



Airbus for AI

A GLOBAL STRATEGY FOR
PUBLIC VALUE CREATION

AUGUST 2025



Public AI

AUTHORS

Joshua Tan is co-founder and research director at Metagov.

Brandon Jackson is a product manager at Metagov.

Robin Berjon is an independent researcher at Supramundane Agency and a senior research fellow at Metagov.

Diane Coyle is Bennett Professor of Public Policy at the University of Cambridge.

EXECUTIVE SUMMARY

Artificial intelligence (AI) is becoming central to economic productivity and national sovereignty, but its current frontier development is concentrated in a few dominant U.S. and Chinese firms. Many middle power countries like Canada, Germany, Japan, Spain, Sweden, Switzerland, and the U.K. have struggled to field competitive AI products despite substantial public and private investment across national labs and national champions (Allen, 2025; Council of the European Union, 2023; Escritt, 2023; Kreps, 2024; Scott, 2024; Stage et al., 2024; Stanford HAI, 2024). Insufficient funding, a talent/compensation gap, regulatory constraints, and fragmented markets all hinder their ability to scale AI initiatives effectively. As a result, directors in several national AI labs have called for more coordinated public interventions (Bengio, 2023; Valero & Crespo, 2024) while policymakers are debating the case for industrial policy interventions to support domestic AI firms (AI Now Institute, 2024; European Commission, 2025b; UK Government, 2025).

If AI is to serve the economic and security interests of all countries, it must be developed in a way that balances public value generation with a credible business model and a pathway to scale in global markets. Any solution must also respond to emerging industrial trends in AI, including the increasing commodification of large language models (LLMs), the rise of open source, and the dramatically increasing salience of national security and geopolitics to trade in AI.

There are precedents for such solutions. In particular, Airbus, originally developed by a set of middle powers to compete in aerospace, provides a successful example of an international public-private collaboration. This policy brief proposes a **Public AI Company** inspired by such models—a coordinated public-private partnership that ensures AI development serves shared national interests and maximizes public value rather than concentrating benefits in a handful of private entities headquartered in global superpowers. Drawing on decades of economic and public policy research as well as a technical analysis of current AI supply chains, this brief outlines why such a global approach is both necessary and achievable, and discusses some of the risks and challenges involved.

THE MIDDLE POWERS' DILEMMA

Middle power countries face structural barriers that prevent them from building globally competitive frontier AI systems.

Scale At every stage of the supply chain—data, compute, model training, deployment, and talent—scale dynamics favor the largest, most vertically-integrated players. These monopoly pressures have stratified the market into a top tier of U.S. and Chinese hyperscalers (such as OpenAI, Google, Microsoft, Alibaba, Anthropic, DeepSeek) and a fragmented landscape of dependent smaller firms. By way of example: ChatGPT has approximately 79.8% of the global consumer chatbot market as of July 2025. Perhaps the best-known “sovereign” competitor, Mistral’s Le Chat, has less than 4% market penetration in France, its home market (Statcounter, 2025).

Limited Markets States’ default approach has been to support the growth of local private firms (Department of Finance, Canada, 2024; European Commission, 2024; Ministry of Economy, Trade and Industry, Japan, 2024; UK Department for Science, Innovation and Technology, 2023). But even such national champions with overt state support struggle to access risk capital or penetrate markets outside their home market—consider Mistral’s performance in France compared to Germany (Dillet, 2025). There is currently insufficient economic incentive for new private sector entrants—companies like Sakana, Lighton, and Cohere—to develop competitive frontier models, especially the large pre-trained models that ground the modern AI stack (Azoulay et al., 2024).

Foreign Capture Public subsidies to private champions risk later rent extraction or foreign expropriation by acquisition—consider for example AMD’s acquisition of Finland’s Silo, which had been supported by Finnish/EU supercomputing resources (Cherney, 2024; LUMI, 2024).

Fragmented Sovereigns Apart from supporting domestic firms, several middle powers have adopted public AI strategies anchored by large public compute investments and domestic production of large language models within national labs (European Commission, 2024; Innovation, Science and Economic Development Canada, 2024; Sastry et al., 2024; Takano et al., 2024). Such investments can produce interesting research outputs, but public AI investments of \$30M for a new public

model or \$500M for a new public data center do not meaningfully compete with the dominance of the leading U.S. and Chinese frontier labs. Fragmentation of these national initiatives also results in duplicated efforts and contributes to scale problems, preventing coordinated advances in AI.

Research, Not Development States often fund national labs to build frontier models, such as Sweden's GPT-SW3, Singapore's SEA-LION, Spain's Salamandra, or the EU's many pan-European research consortia (Ekgren et al., 2024; European Commission, 2025a; Gonzalez-Agirre et al., 2025; Ng et al., 2025). Yet these labs, endowed with some of the world's most powerful supercomputers, struggle with bureaucratic constraints, no access to risk capital, a lack of skills for product development and subsequent commercialisation, high sensitivity to legal risk (e.g. on copyright), and the challenge of competing with private firms with more focus and greater agility. The result is a fragmented landscape of publicly-funded models that are not actually being used.

How does your country access frontier AI?	Is it competitive with frontier labs and models?	Does it generate public value?	Does it protect national sovereignty?
Import private models (e.g. OpenAI, Anthropic)	✓	✗ Most value is captured privately.	✗
Import open models (e.g. Meta, DeepSeek)	⚠ Cheap, but not best-in-class.	✓	⚠ Production still lives abroad.
Back national champions and local startups (e.g. Mistral)	⚠ Not in consumer or pretraining, but some traction in enterprise.	⚠ Spillovers, but no public accountability.	⚠ Threat of foreign acquisition (e.g. DeepMind, Silo).
Fund national AI labs (e.g. Swiss AI)	✗ Note: a bit of traction in public sector use.	✓	✓
Fund research consortia (e.g. OpenEuroLLM, CERN for AI proposals)	✗ Wrong incentives for product. Limited access to capital.	✓	✓

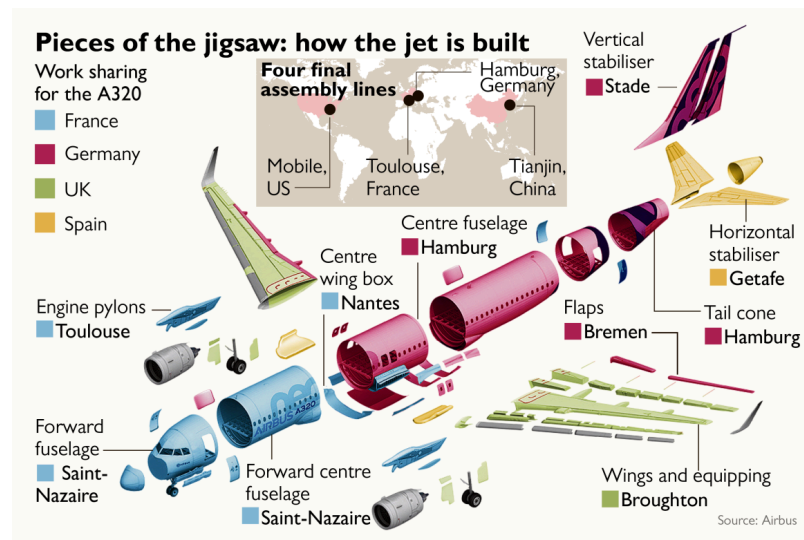
THE AIRBUS MODEL

The Airbus story began in the late 1960s as European governments recognized that their domestic aerospace industries were struggling to compete with dominant American firms Boeing and McDonnell Douglas. The fragmented nature of European aviation, with multiple national firms developing independent aircraft models, led to inefficiencies and an inability to achieve the economies of scale necessary for global competitiveness. In response, France, Germany, and later the U.K. and Spain collaborated to create Airbus, a consortium that pooled resources, expertise, and funding to develop a competitive alternative to American aviation giants.

Government support played a crucial role in Airbus' success. European states provided financial backing and advanced market commitments, ensuring that Airbus had the stability to undertake complex aircraft development projects. This public-private partnership allowed Airbus to invest in cutting-edge research, standardized manufacturing, and coordinated supply chains that spanned multiple countries. The Airbus A300, the consortium's first commercial aircraft, demonstrated the viability of this model by successfully entering the market against well-established and dominant American competitors.

Over time, Airbus refined its approach by embracing **supply chain specialization**. Each participating country focused on different aspects of aircraft production—France on avionics and final assembly, Germany on fuselage construction, the U.K. on wings, and Spain on tail sections. This structured division of labor contributed to efficiency while maintaining national contributions and industrial expertise.

Of course, Airbus has faced challenges—such as the failure of the A380 (Seabright, 2011)—and the 1960s aerospace industry differs in important ways from AI today (see Appendix: The 1960s Aerospace Industry). But Airbus remains one of the most successful examples of coordinated industrial policy in the modern era, and a promising model for global AI development.



AN AIRBUS FOR AI

Inspired by Airbus, we propose that Canada, Germany, Japan, Singapore, South Korea, Spain, Sweden, Switzerland, France, the U.K., and other middle powers unify their **existing sovereign AI efforts** into a **Public AI Company**—a competitive frontier AI lab to build and commercialize AI in the public interest.

How?

First, each participating nation needs to identify or organize at least one **national entity** which will participate in the Public AI Company. Each national entity should have top-level political support and frontier AI capabilities in at least one component of the AI stack. Depending on the country and its sovereign AI strategy, this entity may be an existing private national champion (e.g. France's Mistral, Canada's Cohere), a public national lab (e.g. Spain's BSC, Japan's RIKEN), an existing regional coalition (e.g. New Nordics AI), or a new public-private partnership.

Second, integrate these national entities into an **Airbus-style enterprise** that coordinates key shared activities:

1. **model building**, especially pre-training,
2. **resource sharing**, especially compute and expertise,
3. **market strategy**, including shared branding, and
4. **product strategy**—so one API, not twenty.

Third, each national entity should handle **localization** for language, culture, and in-country data, **deployment** in local data centers for inference, and **public engagement** with citizens and political leadership. Depending on the local commercial and political environment, each national entity may consider adopting a **public utility** business model (see Appendix: AI as a Public Utility).

Why?

Four core arguments:

Collective Action The middle powers' dilemma is fundamentally a coordination problem. Combined, the middle powers have the talent and resources to compete. But private national champions and public national labs in each country are incentivized to compete with each other over markets, funding, and prestige. Nonprofit R&D consortia like the EU's TrustLLM, EuroLLM, AI4EU, and OpenEuroLLM are unlikely to move the market due to a skills and incentives mismatch—any competitor needs to have a

focused product strategy, not 20 separate research agendas defined over a finite pot of funding. Private firms are unlikely to coordinate without government intervention or the prospect of being acquired, and governments are already cautious about consolidation as the most likely acquirers are American or Chinese giants (Benaich & Chalmers, 2024). Government intervention to form a **public-private consortium** solves the collective action problem for firms while sidestepping the zero-sum funding incentives for public labs and nonprofits.

Public Goods are Good Foundation models are arguably already infrastructure for modern applications, but they are not *public infrastructure* like highways, electricity, or the Internet. They should be (Mozilla Foundation, 2024; Public AI Network, 2024; Vanderbilt Policy Accelerator, 2025). Frontier models exhibit strong public good characteristics—they are non-rival, widely re-usable, and generate broad productivity spillovers—implying that the social returns of investing in these technologies can far exceed private returns. Left to the market, they will be undersupplied and over-enclosed relative to the social optimum. A Public AI Company operationalizes these economic arguments into a concrete business—one that is premised on (existing) sovereign AI investments and designed from the ground up to maximize public value generation ahead of private value extraction. Note that this line of reasoning strongly argues for a public utility business model (see Appendix: AI as a Public Utility) or even a public option (Coyle, 2022; Schneier & Sanders, 2023) for the national entities.

“A little known fact, we tried to get the public sector to fund us before we went to the capped profit model. There was no interest. But yeah, I think **if the country were working a different way — I would say a better way — this would be a public sector project.** But it’s not and here we are.”
- Sam Altman (New York Times, 2021)

Autonomy and Sovereignty Airbus markets the planes, but subsidiaries in each country actually make them, and then airlines in each nation purchase, own, and operate them. We imagine a similar market structure: some national entities contributing to the pre-training supply chain while others operate as a system of “flag carriers”—in place of Iberia, Japan Airlines, and SAS, imagine a Spanish AI Company, a Japanese AI Company, and a Nordic AI Company each of whom owns and deploys their own post-trained models. National labs and national champions already underwrite their countries’ digital

sovereignty (see Appendix: Why Sovereignty?), but they face the middle powers' dilemma described earlier. The alternatives are stark: continue investing in sub-scale sovereign AI strategies, import AI from abroad—with clear and growing risks to national security, economic competitiveness, and cultural identity—or put faith in pitches like OpenAI for Countries (OpenAI, 2025) or NVIDIA and Perplexity's sovereignty sandwich (Laurent, 2025; NVIDIA, 2025).

Scale Without Monopoly Narrow commitments to single-nation sovereign AI are unlikely to generate competitive alternatives to the leading closed or open source options. Like aerospace in the 20th century, competing on frontier AI requires both horizontal and vertical integration that neither private national champions or public national labs can manage alone. A consortium model allows countries to achieve scale without centralizing power into a single nation or national entity, which is preferable for many reasons from economic efficiency to AI safety (Bengio, 2023; Future of Life Institute, 2024; Verdegem, 2024). A multi-state setup also naturally incentivizes the production of interoperable AI ecosystems across member countries. Interoperability tends to level the playing field, benefiting startups and fostering a more competitive AI landscape rather than reinforcing existing tech oligopolies (AI Now Institute, 2023).

KEYS TO SUCCESS

Market Strategy Public national labs and private national champions already struggle to compete against dominant hyperscalers like OpenAI and Anthropic, large firms with significant first-mover advantages in distribution, compute, proprietary data, and talent. For example, France's Mistral has delivered technically but has still struggled to find substantial consumer adoption. The Public AI Company's market strategy needs to face this reality, pick a counter-positioned product and brand, and pursue it with conviction. This is critical; Airbus succeeded because it addressed a clear market need. Open source communities will be critical allies in generating demand, as will governments.

Open Source Existing middle power strategies are based on open models, including open weight (just the weights are shared) and open source (all aspects of training and data are shared). There are good reasons for this. But: (1) the world's leading models are not open weight and definitely not open source, putting a significant cap on open source adoption (Bilski, 2025); (2) the leading open models are built by Meta and a Chinese hedge fund; (3) open source is a mode of production, not a business model; and (4) it still takes substantial investment to activate open models, meaning market power still tends to concentrate in large hyperscalers (Tan et al., 2025). To win allies in the open source community, the Public AI Company needs to demonstrate that it can complement open source production with a scalable business model.

Bootstrap the Product The product is clear. Airbus built airplanes, and the Public AI Company will build frontier models and the APIs to serve them, at scale. There are many more decisions that go into the product: what kinds of models, whether multi-modal, domain-specific, consumer or enterprise, how much fine-tuning, etc. But it is also clear that training a \$2 billion "leading frontier model" is not the minimum viable product; for now, the world does not lack good-enough open source and open weight models, including ones from existing national champions and national labs. Early on, therefore, the Public AI Company should focus on generating local demand for existing open models, bootstrapping data flywheels, and plan to launch partner with an upcoming public AI project like Swiss AI or OpenEuroLLM.

Focus on Product Airbus succeeded because public action catalyzed the creation of an innovative product that met a market demand. The Public AI Company must be set up so that its long-term success depends on its ability to build frontier AI models that can compete against dominant alternatives like OpenAI's ChatGPT, Anthropic's Claude, Alibaba's Qwen, and DeepSeek's R1 as use of AI grows. That means empowering leadership to make commercial decisions based on customer needs, not political favor.

Public Support To succeed, a Public AI Company needs to rally public support and public investment. To do so, it needs to build a world-class policy team, find allies in national champions and national labs, and overcome the cold start problem with governments. But public support can mean many things. States can catalyze efforts with advance market commitments, incentivising public services to use its model. They can invest directly, taking equity stakes. They can also grant specific privileges that allow companies to prioritize public value creation over short-term profitability. Just as public utilities receive advantaged access to public markets in exchange for providing universal public access, each national entity within the Public AI Company could receive preferential access to public datasets, public compute, and other regulatory incentives from its sponsoring country. See more in Appendix: AI as a Public Utility.

Good Governance Collective action is hard. The consortium needs to act with the speed and coherence of a single lab. But distributing ownership, accountability, and geographic operations tends to introduce a lot of incentive and coordination problems, especially when the underlying entities of a sovereign AI strategy—private startups, public labs, private multinationals, open-source communities—have structurally different incentives and resources. In particular, there will be considerable political interest among consortium member countries in the decisions made. The Public AI Company needs to take a multipronged approach. First, start small: four labs, not twenty. Second, select carefully the partners in each country, aiming to balance leadership, operational capacity, and political backing. Third, explore and amend the consortium structure: depending on the participating labs, the right structure may be an Airbus-style merger, an early Visa-like co-op, a federated structure granting substantial autonomy to member firms, or most likely something bespoke. Fourth, let the supply chain evolve over time, responding to developing capabilities, with different members

focusing on pre-training, evaluation, fine-tuning, inference, or deployment based on their comparative advantages.

Public Accountability CERN and the European Space Agency demonstrate how transparent governance ensured sustained political and financial commitment. Another example is the Human Genome Project, which coordinated international scientific collaboration through ensuring open data-sharing policies, enabling broad participation in genetic research. It will be important to balance the activities of participating national entities—which link the consortium to public compute, local markets, and national populations—with the international activities of the consortium as a whole.

Pitfalls to Avoid

Governance Challenges The political economy behind structuring a Public AI Company poses particularly challenging questions. Some of these concern the right governance structure. On the one hand, a purely governmental model presents the well-known problems of bureaucratic management. Overly rigid governance and rules could stifle innovation and slow decision-making. For example, the U.S. Department of Defense’s lengthy and complex procurement process has historically delayed the adoption of cutting-edge technologies (Oakley, 2025). On the other hand, complicated public-private governance models could risk slowing down progress by imposing excessive administrative burdens on labs and research institutions.

Beyond these specific organisational design questions, there is a clear danger of national political imperatives and corporate lobbying hindering the potential of a Public AI consortium to develop AI models that are commercially viable in global markets. Member governments will want to support national companies, and recent history offers warning examples of politically-driven but commercially-inefficient projects such as Northvolt and Gaia-X (Goujard & Cerulus, 2021; Milne, 2024). To reduce the chance of political capture, the Public AI Company should emphasize a product mission and ensure that there are no permanent national quotas for contributions, while recognising the political realities. Here, a common public mission can also help align the competing national interests that have undermined many other industrial policies.

Market Distortions Poorly designed subsidies or exclusive public contracting can lead to inefficiencies and crowd out private sector innovation instead of fostering competition. For example, the European Union's early subsidies for semiconductor manufacturing led to dependence on government funding without fostering long-term competitiveness, resulting in firms that struggled to scale without continued financial support. Similarly, public AI funding programs that lock in a handful of providers could prevent a more diverse and competitive ecosystem from emerging. The Public AI Company should be designed with clear milestones and plans for a structured shutdown if it does not scale or deliver.

Lack of Sustained Political Support On the other hand, large-scale PPPs require stable, long-term commitments. Projects like the ITER fusion project have faced delays due to fluctuating political and financial commitments. Another example is the Galileo satellite navigation system, which suffered from repeated budgetary constraints and disagreements among EU member states, leading to prolonged development timelines and inefficiencies. From AI's own history, the U.S. government's inconsistent funding for AI research has resulted in successive AI winters and summers, undermining sustained progress in public sector AI initiatives.

THE TIME IS NOW

An Airbus for AI is achievable.

It can start now. Building a Public AI Company does not require coordinating between 20 different bodies; it can start with a small team and two labs. It also doesn't require high-level political approval to get started. A 10-person team working for 3 months, building on existing collaborations and open-source models, and supplied with about 100k GPU-hours (~150k USD), can deliver a minimum viable product and establish first proof points around product, coordination, and market demand.

It makes financial sense. A huge amount of costly public compute is coming online over the next few years. Collaboration between countries will maximise the return to taxpayers for these investments. At the same time, most private companies are ditching the pre-training game as the existing hyperscalers pull ahead, open-source eats into margins, and distillation undermines even closed source business models. This creates

an opening in the market for a focused pre-training play. A consortium of national entities supported by several middle powers would immediately become a global contender among frontier AI labs.

Politicians are ready. The Overton window has shifted in favour of bolder government actions in many middle powers due to the rapidly-changing geopolitical and security environment. Governments are already making large investments in public compute and are urgently searching for alternatives to the American and Chinese options. Our proposal aligns especially well with Canada's Pan-Canadian AI Strategy, Japan's AI Promotion Act, and Europe's AI Continent Plan and Apply AI Strategy. In particular, an Airbus-style approach would directly complement (and resolve some of the gaps in) the existing EU AI Gigafactories program.

A team is ready. National champions and national labs have built and released foundation models, and they have been doing so for years. Engineers and leaders at these labs already agree on the urgency of collaboration, the insufficiency of existing efforts, and the need for a more competitive product.

CONCLUSION

The Airbus model offers a clear precedent for how nations can collaboratively build advanced, competitive, and publicly accountable technological capabilities. An Airbus for AI can link together model builders across several labs into a sustainable commercial entity, ensuring that AI development reflects democratic values, advances scientific progress, and delivers broad economic benefits rather than concentrating wealth and power in a few dominant overseas firms. With the incoming surge of public investment into AI and increasing political urgency, now is the time to act.

ACKNOWLEDGEMENTS

The authors are grateful to Paul Seabright, Magnus Sahlgren, Aleksei Turobov, Ricardo Baeza-Yates, Albert Cañigueral, Felix Sieker, Sam Klein, Joseph Low, Alex Krasodonski, Jacob Taylor, Jake Hirsch-Allen, Katherine Gorman, Nick Vincent, and Yoshua Bengio for their comments and suggestions. Joshua Tan was supported in the writing of this essay by a grant from the Future of Life Institute.

REFERENCES

- AI Now Institute. (2023, April 11). 2023 Landscape: Confronting Tech Power. *AI Now Institute*.
<https://ainowinstitute.org/publications/research/2023-landscape-confronting-tech-power>
- AI Now Institute. (2024, October 15). Redirecting Europe's AI Industrial Policy. *AI Now Institute*.
<https://ainowinstitute.org/publications/research/redirecting-europes-ai-industrial-policy>
- Allen, M. (2025, January 30). Switzerland caught in middle of US-China race for AI dominance. *SWI Swissinfo.Ch*.
<https://www.swissinfo.ch/eng/science/switzerland-caught-in-middle-of-us-china-race-for-ai-dominance/88804566>
- Azoulay, P., Krieger, J. L., & Nagaraj, A. (2024). *Old Moats for New Models: Openness, Control, and Competition in Generative AI* (Working Paper 32474). National Bureau of Economic Research. <https://doi.org/10.3386/w32474>
- Benaich, N., & Chalmers, A. (2024, April 18). Alchemy is all you need. *Air Street Press*.
<https://press.airstreet.com/p/alchemy-is-all-you-need>
- Bengio, Y. (2023). AI and Catastrophic Risk. *Journal of Democracy*.
<https://www.journalofdemocracy.org/ai-and-catastrophic-risk/>
- Berjon, R. (2025, February 25). Digital Sovereignty: The Missing Manual. *Robin Berjon*.
<https://berjon.com/digital-sovereignty/>
- Bilski, D. (2025, July 31). 2025 Mid-Year LLM Market Update: Foundation Model Landscape + Economics. *Menlo Ventures*.
<https://menlovc.com/perspective/2025-mid-year-llm-market-update/>
- Cherney, M. A. (2024, July 10). AMD to acquire Finnish startup Silo AI for \$665 mln to step up in AI race. *Reuters*.
<https://www.reuters.com/markets/deals/amd-acquire-finnish-startup-silo-ai-665-mln-step-up-ai-race-2024-07-10/>
- Clark, D. (2012). *Control Point Analysis*. SSRN.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2032124
- Cottier, B., Rahman, R., Fattorini, L., Maslej, N., & Owen, D. (2024). *The rising costs of training frontier AI models* (arXiv:2405.21015; Version 1). arXiv.
<https://doi.org/10.48550/arXiv.2405.21015>
- Council of the European Union, A. and R. T. (2023). *ChatGPT in the Public Sector – overhyped or overlooked?*
https://www.consilium.europa.eu/media/63818/art-paper-chatgpt-in-the-public-sector-overhyped-or-overlooked-24-april-2023_ext.pdf
- Coyle, D. (2022). The Public Option. *Royal Institute of Philosophy Supplements*, 91, 39–52.

- <https://doi.org/10.1017/S1358246121000394>
- Crew, M. A., & Kleindorfer, P. R. (1986). *The Economics of Public Utility Regulation*. Springer.
- Department of Finance, Canada. (2024, December 6). *Deputy Prime Minister announces \$240 million for Cohere to scale-up AI compute capacity*.
<https://www.canada.ca/en/departement-finance/news/2024/12/deputy-prime-minister-announces-240-million-for-cohere-to-scale-up-ai-compute-capacity.html>
- Dillet, R. (2025, February 19). Mistral's Le Chat tops 1M downloads in just 14 days. *TechCrunch*.
<https://techcrunch.com/2025/02/19/mistrals-le-chat-tops-1m-downloads-in-just-14-days/>
- Ekgren, A., Gyllensten, A. C., Stollenwerk, F., Öhman, J., Isbister, T., Gogoulou, E., Carlsson, F., Heiman, A., Casademont, J., & Sahlgren, M. (2024). *GPT-SW3: An Autoregressive Language Model for the Nordic Languages* (arXiv:2305.12987). arXiv.
<https://doi.org/10.48550/arXiv.2305.12987>
- Escritt, T. (2023, August 23). Germany plans to double AI funding in race with China, U.S. *Reuters*.
<https://www.reuters.com/technology/germany-plans-double-ai-funding-race-with-china-us-2023-08-23/>
- European Commission. (2024, December 10). *Seven consortia selected to establish AI factories* [Text]. European Commission.
https://ec.europa.eu/commission/presscorner/detail/en/ip_24_6302
- European Commission. (2025a, February 3). *A pioneering AI project awarded for opening Large Language Models to European languages | Shaping Europe's digital future*.
<https://digital-strategy.ec.europa.eu/en/news/pioneering-ai-project-awarded-opening-large-language-models-european-languages>
- European Commission. (2025b, April). *AI Continent Action Plan*.
<https://digital-strategy.ec.europa.eu/en/factpages/ai-continent-action-plan>
- Future of Life Institute. (2024). How to mitigate AI-driven power concentration. *Future of Life Institute*.
<https://futureoflife.org/grant-program/mitigate-ai-driven-power-concentration/>
- Gonzalez-Agirre, A., Pàmies, M., Llop, J., Baucells, I., Dalt, S. D., Tamayo, D., Saiz, J. J., Espuña, F., Prats, J., Aula-Blasco, J., Mina, M., Pikabea, I., Rubio, A., Shvets, A., Sallés, A., Lacunza, I., Palomar, J., Falcão, J., Tormo, L., ... Villegas, M. (2025). *Salamandra Technical Report* (arXiv:2502.08489). arXiv.
<https://doi.org/10.48550/arXiv.2502.08489>
- Goujard, C., & Cerulus, L. (2021, October 26). Inside Gaia-X: How chaos and infighting are killing Europe's grand cloud project. *POLITICO*.
<https://www.politico.eu/article/chaos-and-infighting-are->

- killing-europes-grand-cloud-project/
 Innovation, Science and Economic Development Canada. (2024, December 5). *Canada to drive billions in investments to build domestic AI compute capacity at home*. Government of Canada.
<https://www.canada.ca/en/innovation-science-economic-development/news/2024/12/canada-to-drive-billions-in-investments-to-build-domestic-ai-compute-capacity-at-home.html>
- Jiang, M., & Belli, L. (2024). *Digital Sovereignty in the BRICS Countries*.
- Kellerman, M. (2025). This Chip will Self-Destruct. *False Positive*.
<https://mileskellerman.substack.com/p/this-chip-will-self-destruct>
- Kreps, S. (2024, December 6). *The global AI race: Will US innovation lead or lag?* Brookings.
<https://www.brookings.edu/articles/the-global-ai-race-will-us-innovation-lead-or-lag/>
- Laurent, L. (2025, June 18). Nvidia's 'Sovereign' AI Could Win a Prize for Irony. *Bloomberg*.
<https://www.bloomberg.com/opinion/articles/2025-06-18/nvidia-s-sovereign-ai-could-win-a-prize-for-irony>
- Lehdonvirta, V. (2022). *Cloud empires: How digital platforms are overtaking the state and how we can regain control*. The MIT Press.
- LUMI. (2024, December 11). The successor for LUMI announced—A new pan-European supercomputer and one of the first European Union's AI Factories to be placed in Finland. *LUMI*.
<https://lumi-supercomputer.eu/successor-for-lumi-announced/>
- Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., ... Oak, S. (2025). *Artificial Intelligence Index Report 2025* (arXiv:2504.07139). arXiv.
<https://doi.org/10.48550/arXiv.2504.07139>
- Milne, R. (2024, November 22). 'There was so much promise': How Northvolt tumbled into bankruptcy. *Financial Times*.
- Ministry of Economy, Trade and Industry, Japan. (2024, February 2). AIの に けたプロジェクト「GENIAC」をします.
<https://www.meti.go.jp/press/2023/02/20240202003/20240202003.html>
- Mozilla Foundation. (2024). *Public AI: Building an AI ecosystem that is of, by, and for the people*. Mozilla Foundation.
- Neven, D., Seabright, P., & Grossman, G. M. (1995). European Industrial Policy: The Airbus Case. *Economic Policy*, 10(21), 313. <https://doi.org/10.2307/1344592>
- Ng, R., Nguyen, T. N., Huang, Y., Tai, N. C., Leong, W. Y., Leong, W. Q., Yong, X., Ngui, J. G., Susanto, Y., Cheng, N.,

- Rengarajan, H., Limkonchotiawat, P., Hulagadri, A. V., Teng, K. W., Tong, Y. Y., Siow, B., Teo, W. Y., Lau, W., Tan, C. M., ... Teo, L. (2025). *SEA-LION: Southeast Asian Languages in One Network* (arXiv:2504.05747). arXiv. <https://doi.org/10.48550/arXiv.2504.05747>
- NTIA. (2024). *Dual-Use Foundation Models with Widely Available Model Weights*. National Telecommunications and Information Administration, US Department of Commerce.
- NVIDIA. (2025, June 11). *NVIDIA Partners With Europe Model Builders and Cloud Providers to Accelerate Region's Leap Into AI*. <https://nvidianews.nvidia.com/news/nvidia-partners-with-europe-model-builders-and-cloud-providers-to-accelerate-regions-leap-into-ai>
- Oakley, S. S. (2025). *Defense Acquisition Reform: Persistent Challenges Require New Iterative Approaches to Delivering Capability with Speed*. Government Accountability Office. https://www.gao.gov/assets/gao-25-108528.pdf?utm_source=chatgpt.com
- OpenAI. (2025, May 7). *Introducing OpenAI for Countries*. <https://openai.com/global-affairs/openai-for-countries/>
- Potter, B. (2024, January 31). A Cycle of Misery: The Business of Building Commercial Aircraft. *Construction Physics*. <https://www.construction-physics.com/p/a-cycle-of-misery-the-business-of>
- Public AI Network. (2024). *Public AI: Infrastructure for the Common Good*. <https://publicai.network/whitepaper>
- Sastry, G., Heim, L., Belfield, H., Anderljung, M., Brundage, M., Hazell, J., O'Keefe, C., Hadfield, G. K., Ngo, R., Pilz, K., Gor, G., Bluemke, E., Shoker, S., Egan, J., Trager, R. F., Avin, S., Weller, A., Bengio, Y., & Coyle, D. (2024). *Computing Power and the Governance of Artificial Intelligence* (arXiv:2402.08797; Version 1). arXiv. <https://doi.org/10.48550/arXiv.2402.08797>
- Schmid, J., Sytsma, T., & Shenk, A. (2024). *Evaluating Natural Monopoly Conditions in the AI Foundation Model Market*. RAND Corporation. <https://doi.org/10.7249/RR3415-1>
- Schneier, B., & Sanders, N. E. (2023, June 12). Build AI by the People, for the People. *Foreign Policy*. <https://foreignpolicy.com/2023/06/12/ai-regulation-technology-us-china-eu-governance/>
- Scott, J. (2024, December 20). *In 2024, Canada struggled to find its place in the global AI race | BetaKit*. <https://betakit.com/in-2024-canada-struggled-to-find-its-place-in-the-global-ai-race/>
- Seabright, P. (2011). The Hidden Costs of Political Sponsorship of Industrial Firms. In O. Falck, C. Gollier, & L. Woessmann (Eds.), *Industrial Policy for National Champions*. MIT Press. <https://direct.mit.edu/books/edited-volume/2196/chapter/>

- 58226/The-Hidden-Costs-of-Political-Sponsorship-of Stage, C., Slottnér, E., Tung, K. O., Rydman, W., Sigurbjörnsdóttir, Á. A., Ellemann, K., Rogaczewski, A., Svenson, M., Partanen, K., & Husby, Ø. (2024, October 31). *Nordic Region must be Europe's driving force in AI race*. <https://www.newnordics.ai/nordic-region-europes-driving-force-in-ai-race>
- Stanford HAI. (2024, November 21). *Global AI Power Rankings: Stanford HAI Tool Ranks 36 Countries in AI | Stanford HAI*. <https://hai.stanford.edu/news/global-ai-power-rankings-stanford-hai-tool-ranks-36-countries-in-ai>
- Statcounter. (2025, June). *Statcounter GlobalStats AI Chatbot Market Share*. StatCounter Global Stats. <https://gs.statcounter.com/ai-chatbot-market-share/>
- Takano, R., Takizawa, S., Tanimura, Y., Nakada, H., & Ogawa, H. (2024). *ABCI 3.0: Evolution of the leading AI infrastructure in Japan* (arXiv:2411.09134). arXiv. <https://doi.org/10.48550/arXiv.2411.09134>
- Tan, J., Vincent, N., Elkins, K., & Sahlgren, M. (2025). *If open source is to win, it must go public* (arXiv:2507.09296). arXiv. <https://doi.org/10.48550/arXiv.2507.09296>
- Transcript: Ezra Klein Interviews Sam Altman. (2021, June 11). *The New York Times*. <https://www.nytimes.com/2021/06/11/podcasts/transcript-ezra-klein-interviews-sam-altman.html>
- UK Department for Science, Innovation and Technology. (2023). *AI regulation: A pro-innovation approach*. UK Department for Science, Innovation and Technology. <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach>
- UK Government. (2025, January 13). *AI Opportunities Action Plan*. GOV.UK. <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>
- UK Ministry of Defence. (2025). *Strategic Defence Review 2025 – Making Britain Safer: Secure at home, strong abroad*. UK Ministry of Defence.
- Valero, M., & Crespo, D. (2024, April 11). Es la hora de 'Albus': Por qué Europa debe crear una gran empresa de AI. *El País*. <https://elpais.com/tecnologia/2024-04-11/es-la-hora-de-aibus-por-que-europa-debe-crear-una-gran-empresa-de-ai.html>
- Vanderbilt Policy Accelerator. (2025). *Global Rise of Public AI*. <https://cdn.vanderbilt.edu/vu-URL/wp-content/uploads/sites/412/2025/05/05220054/The-Global-Rise-of-Public-AI.pdf>
- Verdegem, P. (2024). Dismantling AI capitalism: The commons as an alternative to the power concentration of Big Tech. *AI & SOCIETY*, 39(2), 727–737. <https://doi.org/10.1007/s00146-022-01437-8>

APPENDIX: THE 1960S AEROSPACE INDUSTRY

In the 1960s, the global jetliner market (in the Western/OECD bloc) was dominated by effectively two firms in the US: Boeing and McDonnell-Douglas (MDD). This appendix draws out the parallels and distinctions between that era and today's AI industry.

Dimension	1960s aerospace industry	2020s frontier AI industry
Market Concentration	Functional duopoly: Boeing dominant; MDD secondary	Near duopoly: OpenAI dominant, Anthropic/Google secondary
Barriers to Entry	Very high: capex, engineering, regulatory	Very high: compute, talent, data
Product Characteristics	Physical goods (aircraft); high durability	Cloud-based SaaS and APIs; continuous updates
Customer Base	Few, stable: national airlines, militaries	Fragmented: enterprises, governments, developers, consumers
Supply Chain Structure	Nationalized, vertically integrated	Private models are vertically integrated, some completely (e.g. Google); open models are not
Revenue & Pricing	Large contracts, capex-heavy sales	Usage-based APIs, subscriptions; licensing, limited pricing power
Strategic Importance	High: military-industrial complex, infrastructure	High: dual-use, economic and cultural implications
Geopolitical Structure	Bipolar (U.S. vs USSR); closed blocs	Bipolar U.S.-China bifurcation, but fluid alliances
Innovation Cycle	Slow: decades-long product timelines	Rapid: quarterly releases, open source model replication
Public Intervention	High in Europe (Airbus); moderate in U.S.	High everywhere: subsidies, labs, export controls

Key Parallels

Market concentration. For the aerospace industry then, Boeing and McDonnell-Douglas (MDD) were the key firms, though Boeing was by far the more dominant player (Neven et al., 1995). This mirrors today's AI landscape, where a handful of frontier

labs control the cutting edge of foundation model development. OpenAI, analogous to Boeing, holds outsized market share, though there is no clear one firm in second place, with Google's Gemini, Anthropic's Claude, Meta's LLaMa, and DeepSeek's R1 regarded as the leading competitors to OpenAI's ChatGPT across different markets.

Barriers to entry. Building the product in both industries demands massive scale and investment. In aerospace, developing a new airliner was—and remains—extraordinarily expensive. Boeing's 747 program in the late 1960s, for example, cost on the order of \$1–2 billion, approximately \$10–20 billion in today's dollars (Potter, 2024). Such programs often amounted to bet-the-company ventures for manufacturers. Similarly, state-of-the-art AI models require vast compute resources, data, and capital. Training a single frontier model now incurs tens of millions of dollars in compute cost alone, and this figure is rising exponentially. According to one recent analysis, if current trends continue, the largest AI training runs will exceed \$1 billion in cost by 2027 (Cottier et al., 2024).

Strategic importance, both industrial and military. Both industries carry strategic significance well beyond their market size. In the Cold War era, a robust aerospace sector was synonymous with national power: it conferred military advantage, high-tech industrial capabilities, and prestige. Governments treated large civil aircraft as strategic assets, often supporting domestic manufacturers for defense-industrial reasons (Neven et al., 1995). Likewise, frontier AI is now viewed as a critical strategic technology by governments around the world. AI capability has dual civilian-military uses and is seen as pivotal for future economic competitiveness and security (NTIA, 2024).

Key Differences

Open source. Aircraft designs and technologies were and are entirely proprietary. Boeing, MDD, and others closely guarded their engineering; no equivalent of “open-source airplanes” existed. By contrast, today's AI landscape includes a prominent open-source/open-weight movement. Many advanced AI models have their model weights or code openly released, or at least accessible to researchers. For example, Meta's LLaMa large language models were released openly (with certain usage restrictions), allowing other organizations to build on it. This openness means that, unlike an upstart in 1965 who could not

simply copy a 707, an AI entrant in 2025 can potentially take a state-of-the-art model's weights off the shelf. However, open models come with caveats: even when cutting-edge models are released, they remain expensive and complex to activate at scale (Tan et al., 2025), and a persistent gap in performance undermines adoption. This is especially true in enterprise software (Bilski, 2025), where historically open source has performed better (e.g. in cloud and operating systems). In summary, open source has no true parallel in 1960s aerospace, and it both empowers new entrants (by lowering IP barriers) and complicates the business model for companies seeking an open source strategy. Strategically: open source, combined with the ubiquity/speed of global software distribution (see next point) mean that frontier labs (public or private) cannot rely solely on proprietary advantage or captive domestic markets.

Durable goods vs. SaaS. The way value is delivered and monetized diverges greatly. Aerospace companies sell durable capital goods—physical airplanes and components—typically via procurement contracts or sales to airlines. These are high-priced, one-off (or few-off) transactions followed by long service relationships. The revenue model relies on upfront sales and long-term maintenance. In contrast, many AI frontier labs distribute their models as a cloud service via a managed frontend or API calls; customers generally do not take possession of the model itself. This means usage is metered and subscription-based rather than a single capital expenditure. The SaaS model allows continuous updates and iterative improvement, but also means that “sales” are low-friction and usage-dependent. (A caveat: many companies are also pursuing enterprise-style licensing deals that can sometimes mimic the sales and maintenance business cycles of aerospace.) It also means global distribution can happen instantly online. This difference in business model affects everything from cash flow to customer relationships and must inform how closely to apply an Airbus analogy.

Niche vs. mass market customer base. In the 1960s (and to a lesser degree today), demand for jetliners was concentrated within a relatively exclusive set: government-adjacent flag carriers, commercial airlines, and military procurement agencies. These customers were sophisticated and few in number. This arguably made swings in demand harder to forecast (Potter, 2024), but the nationalized character of demand also made it easier for governments to intervene in order to guarantee

demand. By contrast, generative AI's user base is far broader and more diffuse. Modern AI platforms serve hundreds of millions of end-users directly and countless businesses in varied sectors. This makes it harder, though not impossible, to coordinate and aggregate demand. For governments, this means that industrial policy in AI will have to rely on the supply side (investments in compute, talent, foundation models). For firms, this means that competition is more intense and that brand and ecosystem strategy matter more than government relations. Compressed product cycles also mean that any entrant must be structurally agile and innovation-driven—i.e. more research lab than traditional manufacturer.

APPENDIX: AI AS A PUBLIC UTILITY

We do not explicitly call for a public utility model as part of the Airbus for AI proposal, as it is a decision that each national entity needs to make for itself. However, the public utility model, effectively a regulated monopoly model, aligns well with the proposal and its assumptions.

Why a utility lens?

Modern foundation models share the classic signatures of a **natural monopoly**: enormous fixed- and sunk-cost investments (pre-training compute, large proprietary or curated data sets, hyperscale inference) combined with near-zero marginal cost of serving each additional inference. Economic theory predicts under-provision in such markets and increasing returns to scale that favor a handful of dominant providers. Recent empirical work confirms that these conditions already exist in practice, with just a few American and Chinese firms controlling the lion's share of training compute and model releases. RAND's assessment of foundation-model economics goes further, concluding that "foundation models meet every textbook criterion for natural-monopoly regulation" (Schmid et al., 2024).

It bears repeating: a public utility is often just a regulated monopoly.

How conventional utilities are paid

Generally speaking, public utilities do not derive direct profits from the commodity they deliver (water, gas, electrons). Instead, regulators let them earn a regulated rate of return (ROR) on the **regulated asset base (RAB)**: the pipes, poles, wires, and treatment plants (Crew & Kleindorfer, 1986). The model incentivizes long-term infrastructure investment while capping prices to consumers. This regulated business model takes three main forms, depending on who owns and operates the infrastructure:

1. **State-owned enterprise.** Classic public ownership and operation, funded via public capital budgets. For example, municipal water utilities, Scandinavian power grids, many Asian metro systems.
2. **Private regulated monopoly.** Investor-owned and -operated infrastructure where profits are rate-capped; capital earns an approved ROR. Examples include US electric and gas utilities.

3. **Public-private franchise.** Long-term concession of state-owned assets that a private party operates with a service fee based on performance benchmarks, see for example wireless spectrum licenses or UK rail concessions.

For a public AI utility,

- the **regulated asset base** would be the pre-training and inference compute clusters, fundamental research infrastructure, and any related industrial, national, and global data lakes.
- the **universal service obligation** is an affordable, universally accessible inference API.
- the **independent regulator** is to-be-determined, but perhaps some kind of National AI Utility Commission, per country.
- the **financing vehicle** will depend on the particular country / jurisdiction, with possibly different vehicles for different parts of the RAB. Options include long-term loans, equity, and direct public appropriations.

Because any given foundation model's productive life is (far) shorter than that of a dam or pipeline, regulators may want to shift to accelerated depreciation (e.g., two- or three-year recovery) and link new rate-base additions to open-weights releases or independent audits.

The key implication of the model: profits flow from building and maintaining the shared infrastructure, not from vertical capture of downstream applications. This realigns incentives toward safe, high-quality model development instead of data hoarding or rent-seeking "app store"-like lock-ins.

Regulatory architecture

How would a notional National AI Utility Commission govern a public AI utility? Some ideas:

1. **Price-cap or yardstick regulation.** Benchmark the cost per 1k inference tokens to a basket of peer utilities; reward outperformance.
2. **Safety & alignment mandate.** Independent audits of robustness, bias, and misuse risk, funded via a small tariff on gross revenue (analogous to NERC oversight in electricity).

3. **Sunset & innovation clauses.** Allow competitive entrants to bid on specific services (e.g., multimodal extensions) if they can supply at lower regulated cost or higher safety.
4. **Public data dividend.** Require reinvestment of a share of net earnings into expanding public domain datasets, compensating cultural creators, or subsidising civic applications.

Opportunities and Risks

If a given national entity chose to operate as a public utility, it would:

- gain **automatic national champion** status.
- enjoy **stable, long-horizon financing** and profits for pre-training and safety R&D.
- be able to offer **universal, affordable access** to strong models for SMEs, researchers, local governments.
- have **societal alignment** baked into the utility license.

At the same time, not all public utilities are well-run. Some key risks:

- **regulatory capture.** This could be mitigated by rotating commissioners and publishing all cost data.
- **efficiency and innovation drag.** Stable, rate-base profits might crowd out incentives to innovate and compete. This could be mitigated by carve-outs for competitive innovation sandboxes, though arguably innovation might accelerate with a more stable foundation entrant.
- **over- or under-investment.** ROR and RAB rules depend on good forecasting, which is deeply uncertain in AI today. This could be mitigated through more flexible RAB rules or some sort of protocol-based governance mechanism.

More needs to be done to examine the implications and practicalities of a public AI utility strategy. Successful provision would likely require some level of iteration and significant adjustments to existing public utility models.

Why now?

The inflow of speculative capital into frontier model training has begun to taper, while public spending on public AI infrastructure is exploding, with over \$18 billion of confirmed commitments across G7 countries in 2024–25 alone (Maslej et al., 2025). Locking that hardware into a utility-style rate-base gives

taxpayers a clear return on infrastructure and cushions national champions and national labs against the boom-bust cycle in venture funding.

Each participating nation could grant its national entity an **AI Utility License**, entitling preferential access to sovereign data, low-cost public power, or accelerated permitting. In exchange, the national entity would maintain universal access APIs, transparent cost reporting, and alignment audits. To some degree, this is already happening in jurisdictions such as Switzerland, Singapore, and the EU.

APPENDIX: WHY SOVEREIGNTY?

Sovereignty is self-determination over a specific domain, and digital sovereignty can be understood as *the exercise of authority or power over the digital sphere* (Jiang & Belli, 2024).

Tech monopolies are in a position to decide what is relevant, which media will succeed, how retail works, which labor laws apply in entire sectors, what speech rights people have, which business models are acceptable under which condition, whose data is shared, what privacy rules apply, and more. They settle more contractual agreements between parties than governmental judicial systems. They tax businesses that operate on their infrastructure at levels that are multiple integrals of corporate tax rates. In many instances, for those who operate a business in the digital sphere, the rules made by tech monopolies matter more than the laws made by governments (Lehdonvirta, 2022).

This rule-setting power directly limits state sovereignty, and the European Union's experience shows that a purely regulatory approach cannot reclaim it.

The mechanics of power

The power to set rules in the digital sphere comes from uncontested ownership of infrastructural systems and authority over architectural **control points** (Clark, 2012). Sovereignty therefore requires the ability to neutralise the power that actors may have to mobilise these control points whether through regulation (including competition policy) as well as through sovereign alternatives. Both are needed, as regulatory approaches often face limits due to the cost of enforcement and risks stemming from geopolitical constraints. Here, it's important to note how sovereignty depends on maintaining (or removing) control over many control points. Much of the Airbus for AI proposal is derived from the necessity of building a complete AI stack, rather than merely a model, merely compute, or merely data.

A control point may translate into a **kill switch** if control is over an input that is critical to a system's operation. This is particularly true if the input in question is physical (e.g. chips) or strongly entrenched (e.g. a two-sided marketplace) so that it cannot be replaced dynamically in the way that a straightforward software dependency might. As a rule, commercial actors are unlikely to

rely on the use or threat of kill switches (except against direct competitors or commercial threats, e.g. removing a competing app store from their own app store) because overuse threatens a business's viability. However, geopolitics and the weaponisation of interdependencies has to be factored into the equation as a given country may be in a position to direct a company to terminate services or activate a kill switch (in addition to milder interventions such as espionage, degraded performance, and refusal to abide by local law). Kill switches are notably a concern with respect to NVIDIA and AI chips in general (Kellerman, 2025).

The national security case

Loss of sovereignty is a security threat. Placing core physical infrastructure such as energy grids and transportation systems under foreign control puts national defense at the mercy of foreign manipulation. The same logic applies digitally, except that much digital infrastructure is **already captured or at-risk** across the stack (Berjon, 2025). This has two important implications for AI policy:

1. First, **the AI bottleneck**. As AI becomes embedded within computing infrastructure, the lack of sovereign AI for middle powers risks further loss of control over other components of digital infrastructure.
2. Second, **the mutual dependency trap**. The capture of other infrastructure components, for instance B2C and B2B distribution channels or data sources, will in turn severely limit the ability to deploy sovereign AI if and when it is available. This creates a vicious cycle in which AI models have limited data to train on and no distribution outlets, and potential sovereign distribution outlets are thwarted by limited access to sovereign AI.

In general, the interlinkage across multiple architectural control points is why many "sovereign AI" offerings by global hyperscalers amount to **sovereignty-washing**: domestic branding without credible control over all or even most of the relevant architectural control points. Even when given, most often that control is not credible, as it relies on reserve control—control that is not used but *could* be used—or on good behavior that is unlikely to survive geopolitical strain.

Location risk compounds the problem. Countries can and do benefit from AI without training their own models or hosting their own compute. Placing some AI capacity abroad can cut costs

and improve quality. For instance, the Trump administration in May 2025 announced a deal to develop a 5GW AI campus in Abu Dhabi in exchange for the export to the UAE of some of the most advanced semiconductors produced by NVIDIA and AMD. Cheap Gulf energy (and abundant Gulf financing) provides an obvious rationale for such a deal on a cost basis. But in doing so, the US can deter the UAE from cutting off access to the servers through its military power; most middle powers cannot and must find other strategies to secure global supply chains.

Given growing evidence that AI capacity will be critical in future military confrontations (NTIA, 2024; UK Ministry of Defence, 2025), as well as the deteriorating strategic environment for many middle powers, these countries should be willing to pay to ensure a substantial degree of AI sovereignty. The policy question is how to design AI capabilities that will maximise strategic and economic benefits while mitigating costs and dependencies.

The positive case

AI sovereignty isn't only defensive. Domestic capacity delivers its own benefits.

First, there may be positive **locational externalities**: having a technically trained workforce located on the country's territory may create spillovers for other firms; or the implementation of AI in complex corporate settings may be more effective if clients are located physically close to the innovating firms whose technology they are licensing.

Second, tech monopolies introduce a democratic deficit for businesses and citizens where they can no longer resort to the usual political channels in order to change the rules under which they live and operate. Sovereign AI capacity restores the possibility of **democratic oversight** and alignment with local priorities.

Third, frontier AI models trained and governed abroad reflect the cultural biases and commercial imperatives of their origin markets. These may diverge sharply from the aims and values of the societies which import them. Domestic AI development enables models that are **culturally representative**, linguistically inclusive, and responsive to local norms, improving both adoption and trust.

Finally, more **market competition** is good on its own. Tech markets are already prone to high concentration and incumbency advantages; generative AI is a particularly powerful example of the type (Schmid et al., 2024). Public intervention to ensure more domestic capacity can raise the floor of competition.

Taken together, our point is that sovereign AI is not merely an insurance policy. It is also an engine for national competitiveness and sustained innovation.