Where models are released as open weights only (without full training data), this should be explicitly documented.
- Datasets should be released with clear provenance documentation, licensing terms, and data cards describing composition, collection methodology, and known biases. Where datasets cannot be fully released due to privacy or sensitivity constraints (as will often be the case with health or administrative data), the metadata, schema, and documentation should still be published openly, with access to the underlying data governed through appropriate data access agreements.
- Code should be released under OSI-approved licences via accessible repositories such as GitHub, with documentation sufficient to enable reuse.

(b) Apply proportionate requirements to Innovate UK commercial grants. For commercially-oriented grants through Innovate UK, full open release may not always be appropriate. However, the principle should still apply in adapted form. Grant recipients should be required to release at least one of: the trained model under an open licence; the training methodology and evaluation framework; or a demonstrable contribution back to an open source project used in the development. This ensures that even commercially-oriented public funding generates some public return to the ecosystem, without requiring companies to give away their core competitive advantage.

(c) Create a national registry of publicly funded AI assets. The value of open release is undermined if outputs are scattered across individual project websites and institutional repositories with no central discoverability. UKRI should establish (or commission) a searchable registry of publicly funded AI models, datasets, and tools - a "public AI asset catalogue" - so that subsequent researchers and companies can find and build on what has already been funded. This could be integrated with the planned National Data Library infrastructure or hosted on an existing platform such as Hugging Face.

3. Funding

The UK produces world-class AI research and a thriving early-stage startup ecosystem, but consistently fails to retain and scale its most promising companies, specifically open source AI startups. Producing a successful open source AI ecosystem will mean providing incentives for companies to go to market and move beyond being tinkerers.

Europe's research community is a strength that it can leverage to deliver developments in open source AI models. For example, European academic institutions remain powerfully positioned in open source AI, continuing a legacy established in traditional open source software. According to Hugging Face data, Germany ranks second internationally for downloads on models.

As Lucie-Aimée from Hugging Face explains: "In Europe, there is a long tradition of open source AI development and open source software development at large. So in Germany as well as France, there's really strong open source development streams."

This strength reflects deeper roots: as one investor noted, "If you trace back open source, it comes from academia. That's really the roots, that's why the most common license is called the MIT license, literally from MIT. So, probably the fact that we have so many strong engineering universities and research in Europe still, probably aids open source."

Both Linux and MySQL, foundational open source software projects that underpin trillions of dollars of economic activity, originated in Finland. Oxford, Cambridge, and ETH Zurich rank among the top five computer science universities globally. It is from this historic strength in open source software development that this report wishes to explore how open source AI is affecting Europe now – where are Europe's strengths, weaknesses, threats and opportunities in its capabilities and uptake of open source AI.

As Matt Lavergne from Serena VC observed: "There is a very resilient early stage VC market. So if you want to get financing up to 30, or even, 50 millions, you can do it in Europe. But then if you want to go further, you need to talk to US investors." UK ventures rely on foreign capital for over 60% of late-stage funding, and one in three UK AI leaders surveyed by Tech Nation are actively considering relocating to access deeper capital markets abroad.

The government has tried to address this financing gap by asking the British Business Bank to cornerstone later-stage funds, as well as getting pension funds to deploy more capital into venture capital, to support the open source AI startup ecosystem.

Therefore, there must be greater acceleration on this agenda by our largest pension funds, both to plug the later-stage financing gap, but also to ensure pensioners are getting the best returns they can, whilst supporting the open source AI ecosystem.

4. Open Strategic Data Assets

The UK possesses extraordinary data assets.These assets are currently fragmented, poorly documented, and largely inaccessible for AI development. Opening them appropriately could provide UK researchers and companies with competitive advantages unavailable elsewhere.

"If governments, including ours, want to position themselves as world leaders in AI, then investment in supporting the open source community and supporting open data sets is a better investment than the infrastructure," argued Matt Squire from Fuzzy Labs. "It almost feels lazy to say, 'Oh, we want to be world leaders in AI, where does AI come from? Data centres. OK, let's build lots of data centres.' It's completely the wrong way of thinking about it. Because what are we going to train in all these data centres?"

Data investment creates demand for infrastructure rather than the reverse. UK companies with unique data-driven applications will require UK compute capacity. Building infrastructure without data assets is building supply without demand.

Therefore, the Government should prioritise opening public datasets under AI-friendly licenses, focusing on domain-specific data where the UK has distinctive holdings. The development of a National Data Library is a welcome development, yet we wait eagerly to see the structure and potential impact of it.

(a) Integrate data asset opening with the National Data Library. The development of a National Data Library is welcome, but its structure and impact remain unclear. The NDL should be designed from the outset to serve as infrastructure for AI development, not only for traditional research access. This means: supporting bulk data access suitable for model training (not just individual query access); providing API access for programmatic data retrieval; including synthetic data generation capabilities where privacy constraints prevent full data release; and hosting model evaluation datasets that allow developers to benchmark AI systems against UK-specific use cases.

(b) Establish AI-ready data standards for government. Many government datasets exist but are not in formats suitable for AI training or fine-tuning. DSIT should publish AI-ready data standards specifying: minimum documentation requirements (data cards); format and interoperability specifications; licensing terms compatible with AI model training (noting that this requires careful attention to the interaction between Open Government Licence terms and AI training use); and version control and update commitments. These standards should be integrated into the Government Data Quality Framework and the data management requirements that departments already report against.

5. Design Regulation That Rewards Openness

The EU AI Act creates significant unintended barriers to the most promising open source AI development strategies. While intended to ensure transparency and accountability, certain requirements effectively penalise companies building on open source foundations.

The Act's provisions for general-purpose AI models require disclosure of training data composition and computational resources used. For companies training models from scratch, these requirements present compliance challenges but are at least theoretically achievable. For companies building on existing open source models, the resource-efficient approach that could enable European competition may be impossible due to compliance.

While the EU AI Act suggests that AI systems released under 'free and open source licenses' are not subject to some of the regulations it outlines, not many models are completely open source, nor does the act explicitly adopt a clear definition of open source. This is due to the fact that when a company fine-tunes or extends an existing open source model like Qwen (developed in China) or Llama (developed by Meta), they inherit uncertainty about the base model's training data. The original developers may not have disclosed full training data composition, or may have disclosed it inadequately, or may have used data of uncertain provenance. The company, building on this foundation, cannot necessarily certify what they do not know.

The British Opportunity

Having left the EU, the UK is not bound by the AI Act's requirements, or more, its lack of engagement with what constitutes open source AI. The EU AI Act suggests "free and open source AI" can be defined as "AI components that] are made accessible under a free and open-source license" (recital 89) in particular their "parameters, including the weights, the information on the model architecture, and the information on model usage" (recital 102) and open source AI components cover the software, the data and the AI models (including tools, services or processes of an AI system)[1]. However, the Act does not exempt AI systems which are monetised or considered high risk.

This creates an opportunity to design regulatory frameworks that explicitly accommodate and encourage building on open source foundations. The UK could position itself as the jurisdiction where open source AI development is not merely permitted but actively supported by regulatory design. This would attract companies and researchers seeking to build on global open source foundations without facing impossible compliance burdens.

Beyond avoiding penalties for open source approaches, regulation could actively incentivise openness. Transparency requirements that are difficult for closed models to satisfy may be straightforward for fully open ones. If open source releases automatically satisfy certain regulatory obligations while closed models require extensive documentation and certification, companies would face incentives to open their work.

For this to be done, open source AI, and its variations must have robust definitions.

[1] https://legalblogs.wolterskluwer.com/copyright-blog/open-source-ai-definition-and-selected-legal-challenges/

Specific measures:

(a) Adopt a clear, operational definition of open source AI. The UK should publish a government-endorsed definition of open source AI that is clear enough to be used in regulation, procurement, and funding criteria. The definition should be tiered, recognising the spectrum from fully open (weights, data, training code, evaluation) to open-weights-only, and specifying which tier qualifies for which regulatory benefits.

(b) Create regulatory fast-tracks for open source AI. Transparency requirements that are difficult for closed models to satisfy may be straightforward for open models. The UK's emerging regulatory approach — whether through sector regulators, the AI Security Institute, or future legislation — should explicitly include mechanisms by which open source AI providers can satisfy regulatory obligations through the inherent transparency of their release, rather than through duplicative compliance documentation. For example: where a regulator requires evidence of model behaviour, evaluation results, or bias testing, an open source model that publishes its weights, evaluation benchmarks, and documentation on a public repository should be considered to have met a higher transparency standard than a proprietary model that provides the same information only to the regulator.

(c) Consult on incentive structures. DSIT should consult on whether additional regulatory incentives for openness — such as reduced reporting requirements, lighter-touch audit regimes, or explicit recognition in regulatory sandbox programmes — would accelerate adoption of open source approaches. This consultation should include the open source AI community, not only incumbent technology companies.

6. Reform Procurement to Recognise Open Source Providers

The most consequential policy lever available to the government is also the one it has been slowest to reform: procurement. The UK public sector spends billions annually on digital technology, and the structure of that spending — who can bid, what is valued, and how contracts are designed — determines which companies can participate in the market to sell technology to the government.

Procurement processes were designed for a world in which the government buys finished products from large vendors, and open source companies do not, by and large, sell finished products in the way those processes expect. Long and expansive Government procurement evaluations consider companies on turnover thresholds, insurance requirements, track records of contract delivery, and the ability to navigate compliance frameworks that assume a conventional enterprise sales operation. An open source company whose model has been downloaded millions of times, whose code underpins critical infrastructure across multiple governments, and whose community of contributors numbers in the thousands may nonetheless fail to meet the minimum requirements for a typical government contract.

First, framework agreements and approved supplier lists function as gatekeepers that disproportionately favour incumbents. Securing a place on G-Cloud or the Digital Outcomes and Specialists framework requires administrative capacity that bears little relationship to

technical capability. Large systems integrators maintain dedicated bid teams whose sole purpose is navigating these processes.

Second, procurement evaluation criteria consistently undervalue the distinctive advantages of open source. Open source solutions, which may have no licence fee at all, struggle to articulate their value within scoring frameworks designed around proprietary cost structures. The total cost of ownership, including long-term flexibility, avoidance of vendor lock-in, and the ability to modify and redeploy without renegotiating terms, is rarely captured in evaluation models. Government, therefore, routinely selects the option that appears cheaper on a spreadsheet, within that financial year, but proves more expensive in practice.

To change this will require targeted changes that lower barriers for open source providers without dismantling the accountability structures that procurement exists to enforce.

The Government should introduce an open source procurement pathway within existing frameworks, with evaluation criteria adapted to reflect the economics of open source. This means assessing the total cost of ownership over the contract lifecycle rather than upfront licence costs alone.

Specific measures:

(a) Create an "Open Source AI Approved Provider" register. Parallel to but separate from G-Cloud, CCS should maintain a register of open source AI providers that have met basic due diligence requirements (security, data handling, financial stability at an appropriate level) and can be engaged by departments without re-running full procurement processes for each engagement. This register should be designed to be easy for open source companies to join and should recognise non-traditional indicators of quality and reliability — such as community size, contribution history, and deployment track record

7. Invest in Open Source AI Developer Experience and Deployment

The most capable open source AI models in the world will create no value if organisations cannot deploy it. Deployment difficulty, not model quality, can potentially be the primary barrier to open source AI adoption. Raffi Krikorian, CTO of Mozilla, was unequivocal: "In 2026, it's a deployment problem more than anything else." He reports that the overwhelming majority of large enterprises that experiment with open source AI fail to sustain their use, not because the models underperform, but because the tooling, documentation, and operational infrastructure required to move from experimentation to production remain inadequate.

This is a market failure that public investment can address. Government should:

(a) Prioritise open source AI deployment tooling through Innovate UK AI stream. Innovate UK should be open to funding and supporting companies and projects building the infrastructure that makes open source AI easier to deploy — model serving, evaluation, monitoring, fine-tuning pipelines, and integration tooling. This is analogous to investing in roads rather than vehicles: the models exist, but the infrastructure to use them at scale does not.

(b) Establish a government open source AI deployment unit. A small team with practical experience deploying open source AI should be available to advise departments undertaking AI projects, this unit could operate as Subject Matter Specialists within i.AI. This unit would reduce the repeated reinvention of deployment approaches across government, build institutional knowledge, and create a feedback loop between government deployment experience and the open source tooling ecosystem.

Conclusion

Open source AI is not a new opportunity. It is an existing one that the government has systematically failed to recognise, not because the evidence is weak, but because the ways in which it has emerged, not necessarily according to the typical expectations of the existing digital economy.

Open source began as a countercultural movement in California, with programmers who believed that code should be free, that knowledge should be shared, and that collaboration could outcompete hoarding. Over five decades, those ideals became infrastructure. Linux powers the Internet, open source databases store the world's data, and now the most advanced AI models can be downloaded, inspected, and modified by anyone with the technical capability.

Yet the West has largely forgotten the strategic significance of what it created. While European governments treated open source as a technical footnote, China recognised it as an instrument of influence. While American technology giants built empires on open source foundations, they sought to capture the AI layer above. Europe finds itself squeezed between superpowers that have each understood what European policymakers have not.

This is a serious strategic error, and it is becoming a costly one. Beneath what looks like tinkerer culture, commercially viable companies are already emerging, and some are already supplying European governments with critical infrastructure. The Matrix protocol began as exactly the kind of project that a government official might dismiss: a foundation-governed, community-maintained experiment in decentralised encrypted communications. It is now used by the French and German governments for secure internal communications, offering a fully inspectable alternative to proprietary platforms at a moment when the reliability of transatlantic technology partnerships can no longer be assumed.

The good news is that the shift required is not primarily one of resources. It is one of perception. The government does not need to invent a new ecosystem. It needs to see the one that already exists, understand why it has been invisible, and design engagement mechanisms that work with the grain of how open source produces value rather than against it. The builders are already there. Many of them are already trying to work with the government. As Raffi Krikorian puts it, for countries that are not the United States or China, building on open source AI "might be the only viable path forward." The question is not whether open source AI matters. It is whether the government will recognise that the ecosystem producing it already exists, and act accordingly.

STARTUP
C*ALITION