

Who Controls the Stack: Three Strategic Scenarios for AI Governance in the Indo-Pacific

By Mark Bryan Manantan

Over the past few months, the US and China have both demonstrated their capability to wield export control measures over their respective frontier AI developers. On June 12, 2026, the US Department of Commerce (DOC) ordered Anthropic to suspend access to its two most advanced models—Claude Fable 5 and Mythos 5—for any foreign national, including the company’s own employees. With no reliable way to verify nationality in real time, Anthropic disabled both models for its entire global user base within hours. DOC lifted the controls on June 30, in exchange for company commitments on risk detection, standards development, and reporting of malicious activity. Meanwhile, on April 27, 2026, China’s National Development and Reform Commission (NDRC) blocked the USD 2 billion Manus-Meta deal under its Foreign Investment Security Review. The decision extended Beijing’s regulatory toolkit beyond tech transfer, data, and talent to completed overseas transactions.

The reversal does not restore the status quo ante; it confirms that the status quo has changed. For eighteen days, a model widely regarded as the world’s most capable was disabled for every non-US user by a single



Image Credit: Flickr/Ethen Rera

administrative instrument and restored only through a negotiation between one firm and one government in which no other jurisdiction had standing. Access to frontier AI is now a policy variable rather than a commercial one, and the precedent is portable.

Neither move is unprecedented. Over the past decade, both superpowers have refined their export control playbooks. Washington has moved through successive postures: from decoupling to derisking to friendshoring. Beijing has pursued a dual track. It has prioritised self-reliance while expanding research collaboration and diffusing its open-weight models, especially in Southeast Asia. Both converge on the same economic security agenda: who controls the stack, controls the future. Drawing on empirical analysis, this article conducts strategic foresight to outline three scenarios for the Indo-Pacific’s AI governance, to help

scholars, policymakers, and practitioners navigate a more contested and unpredictable region.

Scenario 1: AI Fragmentation

The first and most feasible scenario is AI fragmentation: two competing spheres of influence, one anchored by the US, the other by China. This bifurcation is already visible.

In December 2025, Washington announced Pax Silica, a coalition to streamline the AI supply chain from energy and critical minerals to semiconductors and data centres. Japan, South Korea, and Australia have embraced it, deepening their integration with the US ecosystem.

The window of tech neutrality for Southeast Asia—including close US partners like Singapore and the Philippines—is narrowing. US models remain widely used, but cheaper and strongly multilingual open-weight alternatives from DeepSeek, Alibaba, and Baidu are closing the performance and adoption gaps. The June suspension sharpened this dynamic: Zhipu AI launched its GLM-5.2 flagship the following day. Across the region, enterprises publicly reassessed their dependence on a single American provider. Every disruption in the US stack is an adoption event for the Chinese one. Southeast Asia's deepening reliance on Huawei and ZTE in 5G and data centre's compounds this dynamic.

Both capitals are consolidating jurisdiction. Beijing is expanding its capacity to restrict Chinese firms' overseas investments, transactions, and research partnerships. Washington's AI Diffusion Rule places Singapore and the Philippines under a validated end-user scheme for advanced chips, while the rest of Southeast Asia faces heightened scrutiny and ambiguity.

The Fable–Mythos episode exposed a quieter tiering mechanism operating above the chip layer. Project Glasswing, Anthropic's restricted cybersecurity initiative, expanded in June to around 200 vetted organisations across more than fifteen countries, granting access to vulnerability discovery capabilities unavailable on the open market. Reported participants include NATO, the European Union Agency for Cybersecurity (ENISA), and South Korea's Samsung. Northeast Asia's national champions sit inside the perimeter; Southeast Asia is almost entirely outside it. Frontier cyber defence is now allocated by security vetting and alliance proximity rather than market demand: a validated end-user logic for models rather than chips, with far less regulatory visibility.

In short, both powers can now flex jurisdictional capacity to constrain Southeast Asia's hedging and, at worst, force a choice between Pax Silica and Pax Sinica. Japan and South Korea face the same terrain, exposed to retaliation should they ban Chinese models outright. Partial decoupling results—but interoperability erodes, raising governance

and compliance costs and leaving less secure applications with embedded bias.

Fragmentation and sovereignty assertion are not sequential responses. They are simultaneous. The region's most consequential actors are pursuing both at once.

Scenario 2: AI Sovereignty

The second scenario points to accelerating sovereign AI capabilities. Recognising their exposure to US and Chinese export policy, many countries are investing heavily in building their own.

The Fable–Mythos episode redefined AI sovereignty. For a decade, debates in the Indo-Pacific centred on infrastructure and data residency. The directive to Anthropic restricted access by nationality, irrespective of user location or the data processing jurisdiction. A model can run in an approved region, on compliant infrastructure, over lawfully localised data, and still be unavailable to a government's own officials because of who they are. Data residency is no hedge against eligibility risk—and most regional strategies are built almost entirely around residency.

Northeast Asia is moving fastest, fast-tracking industrial policy to build national champions. Japan's AI Promotion Act, enacted in May 2025, commits nearly JPY 10 trillion to AI and semiconductor infrastructure. South Korea backs its sovereign push with a USD 75 billion

national fund that designates AI data centres as social infrastructure.

Taiwan occupies the most anomalous position. As home to TSMC, it manufactures the chips every sovereign AI programme depends on yet cannot formally join Pax Silica due to PRC sensitivities: indispensable but unrecognised. Taipei has responded with the AI Basic Act, the Tainan Cloud Center, and TAIDE, a sovereign model trained on Taiwanese data to counter the assumptions embedded in Chinese systems. It is building the full stack from below.

Southeast Asia is not idle. Sovereign ambitions are materialising across three layers—legal, infrastructural, and model-level. Vietnam enacted the region's most comprehensive AI law in March 2026, with a risk-based classification system and a National AI Development Fund. Malaysia imposed an AI-only data centre mandate while building ILMU, the region's first sovereign LLM, which outperforms GPT-4o on Malay benchmarks. Indonesia partnered with India on Sahabat-AI, a 70-billion-parameter model supporting Javanese and Sundanese. Southeast Asia is not choosing between Washington and Beijing; it is building a third option.

Scenario 3: AI Resilience

The final scenario is the middle ground. Japan, South Korea, Australia, Taiwan, and ASEAN form ASEAN+Tech Plus—a coalition of small and medium powers pursuing a middle pathway between fragmentation's zero-sum logic and sovereignty's isolation.

Australia, the Philippines, and Indonesia anchor the sourcing and processing layer for critical minerals, while Taiwan, Japan, South Korea, Australia, Taiwan, and ASEAN form ASEAN+Tech Plus—a coalition of small and medium powers pursuing a middle pathway between fragmentation’s zero-sum logic and sovereignty’s isolation. Australia, the Philippines, and Indonesia anchor the sourcing and processing layer for critical minerals, while Taiwan, Japan, South Korea, Malaysia, and Singapore collaborate on advanced chips and data centres. The caveat is structural dependency: US and Chinese frontier models still underpin sovereign LLM development, NVIDIA still holds the most advanced compute, and AWS, Microsoft, Google, Huawei, and Alibaba dominate cloud.

Realising this requires a coordination instrument, and the anticipated ASEAN Digital Economy Framework Agreement (DEFA) provides one. DEFA’s digital trade and data governance architecture can anchor an ASEAN-led AI standards body—with Australia, Japan, and South Korea as full partners and Taiwan as an observer—harmonising sovereign cloud certification and aligning safety benchmarks. The Fable-Mythos resolution makes this urgent. The terms on which the models returned included commitments on pre-release government access, safeguard information sharing, and an industry jailbreak-severity framework drafted with Amazon, Microsoft, and Google. These will function as de facto global standards, and no Indo-Pacific middle power was party to the negotiation that produced them.

Conclusion

The evolving US-China export control rivalry is weaponised interdependence made visible: countries without frontier AI capabilities are structurally exposed to external shocks that threaten both national interest and digital economic development. The three scenarios outlined here represent three different answers to the same underlying question—who controls the stack. Fragmentation accepts the answer as Washington or Beijing. Sovereignty tries to rewrite it nationally. Resilience argues that small and middle powers can collectively redefine it on their own terms. For Northeast Asia’s advanced economies, doubling down on sovereign AI ecosystems is both viable and underway. For Southeast Asia, the capacity gap is real—but not permanent.

The most actionable path forward is deliberate multilateral coordination: diversifying technology partnerships, pooling compute investments, and building the institutional architecture that converts national initiatives into a genuinely interoperable regional AI order. The Fable-Mythos episode offered a preview of how that question gets answered by default: in a private negotiation between one company and one regulator, with every other jurisdiction learning the outcome after the fact. Who controls the stack is still an open question. The window to shape the answer is now.

Dr. Mark Bryan Manantan is research fellow at La Trobe University's Centre for Global Security. He is also the Co-Founder of Balai Labs and a member of UNESCO's AI Experts Without Borders.