Eighteen Days of Dark: The Fable 5 Embargo and India's AI Hedge

July 22, 2026



On June 9, Anthropic released Claude Fable 5 to the general public, incorporating additional safeguards into the Mythos architecture. Meanwhile, the unrestricted Mythos model remained limited to Project Glasswing, i.e., an initiative in which Anthropic granted access to Mythos to organizations that secure critical software, and has included more than 150 organizations across more than 15 countries.

Within days of its release, researchers from Amazon found a way to bypass Fable 5's safeguards, prompting the model to identify what Anthropic described as a handful of previously known, low-severity software vulnerabilities and, in one case, produce code demonstrating how one could be exploited. They also used a series of prompts to get the Mythos-class model to provide restricted information about cyberattacks. Andy Jassy, the CEO of Amazon, first raised such concerns with Treasury Secretary Scott Bessent.

As a result, on 12 June 2026, the US Department of Commerce issued an export control directive barring all foreign nationals, including those working within the US, from using Fable 5 and Mythos. Since Anthropic cannot verify user nationality in real time, the directive required a uniform global ban on both models.

The immediate impact of this ban was felt across sectors in India. Start-ups, banks, and IT firms that had incorporated the model into their regular workflows were unable to adapt because the directive was issued without prior notice and without an exception for Indian users, despite India being Claude's second-largest market after the US. Amid the ban, on 22 June, Japan released its own homegrown AI model, Fugu by Sakana AI, which orchestrates other models rather than competing with them directly.

Finally, after 18 days, during which Anthropic worked with the US Department of Commerce on a new safety classifier, the ban was lifted on 30 June. Fable 5 was re-released globally on 1 July, with additional safeguards that, according to Anthropic, block the bypassing technique in more than 99% of cases.

During the India AI Impact Summit at Bharat Mandapam, New Delhi, from 16 to 21 February, the Director of the White House Office of Science and Technology Policy, Michael Kratsios, defined AI sovereignty as "owning and using best-in-class technology for the benefit of your people, and charting your national destiny in the midst of global transformations."

Based on this definition alone, India's AI strategy is not as sovereign as it appears. Instead of building the AI Infrastructure and its supply chain from scratch, it focuses on importing essential components like GPUs, access to frontier models, and so on. The unilateral ban on Fable 5 exposed key structural problems in this strategy.

The main Frontier AI models are now export-controlled and

are treated like advanced semiconductor chips. This means that access to advanced AI capabilities is neither a commercial purchase nor guaranteed by signed contractual declarations. Unilateral domestic laws prioritize national security while overriding the New Delhi Declaration on 19 February 2026 at the India AI Impact Summit. This incident is likely to accelerate India's efforts to reduce strategic dependence on foreign reliance and invest in indigenous foundational models and open-source alternatives. It is evident that India's AI foundations are rooted in structural dependencies.

The embargo instantly split India's debate into three competing camps.

The first is the "use-case capital" camp, championed by Infosys co-founder Nandan Nilekani, which argues that India's comparative advantage lies in applying AI at massive scale rather than trying to out-train Silicon Valley in raw pretraining. The foundational claim is that India should focus on developing local applications, products, and public-sector tools built on top of imported models. However, the Fable 5 embargo exposed a structural flaw in this approach: if the foreign Department of Commerce can abruptly revoke access to crucial technologies, the domestic "application" layer becomes irrelevant.

The second is the "sovereign-model" camp, which argues that India must train its own foundational models from scratch to achieve AI technological sovereignty and foster indigenous progress. Companies such as Sarvam AI were selected under the IndiaAI Mission, a USD 1.25 billion initiative, and have since released 30B- and 105B-parameter models this year. Sarvam still faces limitations: the embargo demonstrates how these sovereign models lag behind frontier benchmarks and are trained on modest compute, all on imported NVIDIA hardware. This means India's physical AI infrastructure remains as exposed to US export controls as the software APIs relied upon by the first camp.

The middle camp advocates a “do both” strategy, arguing that building applications on top of foreign frontier models while simultaneously funding domestic models is a necessary insurance policy. The harder version of this position—an aggressive pivot to open-source and Chinese open-weight models to avoid sudden vendor shut-offs—belongs with the sovereign model camp; Zoho founder Sridhar Vembu argues that technology is now the substance of national sovereignty. The supply disruption caused by Fable 5’s embargo has triggered caution across the tech industry. Tech founders such as Activate’s Aakrit Vaish weighed in on the sudden export restrictions, commenting in CNBC’s Inside India newsletter that the embargo “completely changes things”, fundamentally altering the risk profile of corporate AI procurement.

Japan’s approach to this problem has been quite distinct. Lacking the enormous resources required to develop a model that matches America’s best, Sakana AI has engineered Fugu. It is a 7-billion-parameter light model, that acts as a sort of dispatcher. Drawing from an ecosystem of other models, it simply reroutes tasks to the most appropriate one.

The company also claims that Fugu Ultra matches or even surpasses frontier models on several benchmarks, including a score of 73.7 on SWE-Bench Pro. Therefore, if an embargo is unexpectedly placed on a particular model, Fugu can redirect the task to the best available alternative, leveraging reliability through redundancy. This strategy is characteristic of Japan’s “soft law” approach to AI governance. As articulated in the Hiroshima AI Process, Japan prioritizes collaborative efforts with minimal interference. This may offer India insight into a more efficient and affordable approach to AI sovereignty.

While India aims for technological sovereignty, its strategy so far has offered little protection in moments like this. Instead of aligning fully with either the US or China, India is likely to hedge. Even as it works toward foundational

models of its own, the Fugu release points to a more attainable middle path.

Given the shortage of semiconductor chips, building such an ecosystem domestically will be challenging and time-consuming. Hence, India is likely to explore middle paths such as Japan’s Fugu, which, despite not being a foundational model, still reduces dependency.

India’s commitment to advancing its AI strategy is visible in the new AI Governance and Economic Group, which exists to operationalize its strategy. Procurement rules are likely to require multi-model architectures, reducing dependence on any single provider. Bilateral commitments on continuity and notice periods for model access, negotiated within the India-US technology dialogue, would harden trust into procedure. Finally, a homegrown orchestration layer, backed as protective infrastructure, could switch among global AI models if access to one were ever cut off.