

From Compute to Access

Inference Sovereignty as the Missing Layer of AI Governance

A measurement framework for the Global South, with case studies from Southeast Asia

Track: Asia · Sub-track: Governance, Policy, and Geopolitics

Global South AI Safety Hackathon, June 2026 · Hub: Ho Chi Minh City

1. Executive summary

Who is really capable of constructing intelligent systems is the central question in the current AI sovereignty debate. In the case of many countries, this issue is already settled. They can no longer manufacture them currently; and do not expect to be able to do so for quite a while. What's perhaps more relevant is who can have reliable, essentially self-providing access to intelligence, whenever they want to use it.

Debate surrounding AI sovereignty remains heavy on ownership of silicon chips alongside various data centers and training kit. Only United States and China can provide all these elements and very few other countries who are capable of doing this. Most countries or states (close to 190) do not develop "frontier" models, so these countries do not develop AI sovereignty at the site where development happened. It happens at an API gateway, the site from which they receive access to intelligence they did not develop.

This paper introduces the idea of inference sovereignty, which refers to a country's ability to have sustained autonomous access to advanced AI capabilities based on reliance on a third-party supplier or an actor who acts on a unilateral basis. The authors additionally propose an Inference Dependence Score (IDS) based on a five-dimensional evaluation to quantify the level of dependence of a state. The paper discusses in detail the IDS score in which the authors apply this score for three countries of Southeast Asia, Vietnam, Indonesia and Singapore. The three countries manage dependency differently. Moreover, they manage dependency through three ways which are Legal-Extraterritorial, Sector-Fragmented, Structural-Normative.

While Vietnam's legally-binding AI regulations were the most legally-aggressive of the three countries in the study, Vietnam had the most dependence measured. In contrast, Singapore was assessed to have a significantly lower reliance rating owing to a voluntary governance model that was backed by more robust technical fallback capabilities and structural market position. The inference from these results is that regulatory bindingness is not a sufficient indicator of sovereignty; equivalently important to law is a state's capacity.

Besides the study results mentioned above, we have also provided a web-based application where a user can enter data regarding any country and calculate its ranking using IDS. In the end, we provide certain practical policies (open weight readiness, API aggregation gateways, local inference mandate, and inference escrow agreement) through which countries from the Global South can start working on their AI sovereignty without having to complete the training phase.

2. Introduction and motivation

Most nations have already long ago settled the matter of which organization could produce intelligence. Therefore, the question that arises is different. The topic of discussion in this paper is the second and definitely even more crucial issue, which involves which organization would be able to gain the necessary intelligence reliably in a certain situation. Figure 1 demonstrates this evolution.




Figure 1. The AI Dependency Stack: The Strategic Bottleneck Moves From Training to Access for Inference.

2.1 Why training infrastructure stopped being the right question

While historically "AI Sovereignty" meant ownership of the inputs necessary for training the models (GPUs, data and architecture design) in a context where there were no technical and institutional barriers to training and deployment, that made some sense. However, in 2026, such a framework holds almost no meaning. Training costs for even a few organizations at the frontier (estimated by GPT-4, ranging from \$78-\$100 million) are out of reach for most. Estimated costs for frontier-level training runs in 2026 fall somewhere between \$200-\$500 million. As Epoch AI itself has noted in its publicly available estimate of the growth rate of training costs, compounded at a rate of around 2.4 times/year since 2016, those will surpass \$1 billion within 2028. At each stage of cost increase, yet another layer of the state becomes impossible as a policy decision for most governments of the world while training could be considered an extremely expensive policy option for most middle powers at \$100 million; at \$1 billion it is not.

2.2 The asymmetry the literature undercounts

Sovereignty in terms of infrastructure for training frontier models, which is the ability to train a frontier model, has become increasingly concentrated in recent years to an extremely limited number of countries; two particular countries, the United States and China, have led in this area, and a handful of others (France, the UAE, Saudi Arabia, and India) have invested in computing power. However, almost all governments, including highly advanced countries like Singapore and South Korea, will obtain access to frontier models in precisely the same way: either by relying on the use of a commercial API that comes from the headquarters of another country or, at best, by the domestic hosting of models trained somewhere else. Sovereignty debates that consider only the small number of governments able to plausibly train a frontier model do not add much value to the discussion.

2.3 What is actually at stake

If a state's critical infrastructure, such as healthcare triage, banking fraud detection, or civil-service case processing, depends on an external Application Programming Interface (API), then it faces the risks associated with those APIs, despite owning just one single chip:

- Geopolitical "kill switches": A state can simply decide to stop offering certain services, because of a shift in foreign policy and/or export controls.
- Muted Model Deprecations: An API provider makes changes or discontinues a particular model version, thereby altering the functionality of a state's system without notifying the state or giving it any options to mitigate the issue.
- Foreign Provider Policy Changes: An API provider can change the pricing terms, the Terms of Service, or the Content Moderation Policies of its service, affecting a dependent state, but not allowing that state to negotiate these changes.

3. Conceptual framework: from infrastructure to inference

Table 1 contrasts the two framings directly.

Dimension	Traditional AI sovereignty	Inference sovereignty
Primary metric	GPUs owned; frontier models trained domestically	API availability, switching options, fine-tuning capacity
Vulnerability focus	Semiconductor supply chain, export controls	API-gateway chokepoints, model deprecation, remote policy change
Strategic goal	Independence in creating intelligence	Resilience and autonomy in applying intelligence

Table 1. Traditional AI sovereignty vs. inference sovereignty.




Figure 2. Inference sovereignty replaces four infrastructure-era questions with five capability-era ones.

4. Methodology: measuring inference dependence

Inference sovereignty is measured through five dimensions, where each dimension is rated from 0 (least risky) to 3 (most risky), and finally all these dimensions are combined to get the total Inference Dependence Score (IDS).

Five dimensions of inference sovereignty include the following:

- Provider Concentration: Dependence on one foreign provider?
- Switching Cost: Technical and Financial costs involved in switching to a different solution?
- Fallback Capability Gap: Availability of viable (domestic/open-weights) solution to substitute a foreign provider in case of loss of access?
- Regulatory Leverage Gap: Ability to make a provider comply with the regulations, or having to request?
- Resiliency Gap: Does an economy have a contingency plan which is tested for maintaining business activity in case of lack of access to critical suppliers?

All these five dimensions of inference sovereignty are given weights and combined to give the IDS.

Formula for IDS calculation is:

$$IDS = 100 \times (\sum w_i \cdot score_i) / (\sum w_i \cdot 3)$$

where w_i is the weight attributed by the analysts to each of the five components. For instance, in case all the five components are weighted equally, the IDS will be ranging between 0 (independence) and 100 (complete dependence). In the event that the IDS is less than 34, the economy's dependence on the

foreign supplier will be regarded as low. An IDS score of more than 66 suggests that the economies dependence on the foreign suppliers is high.

4.1 Scoring each dimension against real-world evidence

It cannot be substituted for such anchoring since in reality each of the dimensions needs to be anchored with a verifiable measure that is observable not only to the public but to the analyst:

- **Provider concentration:** The market share of cloud and AI providers in the relevant country/industry from third party data sources (Synergy Research Group, Gartner, IDC), and an audit of the list of foreign AI providers that register with national digital economy and telecom authorities.
- **Switching costs:** Third party modeled estimates of the difference in performance of the models based on switching between domestic model and alternative model, and if any termination/migration clauses exist in all domestic government procurements of AI solutions.
- **Capability gaps of fall-back options:** Evidence of sovereign computing initiative (GPU cluster), national AI strategy compute requirements, evidence of adoption of open weight models (Llama, Mistral, Qwen etc.), and evidence of credible domestic foundation model initiatives with benchmarks vs the best international alternatives.
- **Regulatory leverage gaps:** Whether the legislation on AI of the country is legally binding or not; whether the legislation on AI of the country is applicable outside the borders of the country for foreign companies who offer services in that country (like Decree 142/2026 of Vietnam); and most importantly, whether there has been an actual enforcement of the law, not just the authorization of its enforcement.
- **Resilience gaps:** The existence of incident reporting framework for the critical sectors where AI is used; plans of the country to keep operating in case of withdrawal of the provider or sanctions; and history of handling previous disruption incidents (cloud outage, changed conditions of service).

The importance of anchoring is seen in the fact that the five dimensions should be re-scored in case of availability of new information. This Index will be a "living document".

4.2 Accompanying tool

The following is how an Interactive Inference Dependence Score (IDS) Tool works according to the Formula. The users can select preloaded country presets or use custom Dimension Scores and Weights to see the results of the Score calculation, Radar Chart, and Risk Tier in real-time.

Live tool: <https://claude.ai/public/artifacts/8b32e2e4-439f-473d-a24c-95035c564139>

The interactive tool allows one to compare different strategies. By keeping constant a country's current scores, the IDS will show two investments that the country can make in order to affect its IDS. These include buying Compute, which changes only the Fallback Capability Dimension, and Securing a National AI Gateway, which changes Switching Cost, Resilience, and Fallback Dimension at once. Figure 3 below shows how the interactive tool works on the example of Vietnam Preset. The Comparison Panel proves the Paper's Policy Argument through numbers instead of mere assertions. Thus, the investment in Compute changes the score from 67 to 63. While a National AI Gateway changes the IDS to 50.




Figure 3. The Inference Dependence Score calculator, shown with Vietnam's preset loaded, including the strategy comparison panel.

4.3 A proposed refinement: the lock-in multiplier

In the linear model, all the dimensions are considered to be additively independent. But there are two dimensions in particular that do not behave independently. For a state experiencing both the dimensions of extreme provider concentration and extreme switching costs, it is not just twice the problem; it is a similar problem, which is even worse because it cannot get out of it—the best available option itself comes with a heavy cost. We present an optional nonlinear extension, which captures the effect in question:

$$M = 1 + \alpha \cdot (PC/3) \cdot (SC/3) \quad IDS_{adj} = \min(100, IDS \times M)$$

where PC and SC refer to the scores of provider-concentration and switching costs respectively, and α is the strength of interaction between the two dimensions (for example, choose $\alpha=0.5$). For cases where none of the dimensions gets a score of zero, M increases from 1 to 1.5.

Application in Vietnam (PC=2.0, SC=2.0): $M = 1.22$, raising the linear IDS of 67 to 82. Application in Indonesia (PC=3.0, SC=2.0): $M = 1.33$, which would raise the adjusted IDS score above 100, and we kept it at 100 to indicate that the near-maximum provider concentration along with the relatively higher switching cost in Indonesia is close to achieving the position of lock-in. We have stuck to using the linear equation for the scores in the table and in the chart for the sake of clarity for our target audience; the multiplier is given as an extra means for those analysts who prefer to see the risk being compounded instead of averaged out.

There are two crucial elements of design in the extension that are worth discussing, as these are the most vulnerable parts of the model. The first one is the nature of the interaction being multiplicative rather than additive because lock-in is a property of previous exposure, rather than additional risk factor: A country having zero prior exposure to foreign providers cannot be 'locked in,' regardless of whatever degree of concentration or switching costs there might be, so M equals 1 if any of PC and SC equals 0; whereas an additive version ($IDS + \alpha \cdot PC \cdot SC$) does not satisfy this requirement, since you add risk to the base index. The second reason we chose the coefficient 0.5 for the same purpose of illustration; not as a result of empirical calibration. There are no publicly available datasets on documented examples of state-level removal or use as a political tool of a provider-country relationship; hence, we did not have anything to calibrate to. Each choice were clearly defined so a reviewer could replace them with another functional form or parameter rather than view the multiplier as a black box.

5. Case studies

5.1 Vietnam: legal-extraterritorial leverage without capability backing

The Vietnamese Law on Artificial Intelligence (No. 134/2025/QH15) came into effect on 1 March 2026, along with its implementing decree, Decree No. 142/2026/ND-CP, on 1 May 2026.¹ The decree is exceptionally proactive in the region, covering all foreign organizations and individuals engaged in business activities related to AI in Vietnam irrespective of their location. The decree divides AI systems into three categories, namely high-risk, medium-risk, and low-risk and mandates the reporting of incidents involving AI systems along with labelling of AI content. The Prime Minister issued an official list of high-risk systems at the time of writing of this paper, and thus far there have been very few, if any, test cases of enforcing these regulations.

¹Law on Artificial Intelligence No. 134/2025/QH15; Decree No. 142/2026/ND-CP. See Tilleke & Gibbins, "A Closer Look at Vietnam's New AI Law," 2026; IAPP, "Vietnam's First Standalone AI Law," 2026.

Unlike Singapore, whose IDS is of 64 (low dependence), Vietnam has IDS of 67, indicating that Vietnam has a moderate dependence on foreign supplier services. While the legal measures employed by Vietnam constitute a much stronger assertion of control over inference based activities than those of Singapore, inability to depend only on domestic sources constitutes the fallback gap and constrains the extent to which Vietnam can become independent from foreign providers. In this regard, legal assertiveness and legal capability constitute two different policy approaches that states can adopt.

The nature of the fallback gap does not entail "there is no Vietnamese frontier model". The frontier models are heavily biased towards English. When applied to the Vietnamese language, as well as other low-resource non-Latin Southeast Asian languages, frontier models tend to have larger token-to-word ratios, latency and reduced accuracy in interpreting local idioms and code-switched sentences. The above observation has been drawn from various tests of the low resource language models.

Consequently, the process of filling the fallback gap will involve much more than training a domestic trillion-parameter foundation model. Instead, low-latency fine tuning of open-weight models like Llama, Qwen and others for Vietnamese in specific high-volume government and corporate use cases such as content moderation, citizen services chatbots and document processing is a realistic capability for a non-training state.

5.2 Indonesia: sectoral fragmentation

As of June 15, 2026, there are no comprehensive laws regulating the usage of AI in Indonesia. There is only one guideline covering ethical uses of AI in Indonesia derived from a Ministerial Circular in April 2023 titled "Ethics of Artificial Intelligence." It is pending the issuance of Presidential Regulation No. 5 of 2019 titled "Safety and Ethics of AI" ² However, there are a number of industries that have some form of "hard" data localization regulations that are compulsory. The banking industry (OJK Reg. 11/2022), the fintech industry (OJK Reg. 3/2024), crypto-asset trading (GR 49/2024) and the healthcare industry (GR 28/2024) all have requirements for having primary and disaster recovery data centers located in-country.³

The result is a hodge-podge of regulations where heavy regulation has been put into place wherever regulators for the financial sector and the health sectors were first to get to the table; but there is a large "voluntary vacuum" with regards to all other uses of AI, especially for any applications using general purpose foundation models.

²See Indonesia Business Post, "Indonesia Prioritizes AI Regulation in 2026," 2026; Pertama Partners, "Indonesia AI Regulations 2026: Complete Compliance Guide," 2026; Global Legal Insights, "AI, Machine Learning & Big Data Laws 2026 — Indonesia," 2026.

³OJK Regulation 11/2022; OJK Regulation 3/2024; Government Regulation 49/2024; Government Regulation 28/2024. See Chambers and Partners, "Data Protection & Privacy 2026 — Indonesia," 2026; HSF Kramer, "AI Tracker: Indonesia," 2026.

5.3 Singapore: structural and normative leverage without a binding general law

Singapore has consciously avoided adopting a mandatory general law on AI and has instead created different versions of a voluntary Model AI Governance Framework for Artificial Intelligence: traditional AI (2020); generative AI (2024) and agenic AI (the first one in the world starting from January 2026), which will be regulated by the Infocomm Media Development Authority (IMDA).⁴

These will work on the top of very stringent binding rules, whenever they need them to be applied compulsorily: The Personal Data Protection Act applies irrespective of where the processing of the data takes place, whereas the Monetary Authority of Singapore Guidelines apply to the financial institutions.⁵

The power which Singapore exercises against foreign providers is structural, and not legislative: The AI Verify testing tool kit and the governance frameworks of Singapore have been used by the regulators all across ASEAN as a non-formalized benchmark for the region. Despite this, or even perhaps due to this, the fact that Singapore does not claim any general mandatory jurisdiction over the foreign providers of the AI, Singapore has the lowest IDS among the three case countries (i.e., 47, moderate reliance) because of the presence of more back up capabilities.

6. Comparative findings

The following is a description of how the countries scored their Inference Dependence Scores (ID S) in relation to their total ID S presented in table one, in table 2 below:

Country	Provider	Switching	Fallback	Leverage	Resilience	Total (IDS)
Singapore	1.5	1.5	1.0	2.0	1.0	47
Vietnam	2.0	2.0	3.0	1.0	2.0	67
Indonesia	3.0	2.0	2.5	2.0	2.0	77

Table 2. Inference Dependence Score components of each country, where equally weighted categories were used (0 = low risk, 3 = high risk).

The scores in the categories mentioned above were given based on a particular quantifiable measurement and not just a subjective measurement since the low scores given in fall back and resilience for Singapore were as a result of AI Verify testing systems and continuity requirements of MAS

⁴Infocomm Media Development Authority, "Singapore Launches New Model AI Governance Framework for Agentic AI," 2026; K&L Gates, "Singapore's New Model AI Governance Framework for Agentic AI (2026)," 2026.

⁵Duane Morris & Selvam, "Singapore's Digital & AI Governance: A Pro-Innovation, Framework-Driven Model," 2026; Pertama Partners, "Singapore AI Governance Framework: What Companies Need to Know," 2026.

for all the financial institutions in Singapore. The biggest gap in Singapore is that of leverage due to the fact that AI regulation is voluntary in the entire country.

Vietnam has received its lowest or strongest score in the leverage gap because of Decree 142/2026 is mandatory even in areas outside Vietnam. But the biggest gap in Vietnam is that of fallback and the gap was given the highest or the weakest rating because of no domestic foundation model projects near the frontier level.

The two highest gaps that were rated were provider concentration and fallback, in that order. It is evident from this that Indonesia has very much depended on hyperscalers for its general purpose AI and also that it has its data localization regulations which are only applicable in the finance and healthcare sectors.




Figure 4. Dimension Scores and Total Inference Dependence Score by Country (Equal Weighting).

First impressions suggest that there is a clear relationship between the degree of regulatory authority, the “assertiveness” of the State, and lower reliance on commercial relationships in order to access services provided abroad. However, such a correlation does not exist. Despite Decree 142/2026 in Vietnam being the most legal ambitious of the three countries studied - both in that it is extraterritorial and it requires incident reporting and content labeling. The IDS score for Vietnam is greater (67) than that of Singapore (47). This is due to the fact that merely having laws which provide for fallback options does not mean that they can replace commercial relationships when implementation is impossible.

Even though Indonesia has a very fragmented regime in relation to the application of sectors, including one which is strong in finance and health and another which is completely absent in general AI, Indonesia's dependency score still stands at 77, which is the highest among the three States. It can thus be argued that having a very fragmented regime of leveraging is even more harmful than both the completely binding and completely voluntary regimes of governance. A fragmented regime of leveraging will give an impression of governance, but fail to take into account some of the biggest dangers involved in uses of general purposes.

Conclusion: while it is essential to enact a legally binding AI statute in order to cut down on inference dependence, it is not enough. Both the separate mechanisms of capability investment and legal assertion need to go hand-in-hand.

7. Policy recommendations

Recommendation 1: National AI gateways

The only possible "lever" for a country without any training data could be either a technological or a financial/legal one. Technological lever consists of creating a sovereign API layer. Such sovereign API layer is analogous to a national payments switch. In such a case, all the important traffic from critical sectors (healthcare, banking, civil services etc) will go through this layer instead of connecting directly to endpoint of a foreign provider.

And here comes the explanation: In case when foreign provider changes its terms of service, unilaterally increases its prices, or blocks access to its data altogether, the API layer could redirect all downstream traffic to other commercial backends, or even to open-source locally hosted weight model without changing all the configuration for every single downstream agencies. Thus, the government could absorb all the volatilities in one single point, API layer, instead of negotiating everything separately with each downstream agency; just like in case when national payment switches were created instead of negotiation of which network cards each bank accepted.

There are two design considerations with regards to the credibility of such gateways, that is, the fact that it does indeed act as a gateway and not just in name only. The first of these is the consideration relating to the Privacy Architecture. In order to enforce the Data Minimization principle for all critical sectors, irrespective of which agency uses it, the gateway has to cleanse all the user supplied Personally Identifiable Information (PII) and mask any fields containing sensitive information in all outgoing connections before the traffic reaches the foreign end point. The second of these design considerations is the trigger condition for failure over at the gateway. The gateway continues to monitor and track the API uptime, response times and where possible, semantic drifts in the model output as a proxy for unannounced model substitution and degradation for all the connected endpoints. In case of any threshold breach, the gateway should automatically failover traffic to a locally hosted Open Weight Fallback Model.

Additional recommendations

- OPEN WEIGHT READINESS: Build sufficient capacity domestically to be able to rapidly test and deploy new open weight models (such as Llama, Mistral, Qwen-Class), and have some form of credible backup in place prior to needing one rather than building one out of necessity.

- **INFERENCE MANDATES FOR LOCAL INFRASTRUCTURE:** Deploy the same thinking behind data residency mandates to the inference layer, in other words mandate that all important sector use case inference happen from infrastructure located on your territory irrespective of where the model itself came from.
- **INFERENCE ESCROWING AGREEMENTS:** Force foreign firms offering their services in critical sectors to escrow the weights of the model and/or any other assurances for future use in case of the provider arbitrarily withdrawing access to its weights.
- **JOINT COLLECTIVE PROCUREMENT & DIVERSIFICATION WITHIN ASEAN JOINTNESS:** Leverage the negotiating power of ASEAN members collectively when it comes to API prices, limits, and assurance of continuing access from the same providers in much the same way as EuroHPC collective bargaining agreements.
- **EXPAND THE REGULATORY JURISDICTION TO THE INFERENCE LAYER AND NOT TO ONLY THE TRAINING PHASE:** The extraterritorial regulation decree of Vietnam is an example that can be replicated by other ASEAN countries in relation to their own jurisdiction, not because of its enforceability but because this is one of the few means by which the non-training country can assert itself.

8. Conclusion

There are some countries which are in a position to control their frontier-scale infrastructure models, but the most prominent ones are China and the US. However, for those countries which are not going to join this elite club in the foreseeable future (case study countries: Vietnam, Indonesia and Singapore) it is more important issue than inference sovereignty. Neither Vietnam nor Indonesia will have its own frontier-scale model any time soon. They will simply keep on using someone else's models.

Here we come to the real question we must ask. We should not worry about getting rid of dependency, which appears to be impossible in the short run. Instead, we should concentrate on how we can assess dependency, negotiate dependency and build resilience of dependency so that our dependency on other countries' models will not stay at the level of unspoken assumption in procurement decision making as it is now. The Inference Dependency Score and its supporting calculator are meant to help with that.

Appendix A: Illustrative extension beyond Southeast Asia

The IDS framework is evaluated for its applicability in three other Global South nations by applying the framework to one nation from each of these three regions – Brazil (Latin America), Kenya (East Africa), and India (South Asia). The ratings are very rough in nature and are presented here as an example of

how the IDS framework can be applied to a new nation. However, these ratings do not reflect the same degree of rigour as in case of rating of Vietnam, Indonesia and Singapore.

IDS framework total

Provider Switching Fallback Leverage Resilience

Country	Provider	Switching	Fallback	Leverage	Resilience	Total (IDS)
India	3.0	2.0	2.0	3.0	2.0	80
Kenya	3.0	2.5	3.0	2.5	2.5	90
Brazil	2.0	2.0	2.5	1.5	2.0	67

Table 3. Illustrative, low-confidence IDS extension to India, Kenya, and Brazil.

Although the figures lack any real predictive capability owing to their very low confidence ratings, there is certainly sufficient direction within them to highlight a number of key insights into the growth of India's AI ecosystem: even though there has been considerable investment by India in building up its digital infrastructure through the IndiaAI initiative, a great deal of the country's commercial AI applications rely on enterprise APIs that belong to firms outside of the nation, and therefore do not have a nationally-binding AI law (compared with sectorial RBI regulations), thereby raising its risk inference immediately. Indeed, this is about the clearest illustration of the central thesis in this paper amongst the six countries being examined. In addition, Kenya comes in last when considering dependency in relation to all six countries that have been analyzed, thus meaning that the country is much more dependent on foreign-owned AI platforms than both Brazil and India. Lastly, Brazil finds itself placed between Vietnam and Kenya, and this is owing to its relatively developed legal system when it comes to the use of human rights-based AI laws (pl 2338/23).

References

Cottier, B. and Rahman, R. et al. "The Rising Costs of Training Frontier AI Models." arXiv:2405.21015, 2024.

Epoch AI. "How Much Does It Cost to Train Frontier AI Models?" epoch.ai, 2024.

Tilleke & Gibbins. "Vietnam Takes Major Step in AI Governance with New Guiding Decree." tilleke.com, 2026.

Tilleke & Gibbins. "A Closer Look at Vietnam's New AI Law: What It Means for AI Businesses." tilleke.com, 2026.

Viet An Law. "Update Decree 142/2026/ND-CP Guiding the Artificial Intelligence Law in Vietnam." vietanlaw.com, 2026.

MLex. “Vietnam Finalizes AI Implementation Decree, Relaxes Incident-Reporting Timelines.” mlex.com, 2026.

Atsumi & Sakai. “Key Highlights of the Decree Implementing the Law on Artificial Intelligence.” aplawjapan.com, 2026.

IAPP. “Vietnam’s First Standalone AI Law: An Overview of Key Provisions, Future Implications.” iapp.org, 2026.

reg-intel.com. “Vietnam AI Law Implementing Decrees Tracker (2026).”

Duane Morris & Selvam. “Singapore’s Digital & AI Governance: A Pro-Innovation, Framework-Driven Model.” blogs.duanemorris.com, 2026.

Infocomm Media Development Authority (IMDA). “Singapore Launches New Model AI Governance Framework for Agentic AI.” imda.gov.sg, 2026.

K&L Gates. “Singapore’s New Model AI Governance Framework for Agentic AI (2026).” klgates.com, 2026.

Pertama Partners. “Singapore AI Governance Framework: What Companies Need to Know.” pertamapartners.com, 2026.

Pertama Partners. “Indonesia AI Regulations 2026: Complete Compliance Guide.” pertamapartners.com, 2026.

Chambers and Partners. “Data Protection & Privacy 2026 — Indonesia.” practiceguides.chambers.com, 2026.

Global Legal Insights. “AI, Machine Learning & Big Data Laws 2026 — Indonesia.” globallegalinsights.com, 2026.

HSF Kramer. “AI Tracker: Indonesia.” hsfkramer.com, 2026.

Indonesia Business Post. “Indonesia Prioritizes AI Regulation in 2026, Says Communications Minister.” indonesiabusinesspost.com, 2026.