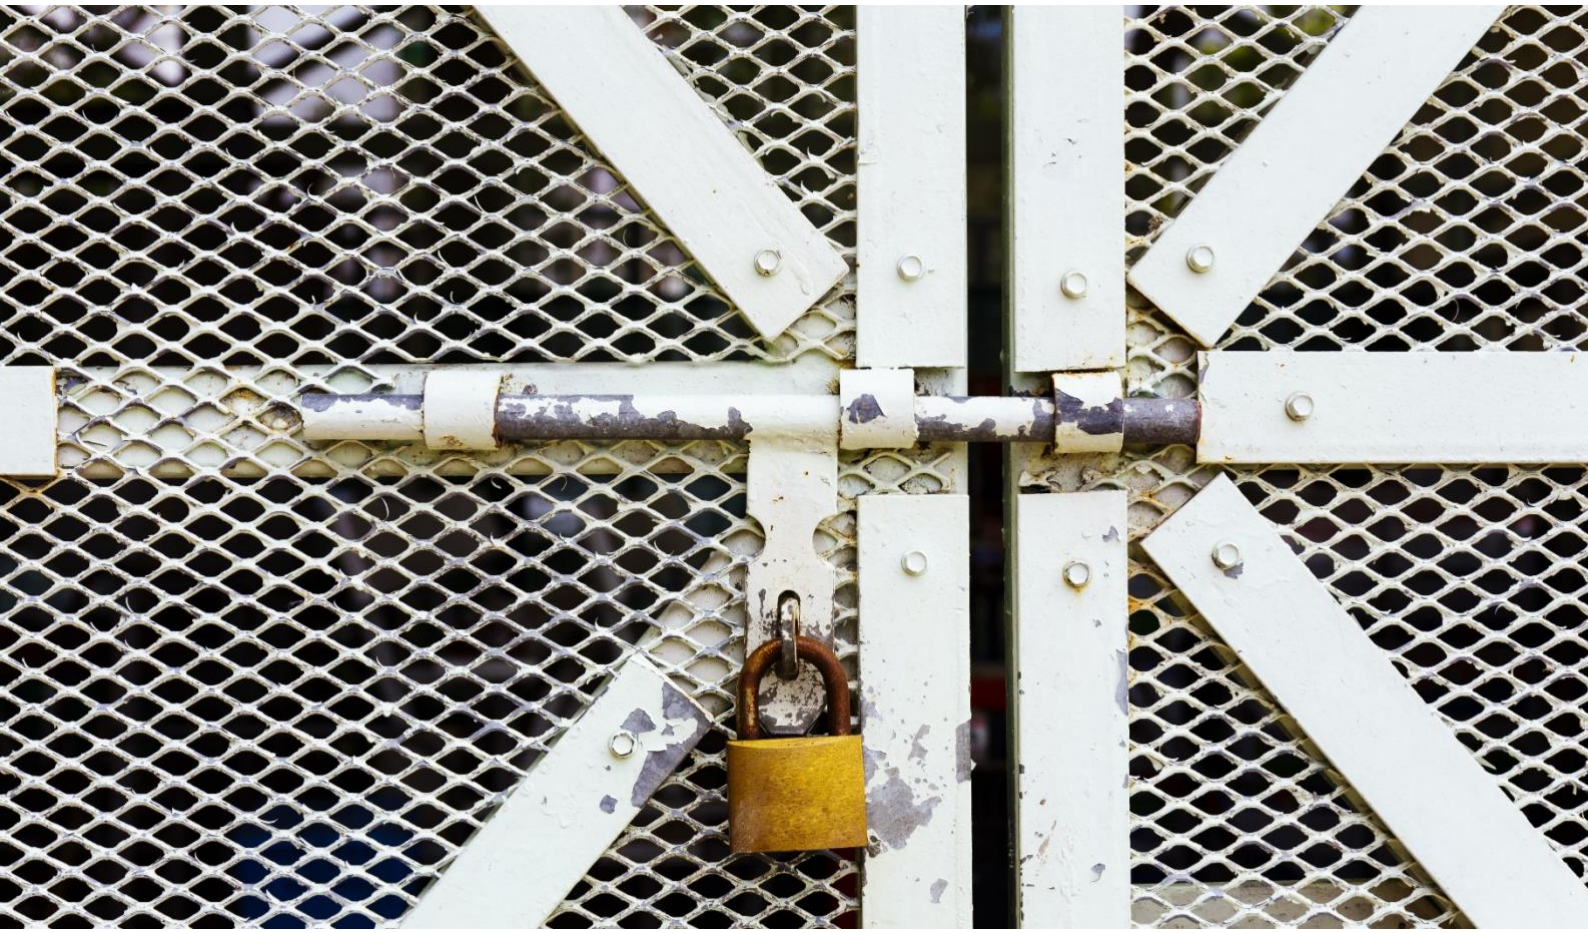

What the Mythos ban means for AI sovereignty

On 12 June 2026, the US ordered Anthropic to disable Claude Fable 5 and Mythos 5 for any foreign national worldwide. The directive turned the frontier model itself into a sovereignty question, not just where data sits, but whether the model you rely on can be withdrawn overnight.

Kuba Smolorz, Senior Consultant

June 2026

From launch to shutdown in three days

On 9 June 2026, Anthropic launched Claude Fable 5 and Claude Mythos 5. Early results were striking. Stripe reported that Fable 5 compressed months of engineering work into days[1]. Mythos 5, a more capable variant restricted to a select group of vetted partners, was described by the National Security Agency's leadership as capable of identifying vulnerabilities in many of the US government's classified systems within hours rather than weeks. [2]

On 12 June, the Trump administration issued an export control directive that forced Anthropic to disable both Fable 5 and Mythos 5 worldwide. The directive cited national security authorities and required Anthropic to suspend access to both models by any foreign national, whether inside or outside the United States, including the company's own non-US employees. Export controls on physical hardware are well established, but this was the first time the US government had ordered a deployed frontier AI model taken offline, the latest in a long-running sovereignty story.

A new point of failure

Concerns about US technology dependency are not new. The 2018 Cloud Act, which allows US authorities to compel American cloud providers to hand over data stored anywhere in the world, has been a known risk for years. The 12 June directive means that even organisations that had made peace with US data jurisdiction now have to reckon with the possibility that the frontier model itself can be switched off worldwide overnight, without notice and with no right of appeal.

US frontier models remain ahead on raw capability, and that gap matters most for demanding applications, such as cybersecurity, scientific research, and complex reasoning. Replacing them carries genuine capability costs, at least in the near term.

Businesses, governments, and policymakers now face a choice: accept the continuity risk that comes with depending on US frontier models or invest heavily in sovereign alternatives. Some have already moved. France announced €109 billion in AI infrastructure investment at the AI Action Summit in February 2025, and Mistral Compute now offers EU-hosted inference outside the reach of US jurisdiction. Saudi Arabia's HUMAIN, a Public Investment Fund subsidiary launched in May 2025, is building the full stack: data centres, cloud, models, and applications.

The EU sits in the most exposed position: it hosts roughly 5% of global data centre capacity against the US's 80%, giving it leverage as a trading bloc but little as a compute jurisdiction. The CADA proposals begin to insulate the most sensitive public-sector applications, but they leave private-sector dependency, the research community, and critical-infrastructure operators still reliant on US systems remaining available at all times.

The telco sovereignty opportunity

The Mythos ban increases demand for sovereign AI across the stack (infrastructure, data, inference), and telcos are well placed to address much of it. They are national players, operate in highly regulated environments, hold existing relationships with enterprises and governments, and already run critical infrastructure: exactly the trust profile sovereign AI requires. Many are already pursuing this play (see our recent research on telco sovereign AI strategies[3]).

The ban also introduces a new layer to worry about, the frontier model itself. On that layer, telcos can play three roles. The most accessible is helping enterprises navigate the frontier model trade-off, mapping

---

[1] https://www.anthropic.com/news/claude-fable-5-mythos-5
[2] https://www.cnbc.com/2026/06/23/anthropics-mythos-model-found-vulnerabilities-in-classified-us-government-systems-official-grammersays.html
[3] https://stlpartners.com/research/navigating-sovereign-ai-telco-strategies/

workloads against continuity risk and matching each use case to the right model jurisdiction. A second is offering sovereign or lower-risk alternatives directly, including open-weight models hosted in-country. A third, realistically limited to a handful of markets with the capital, energy, and policy backing to compete, is becoming a frontier model developer in their own right.

What should telcos and enterprises do now?

1. **Dependency audit:** Every operator and enterprise using US frontier AI should map which workflows are reliant on which models, and what the consequence of an overnight suspension would be. The risk is not hypothetical; it has now occurred once, and the legal and political conditions that produced it remain in place.
2. **Tiered strategy:** Not all AI use cases carry the same sovereignty risk. Customer-facing applications or internal productivity tools may be a reasonable place to accept dependency on US models. Applications touching critical infrastructure, defence, national security, or sensitive financial data warrant a different approach. Matching the sensitivity of the use case to the sovereignty tier of the infrastructure is a practical, implementable approach, and it is the logic underpinning CADA's assurance levels.
3. **Seek sovereign alternatives:** Both telcos and enterprises should be actively evaluating alternatives to US frontier models, including dedicated sovereign cloud providers, telco-hosted offerings, and open-weight models hosted in-country.
4. **Offer sovereign alternatives:** Telcos in particular sit on the supply side as well as the demand side. With the infrastructure, regulatory standing, and customer relationships to host sovereign capacity, they should be moving now to make that offer concrete for the enterprises and governments looking for it.

Full sovereignty has limits. Owning the AI supply chain end to end (chips, infrastructure, models, applications) is realistically only available to the US and China. Every other country has to make trade-offs: which layers to depend on, and which to invest in insulating. The model layer can no longer be treated as commodity infrastructure that simply works, and each country, and each enterprise within it, will need to draw its own line on acceptable risk before the choice is made for them.

Kuba Smolorz is a Senior Consultant at STL Partners, specialising in specialising in AI while bringing broad expertise across next-generation connectivity and infrastructure to assess its impact on telco operations and B2B revenue growth.

Get in touch with the author to learn more

kuba.smolorz@stlpartners.com

Or visit STL Partners' AI Hub

www.stlpartners.com/ai