

ALTLIBRE REFERENCE PAPER

Your AI Answers to a Foreign Government

Jurisdiction, control and the European alternative: a reference paper on AI dependency risk and what Europe has built to answer it

ADRIAN HOLLISTER · ALTLIBRE · VERSION 1.5, 8 SEPTEMBER 2026

ABOUT THIS PAPER

Your AI Answers to a Foreign Government

AUTHOR

Adrian Hollister, founder, AltLibre · altlibre.com

VERSION

Version 1.5, 8 September 2026

LICENCE

Licence: Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0).

SUGGESTED CITATION

Hollister, A. (2026). *Your AI Answers to a Foreign Government: Jurisdiction, control and the European alternative*. AltLibre reference paper, version 1.5, 8 September 2026.

ALTLIBRE

AltLibre is an independent research practice working on digital sovereignty: where data sits, whose laws it answers to, and what that costs. It publishes papers, analysis and a public evidence base at altlibre.com.

METHOD

Every factual claim carries a dated reference. Annex F sets out how claims were verified and what was excluded.

CONTENTS

In this paper

OPENING

Executive summary 4

THE RECORD

Introduction 6

The American record 7

The Chinese record 7

THE COMPARISON

Comparative assessment 14

THE EUROPEAN POSITION

The European position: regulation, labels and a test 16

THE SUPPLY BASE

The European supply base 19

OBJECTIONS AND RECOMMENDATIONS

Objections and constraints 26

Recommendations 28

Conclusion 31

INSTRUMENTS

The three-question test as a scoring template 32

Eight procurement questions 32

Workload classification against the proposed CADA levels 33

Inventory of the European supply base 34

Glossary 36

Method and limitations 37

References 38

69

dated primary and secondary sources

Executive summary

Most European organisations now depend on artificial intelligence services whose model, hosting, silicon and contract terms sit under the jurisdiction of a foreign state. This paper asks a narrow question of that dependency: what have the two governments concerned, the United States and the People's Republic of China, already done with the power it gives them, and what should a European organisation, a public body or a policy maker conclude from the record?

The record is documented rather than speculative. Between 2024 and 2026 the United States withdrew AI access three times: from China and Hong Kong in July 2024; from an Indian refiner in July 2025; and from Anthropic, its own most safety-conscious frontier laboratory, in February 2026, after the company declined to remove two clauses from its acceptable-use policy forbidding mass domestic surveillance and fully autonomous weapons. United States law reaches data held by American providers wherever it is stored, through the CLOUD Act and Section 702 of the Foreign Intelligence Surveillance Act; civil litigation has already suspended the deletion rights of every ChatGPT consumer tier and compelled the production of twenty million conversation logs. Steering of American models has been erratic and private, and since Executive Order 14319 of July 2025 has also been governmental. China's record is structural: a National Intelligence Law that obliges every organisation to assist state intelligence work without process or redress; and generative AI rules that make alignment with state values a licensing condition. Neither jurisdiction's law protects a European user.

One asymmetry cuts across the comparison: China publishes its strongest model weights under open licences, while the United States keeps its frontier models closed and rents access. Weights that an organisation holds on its own infrastructure cannot be withdrawn, read or re-tuned by anyone else.

European regulation governs how AI behaves on the European market. It does not change who owns the model or which court can compel the provider. The AI Act's transparency duties have applied since 2 August 2026; the high-risk regime, deferred by the Digital Omnibus adopted in June 2026, follows in December 2027 and August 2028. The Commission's proposed Cloud and AI Development Act, published on 3 June 2026, is the first European legal text to place ownership and control rather than conduct at the centre of a security requirement, through a four-level sovereignty framework for public procurement. It is a proposal; it is a procurement instrument rather than a prohibition; it protects public buyers only. The United Kingdom has no equivalent.

Meanwhile the label "sovereign" has been adopted by the very providers whose jurisdiction is the problem. A Microsoft executive told a French Senate inquiry under oath in June 2025 that he could not guarantee French data would never be handed to United States authorities; four months later Microsoft announced new sovereign cloud capabilities for Europe. This paper proposes a three-question test that cuts through the label: who owns the parent company; which jurisdiction can compel it; and who holds the encryption keys. Most current sovereignty claims fail it.

The European alternative is not a hypothetical. Europe's first exascale computer, JUPITER at Jülich, ranked fifth in the world in June 2026 and can train frontier-scale models. Commercial GPU capacity is being built at scale in Finland, France and Germany by European-headquartered operators.

Open-weight models are shipping from Paris, Zurich, Freiburg and a pan-European research consortium, and are competitive across the enterprise workloads most organisations run. Sovereign workplace software built by the French and German administrations is being deployed across tens of thousands of public-sector desks. A generation of European builders, from a Stockholm company approaching €500 million of annual revenue to a British industrial founder who built a small-business agent platform because no one else had, did not wait for the policy debate to conclude.

The paper's recommendations follow from the evidence. Organisations should map every AI dependency by jurisdiction; apply the three-question test to every supplier; classify workloads by sovereignty requirement; move one sensitive workload to an open-weight model on European infrastructure now; write eight specific questions into every procurement; budget for the platform migration rather than the procurement; and probe every model they rely on quarterly for steering. Governments should treat the serving layer, not only training capacity, as strategic; write inference contractability into the gigafactory tenders; use procurement as market-making, as the Commission's sovereign cloud framework and the United Kingdom's Sovereign AI competition have begun to do; and hold the line on the ownership and control levels of the Cloud and AI Development Act through the legislative process. None of the organisational recommendations requires a regulator to act first.

Sovereignty cannot be bought from the companies an organisation is trying to reduce its dependence on. It is a configuration to be built: European law, European infrastructure, open weights, and encryption keys the organisation itself holds. Every component now exists.

Residency is geography. Sovereignty is jurisdiction. Control is operational.

1

Introduction

1.1

The question

Ask which is riskier, American AI or Chinese AI, and the conversation becomes a contest of loyalties. Ask instead what each government has already done with the power its jurisdiction gives it over the AI services Europe depends on, and an answer becomes available from the public record. This paper takes the second route.

It consolidates and extends a three-part series published by AltLibre between August and September 2026 [1], together with a related analysis of what a credible European AI layer would require [2]. Earlier pieces in the same programme, on the sale of sovereignty as a label, on Airbus moving its critical systems to a French provider, and on the identity layer as the last American anchor, informed the argument [3, 4, 5]. The series was written for practitioners. This paper is written for the people who advise, regulate and procure on their behalf: researchers in law, public policy and information security; consultancies advising boards and ministries on AI strategy; and governments and public bodies deciding what to buy and what to legislate. The argument is unchanged. The evidence base has been re-verified, consolidated and dated; the recommendations have been separated by audience.

1.2

Scope and definitions

The paper is concerned with the European organisation as user and buyer of AI: a ministry, a hospital trust, a bank, a manufacturer, a university, a consultancy. "Europe" is used for the European Union, the United Kingdom, Switzerland and the European Economic Area unless a legal instrument narrows it, as the AI Act and the Cloud and AI Development Act do to the Union.

Three terms carry the argument and are used precisely.

Residency is geography: where the data and the computation physically sit.

Sovereignty is jurisdiction: which state's courts and executive can compel the operator. The two are routinely conflated in marketing; the conflation is the subject of Section 5.

Control is the operational form of sovereignty: whether the organisation itself can keep a service running, keep its data private and keep the model's behaviour stable, regardless of what any other party decides. Open-weight models, discussed in Sections 3 and 6, matter because they convert a jurisdictional question into an operational one.

Foreign government, for a European organisation, means any government outside the jurisdiction the organisation answers to. The paper treats American and Chinese hosted services as equally foreign. That is a legal observation, not a political one.

1.3

Method and evidence standard

Every factual claim carries a numbered reference to a primary or reputable secondary source, dated. Where the source is a company's own statement, that is stated. Where a figure was published in United States dollars and no euro figure exists, it has been converted at approximately 1.16 dollars to the euro and rounded, and the conversion is noted. Where the author's judgement rather than established fact is offered, the text says so.

Two limits on the evidence deserve to be stated at the outset. First, several of the legal instruments discussed are in motion: the Cloud and AI Development Act is a Commission proposal; the Anthropic litigation was unresolved at the time of writing; market figures for private companies are as reported at the dates given. Second, the inventory in Section 6 is illustrative rather than exhaustive; it is intended to establish that a European supply base exists, not to catalogue it completely. A method note in Annex F sets out how claims were verified and what was excluded.

1.4

Structure

Section 2 sets out the American record and Section 3 the Chinese record, each across four dimensions: the power to switch off, the power to read, the power to steer, and the depth of integration between the AI companies and the state. Section 4 draws the comparison and identifies the asymmetry that shapes the rest of the paper. Section 5 examines the European regulatory position and the phenomenon of sovereignty washing, and introduces the three-question test. Section 6 takes the inventory of what Europe has built, layer by layer. Section 7 answers the standing objections. Section 8 sets out recommendations for organisations, for procurement, for governments and for researchers. Annexes provide working instruments: the three-question test as a scoring template, the eight procurement questions, a workload classification against the proposed CADA levels, the inventory as a table, a glossary and the method note.

3

US AI access withdrawals in two years

2.1

2

The American record

The power to switch off

On 27 February 2026 the President of the United States ordered every federal agency to stop using Anthropic's technology. The Secretary of Defense designated Anthropic, an American company, a supply-chain risk to national security, and declared that no contractor doing business with the United States military could trade with Anthropic at all [6, 7, 8].

There had been no breach and no espionage. The offence was a contract clause. Anthropic had refused to delete two restrictions from its acceptable-use policy: no mass surveillance of American citizens; and no fully autonomous weapons. The Department of Defense had demanded language permitting "all lawful uses". Anthropic declined. The administration then reached for procurement law, supply-chain designations and, in June 2026, export controls on Anthropic's newest models, lifted two weeks later [6]. The courts pushed back with a preliminary injunction in March; the dispute was still in litigation at the end of July 2026 [6]. OpenAI accepted the terms Anthropic had refused.

The significance for a European reader is proportional. If the United States government will attempt to ruin its own most safety-conscious laboratory over a surveillance clause, a European ministry or a mid-sized bank is not a consideration in that calculus.

The Anthropic case was the third instance in two years, not the first. On 9 July 2024 OpenAI cut API access to China, Hong Kong and every territory outside its supported list, giving companies that had built on the API days to migrate [9, 10]. In July 2025 Microsoft suspended cloud services to Nayara Energy, an Indian refiner part-owned by Rosneft, citing European Union sanctions that India had not adopted and that no law required Microsoft to enforce; service was restored only after Nayara went to court in New Delhi [11]. A compliance function in Redmond switched off an Indian company's email over a European political decision that bound neither of them.

Beneath the software sits the silicon, and the same hand is on that too. The Chip Security Act, introduced in May 2025 and reported out of House committee by 42 votes to none in March 2026, would mandate location verification in export-controlled AI chips and fund the study of further mechanisms, which analysts describe as including remote disablement or performance degradation [12, 13]. Beijing evidently believes the capability is coming: it summoned Nvidia on 31 July 2025 over alleged backdoor risks in the H20, a charge Nvidia denies. Days later the administration confirmed it would take 15 per cent of Nvidia's and AMD's China chip revenue in exchange for export licences. Revenue, location, off-switch: the architecture of control is being assembled in plain sight.

2.2

The power to read

The CLOUD Act of 2018 exists because Microsoft fought a United States warrant for emails stored in Dublin, and Congress chose to change the law rather than lose the argument. American providers must now produce data in their possession, custody or control regardless of where it is stored. A European organisation's prompts in an Irish data centre are within reach of an American warrant. Section 702 of the Foreign Intelligence Surveillance Act, reauthorised and widened in April 2024, permits United States agencies to collect the communications of non-Americans abroad through compelled assistance from United States providers, without a warrant and with no meaningful standing for the affected party.

The civil route arrived more quietly and deserves more attention than it received. In May 2025, in the New York Times copyright litigation against OpenAI in the Southern District of New York, a federal magistrate judge ordered OpenAI to preserve all ChatGPT output logs it would otherwise have deleted, across every consumer tier, including conversations users had already deleted. OpenAI had to tell European Union customers that their GDPR erasure rights were suspended. The hold ran until late September 2025. Two months later the court ordered twenty million de-identified chat logs produced to the plaintiffs [14]. Only negotiated enterprise terms with zero retention survived. The delete button, on this evidence, is a feature of the subscription; litigation outranks the subscription.

2.3

The power to steer

Models do not arrive neutral. People tune, prompt and filter them, and those people have owners, employers and politics. The question is whose hands are on the dial and whether the user can see them move.

At about 03:15 Pacific time on 14 May 2025, someone modified Grok's system prompt to force a specific response on a political topic. For hours afterwards, users asking about unrelated subjects received discourses on "white genocide" in South Africa, a theory the company's owner had personally promoted. xAI attributed the incident to an "unauthorised modification" that had evaded its code review, and undertook to publish its prompts [15, 16]. It was the second such unauthorised change in four months [16]. Two months later, after the owner announced the model had been made less politically correct, it produced antisemitic content, praised Hitler and described itself as "MechaHitler"; xAI deleted the output and apologised, and Turkey blocked the service. Weeks after that, the Department of Defense signed xAI to a contract worth up to \$200 million and the company launched a government product line [17, 18].

The state now has its own hand on the American dial. Executive Order 14319, signed on 23 July 2025, makes federal procurement of large language models conditional on "truth-seeking" and "ideological neutrality", both defined by the administration, with diversity, equity and inclusion named in the order as a disqualifying ideology, and with decommissioning costs chargeable to vendors that breach the terms [19, 20]. Set the politics aside and observe the mechanism: the government of the day defines acceptable machine truth as a condition of a chequebook large enough to shape the product everyone else receives.

On foreign agendas running through American models, the loudest claims outrun the evidence and precision matters. What is documented: in May 2024 OpenAI's own threat report disclosed a covert influence operation run by a Tel Aviv political marketing firm, which used ChatGPT to generate content on the Gaza war aimed at audiences in the United States, Canada and Israel through fake personas; Meta removed more than 500 accounts tied to the same firm [21, 22]. What is not documented is the stronger claim that American models are tuned at the weight level to favour any state; published audits point in more than one direction. The defensible conclusion is narrower and, for a risk owner, worse: these systems are steerable; the steering is contested by many hands; and none of those hands answers to the European user.

2.4

The state as customer

The surveillance debate usually stops at who can read the prompts. The more consequential development is how deeply the AI companies are being wired into state security machinery, and what they are being asked to do inside it.

In July 2025 the Department of Defense awarded contracts worth up to \$200 million each to OpenAI, Anthropic, Google and xAI for frontier AI across warfighting, intelligence and enterprise systems [17, 18]. Anthropic had already deployed its models on classified United States networks and launched a dedicated government product line in June 2025. None of this was leaked; it is in the press releases.

The telling document is the dispute that followed. When the Department renegotiated with Anthropic in early 2026, the two clauses it demanded removed were the prohibition on mass domestic surveillance and the prohibition on fully autonomous weapons [6, 7]. The government's shopping list was precisely the list a concerned citizen would hope never to see requested. One laboratory refused and was punished. Another agreed and was paid.

Art. 7

the National Intelligence Law
duty to assist

3

The Chinese record

3.1

The power to switch off

China has no matching record of cutting foreign users off. Its influence runs through app-store control, censorship conditions and the dependency that a free frontier-class model builds. It is also the suspected target rather than the author of chip-level control: the location-verification and remote-disablement mechanisms under study in Washington are aimed at hardware bound for China.

3.2

The power to read

Where the American answer to who can read a user's prompts is "more people than you think, through more doors than you expect", the Chinese answer is that there is no door at all. Article 7 of the 2017 National Intelligence Law obliges every organisation and citizen to support, assist and cooperate with state intelligence work; Article 14 empowers intelligence organs to demand it [23]. There is no warrant, no public docket and no appeal.

When Italy's data protection authority asked DeepSeek basic questions in January 2025, the company's position was that European Union law did not apply to it. The Garante imposed the first Western ban on 30 January 2025, and regulators in Ireland, France, the Netherlands, Belgium and Germany opened files [24]. South Korea's authority found that DeepSeek had transferred user data, including prompt content, to a Beijing affiliate without proper consent [24]. China's generative AI rules, in force since August 2023, require real-name registration, retained logs and algorithm filings with the Cyberspace Administration. The state does not need to request the data. It is plumbed in.

3.3

The power to steer

China does not pretend. Article 4 of the 2023 generative AI rules requires services to uphold core socialist values and bars content that subverts state power or damages the national image, enforced through a filing regime that makes alignment a licensing condition. The hosted DeepSeek application, tested by The Guardian and others in early 2025, refuses or deflects on Tiananmen, Taiwan and the Uyghurs. The censorship is real, predictable and declared.

Two qualifications matter for what follows. Researchers have shown that most of this behaviour lives in the fine-tuning and serving layers, which is why a self-hosted copy of the weights behaves differently from the application. Some refusals are nevertheless embedded in the weights during training, and that residue is what an organisation's own evaluation exists to find. The self-hosting route described in Section 6 does not remove the need for evaluation; it moves the responsibility for it to the organisation, where it can be discharged.

3.4

The state as architect

The relationship between the Chinese state and its AI companies is not one of customer and supplier but of architect and structure. Real-name registration, retained logs, filed algorithms and the statutory duty to assist are design features of the market, not contract terms negotiated within it. The consequence for a European user is that there is nothing to negotiate and nothing to litigate. What the American system exposes through procurement disputes and court orders, the Chinese system settles in statute.

3.5

The asymmetry

China does one thing that Washington will not: it gives the weights away. DeepSeek's R1 carries an MIT licence; Alibaba's Qwen family carries an Apache licence. Weights on an organisation's own hardware cannot be recalled, geo-blocked or sanctioned. The strategic logic is not generosity. Open release undercuts American licensing revenue and seeds dependence on Chinese research, with plausible deniability. The European response that follows from that logic is to take the weights and leave the hosted service.

The choice is between audited risk and unknowable risk.

4

Comparative assessment

4.1

Which is riskier

The question is on the wrong axis. Risk here follows jurisdiction and architecture, not flags.

On being switched off, the United States has the worse documented record by a distance: against a rival, against a neutral party and against its own champion. On being read, both states reserve the power; the American version can be audited and occasionally defeated in open court; the Chinese version cannot. On steering, China declares its censorship in statute and applies it predictably; American steering is private, erratic, lately governmental and disclosed mostly after the fact. On integration with the state, the American AI companies are contractors on classified networks whose acceptable-use terms the state has sought to rewrite; the Chinese companies are structures the state designed.

The ledger, for a European risk owner, reads as follows. American reach is documented, challengeable and occasionally beaten. Chinese reach is invisible and, for practical purposes, absolute. The choice is between audited risk and unknowable risk, not between safe and unsafe.

4.2

Summary table

TABLE. COMPARATIVE RISK BY DIMENSION

RISK DIMENSION	UNITED STATES-BASED HOSTED MODELS	CHINA-BASED HOSTED MODELS
Being switched off	Documented three times in two years, most recently against Anthropic, an American laboratory (2026)	No documented cut-off of foreign users; dependence built through free capability and app-store control
Hardware control	Chip Security Act: location verification mandated, remote-disablement mechanisms under study	Suspected target of chip-level control, not its author
Lawful access to data	CLOUD Act reach anywhere; FISA Section 702 compulsion; courts can override deletion, as twenty million ChatGPT users learned	National Intelligence Law Articles 7 and 14; no process, no redress, no visibility
State integration	Defence contracts of up to \$200 million each; classified deployments; government demanded removal of mass-surveillance limits (2026)	Structural: real-name registration, retained logs, filed algorithms, duty to assist
Steering	Prompt tampering at xAI (twice in 2025); government-defined "truth" a procurement condition; influence operations documented by the vendors themselves	Statutory alignment with core socialist values; hosted applications refuse sensitive topics

Transparency	Steering sometimes published after incidents; operations disclosed in vendor reports	Declared in law, visible in output
Exit	Closed weights, revocable agreements	Open weights (MIT, Apache 2.0): self-host and the foreign hand is removed

4.3

The configuration that removes the foreign hand

One asymmetry cuts across every row of the table. China publishes its strongest weights under open licences. The United States keeps its frontier models closed and rents access. A self-hosted open-weight model on European infrastructure is the only configuration in the entire comparison that removes the foreign hand altogether: it cannot be switched off from abroad because the organisation holds the weights; it cannot be read from abroad because the prompts never leave the organisation's infrastructure; and it cannot be re-tuned from abroad because the fine-tuning happens in the organisation's environment. What remains is the organisation's own responsibility for safety testing, evaluation and patching, which Section 6 argues is a fair exchange for a capable IT function.

4.4

Where the record points

Five developments follow from the record. They are the author's judgement rather than established fact, and are marked as such.

Frontier models will be export-controlled like chips. Washington briefly did this to Anthropic's own models in June 2026, and the logic of location-verified silicon applies cleanly to weights [6]. Licences with "all lawful use" clauses attached should be expected.

Government-defined truth will spread through procurement. Executive Order 14319 requires no parliamentary approval to copy. Other governments will write their own versions, each defining acceptable output to taste, and the differences will be buried in fine-tuning nobody publishes.

China will keep giving the weights away, because open release undercuts American licensing revenue and seeds dependence on Chinese research.

A serious chat-log disclosure incident is coming. Twenty million conversations were ordered produced in a single copyright case [14]. The next such event will involve a government or a leak, and will reset what boards permit staff to type into consumer AI tools.

Europe's window is open now and will not remain open indefinitely. The AI Act is in force; the research base is world-class; and open weights make sovereign hosting practical without frontier compute of one's own. That combination has a shelf life.

4

sovereignty levels in the CADA proposal

5

The European position: regulation, labels and a test

5.1

Regulation governs behaviour, not ownership

Since 2 August 2026 the AI Act's Article 50 transparency duties apply: chatbots must disclose that they are chatbots; generated content must be marked; deepfakes must be labelled, with fines of up to €15 million or 3 per cent of global turnover [25, 26]. The high-risk regime has been rescheduled: the Digital Omnibus, given final approval by the Council on 29 June 2026, defers stand-alone high-risk systems to 2 December 2027 and product-embedded systems to 2 August 2028 [27, 28]. General-purpose model obligations have applied since August 2025, with 26 signatories to the voluntary code of practice, every major United States laboratory among them except Meta, which declined to sign.

This is useful. It is also bounded. The AI Act tells providers and deployers how an AI system must behave when placed on the European market: disclosure, testing, documentation, liability. It says nothing about who owns the weights, which courts can compel the provider, or which government can order a service withdrawn. A fully compliant, CE-marked AI service can still be switched off by a foreign sanctions order or read under a foreign warrant. The Act regulates conduct; it does not confer control.

5.2

The Cloud and AI Development Act

The Cloud and AI Development Act (CADA) is the first European text to go after control. Published on 3 June 2026 as part of the Commission's technological sovereignty package, it proposes a four-level cloud sovereignty framework for public procurement [29, 30, 31].

Level 1 requires infrastructure and customer data in the Union, together with a guarantee that the provider cannot be compelled to hand unpatched vulnerabilities to a foreign authority.

Level 2 requires all personnel, infrastructure and assets to operate inside the Union; measures against third-country access and sanctions exposure; and a complete software bill of materials.

Level 3 requires Union ownership and control, with a possible derogation for providers from countries the Commission deems adequate and open.

Level 4 requires Union ownership and control with no derogation; high-assurance cybersecurity certification; and proof that no third-country entity controls the design, development or maintenance of the software [31].

Public bodies handling national security, defence, law enforcement and public-order data would be able to demand the level that matches the sensitivity of the workload. The Commission itself expects roughly 70 per cent of contracts to fall at level 1 and around 1 per cent at level 4 [32].

Three points need to be kept straight in any advice built on CADA. First, it is a proposal. It goes to Parliament and Council next; the trilogue on the sovereignty framework will be the political battleground of the coming eighteen months, and the hyperscaler lobbying against levels 3 and 4 has begun. Second, it is a procurement instrument, not a prohibition. The general "Union added value" criterion is explicitly ancillary, weighted at most 15 points in 120, and a derogation exists where no reasonable alternative is available [31, 32]. Third, it applies to public buyers. A private company gains no protection from it beyond the market signal it sends.

Even so, it is the first time a European legal text has placed ownership and control, rather than behaviour, at the centre of a security requirement, and the requirement that the most sensitive public workloads keep inference data inside the Union creates the first guaranteed demand a European AI layer has ever had [2]. That is worth having in whatever form it emerges.

For the United Kingdom, none of this applies. Britain has no equivalent framework and no sovereignty tier in public procurement. United Kingdom organisations remain in the position described in Sections 2 to 4, with no equivalent legislation in prospect. The instruments in this paper work without it; Westminster will not apply them on an organisation's behalf.

5.3

Sovereignty washing

While the legislation is argued over, the market for the word "sovereign" has been thriving.

On 15 January 2026 AWS launched its European Sovereign Cloud in Brandenburg, with €7.8 billion of investment behind it and Germany's federal cybersecurity agency present at the launch, describing it as an independent European cloud [33]. It is a serious piece of engineering. It is also a wholly owned subsidiary of Amazon.com, subject to the CLOUD Act, and its legal obligations sit in Seattle whatever the postcode of the data centre.

Microsoft's position deserves particular attention because Microsoft has defined its own limits. In June 2025 Anton Carniaux, director of public and legal affairs at Microsoft France, told a French Senate inquiry under oath that he could not guarantee French citizens' data would never be handed to United States authorities without French consent [34, 35]. Microsoft had earlier told Police Scotland in writing that it could not guarantee data sovereignty for Microsoft 365 [35]. Four months after the Senate testimony, in November 2025, Microsoft announced new sovereign cloud capabilities for Europe [36]. The engineering may well be real. The jurisdiction is unchanged. Only the second of those determines which government can read the data.

5.4

Even the Commission blurs the line. On 20 April 2026 it awarded its €180 million, six-year sovereign cloud framework to four providers: Post Telecom with CleverCloud and OVHcloud, StackIT, Scaleway, and Proximus, whose consortium partners include Mistral and S3NS, a joint venture between Thales and Google Cloud [37, 38]. Google technology can, on this reading, qualify as sovereign given sufficient governance around it. Once the Commission accepts that interpretation, "sovereign" begins to mean "has a French or German operations team".

The three-question test

The label conveys nothing; three questions do.

FIGURE 1. THE THREE-QUESTION TEST

QUESTION	WHAT IT ESTABLISHES
1 Who owns the parent company?	Ultimate beneficial ownership at the top of the group, not the subsidiary that signs.
2 Which jurisdiction can compel it?	The law of the parent's home state: CLOUD Act, FISA 702, National Intelligence Law. Not the postcode of the data centre.
3 Who holds the encryption keys?	Customer-held keys in hardware the provider cannot reach neutralise most of 1 and 2 for data at rest.

A service that fails any question is recorded as resident, never as sovereign.

Who owns the parent company? Ownership determines whose shareholders, whose board and whose national government sit at the top of the chain of command when a conflict arises.

Which jurisdiction can compel it? This is the CLOUD Act question, the Section 702 question and the National Intelligence Law question in one. It is answered by the law of the parent's home state, not by the location of the data centre.

Who holds the encryption keys? If the provider holds them, or holds a copy, or holds the key-management service, the provider can be compelled to use them. If the customer holds them exclusively, in hardware the provider cannot reach, the first two questions lose most of their force for data at rest.

Every current sovereignty claim on the market stands or falls on these three questions, and most fall. Annex A sets the test out as a scoring template for a supplier register. It is the analytical instrument that the rest of this paper applies.

Coordination

EuroStack, CADA, UK Sovereign AI

Builders

Lovable, AIGentHive, the SME base

Workplace

openDesk, Docs, NOVA, Nextcloud

Models

Mistral, Apertus, OpenEuroLLM, Kyutai

Hosting

OVHcloud, Scaleway, IONOS, StackIT

Compute

JUPITER, Alps, Nebius, Schwarz Digits

6.1

6

The European supply base

The persistent myth

The argument most often deployed against a European alternative is that Europe lacks the capacity to provide one. The evidence does not support it. European capacity is smaller, quieter and largely unknown to the people who sign the contracts. This section takes the inventory from the bottom of the stack to the top. Annex D presents it as a table.

FIGURE 2. THE EUROPEAN STACK, WITH NAMED EXAMPLES

LAYER	RUNNING EUROPEAN INSTANCES
Coordination	EuroStack Foundation; Commission sovereign cloud framework; UK Sovereign AI competition; CADA proposal
Builders	Lovable (Stockholm); AIGentHive (UK); the SME base
Workplace	openDesk; Docs (DINUM and ZenDiS); NOVA, 16,000 desks; Schleswig-Holstein; Nextcloud; Lumo
The middle	Where inference runs; identity; logging; version pinning; exit clause
Models	Mistral Medium 3.5; Apertus; OpenEuroLLM; Kyutai; FLUX; PhariaAI; Qwen and DeepSeek weights, self-hosted
Hosting	OVHcloud; Scaleway; IONOS; StackIT; Exoscale; Hetzner; Cloud Temple; Open Telekom Cloud
Compute	JUPITER (Jülich); Alps (Lugano); Nebius (Finland, France); Mistral (Bruyères-le-Châtel); Schwarz Digits (Berlin); BT and Nscale (UK)

Every layer has at least one running European instance. The gap is at the top of the compute layer, frontier training, and in the habit of buying.

6.2

Compute

JUPITER, at the Jülich Supercomputing Centre in North Rhine-Westphalia, became Europe's first exascale computer in November 2025. In June 2026 it ranked fifth in the world on the TOP500 list, with one exaflop at 64-bit precision and more than 40 exaflops at the 8-bit precision that AI training uses, and more than 120 research projects had applied for time on it [39]. It was procured and co-funded through the EuroHPC Joint Undertaking, so the public interest has a say over its use. Europe could not train a frontier-scale model on its own soil in 2024. It can now. Switzerland's Alps system at CSCS in Lugano trained Apertus, the fully open model released by EPFL and ETH Zurich in September 2025 [40], which matters as proof of a complete chain: public compute, public research, public weights, all under European law.

The commercial layer has moved fastest. Nebius, headquartered in Amsterdam and listed on Nasdaq, is tripling its Mäntsälä site in Finland to 75 MW, has announced a 310 MW AI factory at Lappeenranta with first capacity due in 2027 and a 240 MW facility near Lille, and has begun a deployment in Estonia [41, 42]. Finland's attraction is structural: cold air, abundant low-carbon power and a grid with headroom. Mistral has financed its own data centre at Bruyères-le-Châtel near Paris, housing 13,800 Nvidia GB300 GPUs, with roughly €770 million of debt from seven banks led by Bpifrance [43, 44]. Schwarz Digits, the technology arm of the group that owns Lidl and Kaufland, is building an €11 billion data centre near Berlin [45]. OVHcloud, Scaleway, IONOS, Exoscale and StackIT all sell GPU capacity today, and four European providers now sit on the Commission's own sovereign cloud framework [37, 46].

The public programme is slower. InvestAI's €200 billion mobilisation target and the gigafactory programme were announced in February 2025. The formal call for up to seven AI gigafactories opened on 30 July 2026 after two postponements, with about €5 billion from the Union budget, a similar sum from member states and around €20 billion expected from private investors; bids close on 12 November 2026, decisions are expected in early 2027, and the first sites are intended to be operational by mid-2028 [47, 48, 49]. Seventy-six consortia expressed interest before the call opened. Europe's frontier training capacity does not yet exist at the scale envisaged, and the schedule has already slipped once. Planning should assume it will slip again.

One distinction in the gigafactory debate has gone almost unnoticed. It outweighs the headline capacity. A gigafactory is a training story: the money, the grid connections and the consortia are all pointed at building very large models. A hospital trust or a bank with a sensitive workload does not need to train a model. It needs to run one, cheaply and reliably, under European law. Training capacity and serving capacity are different products with different economics. The question to ask of any gigafactory is therefore not how many GPUs it houses but whether inference capacity will be reserved, priced and contractable under Union law, and what happens to the terms when the founding consortia want the machines for their own models [2]. Sovereignty is built at the contract, not at the ribbon-cutting.

Britain is moving on a smaller scale. In April 2026 BT, Nscale and Nvidia announced up to 14 MW of AI data centre capacity across three BT sites, described explicitly as compute under United Kingdom control for United Kingdom organisations [50]. It is modest against the hyperscalers. It is also the first time a British telecommunications incumbent has put its name to sovereign AI compute rather than reselling someone else's.

The caveat is the silicon. Every machine above runs Nvidia hardware. Europe's own processor efforts are years from frontier training parts. But the three-question test asks who owns the operator, which court can compel it and who holds the keys. An Nvidia GPU in a Finnish hall run by a Dutch company under Union law passes all three. Hardware origin is a supply-chain risk to manage over a decade; it is not the sovereignty problem of the coming quarter.

6.3

The hosting market

The market-share figures are the sobering part of the inventory. The Commission's own impact assessment puts three United States firms at more than 70 per cent of the European cloud market [32]; Synergy Research has European providers at around 15 per cent of their home market for several consecutive years [51]. That is the dependency in one statistic. It is also evidence that capacity exists and is under-used.

Price is not the barrier. Like-for-like instances at Scaleway, IONOS or Hetzner generally come in below AWS Frankfurt, and most European providers do not charge for egress [52]. OVHcloud holds France's SecNumCloud qualification for its hosted private cloud; IONOS holds Germany's BSI C5 attestation; StackIT is Schwarz Group's enterprise cloud; Cloud Temple and Open Telekom Cloud serve regulated and public-sector customers [46].

The barrier is technical. Choosing a European provider takes a week. Moving deployment pipelines, identity, secrets and monitoring takes a year. Sovereignty migrations fail at the platform layer, not at procurement, and any plan that budgets only for the procurement will fail on schedule. The identity layer deserves specific mention: it is the one system every other system trusts, and the enterprise market for it belongs almost entirely to two American companies, which is why an organisation can move its systems to a European cloud and still hand the keys to the old landlord [5].

6.4

Models

Here the European habit is at its most visible. American frontier laboratories closed their weights as they scaled. European laboratories kept publishing.

Mistral, Europe's most valuable AI company, released Medium 3.5 in April 2026 under a modified MIT licence with downloadable weights, positioned for agentic and coding work [53]. The company reached a €11.7 billion valuation after its September 2025 Series C led by ASML and was reported in June 2026 to be in talks to raise a further €3 billion at around €20 billion [54]. When Microsoft expanded its partnership with Mistral in July 2026, the agreement placed those open weights into Azure's sovereign and fully disconnected environments and committed Microsoft to buying capacity from Mistral's own European GPU estate [55]. The direction of travel deserves careful reading: the hyperscaler is buying European compute and distributing a European open model. Two caveats stand. Mistral trains on American silicon because no European alternative exists at that scale, and its investors include Nvidia, Andreessen Horowitz and DST Global alongside ASML and General Catalyst. It is a European champion running American hardware with partly American money, and its deepening entanglement with Microsoft means its direction is not entirely Europe's to set. The only guarantee in that relationship that cannot be revised in a boardroom is the open-weight commitment: a laboratory can be bought, but weights already downloaded cannot be repossessed, repriced or switched off [2].

Apertus is the fully open option. Weights, training data, methods and checkpoints are all published under Apache 2.0; it was trained across more than a thousand languages and designed from the outset to meet the AI Act's transparency duties [40]. It does not claim to be a frontier model. It is something rarer: a complete, auditable European AI supply chain that exists today.

OpenEuroLLM is the consortium behind the next wave: Charles University, Tübingen, Helsinki, Oslo and Eindhoven among the universities, the HPC centres CINECA, CSC and SURF, and companies including Aleph Alpha, Silo AI and LightOn, co-funded through the Digital Europe Programme to build open models across the Union's official languages [56]. Around it sit Spain's Salamandra family, France's Pleias, trained on openly licensed text, and Germany's Teuken and Soofi projects [57]. Multilingualism is not a feature these teams added; it is the reason they exist, and it is the one dimension on which European models routinely outperform American ones.

The creativity is most visible in the smaller laboratories. Kyutai in Paris, funded by Iliad, CMA CGM and Schmidt Sciences, has shipped an entire open voice stack in under two years: Moshi for real-time conversation, Hibiki for simultaneous speech translation, streaming speech-to-text and text-to-speech, and in 2026 Pocket TTS, a voice-cloning text-to-speech model that runs on a CPU and covers six European languages [58]. Black Forest Labs in Freiburg, whose FLUX image models are the open-weight standard in their field, announced FLUX 3 in July 2026 as a multimodal model across image, video, audio and action, with open-weight versions promised later in the year; the company has raised more than €390 million at a valuation of about €2.8 billion [59].

Aleph Alpha in Heidelberg took a different route. In April 2026 it merged into Cohere, with Schwarz Digits committing around €515 million and the German government signalling that it would be an anchor customer for sovereign AI procurement [45]. Cohere is Canadian, so the parent sits outside the Union in an adequacy jurisdiction rather than a CLOUD Act one, and PhariaAI, the deployment layer Aleph Alpha built, is designed to run on the customer's own infrastructure. It is a hybrid answer and should be described as one.

Finally, the least preferable option, which Sections 3 and 4 make rational: Chinese open weights on European infrastructure. A DeepSeek or Qwen model that an organisation has downloaded, security-reviewed and fine-tuned, running on OVHcloud or Scaleway hardware, cannot be switched off from Hangzhou or Beijing. Nobody can disable it, because the organisation holds the weights. Nobody can read the prompts, because they never leave the organisation's infrastructure. The organisation decides how the model is steered, because the fine-tuning happens in its environment. The residual training-baked refusals noted in Section 3.3 are what the organisation's own evaluation exists to detect. The responsibility moves to the organisation: safety testing, red-teaming and patching become its obligations. For a capable IT function that is a fair exchange, and European public bodies are already making it.

6.5

The unglamorous middle

Between the model and the user sits the layer where sovereignty is decided in practice. It comes down to five decisions: where the inference runs; what identity system stands in front of it; what is logged and where the logs live; which model version is pinned and who decides when it changes; and what the exit clause says [2]. Leaderboards measure none of these. Auditors will measure all of them.

The version question is the one most contracts miss. A hosted model can be retired or repriced on the vendor's schedule, and an exit clause that ignores versioning covers a product that may not exist when the clause is needed. This is also where the good news lives, because the middle layer needs no gigafactory and no ten-billion-euro tender. Serving, orchestration, logging and identity are skills that competent infrastructure teams already have. What is scarce is the decision to treat the work as sovereignty rather than plumbing.

6.6

The workplace

Models and GPUs are invisible to most staff. The office suite is not, and this is where European public bodies have done the least glamorous and most consequential work.

6.7

Schleswig-Holstein has moved nearly 80 per cent of its state government workstations to LibreOffice, is completing the remainder in 2026, and puts the licence saving at more than €15 million a year; the stated reason is the CLOUD Act [60]. Germany's Centre for Digital Sovereignty in Public Administration, ZenDiS, publishes openDesk, an open-source workplace assembled from Nextcloud, Open-Xchange, Collabora and Univention's Nubus identity layer [61]. The Municipal Data Processing Centre Niederrhein, which serves 46 municipalities, completed the test environment for its NOVA workplace on the openDesk standard in August 2026 and is moving to production for more than 16,000 workstations [62].

France's DINUM and Germany's ZenDiS have gone further and built together. Docs, a real-time collaborative editor released under the MIT licence, was announced in March 2025, is in use by French civil servants through the ProConnect identity service, and is being onboarded in the Netherlands [63]. Two national administrations co-developing a replacement for Google Docs and publishing the code is a thing that did not happen before 2025. On the consumer side, Proton in Geneva launched Lumo, an AI assistant with encrypted conversations and no logs, in July 2025 and shipped a second major version within the year [64].

The builders

The strongest signal is the number of people who started without waiting for the policy debate.

Lovable, in Stockholm, lets non-programmers build software by describing it. Founded in November 2024, it raised around €345 million in August 2026 at a valuation near €11.4 billion, with the EQT-managed Scaleup Europe Fund co-leading, and is tracking toward roughly €500 million of annual recurring revenue [65]. It runs on other companies' models, which is a sovereignty gap it should be pressed on. It is also proof that a European team can build a category-defining AI product from a standing start and keep it headquartered in Europe.

A smaller example illustrates the pattern more directly. In the discussion under the first article in this series, Brian Stack, founder of Digital Twin AI, described watching British industrial teams assume a supplier would always be there until a missing part or platform turned a delay into a shutdown, and drew the line straight to AI: if the model, routing layer or hosting terms sit under someone else's jurisdiction, control is conditional. His response was to build AIgentHive, a platform giving small and growing United Kingdom businesses specialist AI agents with data and decision boundaries kept in the United Kingdom and no third-party large language model in the loop [66]. This paper has not audited the platform and does not endorse it. It records the instinct: a founder outside the AI establishment, in a sector that understands supply-chain failure, concluded that the sovereign version was worth building and built it. Multiplied across Europe's small and medium-sized enterprise base, that instinct is the real market.

6.8

That market is being formed deliberately. The EuroStack Foundation, founded in 2026 with Nextcloud's Frank Karlitschek, Proton's Andy Yen, IONOS's Achim Weiss and Ecosia among its founders and more than 300 chief executives behind it, has three pillars that read like a procurement manual: buy European, sell European, fund European [67]. The Commission's sovereign cloud framework and the CADA proposal are the public-sector end of the same movement.

Britain has, on this occasion, matched the direction. On 31 August 2026 the government launched a £100 million competition through its Sovereign AI unit in which the state acts as first customer rather than grant-giver: British AI companies are paid to build demonstrators for NHS productivity, compute efficiency, defence integration and agent security, retain their intellectual property, and are not screened out by turnover thresholds [68]. Procurement as market-making is precisely the mechanism this paper recommends.

The pattern

Taken together, the inventory forms a pattern that is distinctively European and worth naming. Many mid-sized builders rather than three giants. Open by default, because the research culture and the public funding both push that way. Multilingual by necessity. Public bodies acting as first customers. Energy treated as a design constraint rather than an afterthought, which is why the GPUs are going to Finland. None of it is finished. All of it is real.

7

standing objections, each answered

7

Objections and constraints

Seven objections recur in board rooms and ministries. Each has a direct answer, and none is a reason to wait.

7.1

None of it is frontier

Correct, for now, if frontier means the single largest model on a leaderboard. But the leaderboard is not where most organisations operate. Most enterprise AI work is document handling, classification, retrieval, translation, summarisation, voice and code assistance, and European open models are already competitive across that range and frequently superior in European languages. Most sensitive workloads do not need the frontier; they need a version that does not change underneath them. Frontier capability is also arriving as open weights from several directions at once, and Europe's advantage is that its infrastructure can run any of them.

7.2

The compute gap is real

It is, at the training end. Europe's frontier training capacity does not yet exist at the scale envisaged; the gigafactory schedule has slipped once and should be assumed to slip again; and the public share of the programme has been trimmed from the original ambition. Two things follow. First, the serving half of the market, which is what most organisations need, is a different product with different economics and is already available from European operators. Second, the training gap is a reason to write inference contractability into the gigafactory tenders, not a reason to wait for the machines.

7.3

It still runs on Nvidia

Yes. Hardware supply is a risk to hold on the register and to hedge over a decade. It does not decide who can read an organisation's data or switch off its service next quarter. Jurisdiction does. A supply-chain risk and a sovereignty risk are different kinds of risk and should not be conflated to defer action on the second because the first is unsolved.

7.4

Europe is fragmented

Also correct, and also the point. Three American hyperscalers is a concentration risk; thirty European providers is a market. The cost of fragmentation is real: SecNumCloud in France, C5 in Germany, the contested EUCS scheme and CADA's four proposed levels overlap without aligning; a provider operating across Europe must satisfy several regimes at once; and buyers need legal advice to compare them [31, 69]. Until CADA or its successor collapses this into one framework, the complexity favours the incumbents, because AWS only has to tell one story. Annex B gives the eight procurement questions that manage the cost in the meantime. What EuroStack, openDesk and CADA are each attempting is to make the fragments interoperate without collapsing them into a single new dependency.

7.5

Gaia-X shows frameworks fail

Gaia-X, launched in 2019 to federate European cloud through a common framework, has become a standards body that publishes labels. It cannot host a single container [52]. The lesson is not that European coordination fails but that frameworks do not run workloads. Anyone proposing another framework should be asked where the workloads would run. The initiatives that have moved, JUPITER, the sovereign cloud tender, openDesk, Docs, are the ones that produced a running artefact rather than a label.

7.6

It is subsidised

Some of it is. JUPITER, OpenEuroLLM and the gigafactory programme are public money. So were the internet, satellite positioning and the American semiconductor industry. What matters is whether the public money produces open assets that anyone can use, and in Europe's case it demonstrably does: the weights, the data, the code and the documents are published.

7.7

The sovereign option carries a premium

It does, in some configurations, and the premium should be decided in advance as any resilience measure is, and treated as the price of keeping an exit open [2]. It is also smaller than assumed at the infrastructure layer, where European providers generally undercut AWS Frankfurt and do not charge for egress [52]. The real cost sits in the platform migration, which is why Section 8 recommends budgeting for it rather than for the procurement.

8

questions for every procurement

8

Recommendations

The recommendations are separated by audience. Those for organisations and procurement can be acted on this quarter without any change in law. Those for governments concern the instruments now in the legislative and budgetary process. Those for researchers identify the questions this paper could not settle.

8.1

For organisations

Map every AI dependency by jurisdiction. Model, host, silicon, payment rail, subcontractor. Record which foreign state holds a lever over each layer, in the risk register, by name. Include the AI features that arrived inside products the organisation already owns.

Apply the three-question test to every AI and cloud supplier: who owns the parent; which jurisdiction can compel it; who holds the keys. Record the answers against each supplier. Marketing claims of sovereignty are not evidence; the Carniaux testimony is what evidence looks like [34].

Classify workloads by sovereignty requirement rather than by habit. CADA's four levels are a usable template before they become law: public, non-personal workloads can stay where they are; personal and commercially sensitive data moves to Union-owned providers; the most sensitive workloads are self-hosted or placed with a level 3 equivalent. Annex C sets this out.

Start at the model layer now. Select one sensitive use case and deploy an open-weight model, whether Mistral, Apertus, Qwen or DeepSeek, on European GPU infrastructure. The first deployment teaches more than a year of strategy papers and builds the internal skill every later decision depends on.

Keep sensitive material out of consumer tiers. Litigation holds override the delete button, and the prompt box is now a control surface. Zero-retention enterprise terms were the only carve-out that held in 2025.

Contract for exit: portable fine-tunes; exportable data; notice periods; a tested migration path per supplier; and an exit clause that is version-aware. Nayara's experience shows that terms of service yield to politics; the exit plan is the control that survives.

Budget for the platform migration, not the procurement. Choosing a provider takes a week. Rebuilding deployment pipelines, identity and observability is the substantive work, and starving that phase is how sovereignty programmes fail at the first budget review. Split the estate by data classification so that teams are not blocked waiting for a single large move.

Run the dial test quarterly. Probe every model the organisation relies on with the contested questions in its sector and log the answers. Drift is the early warning of steering. This matters more, not less, once the weights are on the organisation's own hardware, because from that point the steering risk is its own.

Treat American and Chinese hosted services as equally foreign. Neither jurisdiction's law protects the organisation, and neither's politics is its own.

8.2

For procurement

Put three columns into every scoring sheet for anything that touches sensitive data: who hosts the model; where the inference runs; who provides the identity [2]. A question asked in fifty tenders becomes a supplier roadmap.

Put eight questions in every request for proposals: egress and switching charges under the Data Act; the full sub-processor list; key custody and who can access keys; GPU quotas and their guarantees; exit assistance and its cost; recovery objectives for managed databases; the upgrade and deprecation schedule, including model-version pinning; and incident transparency obligations. A provider that cannot answer clearly has answered. Annex B sets these out with the reason each is asked.

Ask for open-weight options in every AI tender. Most tenders will receive none at first. Ask anyway. Suppliers read tenders the way the rest of us read the news.

Write the organisation's own version of the Sovereign AI competition. Even a small organisation can take one problem, invite three European or British builders to solve it, and pay for the demonstrator. The state has shown the model; nothing prevents a company or a council from applying it at its own scale.

8.3

For governments and public bodies

Treat the serving layer as strategic. The gigafactory programme is a training story. Write reserved, priced and contractable inference capacity under Union law into the tenders, and specify what happens to the terms when founding consortia want the machines for their own models. Without that clause the programme builds capacity that public buyers cannot rely on.

Hold the ownership and control levels of CADA through the legislative process. Levels 3 and 4 are the first European legal expression of the difference between residency and sovereignty. The lobbying against them will be intense and will be conducted in the language of choice and competition. The answer is that a market with thirty providers is more competitive than one with three.

Use procurement as market-making. The Commission's €180 million sovereign cloud framework and the United Kingdom's £100 million Sovereign AI competition both place the state in the role of first customer rather than grant-giver, with intellectual property retained by the supplier and turnover thresholds removed [37, 68]. The model scales and should be replicated at member-state, devolved and municipal level.

Consolidate the certification regimes. SecNumCloud, C5, EUCS and the CADA levels should converge into one framework with mutual recognition. Until they do, fragmentation is a subsidy to the incumbents.

Take the frontier-laboratory question seriously. Europe has solved this problem once, at considerable cost: Airbus exists because governments decided aerospace was strategic and then paid for that decision for decades. No equivalent decision has been taken about a frontier laboratory, and until one is, the only guarantee that cannot be revised in a boardroom is open weights [2]. Public funding conditions should require open release of weights, data and evaluation for anything trained on public compute.

For the United Kingdom specifically: adopt a sovereignty tier in public procurement equivalent to CADA's levels, and extend the Sovereign AI procurement model from demonstrators to production contracts. The dependency arithmetic is the same as the Union's, with fewer instruments to address it.

8.4

For researchers

Four questions this paper could not settle would benefit from independent work. First, the empirical performance gap between European open-weight models and American closed models on the enterprise workloads that dominate actual use, measured in European languages rather than English benchmarks. Second, the residual training-time alignment in Chinese open-weight models after self-hosting and fine-tuning, measured systematically rather than anecdotally. Third, the total cost of a sovereign configuration against a hyperscaler configuration over a five-year horizon, including migration, with the platform-layer costs made explicit. Fourth, the effect of procurement scoring on supplier behaviour: whether the three-column and eight-question approach measurably changes what is offered, and how quickly.

European law, European infrastructure, open weights, your own keys.

9

Conclusion

The comfortable European story said that American AI was the safe default and Chinese AI the risky choice. The record of the last two years does not support it. Washington has demonstrated the switch against a rival, a neutral party and its own laboratory, written extraterritorial reach into statute and discovery practice, and put government-defined truth into procurement. Beijing has declared the reach and the censorship in law while distributing the one artefact that neutralises the switch for everyone.

European regulation governs behaviour and, if the Cloud and AI Development Act survives the legislature intact, will begin to govern ownership in public procurement. It does not, and cannot, change who owns the models Europe rents. What can is a configuration: European law, European infrastructure, open weights, and encryption keys the organisation holds itself.

Every component of that configuration exists. An exascale machine in Jülich. GPU halls in Finland and France run by European companies. Open weights from Paris, Zurich and Freiburg and a pan-European consortium. A document editor built by two governments and deployed on sixteen thousand desks. A Swedish company doubling every eight months, and a British founder who built the sovereign version because nobody else had. It is a supply base, not a brochure, and it is open for business.

Sovereignty cannot be bought from the companies an organisation is trying to reduce its dependence on. It is built, at the contract and in the configuration, by people who stop asking which superpower to trust and start asking which dependencies they could survive losing. Nothing in this paper's organisational recommendations requires a regulator to move first, a gigafactory to be built, or permission from anyone.

Your data. Your rules. Your continent.

Score the evidence, not the marketing.

ANNEX A

The three-question test as a scoring template

The test is applied to every supplier whose service touches organisational data, and recorded in the supplier register. Each question is scored on the evidence the supplier provides, not on its marketing.

TABLE. THREE-QUESTION TEST

QUESTION	WHAT TO ESTABLISH	PASSES WHEN	FAILS WHEN
Who owns the parent company?	Ultimate beneficial ownership and the jurisdiction of incorporation of the top of the group	The ultimate parent is incorporated and controlled in the organisation's own jurisdiction or a jurisdiction with equivalent protections	The ultimate parent is subject to a foreign state's extraterritorial disclosure or sanctions law, regardless of subsidiary structure
Which jurisdiction can compel it?	The laws that can require the provider to disclose data, suspend service or alter behaviour, and the process and redress available	Compulsion requires a domestic legal process the organisation can see and contest	Compulsion can occur under foreign law without notice, process or standing for the organisation (CLOUD Act, FISA 702, National Intelligence Law)
Who holds the encryption keys?	Where keys are generated, stored and used, and who can access the key-management service	The organisation holds keys exclusively in hardware or a service the provider cannot reach	The provider holds keys, holds a copy, or operates the key-management service

A supplier that fails any question is recorded as "resident" at best, never as "sovereign". A supplier that passes all three is a candidate for the most sensitive workloads.

ANNEX B

Eight procurement questions

Each question is asked in the request for proposals and the answer is scored. The reason each is asked is given so that evaluators can judge the quality of an answer, not merely its presence.

TABLE. PROCUREMENT QUESTIONS

QUESTION	WHY IT IS ASKED
----------	-----------------

8

procurement questions, with the reason each is asked

Egress and switching charges under the Data Act	The Data Act phases out switching charges; a provider that still levies them is signalling how it views exit
Full sub-processor list, with jurisdictions	The three-question test must be applied to every party in the chain, not only the prime contractor
Key custody: where keys are held and who can access them	The third question of the test; the single most important control for data at rest
GPU quotas and their guarantees	Sovereign capacity that cannot be reserved is not capacity the organisation can plan on
Exit assistance and its cost	The exit plan is the control that survives a political decision; its cost must be known before it is needed
Recovery objectives for managed databases	Recovery time and recovery point objectives determine whether a migration under pressure is survivable
Upgrade and deprecation schedule, including model-version pinning	A hosted model can be retired or repriced on the vendor's schedule; an exit clause that ignores versioning covers a product that may not exist
Incident transparency obligations	What the provider must disclose, to whom and within what time, when it is compelled or compromised

Four workload classes against the proposed CADA levels.

ANNEX C

Workload classification against the proposed CADA levels

The classification uses the four levels of the Commission's proposed Cloud and AI Development Act as a template. It can be applied before the Act becomes law and adjusted if the levels change in the legislative process.

TABLE. WORKLOAD CLASSIFICATION

WORKLOAD CLASS	EXAMPLES	MINIMUM PLACEMENT	PROPOSED CADA EQUIVALENT
Public, non-personal	Published documents, marketing content, open data, public code	Any provider that passes basic residency and contractual requirements	Level 1
Internal, personal or commercially sensitive	Staff records, customer data, contracts, financial systems, most enterprise AI assistants	Union-owned and controlled provider; or Union-resident provider with customer-held keys, recorded as resident not sovereign	Level 2 to 3
Sensitive: regulated,	Health records, legal privilege, security	Self-hosted, or Union-owned and controlled	Level 3

OPENING
THE RECORD
THE COMPARISON
THE EUROPEAN POSITION
THE SUPPLY BASE
OBJECTIONS AND RECOMMENDATIONS
INSTRUMENTS

27

running instances in the inventory

safety-critical or reputationally critical	operations, board and executive material, source code of core systems	provider with customer-held keys	
Most sensitive: national security, defence, law enforcement, public order	As defined by the relevant public authority	Self-hosted on Union-owned infrastructure; open-weight models only; no third-country control of any layer	Level 4

ANNEX D

Inventory of the European supply base

Illustrative, not exhaustive. Status as at the dates in the referenced sources.

TABLE. INVENTORY

LAYER	NAME	LOCATION	WHAT IT IS	STATUS AND LICENCE	SOURCE
Compute	JUPITER	Jülich, Germany	Europe's first exascale supercomputer, EuroHPC	Operational; fifth on TOP500, June 2026	[39]
Compute	Alps	Lugano, Switzerland	CSCS supercomputer that trained Apertus	Operational	[40]
Compute	Nebius	Amsterdam HQ; Finland, France, Estonia	Nasdaq-listed AI cloud operator	75 MW Mäntsälä; 310 MW Lappeenranta due 2027; 240 MW Lille	[41, 42]
Compute	Mistral data centre	Bruyères-le-Châtel, France	13,800 Nvidia GB300 GPUs, debt-financed	Under construction, March 2026	[44]
Compute	Schwarz Digits	Near Berlin, Germany	€11 billion data centre	Announced	[45]
Compute	BT, Nscale, Nvidia	United Kingdom	Up to 14 MW sovereign AI capacity	Announced April 2026	[50]
Hosting	OVHcloud	Roubaix, France	Cloud provider; SecNumCloud	Operational; Commission	[37, 46]

OPENING
THE RECORD
THE COMPARISON
THE EUROPEAN POSITION
THE SUPPLY BASE
OBJECTIONS AND RECOMMENDATIONS
INSTRUMENTS

			qualified	framework	
Hosting	Scaleway	Paris, France	Iliad group cloud; GPU capacity	Operational; Commission framework	[37]
Hosting	StackIT	Neckarsulm, Germany	Schwarz Group enterprise cloud	Operational; Commission framework	[37]
Hosting	IONOS	Karlsruhe, Germany	Cloud provider; BSI C5 attested	Operational	[46]
Hosting	Exoscale	Lausanne, Switzerland	Cloud provider with GPU capacity	Operational	[46]
Model	Mistral Medium 3.5	Paris, France	Frontier-class multimodal model	Open weights, modified MIT, April 2026	[53]
Model	Apertus	Zurich and Lausanne, Switzerland	Fully open multilingual model	Apache 2.0, September 2025	[40]
Model	OpenEuroLLM	Pan-European consortium	Open models across Union languages	In development; Digital Europe Programme	[56]
LAYER	NAME	LOCATION	WHAT IT IS	STATUS AND LICENCE	SOURCE
Model	Kyutai voice stack	Paris, France	Moshi, Hibiki, STT, TTS, Pocket TTS	Open releases 2025 to 2026	[58]
Model	FLUX / FLUX 3	Freiburg, Germany	Image and multimodal models	FLUX open weights; FLUX 3 open weights promised 2026	[59]
Model	Aleph Alpha / PhariaAI	Heidelberg, Germany	Deployment layer, merged into Cohere	Hybrid: Canadian parent, on-premises deployment	[45]
Workplace	openDesk	ZenDiS, Germany	Open-source public-sector workplace	Published; in deployment	[61]
Workplace	NOVA (KRZN)	Lower Rhine, Germany	openDesk-standard workplace, 46 municipalities	Test complete August 2026; 16,000 desks to production	[62]
Workplace	Schleswig-Holstein	Kiel, Germany	State-wide LibreOffice migration	~80 per cent complete; €15 million a year	[60]

saved					
Workplace	Docs	DINUM and ZenDiS	Collaborative editor	MIT; in use in France and the Netherlands	[63]
Workplace	Lumo	Proton, Geneva	Private AI assistant	Live; version 2.0	[64]
Builder	Lovable	Stockholm, Sweden	Natural-language software builder	~€11.4 billion valuation, August 2026	[65]
Builder	AlgentHive	United Kingdom	SME agent platform, UK data boundaries	Live; not audited by this paper	[66]
Coordination	EuroStack Foundation	Berlin	Industry sovereignty initiative	Founded 2026; 300+ chief executives	[67]
Coordination	Sovereign AI (UK)	United Kingdom	£100 million procurement competition	Launched 31 August 2026	[68]
Coordination	Commission sovereign cloud framework	Brussels	€180 million, six-year framework	Awarded 20 April 2026	[37]

Terms as used in this paper.

ANNEX E

Glossary

TABLE. GLOSSARY

TERM	MEANING IN THIS PAPER
AI Act	Regulation (EU) 2024/1689, the European Union's horizontal regulation of artificial intelligence, as amended by the Digital Omnibus
CADA	The Commission's proposed Cloud and AI Development Act, published 3 June 2026
CLOUD Act	United States Clarifying Lawful Overseas Use of Data Act 2018, requiring United States providers to produce data in their possession, custody or control wherever stored
Control	The operational form of sovereignty: the organisation's ability to keep a service running, private and stable regardless of any other party's decisions
Digital Omnibus	The 2026 European Union package that, among other changes, deferred the AI Act's high-risk obligations

FISA Section 702	The provision of the United States Foreign Intelligence Surveillance Act permitting collection of non-Americans' communications abroad through compelled assistance from United States providers
Inference, serving	Running a trained model to answer requests, as distinct from training it
Open weights	A model whose trained parameters are published under a licence permitting download, self-hosting and modification
Residency	Geography: where data and computation physically sit
Sovereignty	Jurisdiction: which state's courts and executive can compel the operator of a service
Sovereignty washing	Marketing a service as sovereign on the basis of residency, governance layers or local staff while the parent and jurisdiction are foreign
Steering	Changing a model's outputs by prompt, fine-tuning, filtering or policy, whether by the vendor, a third party or a state
Three- question test	Who owns the parent, which jurisdiction can compel it, who holds the keys

How claims were verified, and what was excluded.

ANNEX F

Method and limitations

Sources. Every factual claim in this paper is referenced to a dated source. Primary sources (legislation, court orders, official press releases, company statements and model cards) were preferred. Where a secondary source is used, it is a specialist legal, technical or trade publication. Vendor blogs are used only for observations about pricing and migration and are labelled as such in the reference list.

Verification. All claims carried over from the three-part series were re-checked against their sources in the first week of September 2026. Three claims in the second part had been overtaken by events between drafting and publication and were corrected: the status of the Digital Omnibus, the opening of the gigafactory call, and the publication date and content of the Cloud and AI Development Act proposal. The description of the CADA levels in this paper is taken from legal analyses of the proposal text rather than from press coverage.

Currency. Figures first reported in United States dollars have been converted to euros at approximately 1.16 dollars to the euro and rounded, except for United States government contract values, which are given in dollars because no euro figure exists.

Exclusions. The paper does not assess model quality, safety or bias; it is concerned with jurisdiction and control. It does not assess the internal security of any named product. The inventory in Section 6 and Annex D is illustrative and omits many providers; omission is not a judgement. The paper does not audit any company's sovereignty claims beyond the public record, and says so where a claim is reported rather than verified.

Judgement. Where the author's judgement is offered rather than established fact, principally in Section 4.4 and Section 8, the text says so.

Series. The paper consolidates three articles published by AltLibre: "Your AI Answers to a Foreign Government. Part 1: China vs America" (25 August 2026), "Part 2: The European Answer" and "Part 3: What Europe Is Already Building" (September 2026), together with "What a Credible European AI Layer Would Take" (2 September 2026) [1, 2].

Licence. This paper is published under the Creative Commons Attribution-ShareAlike 4.0 International licence. It may be shared, republished and translated with credit to Adrian Hollister and a link to the original, and anything built on it must be released under the same terms.

69

dated sources, re-verified
September 2026

References

1. Hollister, A., "Your AI Answers to a Foreign Government. Part 1: China vs America", AltLibre, 25 August 2026. <https://altlibre.com/blog/us-vs-chinese-ai-foreign-government-risk>
2. Hollister, A., "What a Credible European AI Layer Would Take", AltLibre, 2 September 2026. <https://altlibre.com/blog/credible-european-ai-layer>
3. Hollister, A., "Sovereignty, Sold Separately", AltLibre, 27 July 2026. <https://altlibre.com/blog>
4. Hollister, A., "Airbus Just Pulled Its Crown Jewels Out of the American Cloud", AltLibre, 21 July 2026. <https://altlibre.com/blog>
5. Hollister, A., "The Last American Anchor: The Login", AltLibre, 4 August 2026. <https://altlibre.com/blog>
6. Congressional Research Service, "Federal Government and Anthropic: Considerations for AI Innovation and Competition", IF13217, 31 July 2026.
7. Mayer Brown, "Pentagon Designates Anthropic a Supply Chain Risk: What Government Contractors Need to Know", 2 March 2026.
8. Lawfare, "Pentagon's Anthropic Designation Won't Survive First Contact with Legal System", 2 March 2026.

9. Rest of World, "OpenAI cuts its last and most important link to China", 27 June 2024.
10. ThinkChina, "OpenAI's block in Hong Kong and mainland China deepens China-US divide", 10 July 2024.
11. BankInfoSecurity, "Nayara Energy vs. Microsoft and Compliance-Driven Lockouts", 13 August 2025.
12. Office of Senator Tom Cotton, "Cotton Introduces Bill to Prevent Diversion of Advanced Chips to America's Adversaries and Protect US Product Integrity", 8 May 2025.
13. Legis1, "In a Fractured Congress, a Chip Security Bill Finds Rare Common Ground", 30 March 2026.
14. Terms.law, "OpenAI v New York Times: When Your ChatGPT Logs Become Evidence", 12 November 2025.
15. xAI, statement on the Grok system prompt incident, via X, 16 May 2025.
16. The Decoder, "xAI reverses Grok's prompt changes after racist responses", 21 May 2025.
17. DefenseScoop, "Pentagon awards mega contracts to Musk-owned company, other firms for new frontier AI projects", 14 July 2025.
18. Breaking Defense, "Anthropic, Google and xAI win \$200M each from Pentagon AI chief for agentic AI", 14 July 2025.
19. The White House, Executive Order 14319, "Preventing Woke AI in the Federal Government", 23 July 2025.
20. Paul Hastings, "President Trump Signs Three Executive Orders Relating to Artificial Intelligence", 8 August 2025.
21. Time, "OpenAI: Russia, China, Israel Use It for Influence Campaigns", 30 May 2024.
22. SC World, "OpenAI report reveals threat actors using ChatGPT in influence operations", 31 May 2024.
23. National Intelligence Law of the People's Republic of China, 2017, Articles 7 and 14.
24. AI-Regulation.com, "DeepSeek One Year Later: Regulatory Storm, Global Surge", 28 January 2026.
25. Cooley, "EU AI Act: Transparency Obligations Take Effect 2 August 2026", 3 August 2026. <https://www.cooley.com/news/insight/2026/2026-08-03-eu-ai-act-transparency-obligations-take-effect-2-august-2026>

26. Morgan Lewis, "EU AI Act's Transparency Rules: What Went Into Effect on 2 August", 12 August 2026. <https://www.morganlewis.com/blogs/sourcingatmorganlewis/2026/08/eu-ai-acts-transparency-rules-what-went-into-effect-on-2-august>
27. Council of the European Union, "Artificial Intelligence: Council gives final green light to simplify and streamline rules", press release, 29 June 2026. <https://www.consilium.europa.eu/en/press/press-releases/2026/06/29/artificial-intelligence-council-gives-final-green-light-to-simplify-and-streamline-rules/>
28. Cloud Security Alliance, "EU AI Act's High-Risk Deadline: Deferred, Not Cancelled", 2026. <https://labs.cloudsecurityalliance.org/research/csa-research-note-eu-ai-act-high-risk-deadline-omnibus-20260/>
29. European Commission, "Commission proposes tech sovereignty package to strengthen Europe's digital autonomy and resilience", press release IP/26/1187, 3 June 2026. https://ec.europa.eu/commission/presscorner/detail/en/ip_26_1187
30. European Commission, "Proposal for the Cloud and AI Development Act (CADA)", 3 June 2026. <https://digital-strategy.ec.europa.eu/en/library/proposal-cloud-and-ai-development-act-cada>
31. Covington & Burling, Inside Global Tech, "The EU Cloud and AI Development Act in Depth", 11 June 2026. <https://www.insideglobaltech.com/2026/06/11/the-eu-cloud-and-ai-development-act-in-depth/>
32. Lawfare, "The EU Cloud and AI Development Act", 2026. <https://www.lawfaremedia.org/article/the-eu-cloud-and-ai-development-act>
33. Amazon, "AWS Launches AWS European Sovereign Cloud and Announces Expansion Across Europe", 15 January 2026. <https://press.aboutamazon.com/aws/2026/1/aws-launches-aws-european-sovereign-cloud-and-announces-expansion-across-europe>
34. Forbes, "Microsoft Can't Keep EU Data Safe From US Authorities", 22 July 2025. <https://www.forbes.com/sites/emmawoollacott/2025/07/22/microsoft-cant-keep-eu-data-safe-from-us-authorities/>
35. The Register, "Microsoft admits it 'cannot guarantee' data sovereignty", 25 July 2025. <https://www.theregister.com/off-prem/2025/07/25/microsoft-exec-admits-it-cannot-guarantee-data-sovereignty/>
36. Microsoft Azure blog, "Microsoft strengthens sovereign cloud capabilities with new services", 4 November 2025. <https://azure.microsoft.com/en-us/blog/microsoft-strengthens-sovereign-cloud-capabilities-with-new-services/>
37. European Commission, "Commission awards €180 million tender for sovereign cloud to four European providers", press release IP/26/833, 20 April 2026. https://ec.europa.eu/commission/presscorner/detail/en/ip_26_833

38. Help Net Security, "EU pushes for stronger cloud sovereignty, awards €180 million to four providers", 20 April 2026. <https://www.helpnetsecurity.com/2026/04/20/eu-sovereign-cloud-tender-180-million-eu/>
39. HPCwire, "JUPITER Ranked 5th on TOP500, Powers More Than 120 Research Projects", 24 June 2026. <https://www.hpcwire.com/off-the-wire/jupiter-ranked-5th-on-top500-powers-more-than-120-research-projects/>
40. ETH Zurich, "Apertus: a fully open, transparent, multilingual language model", press release, 2 September 2025. <https://ethz.ch/en/news-and-events/eth-news/news/2025/09/press-release-apertus-a-fully-open-transparent-multilingual-language-model.html>
41. Nebius, "Nebius to construct 310 MW AI factory in Finland", 31 March 2026. <https://nebius.com/newsroom/nebius-to-construct-310-mw-ai-factory-in-finland>
42. Data Center Dynamics, "Nebius expands European presence, announces deployment in Estonia and second data center in Mäntsälä, Finland", August 2026. <https://www.datacenterdynamics.com/en/news/nebius-expands-european-presence-announces-deployment-in-estonia-and-second-data-center-in-m%C3%A4nts%C3%A4l%C3%A4-finland/>
43. TechCrunch, "Mistral AI raises \$830M in debt to set up a data center near Paris", 30 March 2026. <https://techcrunch.com/2026/03/30/mistral-ai-raises-830m-in-debt-to-set-up-a-data-center-near-paris/>
44. Data Center Dynamics, "Mistral AI raises \$830m in debt financing for data center in Paris, France", 30 March 2026. <https://www.datacenterdynamics.com/en/news/mistral-ai-raises-830m-in-debt-financing-for-data-center-in-paris-france/>
45. Futurum Group, "Cohere Acquires Aleph Alpha: Deal Born of Sovereignty, Necessity", 2026. <https://futurumgroup.com/insights/cohere-acquires-aleph-alpha-a-deal-born-of-sovereignty-necessity/>
46. Exoscale, "Best European Cloud Provider for Sovereign Hosting", 27 November 2025. <https://www.exoscale.com/blog/european-cloud-providers/>
47. Euronews, "EU opens call for seven 'gigafactories' to train next-generation AI technologies", 30 July 2026. <https://www.euronews.com/my-europe/2026/07/30/eu-opens-call-for-seven-gigafactories-to-train-next-generation-ai-technologies>
48. EU Perspectives, "Europe's AI gigafactory race is on: who will win it?", July 2026. <https://euperspectives.eu/2026/07/europes-ai-gigafactory-race-is-on-who-will-win-it/>
49. STL Partners, "The EU's AI Gigafactory Initiative: What it means for digital infrastructure?", 18 May 2026. <https://stlpartners.com/articles/data-centres/eu-ai-gigafactory-initiative/>

50. The Next Web, "BT, Nscale, and Nvidia announce UK sovereign AI partnership", 23 April 2026. <https://thenextweb.com/news/bt-nscale-nvidia-14mw-sovereign-ai-data-centres-uk>
51. Synergy Research Group, "European Cloud Providers' Local Market Share Now Holds Steady at 15%". <https://www.srgresearch.com/articles/european-cloud-providers-local-market-share-now-holds-steady-at-15>
52. Qovery, "Top EU Sovereign Cloud Platforms in 2026", August 2026. Vendor blog; used for pricing, egress and migration observations only. <https://www.qovery.com/blog/top-eu-sovereign-cloud-platforms-migration-guide>
53. Mistral AI, model card, "Mistral Medium 3.5", 28 April 2026. <https://docs.mistral.ai/models/model-cards/mistral-medium-3-5-26-04>
54. Bloomberg, "France's Mistral in Funding Talks at About €20 Billion Valuation", 12 June 2026. <https://www.bloomberg.com/news/articles/2026-06-12/france-s-mistral-in-funding-talks-at-about-20-billion-valuation>
55. Microsoft, "Microsoft and Mistral expand strategic partnership to give enterprises and regulated industries frontier AI they can control", 21 July 2026. <https://news.microsoft.com/source/2026/07/21/microsoft-and-mistral-expand-strategic-partnership-to-give-enterprises-and-regulated-industries-frontier-ai-they-can-control/>
56. OpenEuroLLM, project site and consortium list. <https://openeurollm.eu/>
57. OpenLLM-Europe, "European Open Source LLM Projects", living catalogue. <https://github.com/OpenLLM-Europe/European-OpenLLM-Projects>
58. Kyutai, blog and release history, 2025 to 2026. <https://kyutai.org/blog/>
59. Black Forest Labs via GlobeNewswire, "Black Forest Labs Unveils FLUX 3, A New Multimodal Frontier Model For Visual Intelligence", 23 July 2026. <https://www.globenewswire.com/news-release/2026/07/23/3332364/0/en/black-forest-labs-unveils-flux-3-a-new-multimodal-frontier-model-for-visual-intelligence.html>
60. Cybernews, "German state replaces Microsoft with open source, saves millions each year", December 2025. <https://cybernews.com/news/schleswig-holstein-germany-microsoft-open-source/>
61. ZenDiS, "openDesk", project site. <https://www.opendesk.eu/en/about>
62. Univention, "NOVA: Sovereign Open-Source Workplace for 16,000 Municipal Workstations", August 2026. <https://www.univention.com/blog-en/2026/08/krzn-nova-open-source-administrative-workplace-nubus/>
63. Wikipedia, "Docs (software)"; La Suite numérique (DINUM), "Docs". [https://en.wikipedia.org/wiki/Docs_\(software\)](https://en.wikipedia.org/wiki/Docs_(software)) ; <https://lasuite.numerique.gouv.fr/produits/docs>
64. Proton, "Introducing Lumo: AI where every conversation is confidential", 23 July 2025; "Lumo 2.0". <https://proton.me/blog/lumo-ai> ; <https://proton.me/blog/lumo-2>

65. Tech Startups, "Swedish vibe-coding startup Lovable raises \$400 million at \$13.3 billion valuation as ARR nears \$600 million", 12 August 2026. <https://techstartups.com/2026/08/12/swedish-vibe-coding-startup-lovable-raises-400-million-at-13-3-billion-valuation-as-arr-nears-600-million/>

66. Brian Stack (Digital Twin AI), comment on "Your AI answers to a foreign government: China vs America", LinkedIn, August 2026; AlgentHive, product site. <https://www.aigenthive.co.uk>

67. heise online, "EuroStack Foundation: Europe's Tech Industry Finds Sovereignty Initiative", 2026. <https://www.heise.de/en/news/EuroStack-Foundation-Europe-s-Tech-Industry-Finds-Sovereignty-Initiative-10964468.html>

68. GOV.UK, "£100 million competition to back British AI companies to fix public services", 31 August 2026. <https://www.gov.uk/government/news/100-million-competition-to-back-british-ai-companies-to-fix-public-services>

69. Cloud Security Alliance, "EU Tech Sovereignty: Cloud Concentration Risk and the Compliance Cascade", 7 June 2026. <https://labs.cloudsecurityalliance.org/research/eu-tech-sovereignty-cloud-ai-enterprise-risk-v1-0-csa-styled/>