AI & Society Institute

PSL ENS Dauphine Université Paris

Europe and Open Source AI

Beyond Licenses

Mitchell Baker, Co-founder of the Mozilla project · Expert Fellow at the AI and Society Institute

Mitchell Baker is the co-founder of the Mozilla project and led Mozilla – the organization behind the Firefox web browser – for more than 25 years, building it into one of the largest and most influential open source efforts in the world. Her career has centered on the questions this paper addresses: how open source projects are governed, funded and turned into products people actually use, and how a public-benefit mission responds to market forces. Mitchell is a member of the American Academy of Arts and Sciences, Advisory Boards of the Oxford Internet Institute, the Center for Machine Intelligence at the University of Sheffield and the G1 Silicon Valley. She is a board member of OpenMRS, Inc., which develops open source medical records systems for use in resource-constrained environments.

The **AI and Society Institute** fosters the responsible development and use of AI by studying its interactions with society. It produces analyses grounded in scientific methods to inform policy in France and Europe. It is led by **Hugo Mercier** (Scientific Director) and **Constance de Leusse** (Executive Director).

Published in July 2026.

Table of Contents

Key Messages

1. Europe is at a strategic turning point for open source AI. The European Technological Sovereignty Package and its Open Source Strategy create a rare opportunity to shape a sovereign AI future. The critical question is no longer whether Europe should support open source AI, but rather how it builds the structures and capabilities to succeed. Outcomes will depend as much on implementation as on the commitments themselves.

2. Open source offers a multi-layer stack, not just a license. Creating a new degree of technical sovereignty depends on engaging multiple layers: license terms, the technology, the team, collaborative development and governance, product strategy, and organizational design.

3. The definition of "open source AI" is unsettled, and Europe should build anyway. The European Union (EU) can act today: adopt the Open Source Initiative's Open Source AI Definition as a funding floor, reward deeper openness through the Linux Foundation's Model Openness Framework and the DPGA's AI Standard, and earn definitional influence by actively building and deploying.

4. Adopting open source models rather than building from scratch alleviates funding and capacity problems. But access alone does not deliver the knowledge, the ability to steer or replace dependencies, trust in the model, or capabilities that come from building and participating. If adopted, open source models without a collaborative development process should entail an exit plan.

5. Sustainable open source AI also depends on capabilities, people, and governance. European investments should include the expertise, leadership, and institutional capacity needed to build, govern, and maintain open source AI over time. Further, publicly funded assets need safeguards to remain public-benefit despite any acquisition or restructuring.

6. Interoperability retains its importance. The EU should support its teams to identify key interoperability points as AI technology develops and participate in related standards work.

7. Research and technology development must be paired with a focus on products people use. Success metrics should include voluntary adoption and satisfaction metrics.

8. The idea that "Secret is safer" is a recurring problem. Closed systems hide mistakes, frustrate regulation, and make security a matter of faith in a brand rather than something society can inspect. The EU needs the institutional reflex to recognize and interrogate this argument rather than retreat into secrecy.

Introduction: Europe and Open Source AI

In early June 2026, the European Commission announced the European Technological Sovereignty Package – an ambitious set of policies and legislative proposals to boost European competitiveness and strategic autonomy in the AI sector. A central piece of this package is the Open Source Strategy, which promotes the use of open source technologies by public administrations and open source components in critical infrastructure, and seeks to encourage a pan-European ecosystem around open source.

AI is emerging as a general-purpose technology likely to shape many aspects of life. If these capabilities become concentrated within a small number of actors and jurisdictions, Europe risks becoming dependent on external infrastructures that influence not only markets but also the conditions under which citizens access knowledge, services, and opportunities. This is not hypothetical. In June 2026, within days of the Open Source Strategy's announcement, the US government suspended access to Anthropic's most capable models on national-security grounds; OpenAI was persuaded to withhold its latest release, and Reuters reported that China may consider limits of its own. The controls were lifted within weeks, but the point stands: access to frontier capabilities can be granted and withdrawn by decisions made outside Europe. Open source AI should therefore be understood as a strategic priority for Europe's ability to shape its own future.

European institutions bring to AI something valuable and hard-won: a deep understanding of how technology lands in society (including its effects on citizens, on markets, and on rights) and the will to act on that understanding. What remains scarce, in Europe as everywhere, is the complementary experience of building and leading large scale open-source projects. This includes activities such as assembling the team, building a community and ecosystem, governing its conflicts, and carrying the work from promising technology to a product that millions of people choose to use. I write from inside that experience. Many of its lessons are encouraging, as reflected in the Open Source Strategy itself. Some are discouraging. But nearly all reflect processes and approaches that are subtly or significantly different from traditional development processes. The full power of open source is best unleashed when action is cultivated, resourced and protected, rather than simply mandated. This requires institutional processes focused on "what conditions do we create" as well as "what do we require." The success of the Open Source Strategy depends on how well Europe integrates its institutional strengths with the experience of those who have built and led such projects. This should be done at the working level, as part of how its projects are designed and run.

The goal of this paper is to help the Open Source Strategy succeed in producing projects that endure, products that citizens choose, and an ecosystem that strengthens Europe's hand. Thirty years of open-source practice says that this kind of success is decided at multiple layers: the license is the entry point, and above it sits the technology, the team, the collaborative process, the product, and the organization. To be successful it is critical that Europe implement open source as an approach that addresses each of these layers. This paper takes the layers in turn, drawing on projects that got them right and some that did not, and asks at each one: what would it take for Europe's projects to succeed here?

Multiple Layers of Open Source

An asset is "open source" or not based on its license terms. An open source asset can have a wide variety of impacts: public benefit or private enrichment; sustainable or a "flash in the pan;" research achievements or products that people use, to name a few. An open source asset's impact depends on other layers of activity. There is no canonical artifact that sets out these layers of activity; the framework below is distilled from three decades of practice.

The first layer is the license, which determines whether or not something is actually "open source." The definition of "open source AI" is under discussion, so the License Terms section below recommends a way to move forward despite the ambiguity. A second layer is the technology itself – is there a clear enough vision of what a project aims to build and the success criteria? A third layer is the team, which includes the people who build the project, the reasons they choose to participate, and what they need to give their best. A fourth layer is the collaborative process by which that technology is developed – is it a process in which multiple stakeholders in the EU share a vision, a community sense of ownership and a commitment to the whole? A fifth layer is the finished product – can the research and development advancements be turned into products that are used regularly and productively? A sixth layer is the organization that holds responsibility for the technology and/or products.

Two things about frontier AI are unlike the open source software experience this paper draws on: the capital required to build the largest models, and the structural challenges of the underlying hardware, chips, and compute.[1] How frontier-scale training is financed in a public-benefit model is a genuinely open question, and this paper does not resolve it. The focus of this paper is different: whatever methods and amounts of funding are allocated, the likelihood of bringing the open source AI projects to successful impact depends on the variables discussed here. This paper will limit itself to the objectives for open source software and AI.

Layer 1: License Terms

At the licensing layer, the term "open source" does not apply exactly to significant aspects of AI. It applies to relevant software. The architecture of AI requires other elements, in particular training data and model weights. As a result, there is significant discussion underway about what the technical definition of "open source AI" should be. The lack of a consensus is largely about how the training data is treated, which can be seen in the response to a definition of open source AI proposed by the Open Source Initiative, which stewards open source licenses for software. One view holds that the training data must all be released for "open source" to have real meaning. Another view holds that this is impractical – due to copyright issues, privacy issues, data protection issues, etc. – so that very few if any models would meet this bar today. Another approach is to describe different "gradations" of open. The Linux Foundation has produced a Model Openness Framework using this approach.

[1] These involve additional complex policy considerations, being addressed through the AI Continent Action Plan, the proposed Chips Act 2.0, and the Cloud and AI Development Act.

The Digital Public Good Alliance ("DPGA") has created an AI standard to determine what is required for an AI system to be eligible for status as a digital public good. This standard requires all data used in a system to be available, and also includes "responsible AI considerations." These deep divisions are further complicated by the need for high performance computing facilities which sets up another barrier to the access, reproducibility and use rights that form the heart of open source software.

We should expect the exact definition to be in flux for some time. It will settle when enough AI practitioners and policy makers have enough real-world experience building artifacts to coalesce on a definition that is powerful enough to matter and practical enough to support an effective development process. A definition cannot be agreed upon as a research or policy process alone. It must include the real-world experience gained from building and deploying technical solutions. With open source software, these real-world experiences have led to multiple licenses that are quite different in practice and lived experience, while still meeting the required freedoms of use to be approved by the relevant organization (the Open Source Initiative). This may well be the case in the AI space as well.

As a result, the EU initiatives to build an open source AI future will need to begin without definitional certainty. This creates challenges, such as what the EU should mandate in its funding requirements. It also creates opportunities: leadership earned through participation is a key tenet in many open source projects. EU projects that are actively building and deploying AI can participate in the definitional process, broadening the EU's voice.

As the definition of Open Source AI is being determined, there are steps the EU can take today to apply structure and consistency to the selection, funding and access to compute of projects based on their openness, without waiting for the canonical definition of "Open Source AI" to be settled. The Open Source AI Definition provided by the Open Source Initiative is a widely understood document. It is controversial because many believe it does not go far enough. This definition can be the floor – open source AI projects seeking public funding must meet this definition at a minimum. Because this definition is already widely understood, adopting it as a baseline lets the EU bring structure and consistency to funding decisions today and ensures public money flows only to projects that clear a recognized bar for openness.

In addition, the Linux Foundation's Model Openness Framework and the DPGA's AI standard provide consistent mechanisms for evaluating how far above the bare minimum of "open sourceness" a project goes. Public funding and access to compute can increase for projects in higher tiers of the Model Openness Framework and for meeting the DPGA's AI standard. This rewards deeper openness across data, code, and documentation, rather than just the minimum. It also gives projects a consistent, transparent incentive to move up the tiers.

Finally, task an appropriate entity (perhaps the AI Office and/or the EC Open Source Program Office network) with guidance ensuring that systems claiming the AI Act's open source exemptions[2] meet published criteria in order to avoid "openwashing" concerns related to the open source exemptions.

Layer 2: Technology

The technology layer holds a myriad of specific decisions: which models to build and which to adopt, which components beyond the model matter most, how to weigh quick wins against long, expensive bets; where public funding and public infrastructure are most important, among other aspects. This section applies an open source lens to these topics; it does not seek to provide the strategy and roadmap document.

Model-Specific or Broader Scope

A key decision is whether the focus will be almost entirely on models, or a broader range of AI capabilities and products. And if models: frontier models only, or also smaller, domain-specific, and exploratory models?

To the extent that the focus is developing frontier models that require large amounts of capital and hardware, certain forms of open source participation will be limited. That said, there are a host of decisions to be made in this process, and the open source practices described in this document can serve as an organizing structure for member-state participants.

To the extent the scope is broader than frontier models, the opportunities for open source participation expand. The Llama example described below demonstrates how open models can unleash new forms of participation and development.

Build or Adopt

Since the "DeepSeek moment" in early 2025 China has become a leading actor in the global AI arena through its open models. As of 14 July 2026 the six most heavily used models on OpenRouter, a major routing platform for AI usage, were all open models from Chinese firms. Data from Hugging Face show that in the past year Chinese open-weight models rose to a plurality of 41 percent of model downloads on the platform, surpassing the U.S. in monthly downloads and overall downloads.

In mid-July 2026 China's President Xi Jinping laid out China's goal to build an international AI ecosystem with particular relevance to the Global South, including a call to encourage open source. At the World AI Conference in Shanghai, China also launched the World Artificial Intelligence Cooperation Organisation (WAICO), with

[2] Specifically, Articles 2(12) and 53(2) of the EU AI Act. Providers of general-purpose AI models who release each such model under "a free and open-source licence that allows for the access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available" are not required maintain technical documentation nor share certain information on capabilities and limitations of these models with downstream users.

twenty-nine founding member states, most from the Global South, a headquarters in Shanghai, and programs to build AI capacity in developing countries.

Another key technology consideration is whether, and how, to adopt open source models available from Chinese labs. For this discussion let's assume the Chinese models have, and continue to have, terms that meet an open source definition. An important next evaluation will be the degree to which these model providers engage in collaborative development and/or have an organizational structure that provides long-term reliability. If not, then using these models will not deliver a number of things that are generally viewed as important in open source evaluation, even with technical fixes such as self-hosting these models on European infrastructure:

- knowledge of the decisions involved in creating the model, including training data and the nuances of pre- and post-training;
- the ability to steer, maintain, or replace one's dependencies;
- the knowledge gained through participation;
- provenance, auditability, and reproducibility; insight into the project's decision-making processes.

In an ideal world, adopting a model would deliver both: the capability of the model itself, and the additional capabilities listed above immediately. Given that's not currently possible, the importance of these gaps can be evaluated based on scope, which is a strategy question the EU strategy will need to address. One possible scope is an alternative for government services themselves, so that EU member-states aren't relying on providers outside their control for the services they deliver. Some of these services, national security among them, might well require understanding every aspect of the model. At the other end of the spectrum is a scope covering all the actions citizens might take, where different tradeoffs might make sense.

Analysis of adopting open source models which lack collaborative development should be paired with an exit

plan. In a collaborative development environment the ability to see and participate in the development gives one a credible capacity to sustain the software if one utilizes the right to fork. In cases where an open source model lacks an open, collaborative development process then the right to fork remains but the experience base to use it will be limited.

Decisions about if and when to adopt Chinese models can also be evaluated on the degree to which Europe embraces international interest in joining with it. China is emerging as a strong proponent for open source AI. This can be both a challenge and an opportunity for Europe. Countries uncomfortable with China's plans may seek an alternative. Countries working with China may also prefer a second option. Europe can offer jurisdictional trust, GDPR-compatible provenance, auditability, public-sector deployability, co-development, and Horizon Europe association. More than twenty non-EU countries already participate in Horizon Europe, and the EU can choose to extend the same welcome to open source AI projects.

Unexpected Opportunities

Ground-breaking ideas arrive from unexpected places, and development processes need room for them. When Meta's Llama model became widely available in 2023 there was a frenzied period as individuals made unexpected progress in making the models work on smaller and smaller devices. The lack of resources available to individuals was a motivator to pursue a development direction that wasn't centrally funded by the big model provider. We experienced this phenomenon in creating Firefox as well. Early versions of Firefox offered features that were new to consumers and dramatically improved the experience. One of them was tabbed browsing. This feature was created by a volunteer. There was no central planning process, no one at the core Mozilla organization to say yes or no. Mozilla's open development system let people try ideas and see whether others liked them; people liked tabs so much they became standard in Firefox and, from there, in browsers generally. Both examples demonstrate that carefully planned work and unplanned exploration can proceed at the same time, and projects that allow both tend to capture ideas the plan would never have produced. To help enable this the EU should reserve a fixed share of public compute capacity for projects that meet the openness tiers in Recommendation 1 (see recommendations section below).

Interoperability

One class of technology decision deserves particular attention, because it shapes market structure. This is the question of where interoperability should be required. A good current example is the agentic harness that wraps around a core model. The harness supplies memory between interactions, tools for taking action beyond text generation, and an execution environment that orchestrates the work. Harnesses are available today from independent organizations and from the model providers themselves. There is concern that model providers will vertically integrate: folding the harness into the core offering and making independent alternatives irrelevant, much like Microsoft's practices of absorbing categories of products into Windows. Anthropic's Model Context Protocol ("MCP"), released in 2024, created a standard interface for how models connect to external resources, providing a level of interoperability. Today the harness integration concern involves tight coupling between the model and the harness, and a new set of interoperability questions has arisen here, despite the MCP layer. The EU's open source projects should be actively involved in identifying the interoperability points that matter and actively involved in standards development for interoperability. A proactive approach will help avoid lock-in and reduce the need to rely on post-facto regulation.

Layer 3: Team

Implementation of the Open Source Strategy should put a heavy focus on attracting, retaining and supporting the teams of people who create the technology and products. Motivating teams is a distinctly different mindset than managing the outcome and should be intentionally built into the operating plans of open source AI projects.

This is true in many settings, and especially important here. It will be a challenge to build an alternative open source AI ecosystem. Many of the possible participants will have other enticing options. The lure of market compensation cannot be eliminated. Some people today will be offered packages that are impossible to match and perhaps near-impossible to turn down. Open source does not fix this. However, I was told many times that open source software could not succeed because the tech jobs of the day were too good. That turned out not to be the case.

Fortunately, building open source public benefit technology is a powerful pull for talented and dedicated people. We experienced this at Mozilla. People sought us out, they worked for less compensation than elsewhere, they volunteered at a massive scale because Mozilla stood for something they wanted to exist in the world. At Mozilla we thought of this as our "community"– people bound by the shared commitment to open source public benefit technology, as expressed in the Mozilla Manifesto. The clearly stated mission and the fiercely committed community that developed around it is one of the keys that has allowed Mozilla to remain relevant and compete against vastly wealthier organizations for so long.

I experience a similar interest among a subset of AI practitioners. The pull of the mission can be deliberately strengthened. The starting points include: Articulate the public mission of European open source AI in terms people can sign up for. Seek out the people for whom mission outranks personal compensation, and people at the right phase of their careers, often before or after private-sector chapters (some come intending to leave but never do). Convene the practitioners, including researchers, product people, engineers in for-profit jobs who will contribute alongside their day jobs. Give them places to meet, prototype, and find each other. Create fellowships that let strong people work on public AI infrastructure full-time, and on-ramps that harness the time and attention part-time contributors do have. Support talented people from unconventional places. Remain stable when tech slumps come. Make participation visible and prestigious. Intentionally build teams with global scale.

It was enough to build the open source software movement. The open source AI world may well need more, which is one more reason to gather the people who want the Open Source Strategy to succeed and create the space for further solutions to emerge.

Layer 4: Collaborative Development and Governance

This layer focuses on how the technology is built. Is it a truly collaborative development process? The more this is true, the greater the chance of European sovereign AI that also serves the public interest. This also aligns with the growing focus on digital commons as a backbone of European tech sovereignty, evident in the newly constituted European Digital Infrastructure Consortium for Digital Commons.

We are in the early days; the underlying technology will be developing for quite some time. This means that the EU and its member states will want a process that is sustainable. Designing effective participation models,

transparent decision-making, and overall project governance will be key. A few key elements where I have seen projects and contributors struggle to understand and implement are described below.

Decision–making for the benefit of the project as a whole.

In collaborative development projects contributors may be employed by many different organizations. This leads to the question: to what degree do technical and product leadership make decisions based on what they believe is best for the project as a whole and all its participants? And to what degree do they make decisions that benefit the entity that pays them or sponsors their participation? Open source projects such as Apache, Mozilla and the Eclipse Foundation have answered clearly: decision-making should be for the benefit of the project, even if the decision is not what the decision-maker's employer believes is uniquely best for it. At Mozilla we described this as "wearing two hats." It was perfectly acceptable to say something like: "Wearing my hat as an authorized decision-maker for the project, I'm deciding we will do X. Wearing my employee hat, I'm going to spend a significant amount of my time implementing what my employer needs for X to be valuable to it." Apache takes a similar approach, requiring that their Project Management Committees act in the best interest of the project as a whole. The Eclipse Foundation, focused on open source collaboration in industry from its headquarters in Belgium, also requires vendor neutrality. The EU will want formalized mechanisms for managing these multiple roles. EU funding programs such as Horizon Europe and Digital Europe Program can require analogous mechanisms as part of their processes.

A robust collaborative process will also include guidelines for organizations addressing the interaction of these open source processes with their internal processes. This includes topics such as the treatment of company material, which information needed to be public for the project, interaction with company product planning and ROI planning, and practical operational topics such as how company contributions would be evaluated.

The remit of open source project leadership

One foundational question is the degree to which project leadership makes the major decisions, and the degree to which the funding organization makes (or designates another political or regulatory entity to make) technical and product decisions. In my experience, it's nearly impossible for people not responsible for building the technology and product to make good decisions in these areas. Project leadership should be chartered to make independent decisions over technical, architectural and product matters – they will need room to explore, prototype, adjust and iterate.

Forking and the cost of fragmentation

One topic that could come up is the temptation for EU member states involved in open source projects to pull the project into a private development environment and continue developing it differently from the original. In open source software this is known as "forking" – like taking a different fork in the road. However, there can be a severe forking "tax" in the form of long-term incompatibility and fragmented communities. In general, forking is best understood as (1) a credible exit plan whose existence disciplines project governance and (2) a strategy for cases of ongoing, unresolvable conflict over the direction of a project. It is a guarantee of sovereignty and control

over the technology. However, EU member states should default to contributing their improvements to the collaborative development process, utilizing the right to fork only when absolutely necessary.

Layer 5: Products

"Open source" often leads people to focus exclusively on raw technology, such as frontier models, agentic frameworks, or base code. This is necessary but insufficient. Success will rest in large part on the ultimate products that citizens and organizations use.

Consider historical precedents: the Android operating system used Linux as a starting point, and early Amazon and Meta server capabilities relied heavily on Linux. Parts of that infrastructure may still be technically open source software, yet the layers of commercial infrastructure and private commercial products obscure this. Similarly, an open source frontier model wrapped in as many layers of privatization, commercial agreement, and vendor control as today's Linux-based systems would be only a partial success for the Open Source Strategy. Broader success is a robust, open alternative AI future for citizens, with real choices as opposed to lock-ins. Turning great technology into an intuitive, problem-solving product is difficult, and it requires a dedicated organizational focus. This is a focus that EU projects will be called on to develop.

Mozilla started out as a technology organization and evolved into a product organization as well. The difference is huge. It was a fraught transition – the technology-focused and product-focused groups wanted different things, and resolving that tension took sustained attention from my co-founder and me. Technology organizations and product organizations may be close siblings, but they are not the same. In a technology organization, the core focus is on the technology: what can it do, and how far can it evolve? In a product organization, the core focus is on the customer: what kind of experience will they have, what problems are being solved for them, what new capability will they enjoy? Technology organizations are renowned for producing products with advanced technology that don't quite fit what consumers want. Mozilla's first product (the "Mozilla Application Suite") is an example – well regarded by engineers but not consumer-friendly. Mozilla Firefox was built on the same technology with the opposite priority, and was widely credited with changing the internet landscape following its release in 2004. This shift requires as much focus on the intended end user as on the underlying technology.

The Firefox story must include a discussion of the Chrome era. Google released Chrome in 2008 and put its distribution power behind it. This included activities such as advertising on Google's highly visited webpages, tuning their other products for Chrome, and setting the default browser on Android phones. Firefox's share of the market declined over the years, and it remains small today. Product against product, Firefox can be competitive. But competing against the entirety of Google and its resources is a different matter, and one that few attempt. Both Microsoft and Apple have decided to rely on Google's might instead. Microsoft decided to use Google's open source Chromium offering as the basis of Microsoft's browser, and Apple uses Google's search feature in Safari. (Mozilla does as well, though not on an exclusive basis.)

The Firefox story is two things at once: proof that an open source alternative can be built and make a difference, and a cautionary tale about the power of large integrated technology companies.

The EU has addressed similar concerns with the Digital Markets Act ("DMA"). The DMA and competition enforcement constrain the distribution and default-setting conduct that independent producers struggle to overcome on their own. Europe's open source AI products will at least enter a different environment than Mozilla did. How much this new environment changes outcomes will only be learned by building. Applying the levers to the AI world is a necessary complement to the open source product development process.

There are additional ways the EU can support product development. One is to fund productization and user experience goals as well as research and infrastructure, learning from the work of the Next Generation Internet programs. Another is to require explicit openness scoring in procurement tenders and through the Open Source Strategy's public-administration objective. A third is to focus success metrics on voluntary adoption – the number of people and organizations that find products useful. Avoid success metrics focused on shipping because they don't measure the impact generated by the artifact. If people are required to use a product (for example, to access government services), include satisfaction criteria in the metrics. This is to avoid settings where a project appears successful even though it is forcing people to use products that perform poorly.

Layer 6: Organizations

Open source projects can exist in a wide variety of organizations, from venture-backed companies, to government funded creation of public goods, to a myriad of hybrid models. Different models produce different results. A robust open source AI ecosystem might include all of them. My experience suggests core lessons.

First, when products are successful at scale, there is usually a formal organization that takes responsibility for them. I experienced this directly at Mozilla and at OpenMRS, an open source project for electronic medical record systems. Mozilla has existed since 1998 and was a successful open source project for years before the Mozilla Foundation was formed in 2003. Despite this, Mozilla's work and our products were seen in a very different light after the Mozilla Foundation was established. Organizations were much more comfortable engaging with me as the President of Mozilla than they had been engaging with me as the acknowledged leader of Mozilla as an open source project. Within weeks of the Mozilla Foundation's launch we received expressions of institutional interest that had not come our way before. OpenMRS's legal entity has similarly made it easier for Ministries of Health to engage with the project.

Second, public infrastructure requires organizational forms that are focused on protecting the public assets, especially if the project is also trying to generate revenue to sustain itself. Publicly funded assets should be subject to mechanisms that ensure they remain public-benefit despite any acquisition or restructuring.

Third, public benefit projects that want to compete with proprietary technology need tools in addition to the open source layers this paper describes. These needs include distribution at scale, a way to challenge the vertical

integration and market concentration of the incumbents, and an alternative to the financial resources the public markets bring. Open source can do a shockingly good job at building alternatives, but building alone doesn't solve these problems. Firefox managed distribution and a genuine market challenge for some years, but not indefinitely. Projects now operate alongside levers the EU already applies to technology it deems strategic, such as the DMA as a constraint on incumbents, public procurement, and government funding.

The second lesson is playing out in real time. The company "OpenAI" has undergone organizational change, moving its operating entity from a non-profit to a business ("OpenAI PBC"). Both OpenAI and Mozilla have operating companies as subsidiaries. In OpenAI's case the subsidiary company has investors, and a mechanism for private financial return. Mozilla's subsidiaries were set up as wholly-owned subsidiaries and have historically not accepted investment to maintain focus on the public mission. The two structures answer the same question differently: whose interests the organization must weigh when mission and financial return diverge.

MySQL / MariaDB is another example that demonstrates the importance of an organization committed to the public benefit nature of the asset. MySQL was an open source database project founded in Scandinavia that became widely used globally. It made its product available as open source under a restrictive open source license (technically known as "a copyleft under dual licensing") for embedding in open source applications, and also charged customers who wanted to use MySQL in proprietary applications. MySQL operated as a traditional profit-seeking company while also satisfying its open source user base for many years. Then it sold itself to a larger company which in turn was bought by Oracle. At this point the open source founders became concerned about the direction of the project under Oracle's leadership. They exercised the right to fork and formed the MariaDB Foundation which successfully builds an open source version under the name of MariaDB. A second commercial entity based on MariaDB code was then formed. It struggled and was eventually acquired by a private equity firm. Because the MariaDB Foundation is separate and stewards the open source code, the open source project continues.

The Eclipse Foundation has also changed its organizational structure. It was originally chartered as a US 501(c)(6) non-profit organization. In 2021 it reorganized as the Eclipse Foundation AISBL, a Belgian international non-profit, membership-based association. It currently hosts over 400 projects with nearly 100 employees. The organizational form changed; the stewardship of the projects continued.

None of the organizational forms described here provides, on its own, distribution at scale, market power, or an alternative to the financial resources the public markets bring. What an organization secures is durability and commitment to the public mission through changes no one can predict, forming the base the other levers work from.

Context: The "Secret is Safer" Reprise

The argument that "secret is safer" surfaces repeatedly as a reason to restrict open source systems and components. It cuts across the layers of open source – it has appeared around new technical capabilities, around

training data, and around product assets that institutions have opened. We should expect it to resurface so I address it here.

We can stipulate that secret technology is certainly safer for the product company. Secrecy lets an organization act in ways people don't know or expect, hide its mistakes, and tune its technology for its own benefit without others understanding what that means. It makes good regulation harder and noncompliance easier (good-faith mistakes and intentional avoidance alike). And it leaves society unable to see the risks at all: the dangers of closed technology are invisible from outside, so security becomes a matter of faith in the brand. As a long-term plan it is neither robust nor resilient.

In Mozilla's early days, people routinely questioned us about whether our open process made our products less secure. We answered by shipping Firefox, which set and then raised the security bar for all browsers. We did this with an open bug process, an open development process, the participation of volunteers around the globe, and a Bug Bounty program that rewarded people for finding problems and working with our developers to fix and deploy quickly. Because browsers then touched everything on the internet, I also saw the internal security practices of some of the era's biggest software companies. Those practices put convenience and profitability well before the safety of customers, at brands that remained trusted precisely because no one outside could see how they worked.

There are times when a new question arises and it's important to rethink whether "secret is safer" might be true in that setting. When the question of whether open weights were dangerous and should be restricted was first raised there was a period of intense evaluation. Ultimately the US National Telecommunications and Information Administration (NTIA) concluded that the evidence of incremental increase in risk did not warrant restricting open weight models.

The secrecy discussion arises in connection with training data. Making training data open makes copyright and data-protection issues visible, and visibility has brought complaints and litigation. One available response is to keep the data secret and avoid the lawsuits. But secrecy does not stop the harm; it only rewards those who keep the harm hidden. The policy response is not easy, but rewarding secret harm cannot be a good answer. One direction worth exploring is training-data documentation requirements with mechanisms that favor documented, remediable open datasets over undisclosed ones.

We are also seeing the "secret is safer" discussion play out as a reaction to the enhanced ability to find software vulnerabilities demonstrated by Anthropic's Mythos model. The British National Health Service made its code repositories private by default while it adjusted to the new reality. Days later, the UK Government Digital Service published guidance on AI, open code, and vulnerability risk in the public sector. This guidance is explicit that open should remain the default, and includes the analysis underlying that recommendation. The sequence demonstrates two things at once: how quickly "secret is safer" can seem the obvious solution, and that institutions willing to stop and interrogate it can reach the opposite conclusion.

Looking forward, the same capabilities that find vulnerabilities can be used to test and harden software before it is deployed. Whether that happens depends on who controls the capability, who is granted access, and at what cost. These are market structure and access questions rather than questions of open versus closed.

The EU will benefit enormously from an institutional reflex with three parts: recognize the "secret is safer" argument even when it arrives disguised as a practical reaction to the crisis of the moment; interrogate its assumptions precisely; and craft responses that weigh the harms of secret infrastructure alongside the perceived risks of openness. In addition, the EU will benefit from a funded program that hardens open source projects before release, coordinates vulnerability-disclosure safe harbors under applicable regulation, and supports documented, remediable open datasets. Put simply, the EU should treat openness as a security method rather than a security risk.

Recommendations

This analysis points toward a set of practical priorities for European leaders and policymakers. If Europe wants open source AI to strengthen technological sovereignty and deliver public benefit, its interventions must engage the full stack of open source rather than licensing or technology alone.

License terms

1. ***Set the Open Source AI Definition as the funding floor.*** The EU should require every AI model project seeking public funding to meet, at a minimum, the Open Source Initiative's Open Source AI Definition. Because this definition is already widely understood, adopting it as a baseline brings structure and consistency to funding decisions today, without waiting for the definition of "open source AI" to be settled. On training data, the definition requires detailed disclosure of what the data is and how it was prepared; it does not require release of the data itself.

2. ***Scale funding and compute to tiers of the Model Openness Framework. Above that floor, use the Linux Foundation's Model Openness Framework and the DPGA's AI Standard to measure how*** far a project's openness extends beyond the minimum, and increase support accordingly. Projects reaching higher tiers should receive greater public funding and expanded access to compute. This rewards deeper openness across data, code, and documentation, and gives projects a consistent, transparent incentive to move up the tiers.

3. ***Issue guidance on the AI Act's open source exemptions.*** Task an appropriate entity (for example the AI Office and/or the EC Open Source Program Office network) with issuing guidance so that systems claiming the AI Act's open source exemptions meet published criteria.

Technology

4. *Reserve compute for openness.* Convert the precedent of reserving capacity for open projects into standing policy: a fixed share of public compute capacity (for example, EuroHPC and the planned AI Gigafactories) reserved for projects that meet the openness tiers of Recommendation 2.

5. *Evaluate open source models for possible adoption on criteria beyond the license terms.* Require that models meet an open source definition, then assess the degree of collaborative development or other mechanisms for building capacity. Repeat this evaluation as the model capabilities grow. Pair adopted models with an exit plan.

6. *Interoperability continues to matter.* Charter and fund EU-funded projects to actively identify the interoperability points that matter most. Back their participation in standards work that results.

Team

7. *Invest in participation.* Fund training programs, fellowships, and on-ramps that let others contribute alongside their day jobs. Convene the practitioners: support recurring gatherings that connect researchers, product people, and contributors from for-profit settings, whether or not they work on EU-funded projects. Build this community on a global basis, welcoming participants regardless of geography. Horizon Europe already welcomes non-EU countries, and the same welcome can extend to open source AI.

Collaborative development and governance

8. *Institutionalize decision making for the benefit of the entire project*. Require EU-funded programs to build in formal mechanisms for decision-making "for the benefit of the project as a whole," including clear guidance for managing the dual roles contributors hold.

9. *Make governance quality a condition of public funding*. Projects seeking public money should demonstrate sound open governance (documented decision-making, open contribution and maintainer structures, and safeguards against capture by any single vendor or interest) and governance quality should be revisited over the life of the funding. This review should stay on governance: it assesses these structures, not the technical and product decisions that project leadership is chartered to make.

Products

10. *Fund the full product lifecycle.* Public funding should support the entire lifecycle of open source AI. This includes development, usability, productization, deployment, product innovation, long-term maintenance, and open source stewardship, so that European investment yields products people choose rather than isolated technical artifacts. Require explicit openness scoring in procurement, through the Interoperable

Europe Act and the forthcoming Cloud and AI Development Act[3] and the Open Source Strategy's public-administration objective.

11. ***Anchor success metrics in product adoption rather than shipping***. Measure voluntary adoption and required adoption (e.g. ‚required for use of government services) separately. Where use is required, measure satisfaction criteria.

Organizations

12. ***Build durable, independent stewards, and lock in public benefit.*** Implement the Cyber Resilience Act's "open source software steward" category with real resourcing. Designate a body that defines the conditions attached to public funding, supplies tested governance templates for funded projects, and insulates their governance from vendor and political capture. Ensure that publicly funded assets carry mechanisms to remain public-benefit through any later acquisition, restructuring, or change of corporate control.

Security

13. ***Treat openness as a security method.*** Build the institutional reflex to recognize and interrogate the "secret is safer" argument wherever it surfaces. Fund a pre-deployment hardening program that lets open source projects use AI vulnerability-finding capabilities before release. Coordinate vulnerability-disclosure safe harbors under applicable regulation. Favor openness on training data by rewarding documented, remediable open datasets over undisclosed ones.

[3] European Commission, Proposal for the Cloud and AI Development Act (CADA), Shaping Europe's Digital Future, https://digital-strategy.ec.europa.eu/en/library/proposal-cloud-and-ai-development-act-cada.