LAB 45 — TELAAH KEBIJAKAN

No. 049 - 24 April 2026

Telaah Kebijakan LAB 45 adalah wadah yang dirancang untuk menyampaikan pandangan kritis dan analisis terkini dari para peneliti serta analis kebijakan terkait berbagai isu strategis seputar politik keamanan, ekonomi politik, politik media, dan gender. Platform ini bertujuan untuk memberikan wawasan mendalam sekaligus menawarkan gagasan inovatif dalam menghadapi tantangan lokal ataupun global. Pendapat yang tercantum dalam setiap komentar merupakan tanggung jawab penulis sepenuhnya dan tidak merefleksikan posisi resmi LAB 45. Jika Anda memiliki pertanyaan atau memerlukan informasi lebih lanjut, silakan menghubungi tim kami melalui lab45@lab45.id.

Indonesia's AI Architecture Was Designed for a Different World

Christian Guntur Lebang
Koordinator Analis, Laboratorium Indonesia 2045

On April 7, 2026, Anthropic announced that its newest AI model, Claude Mythos Preview, had discovered thousands of previously unknown critical vulnerabilities in every major operating system and every major web browser in the world. Some of these flaws had gone undetected for decades. The model didn't just find the holes. It demonstrated the ability to exploit them, chaining multiple vulnerabilities together into attack sequences of a sophistication previously associated with elite state-sponsored hacking units.

Anthropic chose not to release the model publicly. Instead, it convened Project Glasswing, a coalition of twelve companies, including Amazon, Apple, Microsoft, Google, and NVIDIA, plus roughly forty additional organizations, to use Mythos for defensive purposes: scanning their own systems, patching vulnerabilities, and preparing for a world in which models with similar capabilities will proliferate within months. Anthropic committed $100 million to the initiative.

No government authorized this process. Not even the United States, where Anthropic is headquartered and currently locked in a legal battle with the Pentagon over the military use of AI, had a seat at the table when these decisions were made. No treaty governed the disclosure timeline. No multilateral body determined who would be included or excluded. A private company discovered that the world's shared digital infrastructure is critically compromised, and then it decided, on its own authority and according to its own judgment, who gets told, in what order, and on what timeline.

Indonesia's government agencies, financial institutions, and critical infrastructure run on the same operating systems and browsers that Mythos found riddled with vulnerabilities. The question is whether anyone in Jakarta has registered what this means for sovereignty itself.

The vacuum

The Mythos episode is the sharpest illustration to date of a problem that has been building for years: AI capabilities have outpaced every

institution meant to govern them.

Nuclear technology, the comparison that AI developers themselves keep reaching for, produced the Non-Proliferation Treaty, the IAEA, international inspection regimes. Decades of diplomatic effort went into ensuring that capabilities powerful enough to threaten entire nations would not be left in the hands of individual actors without oversight. AI has been compared to "a factory that produces cars, micro scooters, animals, and nuclear weapons all at the same time" (Clark, 2026), yet has produced nothing equivalent. There is no treaty. There is no agency. When a capability this powerful emerges, the world has no agreed-upon process for deciding who gets access, who gets warned, or who is responsible for the consequences. In the absence of any framework, those decisions default to the company that built the model.

Both Washington and Beijing are racing to capture AI for state purposes, yet neither has built the governance frameworks the technology demands (Bremmer, 2025). As one observer put it, America is "barrelling toward a future in which nobody claims responsibility for AI" (Wong, 2026). The reason is structural. Governments hesitate to constrain companies they view as engines of economic growth, national competitiveness, and military advantage. The result is the rise of "silicon sovereigns": private firms exercising authority that once belonged to states, while public power is steadily hollowed out (Chesterman, 2026).

Jakarta has not stepped into that vacuum. Its political leadership treats AI as an economic opportunity, not a governance question, while rapidly building its digital economy on the infrastructure of others. This governance challenge is not a theoretical risk for Jakarta. It is already the operating reality.

What Indonesia sees when it looks at AI

To understand how Indonesia's political leadership thinks about AI, start with what is most visible. Vice President Gibran Rakabuming Raka has made AI arguably his signature issue. He visits high school AI workshops, judges student competitions, posts AI-generated content on social media, and has pushed for AI and coding to enter the national curriculum. At the G20 Summit in Johannesburg in November 2025, he was tasked with delivering Indonesia's position on AI governance and critical minerals. His framing, when he does speak substantively, centres on youth empowerment, digital literacy, and economic fairness. These are not trivial concerns. But they are also not a strategic response to what AI has become.

In February 2025, weeks after DeepSeek's open-source model went viral, Luhut Binsar Pandjaitan, chair of the National Economic Council, announced that Indonesia would develop its own rival to DeepSeek and ChatGPT. The justification was national competitiveness: Indonesia, he argued, had no reason to cede this ground to the United States and China. The initiative bore no visible connection to Indonesia's existing National AI Strategy, an instance of what has been characterised as governmental fear of missing out: a reactive response to a viral moment dressed up as industrial policy (Alkaf, 2025).

The gap between this kind of rhetoric and concrete, investable AI policy has consequences. When NVIDIA CEO Jensen Huang visited Jakarta in November 2024, a $200 million AI centre in Solo, Central Java was announced as a joint NVIDIA-Indosat investment. Nearly two years later, no public reporting has documented construction progress at the site. What materialized instead was a showroom with demonstration robots and a photobooth. Three weeks after his Jakarta visit, Huang flew to Hanoi and committed a genuine research and development centre to Vietnam. Indonesia's own

investment agency acknowledged that Malaysia and Vietnam were preferred for major AI investments, citing talent shortages and regulatory complexity (VOI, 2026). Malaysia has attracted roughly four times Indonesia's committed AI and data centre investment despite having one-eighth of the population (Nikkei Asia, 2026).

This pattern of grand announcements without sustained follow-through is not unique to AI. Indonesian foreign policy in the past decade has been described as "reference-point diplomacy," where the production of initiatives and documents substitutes for shaping outcomes (Laksmana, 2024). AI policy is not foreign policy, but the institutional habit carries over.

Indonesia's discourse on digital sovereignty follows a similar logic. In practice, it typically means not relying too heavily on any single technology source: diversifying suppliers, buying from both American and Chinese vendors, hedging between cloud platforms. This amounts to what has been called "sovereignty as a service": US and Chinese firms bundle chips, cloud infrastructure, and AI platforms into packaged deals that nations adopt under the banner of sovereignty while locking themselves into long-term dependencies (Popko, 2025). When a single AI model can discover thousands of critical vulnerabilities in the operating systems that all your vendors run on, diversifying who you buy from protects nothing.

What is striking is the absence of AI from Indonesia's strategic and security conversation. Defence officials have discussed AI only in the narrow context of autonomous weapons and military platforms, acknowledging it as unavoidable while citing precision concerns and funding constraints. BSSN, the national cyber agency, frames AI primarily as a threat vector, focused on phishing and deepfakes, not as a strategic capability reshaping global power. Deputy Minister of Communication and Digital Nezar Patria has spoken about geopolitical factors shaping AI policy and called for "sovereign technology," but his influence on the administration's overall direction remains limited. This silence at the strategic level persists even as restricted NVIDIA AI chips have been rerouted through Indonesian cloud infrastructure, drawing international scrutiny that went publicly unaddressed (Wall Street Journal, 2025).

The most revealing gap, however, lies not in what officials fail to say but in what they say and then fail to act on. In a public presentation on the AI roadmap, Komdigi's Director General for Digital Ecosystem, Edwin Hidayat Abdullah, explicitly cited the World Economic Forum's Global Risk Report and identified geoeconomic confrontation as the number one global risk associated with AI. He walked through Jensen Huang's five-layer cake of AI infrastructure, connecting each layer to geopolitical competition: energy, semiconductors, connectivity, large language models, applications. Then the presentation moved to the actual policy response: ten priority sectors, quick wins tied to Prabowo's flagship programs, a four-pillar framework of institutional collaboration, innovation, talent development, and risk mitigation. The geopolitical diagnosis was acknowledged and then set aside. It shaped no part of the policy architecture that followed.

Within ASEAN, the comparison is unflattering. Singapore has launched 25 AI governance initiatives and Vietnam passed an actual AI law in December 2025. Indonesia has produced one initiative and its presidential regulation still awaits the President's signature as of April 2026. Billions in foreign AI infrastructure investment are flowing into Indonesia, and the regulation meant to govern it has not yet been signed.

What a serious approach would require

Indonesia cannot build its own Mythos. A country with 270 million people, pressing development needs, and limited fiscal space will not compete at the frontier of AI capability, nor should it try. AI can genuinely improve public services, agricultural productivity, and economic inclusion, and the impulse behind the government's education and sovereignty rhetoric reflects real needs. What is missing is not ambition but recognition: the development agenda, pursued in isolation, does not constitute a complete response to what AI has become.

The United States, with vastly more institutional capacity, held a congressional hearing in April 2026 on the theft of AI capabilities by China. The session exposed a tension American policymakers themselves have not resolved: whether AI is fundamentally a security enforcement problem or an innovation leadership problem. If the country that produces most of the world's frontier models cannot settle this question, Indonesia's single-track framing of AI as a development opportunity looks less like a strategic choice than a gap no one has gotten around to filling. Even the most optimistic assessments of AI and cybersecurity assume that defenders have the institutional capacity and resources to act. Indonesia has not built either.

Two shifts would mark the beginning of a more serious approach.

First, the AI presidential regulation currently awaiting Prabowo's signature should close the gap between the government's own geopolitical diagnosis and its policy response. As it stands, the draft framework treats AI as a collection of sectoral applications without distinguishing between development tools and strategic dependencies. The Komdigi Director General's own presentation identified geoeconomic confrontation as the number one global AI risk. The regulation should reflect that diagnosis, not set it aside. The fact that the regulation is still being finalized is precisely why this is the moment to ensure the strategic dimension is not lost.

Second, the administration should recognize that the current mode of international engagement on AI is insufficient. President Prabowo has been to London, Davos, Moscow, and multiple summits. He personally witnessed the Danantara-Arm chip design MoU. He called Jensen Huang during Indonesia AI Day. Yet none of this engagement has been directed toward the governance conversations that will determine who controls the AI infrastructure Indonesia is building its economy on.

In the same week that Indonesia's defence minister met his American counterpart at the Pentagon, three American AI companies quietly formed a private coalition to govern what they call "distillation attacks," a governance decision with global consequences made in three boardrooms. Two days before that, Congress held a hearing on AI chip smuggling that directly implicated Indonesian cloud infrastructure. Indonesia had no presence in any of these spaces, and no apparent awareness that it should. The Danantara-Arm MoU and the summit appearances follow the same pattern the piece has described throughout: transactional engagements focused on investment attraction, not on shaping the frameworks that will govern what those investments mean. Countries that remain outside these conversations risk becoming permanent renters of external infrastructure, locked into the application layer and subject to decisions made without their input (Taylor & Tan, 2025).

None of this requires matching American or Chinese spending. It requires recognizing that the world Indonesia is building its digital economy into has changed, and that the policy architecture meant to navigate it was designed for a world that no longer exists.

Referensi

Alkaf, A. M. (2025, June 25). Indonesia's AI FOMO threatens real progress. *East Asia Forum*. https://eastasiaforum.org/2025/06/25/indonesias-ai-fomo-threatens-real-progress/

Bremmer, I. (2025, May 13). The technopolar paradox: The frightening fusion of tech power and state power. *Foreign Affairs*. https://www.foreignaffairs.com/united-states/technopolar-paradox-ian-bremmer-fusion-tech-state-power

Chesterman, S. (2026). Silicon sovereigns: Artificial intelligence, international law, and the tech-industrial complex. *American Journal of International Law,* 120(1), 44.

Clark, J. (2026, March 27). Anthropic thinks AI might destroy the economy. It's building it anyway [Audio podcast episode]. In D. Thompson (Host), *Plain English*. *The Ringer*. https://www.theringer.com/podcasts/plain-english-with-derek-thompson/2026/03/27/anthropic-thinks-ai-might-destroy-the-economy-its-building-it-anyway

Laksmana, E. A. (2024). Indonesia's reference-point diplomacy decade under Jokowi. *IISS Online Analysis*. https://www.iiss.org/online-analysis/online-analysis/2024/10/indonesias-reference-point-diplomacy-decade-under-jokowi/

Nikkei Asia. (2026, January 6). Malaysia's data center boom: An inside look at Asia's battle for AI supremacy. *KrASIA*. https://kr-asia.com/malaysias-data-center-boom-an-inside-look-at-asias-battle-for-ai-supremacy

Popko, J. (2025, September 12). Chinese and U.S. tech keeps countries dependent on foreign AI. *Rest of World*. https://restofworld.org/2025/chinese-us-tech-foreign-ai-dependence/

Taylor, J., & Tan, J. (2025, April 13). Asia needs an AI third way. *East Asia Forum*. https://eastasiaforum.org/2025/04/13/asia-needs-an-ai-third-way/

VOI. (2026, January 9). Malaysia dan Vietnam lebih diminati untuk investasi AI besar. VOI. https://voi.id/en/economy/550114

Wall Street Journal. (2025, November 12). Chinese startup accessed banned Nvidia chips through Indonesia. *The Wall Street Journal*. https://www.wsj.com/tech/ai/china-ai-nvidia-chip-access-6a4fa63d

Wong, M. (2026, March). Pentagon and Anthropic dispute [Article]. *The Atlantic*. https://www.theatlantic.com/technology/2026/03/pentagon-anthropic-dispute/686307/

Jalan Mabes Hankam No. T65, Bambu Apus, Cilangkap, DKI Jakarta 13890
Email: lab45@lab45.id | Telpon: +62811452045

Silahkan hubungi tim editorial untuk pertanyaan melalui lab45@lab45.id