Abilian
ECON LAB

# The EU's AI money is aimed one layer too low

Reading Mozilla's State of Open Source AI (July 2026) against the Commission's AI technology area report (June 2026)

**Stefane Fermigier, Abilian Econ Lab · 18 July 2026 · DRAFT WORKING PAPER**

## Executive summary

EU investment in AI for the 2028-2034 funding period is being planned against a diagnosis of scarce compute, scarce capital and scarce skills. Two assessments published this summer, one of them the Commission's own, locate Europe's failure elsewhere: at deployment, where pilots fail to become production systems. Annex 4 on Artificial Intelligence, prepared by Technopolis Group for DG CNECT and published on 23 June 2026 as part of the study on EU critical digital capacities beyond 2027, is explicitly an input to the next Multiannual Financial Framework. Mozilla's *State of Open Source AI* reaches the same reading of the constraint from market data and places it in a specific layer, the software that sits between a trained model and a working application, which is also where it locates the commercial value; the funding conclusion drawn from the two documents is this note's own.

**Teams that start building with AI and never get past the pilot stage are not held back by the quality of the models available to them.** In Mozilla's global developer survey, 63% of teams using closed models got something into live production and 51% of teams using open models did: between a third and a half of teams stall short of production whichever way they buy, including those buying the most capable models with the full vendor bundle attached. Annex 4 reports the same failure for Europe in qualitative terms, citing absent return on investment, failed technology transfer, and the difficulty even well-funded teams have turning experiments into reliable production systems. Capable models are abundant and cheap. At the model layer Europe lacks European models, a separate problem with a separate remedy, and nothing here should be read as denying it. Annex 4 nonetheless directs its investment needs to compute, chips, data centres, energy, connectivity and skills.

**The layer where value is moving has no cell in the Commission's value chain.** Annex 4 reproduces a McKinsey eight-segment value chain that ends at applications and services, with no segment for orchestration, memory, tool access or agent permissions. The text names the gap twice, once as missing out on commercialising agentic AI and once as a need to invest in AI middleware, and attaches no figure to either.

**The layer moves on market time and the money moves on framework time.** Annex 4, as an input to the next framework must, works in actions by 2028 and actions for 2036-2040. Mozilla dates vertical integration of model and orchestration layer to the current year, and forecasts the expiry of introductory API pricing around 2027-2028. The mismatch is between the instrument and the

object, with the two documents in agreement on the substance, and it argues for procurement and standardisation ahead of the framework.

**Europe's public money backs open supply while its developers buy closed.** Mozilla's regional survey finds Western Europe to be one of only two regions where closed-model adoption exceeds open.[1] Annex 4 notes in passing that public procurement for AI is under-utilised and hard for suppliers to navigate; no recommendation in the document develops it as a demand instrument.

**Deployment autonomy bought with foreign open weights is a stock, and the capability that replenishes it is a flow Europe does not control.** Weights already held cannot be revoked and do not phone home. The pipeline producing the next generation is foreign and discretionary, and a stock of weights loses value at the rate the frontier moves. Mozilla's own figures make that flow look fragile: on OpenRouter over May to September 2025, open models carried a fifth of token volume and captured about 4% of revenue; the token share has since passed one half, and two of the leading open-weight labs are now listed companies answerable to public markets.

None of this settles whether Europe should hold sovereign training capacity, and the supply argument above strengthens that case. The claim here is about which investment to make first under uncertainty. **The deployment layer pays off whichever models Europe ends up using: it serves foreign open weights today and would serve a European frontier model tomorrow.** Training capacity pays off on the branch where Europe produces competitive models, and even that branch needs the deployment layer to use them. That is the argument for the marginal euro reaching the deployment layer first. The comparison is between a payoff on every branch and a larger payoff on one branch, so a reader who rates Europe's chances at the frontier higher than this note does should weigh it differently.

Five recommendations follow: give the deployment layer a named budget line costed to its instruments; fund the contracting counterparty that open weights lack; use procurement as a demand instrument on portability terms; take the agent permission standard through CEN-CENELEC and ETSI; and treat plurality of open-weight supply, with the verification capacity openness does not supply, as objectives in their own right.

# 1. What both documents establish

## 1.1 Open weights are a European strength on paper

Annex 4 states that Europe, with its long tradition of open-source innovation, is well positioned to lead on open-source and open-weight AI. It draws on the Commission's own *Europe's open source AI landscape* study for the argument that open source lowers licensing costs, reduces vendor lock-in and enables deployment on sovereign infrastructure.

Mozilla supplies the market evidence. A majority of tokens routed through OpenRouter now go to open-weight models, the five models carrying the most token volume on that platform are all open, and inference for GPT-4-class capability has fallen from roughly $20 to $0.40 per million tokens over thirty-six months. **Where developers choose on price, the token volume goes to open models.**

## 1.2 Europe's position in the model layer is real and narrow

Annex 4 records 30 notable European models out of 2,900 tracked by Epoch AI, against more than half from the United States and a quarter from China, with the Chinese share rising. On compute, it

[1] Reported without sample sizes or confidence intervals; treated throughout as a signal rather than a measurement.

cites Epoch's GPU cluster data showing the United States holding about three quarters of global cluster performance as of May 2025, with Germany, Japan and France playing marginal roles in the AI cluster landscape.

Mozilla's funding table places Mistral third among the open-weight developers it tracks by disclosed funding, at $3.05bn and approximately $400m ARR, behind DeepSeek at $7.4bn and Moonshot at $3.9bn; Zhipu AI and MiniMax, which listed in Hong Kong in 2026, disclose no totals. Annex 4 records the €1.7bn Series C at €11.7bn post-money led by ASML. **One credible European champion exists, and it sits alongside two Chinese firms with more capital and a state strategy at their back.**

### 1.3 Europe's position below the model layer is close to absent

Annex 4′s value-chain table gives European cloud providers less than 5% market share against roughly 85% for US hyperscalers; its investment chapter adds that 74% of Member States rely at least partially on US cloud providers while only 14% use European ones, and that over 80% of Europe's technology stack is imported. On Mozilla's funding map, the companies listed under the inference and tooling layers are American, with no European entry at either layer. The two dependencies compound: the EuroStack initiative observes that frontier models arrive bundled with hyperscaler clouds, so the model decision and the cloud decision reach a European buyer as one, in exactly the enterprise channel that OpenRouter cannot see. The portability terms of the third recommendation below are written for that bundle.

## 2. Where the money parts from the diagnosis

### 2.1 The investment needs answer input scarcity

Annex 4′s investment needs are dominated by compute access, chips, data centre capacity, energy supply, connectivity, data interoperability and skills. The financial profile follows: €1,321m from DIGITAL and CEF Digital for AI over 2021-2027, of which €594m goes to testing facilities and technology infrastructure and €272m to deployment proper, alongside the AI Continent Action Plan and InvestAI, which mobilises €200bn with €50bn of EU public funding and €20bn earmarked for gigafactories. The framing responds to the European Court of Auditors' 2024 finding that EU investment lacked ambition against global leaders, a judgement about aggregate volume rather than about allocation across layers.

Mozilla's evidence suggests that the inputs at the centre of that list, capable models and the compute to train them, are no longer the binding constraint. On Chatbot Arena's preference ratings, the best open models sit 3.3% behind the top closed models, at or near parity on coding, instruction-following and general knowledge, with the residue concentrated in reasoning, long context and agentic tasks. Cost has fallen by a factor of fifty. On a neutral third-party run of Terminal-Bench 2.1, one harness driving every model, the best-scoring open model reaches 67.8% against 71.9% for the leading closed model, at roughly one fifth the cost per task. **A European buyer in 2026 does not fail to deploy because capable models are scarce or hosted inference overpriced.** The obstacle open-model teams rank at or near the top in every region of Mozilla's survey is infrastructure and compute cost: the bill for running models oneself, an operations cost that no amount of model quality removes.

### 2.2 The stated open-source strength ranks second tier

**Annex 4′s prose and its tables disagree about how much open source matters.** The narrative twice presents promotion of open-source AI as one of the high-priority trends identified for deployment by 2028. Both ranking tables place it lower: Table 2 lists "Open-source AI" in the medium-priority column alongside neurosymbolic AI, embodied AI and neuromorphic computing, and Table 3 marks the row (M). The high-priority column holds foundation models, AI agents and multi-agent systems and explainable

and trustworthy AI, among other entries, with HPC for AI, Edge AI and AIoT as cross-technology priorities. **A trend described in the text as a European strength is ranked in the tables as a second-tier deployment priority, and the document never reconciles the two.** The 2028 action column does carry the open-source measures forward despite the ranking; the investment needs and the figures sit with compute and infrastructure.

### 2.3 The value chain has no cell for the layer where value is moving

The eight segments of the value chain Annex 4 reproduces run from raw materials through semiconductors and cloud to foundation models, applications and services. Nothing in that map corresponds to orchestration, memory, tool access or agent permissions. **The value-chain table carries 2023 market-share data into a document allocating investment against a 2026 problem.**

Mozilla puts that layer at the centre and calls it the agentic harness. Its Terminal-Bench sequence is the sharpest evidence in either document. In May 2026, on the benchmark's 2.0 release, an independent scaffold drove Anthropic's model to 79.8% while the vendor's own harness managed 58.0% on the same model, a spread of 21.8 points that made the orchestration layer worth more than the model. Eight weeks later, on the 2.1 revision, the frontier labs had internalised the harness: the lab harnesses now win wherever both appear, the spread at the top has compressed to roughly three points, and no open model appears in the verified top tier, because none has a first-party harness to answer with.[2] The direction is consistent with the commercial logic: **once models commoditise, pricing power returns to the integrated harness.**

## 3. Public money backs open supply while developers buy closed

Mozilla's regional survey finds Western Europe and South America to be the only two regions where closed-model adoption runs ahead of open.[3] The measure is the share of surveyed developers using each kind of model at all, so the categories overlap: 89% of developers in Greater China and East Asia use open models against 70% in Western Europe, where closed use is higher still.

The instruments named in Annex 4′s open-source paragraph act on supply: a fund for open LLMs, multilingual model programmes, an EIF co-investment window for open-weight startups. **None of them changes what European public bodies buy.** Annex 4 does record that public procurement mechanisms for AI are under-utilised, and that companies moving from lab to market find public procurement hard to navigate; its one demand-side action for 2028 supports the public sector as a deployer of AI, and procurement itself appears in no action line, left as a barrier for sellers to navigate rather than as leverage over European purchasing.

## 4. The instrument and the layer run on different clocks

This is a difference of constraints rather than of diagnosis. Annex 4 is an input to a multiannual framework and works in framework time by construction: actions by 2028, actions for 2036-2040, first disbursement in 2028. Mozilla describes a market moving on its own clock: vertical integration of model and orchestration layer forming this year, the metered-pricing breakpoint arriving around 2027-2028 as providers list publicly and introductory discounts expire. Each clock is proper to its object, and the mismatch is

[2] Two benchmark snapshots are thin evidence for a structural claim.

[3] Mozilla reports no sample size for the regional adoption exhibit and no confidence interval anywhere in the survey, so an ordering claim of this kind is a signal rather than a measurement.

the point: **by the time the framework's first action column opens, the standards and integration choices that give the layer its shape will already have been made.** The Annex itself strains against its own grid, with Edge AI in the 2036-2040 investment column of the same document that ranks it a high priority for deployment by 2028 and neurosymbolic AI beside it, while the agent permission standard is being decided in 2026. Redirecting settled choices later remains possible, as the Union showed with data protection and the common charger, but joining a standards process early is cheap and redirecting a settled layer is expensive and uncertain even for a regulator able to mandate. The argument is therefore for using the instruments the Union already holds, above all procurement and standardisation, and not for waiting on the next framework. The framework money still matters; the first recommendation says what it should buy once it moves.

## 5. What open weights do and do not buy

Both accounts recommend open weights as a route to sovereignty, and Mozilla's data shows the supply is increasingly Chinese: Alibaba's Qwen at 942m cumulative Hugging Face downloads against 476m for Llama, Chinese models outrunning US models by more than three to one in token volume across OpenRouter's nine most-used models, DeepSeek reporting more than 26,000 enterprise accounts. Annex 4, whose manuscript closed in May 2026, credits Chinese state industrial policy for the advance in general, files DeepSeek under efficiency-driven responses to US export controls, and says nothing about the open-weight character of the releases. Mozilla reads the open-weighting itself as deliberate industrial policy, and the 15th Five-Year Plan, adopted on 13 March 2026, supports that reading with an explicit objective that Chinese models lead the global open-source ecosystem. The State Council's Opinions on Deepening the Implementation of the "AI Plus" Initiative of 26 August 2025 supply the demand side, setting penetration targets for next-generation intelligent terminals and agents in China of 70% by 2027 and 90% by 2030. Public release also functions as a hedge against semiconductor restrictions, since it moves global inference onto hardware users already own.

### 5.1 A downloaded weight file gives control

It does not phone home, cannot be revoked, and does not place the user under a foreign regulator's discretion at runtime. Mozilla's account of the June 2026 episode illustrates the property: a single export order obliged one American provider to cut access for every foreign national with no notice and no possibility of selective compliance. **It holds equally for weights of Chinese origin, which is the strongest thing that can be said for the current supply.**

### 5.2 It gives no assurance about the artefact

Weights are not training data, training code or evaluation. No practical audit exists for implanted triggers, and the backdoor literature is discouraging: Hubinger et al. (2024) construct models whose harmful behaviour survives supervised fine-tuning, reinforcement learning and adversarial training, and find that adversarial training can teach a model to recognise its own trigger and so conceal the behaviour more effectively, producing a false impression of safety. For a public body deploying in health, justice or defence, the openness of the weights establishes very little about whether the model can be trusted.

**Openness is a property of control, and verification is a separate question that openness does not answer.** Treating the licence as evidence of provenance is a market for lemons: the label is cheap and observable, the property it is taken to certify is neither. That gap narrows with verification infrastructure of the kind the backdoor result dictates: provenance rather than inspection, meaning training data, training code, reproducible evaluation and attested builds, since the weights themselves cannot be audited into trust. A licence supplies none of it.

### 5.3 Weights are a stock, and capability is a flow

Europe can hold the stock, and it cannot be taken back. The capability to produce the next generation is a flow, it is foreign, it is discretionary, and models depreciate fast enough that autonomy resting on a stock erodes at the rate the frontier moves. **Sovereignty gained at the deployment layer is real, and it does not extend backwards up the production chain.**

Whether the flow persists is a policy question with an economic answer, and the evidence in these documents points to fragility. Open release is a stated strategy, and strategies are revisable. Three routes lead away from it: the calculus shifts as export controls ease; a laboratory establishes a defensible lead and monetises it closed; or listed open-weight companies come under pressure to close the revenue asymmetry. The third is already in motion. Zhipu AI and MiniMax listed in Hong Kong in 2026, and DeepSeek, still private, has raised at a valuation above $50bn against roughly $220m of annual revenue. The OpenRouter split from May to September 2025, a fifth of token volume against 4% of revenue, illustrates the underlying economics without being any lab's income statement: weights given away and served at commodity prices earn their revenue mostly for whoever hosts them. Public-market ownership and thin monetisation pull in the same direction.

**A capability plan for 2028 to 2034 resting on a revisable foreign policy choice is an exposure, and the word "open" does not remove it.** Plurality of origin therefore belongs in the core of the programme, together with the verification capacity that openness does not supply. Annex 4 treats neither as a policy objective; Mozilla's closing recommendation to funders argues for the first.

## 6. Recommendations

### 6.1 Give the deployment layer a named budget line

**The evidence assembled here does not settle whether Europe should hold sovereign training capacity.** That case rests on optionality and industrial capability, and deployment data bears on it from both directions without deciding it. The claim in this section is narrower, and it turns on complementarity: any model needs the deployment layer to be useful, European or otherwise.

The Commission files the €20bn earmarked under InvestAI for AI gigafactories under training: the facilities, up to five of them, are presented as enabling the training of complex, large-scale AI models at unprecedented scale. The label does not survive the arithmetic of Cottier et al. (2024), the one systematic cost series for frontier training. On their preferred measure, the most expensive disclosed training runs cost $40m in amortized hardware and energy, on a 2.4-times-per-year trend that reaches $1bn per run around the start of 2027; the up-front acquisition cost of the hardware sits one to two orders of magnitude higher, roughly $800m for the cluster behind GPT-4. Against that series, €20bn split over five sites fits no training strategy. It is orders of magnitude above what the efficient path needs, the path Annex 4 itself credits to DeepSeek; on the frontier path, the same trend puts the hardware behind a single 2027-class run in the tens of billions, a league that five facilities of roughly €4bn do not enter and that a one-off fund cannot keep pace with. The European training demand to fill the facilities does not exist, in a market where Europe fields one frontier-adjacent laboratory; filled at all, they will fill with inference and serving. The capex also leaks on arrival: the dominant cost items in an AI data centre are accelerators and high-bandwidth memory, with no European supplier at the required level, so a gigafactory euro buys capacity located in Europe and supplier value booked abroad. The programme ends up owning general accelerator compute, and how much of it serves deployment workloads is an access-design choice it has not yet made. Annex 4's wider compute argument points the same way, covering access to inference as well as training, and noting that European resources are often optimised for scientific simulation rather than for either. **The evidence supports a narrower claim: adding training compute will not close the production gap,**

**because the obstacles open-model teams report are operational and capable models are already abundant.**

The layer that would move it sits between the model and the application. Annex 4 approaches it as AI middleware, covering model deployment, monitoring, security, compliance and interoperability. Mozilla approaches it as the agentic harness, covering the orchestration loop, tools, memory, sandboxing and the permission model. The two overlap without coinciding, and European capability is absent from the overlap. That layer currently holds one bullet point in Annex 4's infrastructure row and no figure anywhere in the programme. The right size for a named budget line comes from costing its instruments at working strength: integrator accreditation in every Member State, portability tooling, standards participation, and inference and operations access for the teams doing the deploying, including designed-in deployer access to the factory and gigafactory capacity as it comes online. A first-pass costing lands in the low billions over the framework's seven years, with the integrator networks and deployers' inference access as the large items: a tenth of the €20bn compute commitment and no rival to it, and far above the current programme's allocation to the layer itself, which is nothing. The €272m deployment line under DIGITAL funds demonstration and roll-out projects, and no line funds the software layer they depend on.

The timing argument made earlier divides the work between the Union's instruments. The standards and integration window closes in 2026 and 2027; the procurement and standardisation instruments below, which the Union already holds, answer on that timescale. A budget line opening in 2028 buys capacity that does not expire with the window: integrators, portability, operations, verification.

## 6.2 Fund the contracting counterparty alongside the technology

The most useful number in either document is Mozilla's breakdown of production success by company size. Among companies of up to fifty people, teams building on closed and on open models reach live production at nearly the same rate: 54% against 53%. The two curves then separate as companies grow, and at enterprises above a thousand employees 73% of closed-model teams have something in production against 57% for open. Everything a large firm has more of, engineers included, lifts closed deployment and barely lifts open, which is hard to square with a shortage of in-house engineering as the explanation for where open stalls.

**A reading consistent with both curves is that firm size brings procurement capacity: the ability to buy an indemnity, a service level agreement, a support contract and a named party to hold liable.** Closed vendors sell that bundle with the model; open weights arrive with no counterparty attached, and a buyer must assemble one. Mozilla's own reading is adjacent, with enterprises buying their way through closed deployment while open deployment waits on operational tooling nobody has finished. The survey's blocker table accommodates both readings: infrastructure cost sits at or near the top in every region, Western European teams cite maintenance and deployment complexity next, and security and support concerns weigh heaviest in South Asia. Tooling and counterparty point at the same remedy, an accountable party that packages the deployment work, the support and the liability.

Annex 4 already contains the right instrument, filed under ecosystem building as a 2028 action: trusted AI integrator networks accrediting 10 to 20 SME-focused providers per Member State through the European Digital Innovation Hubs. That instrument belongs at the centre of the programme, resourced accordingly, with accreditation conditioned on the integrator carrying contractual liability for the deployment itself: service levels, security of the integration, maintenance. Liability for the artefact is a different matter; no small integrator can insure an unauditable model, and closed vendors' contracts do not insure their models' behaviour either. Behind the artefact stands the verification infrastructure of the final recommendation.

### 6.3 Use procurement as the demand instrument, on portability terms

Europe's adoption gap has a demand side, and the Union controls a large share of the demand. The tractable requirement is portability rather than origin: exit-cost disclosure at contract award, second-sourcing capability for any model-dependent system above a threshold, and portable formats for the accumulated memory and evaluation data a deployment generates. Requiring that a system can be moved is enforceable at procurement, benefits European suppliers without discriminating in their favour, and preserves an option that a metered contract removes. Pooled public demand also acts on supply. EuroStack names the coordination failure on the inference side: Europe has not shown enough inference demand to motivate building European capacity, and the hyperscalers' fleets are already secured. Aggregated public purchasing of inference and operations, on the portability terms above, breaks that circle and connects this recommendation to the budget line of the first. **The timing argument is specific: second-sourcing is cheap while introductory pricing holds, and Mozilla's forecast, hedged but explicit, puts the expiry around 2027-2028.**

### 6.4 Take the agent permission standard

Mozilla reports that no portable specification exists for which write actions an agent may perform unattended across a Model Context Protocol host, an A2A peer, a direct tool call and a framework boundary, and identifies this as the largest unresolved gap in the layer. On its account both protocols settled authentication and left authorisation open, with governance of both now sitting under the Linux Foundation's Agentic AI Foundation.[4] **Knowing which agent is acting says nothing about what it is permitted to do.**

This is a standard-setting opening of the kind the Union has taken before. It maps onto machinery the AI Act already establishes, since Article 40 routes conformity for high-risk systems through harmonised standards requested from the European standardisation organisations, and it sits at the exact point where a high-risk system becomes operationally dangerous. Work through CEN-CENELEC and ETSI should start, or join what already exists, on a timescale measured in months. If the de facto answer ships inside a closed platform first, the layer closes on terms Europe does not write. The standard is also the one instrument here aimed at the layer's pricing power: a portable permission model keeps the harness contestable rather than welded to a single vendor's stack.

### 6.5 Make plurality of origin and verification capacity explicit objectives

Support for open-weight adoption should be conditioned on maintaining more than one viable supply origin, including European ones, and on demonstrated ability to substitute between them. Otherwise the open-weight strategy delivers deployment autonomy on a supply base Europe does not influence. Mozilla's one recommendation addressed to public money points the same way: sustained funding for plural open-model supply through the period when open-lab economics are weakest, so that the open default does not settle into a single origin by absence. Plurality spending is model-layer spending, and it is compatible with the ordering above: the deployment layer holds the first claim on the marginal euro, and plural supply is the insurance that keeps a deployment strategy built on foreign weights safe to run.

Verification is the second objective, and the backdoor evidence fixes its form: provenance. Public deployment programmes should fund, and progressively require, disclosed training data and code, reproducible evaluation and attested builds for the models they support. That is the infrastructure that narrows the lemons gap a licence leaves open; it would also let the integrators of the second recommendation extend their liability from the deployment toward the artefact at a price a small firm can bear.

[4] The claim is a negative about the standards landscape and should be checked against current CEN-CENELEC and ETSI work before it is relied on in a formal recommendation.

## References

Cottier, B., Rahman, R., Fattorini, L., Maslej, N., Besiroglu, T., Owen, D. (2024). *The rising costs of training frontier AI models*. arXiv:2405.21015.

Epoch AI. *AI Models* dataset, epoch.ai/data/ai-models.

Epoch AI. *GPU Clusters* dataset, epoch.ai/data/gpu-clusters.

EU Artificial Intelligence Act, Article 40 (harmonised standards and standardisation deliverables), artificialintelligenceact.eu/article/40.

European Commission (2025). *AI Continent Action Plan*.

European Commission (2025). *EU launches InvestAI initiative to mobilise €200 billion of investment in artificial intelligence*. Press release IP/25/467, 11 February 2025.

European Commission, DG CNECT (2026). *Boosting tech deployment beyond 2027: new study highlights digital opportunities and challenges for the EU*. Study landing page, 23 June 2026, covering 13 technology areas and informing the 2028-2034 funding period.

European Commission, DG CNECT. *Europe's open source AI landscape: a lever for innovation and sovereignty*.

European Court of Auditors (2024). *Special report 08/2024: EU artificial intelligence ambition. Stronger governance and increased, more focused investment essential going forward*, eca.europa.eu/en/publications/SR-2024-08.

EuroStack (2026). *European AI maximalists and their masters worry about EuroStack*. Blog post, 9 July 2026.

Directive (EU) 2022/2380 amending Directive 2014/53/EU on radio equipment (common charger), eur-lex.europa.eu/eli/dir/2022/2380/oj.

General Data Protection Regulation, Regulation (EU) 2016/679, eur-lex.europa.eu/eli/reg/2016/679/oj.

Hubinger, E., Denison, C., Mu, J., Lambert, M., Tong, M., MacDiarmid, M., Schiefer, N. et al. (2024). *Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training*. arXiv:2401.05566.

McKinsey (2024). *Time to place our bets: Europe's AI opportunity*.

Mistral AI (2025). *Mistral AI raises €1.7bn to accelerate technological progress with AI*.

Mozilla (2026). *The State of Open Source AI, V1.0, July 2026*, stateofopensource.ai. Also the full report PDF and the launch announcement.

Stanford DigiChina (2026). *Forum: Technology in China's 15th Five-Year Plan*. Analysis of the plan adopted by the National People's Congress on 13 March 2026.

State Council of the People's Republic of China (2025). *Opinions on Deepening the Implementation of the "Artificial Intelligence Plus" Initiative*, 26 August 2025. English summary.

Strauka, O., Timan, T., Denais, A. (Technopolis Group) (2026). *Study on the EU's critical digital capacities deployment beyond 2027. Annex 4: Technology Area Report, Artificial Intelligence*. European Commission, DG CNECT. Manuscript completed May 2026, published 23 June 2026. ISBN 978-92-68-40090-6, data.europa.eu/doi/10.2759/1036940.

**Technical references:** Agentic AI Foundation (Linux Foundation), agenticaifoundation.org · A2A Protocol, a2a-protocol.org · Model Context Protocol, modelcontextprotocol.io · OpenRouter, openrouter.ai · Terminal-Bench, tbench.ai

## Caveats on the underlying evidence

Mozilla is an interested party and the report is written as advocacy, though it discloses its weak points. Four measurement issues affect how its figures should be quoted.

All usage and revenue shares derive from OpenRouter, a price-comparison router whose users are selected on cost sensitivity, which is the variable of interest. Direct enterprise consumption through Azure, Bedrock and vendor contracts is invisible in it. Closed providers still lead the platform by request count; the open lead is a lead in token volume, concentrated in coding and agentic workloads.

The token-volume and revenue shares quoted together are OpenRouter measurements over the same May to September 2025 window. The majority-of-tokens figure refers to mid-2026 and has no matching revenue measurement.

The 3.3% capability gap is a percentage difference of Chatbot Arena scores, preference-based ratings that carry no cardinal interpretation.

The adoption and production figures are global. Mozilla breaks production down by company size and not by region, so reading the 51% and 63% rates as European assumes European teams are not systematically different from others. Mozilla's regional table of blockers, in which the same four operational obstacles dominate every region in shifting order, is the basis for that assumption. The overall open production rate of 51% sits below every band of the firm-size breakdown, 53% to 57%, a discrepancy the report does not reconcile.

On the Commission side, Annex 4 reproduces a McKinsey value chain built on 2023 market-share data and a 2024 Court of Auditors judgement on investment scale, both of which predate the developments this note addresses.

## About the Abilian Econ Lab - econ.lab.abilian.com

The Econ Lab is Abilian's economics research laboratory. It works on a single research programme, the economics of European digital sovereignty, which asks through which mechanisms dependence on foreign technology providers translates into economic and political exposure for Europe, and which interventions actually reduce it. It answers with measured dependency graphs, formal theory, game-theoretic mechanism analysis, service-level architecture measurement, and the economics of open-source adoption and supply, in a field where most public debate runs on anecdote.

**Author.** Stefane Fermigier (sf@abilian.com) holds a PhD in mathematics and is adjunct professor (*professeur associé*) of open-source software at Université Paris Cité. He is the founder and CEO of Abilian, a French open-source software company, and has spent three decades building and organising the European open-source industry: he founded Nuxeo, and co-founded and co-chairs the CNLL, the French open-source business association.