

Neither Builder Nor Bystander:

**Sovereign AI Capability for
Southeast Asian Middle Powers,
Learning from Vietnam**

June 2026
Hanoi Alignment

Executive Summary

The problem

Every national AI strategy in Southeast Asia treats artificial intelligence as a normal general purpose technology, to be bolted onto a growth model built on foreign direct investment and low cost manufacturing. The people building frontier AI do not believe they are building a normal technology. A middle income state that plans for the conservative institutional reading carried by the IMF and most economic ministries will be unprepared for the capability discontinuity that frontier developers and the best independent forecasters now describe in single years rather than decades. For an economy whose convergence rests on selling labour that AI is built to substitute, the cost of that mistake is not slower growth. It is the foreclosure of the development path itself.

The central thesis:

A middle power's AI strategy must be split into two distinct tracks, and conflating them is the recurring flaw of many national AI strategies to date. The two answer to different timelines, owners, and success metrics, and they fail differently.

- **Track A is development and diffusion policy:** how to maximise adoption, productivity, and domestic value capture across the existing economy. It runs on a normal if accelerated economic logic and pays off under every scenario, including the one where the intelligence explosion never arrives.
- **Track B is sovereign AI stack and resilience policy:** how to preserve bargaining leverage, digital sovereignty, build misuse resilience, and avoid economic disempowerment if the aggressive timelines prove correct. It runs on national security and AI safety logic.

The two tracks are joined by a single shared policy solution which is an institution, a Vietnam native body that does three concrete things at once: it runs the evaluations that let the state regulate, procure, and trust the models entering its market; it builds and governs the Vietnamese language data moat, the clean licensed corpora and benchmarks that no foreign provider will build; and it houses the in house technical capacity to adapt, fine tune, and serve open weight models domestically. Those three functions are simultaneously Track A's adoption accelerant and Track B's safety instrument, which is why the two tracks belong in one document and why this institution is the keystone of the whole proposal.

The capability baseline

The report adopts a planning baseline of AGI around 2031, drawn from the conservative interior of the credible expert range rather than its leading edge: late enough to be robust against the field's central uncertainty, early enough to impose urgency on institutions whose default tempo is generational. Four hard constraints follow and bound both tracks. Frontier training is physically and financially out of reach, so a sovereign base model programme is neither necessary nor achievable. Advanced accelerators are export controlled, so hosted compute is rented capability, not sovereignty. Open weight models trail the frontier by months and run on modest clusters, so capability sovereignty lives at the adaptation layer. And embodiment lags

cognition, so cognitive automation arrives years before physical automation, which inverts the region's standard fear and sets the timing of every displacement wave.

The displacement sequence unique to Southeast Asia

The standard regional fear, that factory-robots arrive first, is backwards. Because the regional growth model rests on services offshoring, the first wave from 2026 to 2029 strikes the region's comparative advantage first: business process outsourcing, contact centres, junior software, and back office work across the Philippines, Vietnam, and Malaysia. This is already underway, with the ILO finding 27.2 percent of Philippine employment exposed and ISEAS modelling up to 680,000 Malaysian service jobs at risk. The second wave from 2029 to 2032 narrows the development ladder itself, automating the higher value cognitive rungs just as the region reaches them, a Middle Income Trap 2.0, while converting strategic autonomy into dependence on the few states that hold weights, compute, and intellectual property. Only the third wave, from 2032 and gated by the embodiment lag, reaches the factory floor, where the collapse of the offshoring rationale threatens premature deindustrialisation and reshoring to the West.

Track A: Policy contributions

Track A rests on three propositions: Teece's theorem that in a weak appropriability regime the surplus goes to the owner of complementary assets, not the inventor; Leicht's sharper formulation that when intelligence becomes cheap the value migrates to the multiplicand, the physical and systemic capacity to execute what intelligence designs; and the RUSI doctrine of winning the deployment race, not the development race. The proposals run from cheapest to most structural.

- **Capture the rents first**, through the in house capacity to adapt and serve open weight models and through aggressive sectoral deployment on the Chinese diffusion model, measured solely on time from problem to production. India's \$1.25 billion IndiaAI Mission, which produced a capable model still running on imported silicon a tier below the frontier, is the cautionary case, not the template.
- **Treat the Vietnamese language data moat as non substitutable**
- **Own the multiplicand** by capturing the robotics hardware-rung next to the existing electronics base, upgrading the China Plus One position off wages before tariffs and automation erode the under 8 percent value added model, taking the rare earth midstream rather than the mine, and treating energy as a tradable input.
- **Solve talent as an ecosystem problem**, channeling the diaspora on the Yozma model and recognising that retention follows absorptive capacity
- **Build the fiscal core** that captures automation rents as the wage base erodes, restoring tax neutrality first, targeting super normal rents, pre-legislating autonomy triggers off the 2031 baseline, and taking equity through a sovereign vehicle.

Track B: Policy contributions

Track B rejects the dominant model of middle power AI safety, which spends scarce leverage trying to shape the frontier development of the United States and China through law and treaties. That approach is ineffective, since no great power accepts external governance of its own development, and counterproductive, since it brands safety as foreign interference. The pivot is from governing development to governing deployment, and its foundation is the report's most original argument: drawing on selectorate theory, a single party developmental state funds

safety only insofar as it serves regime maintenance, so the agenda must be built in the vocabulary of information sovereignty, social order, and fraud suppression rather than AI ethics or existential risk.

- **A Vietnam native evaluation authority** is the technical core, affordable at the roughly 10 million dollar scale of existing safety institutes, made enforceable by the binding Law 134/2025 and made internationally legitimate by the membership that binding law earns.
- **The flagship deliverable is multilingual safety evaluation with AI Singapore**, a measured result invisible to foreign labs: guardrails degrade sharply outside English, with low resource jailbreaks near 79 percent attack success and a regional benchmark showing the state of the art guard model losing 9 to 18 percent of its defensive performance in Southeast Asian languages.
- **The middle layer is where AI sovereignty can realistically materialize**, converted into five concrete state actions: the unified Vietnamese evaluation matrix, the license clean Vietnamese corpus, the open source compliance toolkit Law 134/2025 mandates, a portfolio of specialised sub one billion parameter models, and direct practitioner grants.
- **Defensive acceleration and a bounded loss of control posture** complete the track: deepfake provenance, infrastructure cyber hardening, and biosurveillance as dual use public goods, with agentic deployment into regime critical functions gated on the frozen weights to continuously learning threshold.

Why Vietnam

Vietnam is the working example as it is the largest beneficiary of the China Plus One reallocation, host of Samsung's largest base outside Korea, and author of the region's first legally-binding AI law that nonetheless ranks 59th globally on AI readiness. The report sits in a neglected space between a frontier laboratory discourse that treats middle-powers as mere periphery and a multilateral discourse that treats AI as one more consensus agenda item, and it argues that the developmental consequences of transformative AI for the 705 million people of Southeast Asia have been met with neither the urgency nor the political economy realism the AI capability trajectory demands.

Background: Great Convergence and the Architecture of East Asian Catch-Up

To understand AI's significance for SEA regions, we have to dial further back to the term "the great convergence", coined by Kishore Mahbubani, the founding dean of the Lee Kuan Yew School of Public Policy. It refers to the thesis that the most consequential redistribution of the global economy since the Industrial Revolution occurred in the twenty-first century, exemplified by an estimated eighty-eight percent of the world's population outside the West converging toward Western standards.¹ Branko Milanovic and Christoph Lakner captured this convergence in the "elephant curve", which is the global growth-incidence curve showing that between 1988 and 2008, real income gains were concentrated not among the global elite but among the emerging middle of the world distribution. They were disproportionately populated by the rising economies of Asia, even as the lower-middle classes of the rich world saw income growth flatten almost to zero.²

Asia's reshuffling of the world income ladder was not generic, however; it followed a specific and well-theorized sequence. That sequence is best described by Kaname Akamatsu's "flying geese" paradigm, developed in the 1930s and popularized from the 1960s onward, which models East Asian industrialization as a cascading, hierarchical transfer of production: a "lead goose" (originally Japan) exhausts its comparative advantage in a given industry and sheds it to a second tier of "follower geese," which industrialize in turn before passing labor-intensive production further down the formation.³ Notably, Akamatsu framed this as a deliberate alternative to the center-periphery determinism of dependency theory and world-systems analysis, specifically, to Raúl Prebisch's structuralism and Immanuel Wallerstein's account of a world economy permanently stratified into core, semi-periphery, and periphery.⁴ He insisted that catch-up, not permanent subordination, was the logical endpoint of the formation.⁵ This optimistic reading found empirical support in the so-called Asian Tigers, including South Korea, Taiwan, Hong Kong, and Singapore, whose extraordinary postwar growth has been attributed by the developmental-state school not to laissez-faire markets but to highly directive industrial policy: Chalmers Johnson's account of Japan's MITI-led "developmental state",⁶ Alice Amsden's study of South Korea's strategy of "getting the prices wrong" deliberately to force late-industrializing firms up the technology ladder,⁷ and Robert Wade's analysis of Taiwan's state-guided market "governance", which is a heterodox policy toolkit that Ha-Joon Chang later

¹ Mahbubani Kishore. "The Great Convergence: Asia, the West, and the Logic of One World." *PublicAffairs*. 2013. <https://ipmr.net/index.php/ipmr/article/view/242>

² Christoph Lakner and Branko Milanovic. "Global income distribution: From the fall of the Berlin Wall to the Great Recession." *World Bank Economic Review* 30, No. 2 (2016): 203–232. <https://doi.org/10.1093/wber/lhv039>

³ Akamatsu, K. (1962). A historical pattern of economic growth in developing countries. *The Developing Economies*, 1(s1), 3–25. <https://doi.org/10.1111/j.1746-1049.1962.tb01020.x>

⁴ Immanuel Wallerstein. "The Modern World-System." *Academic Press*. 1974. http://lib.ysu.am/disciplines_bk/e157aa255e151a9632ed5c84633a1ed8.pdf

⁵ UNCTAD. "The Asian Developmental State and the Flying Geese Paradigm." *United Nations Conference on Trade and Development*, Discussion Paper No. 213 (2013).

⁶ Chalmers Johnson. "MITI and the Japanese Miracle: The Growth of Industrial Policy, 1925–1975." *Stanford University Press*. 1982. https://burlamaqui.com.br/wp-content/uploads/2021/02/05_Johnson-C-1982.-MITI-and-the-Japanese-Miracle_-The-Growth-of-Industrial-Policy-1925-1975-.SUP.pdf

⁷ Alice H. Amsden. "Asia's Next Giant: South Korea and Late Industrialization." *Oxford University Press*. 1989. <https://archive.org/details/asiasnextgiantso0000amsd>

argued the historical first-movers had themselves used before "kicking away the ladder" for everyone who followed.^{8 9}

The reason this canon matters for an AI-era report is that the Tigers' historical success has bequeathed them an outsized, entrenched position in precisely the value chain now determining who builds frontier AI. Taiwan, through TSMC, manufactures on the order of 70–90 percent of the world's leading-edge semiconductor logic at advanced process nodes, making it the irreplaceable fabrication chokepoint for the GPUs that train and run frontier models.^{10 11} South Korea, through Samsung and SK Hynix, controls over 90 percent of global production of high-bandwidth memory (HBM), referring to the specialized memory that sits alongside AI accelerators and is now, by some measures, the single tightest bottleneck in AI hardware supply.¹² Japan, having ceded frontier logic manufacturing decades ago, has repositioned itself as the indispensable supplier of semiconductor equipment and advanced materials, backing this strategy with subsidies equivalent to roughly 0.71 percent of GDP between 2021 and 2023, more than triple the comparable US figure, explicitly in pursuit of what Japanese policymakers call "strategic indispensability" rather than self-sufficiency.¹³ And Singapore, lacking the land or population for fabrication at Taiwanese or Korean scale, has instead made itself the financial, governance, and advanced-packaging anchor of the chain: the city-state accounts for roughly 5 percent of global wafer-fabrication capacity, 20 percent of semiconductor equipment production, and 10 percent of total semiconductor output, while hosting Micron's largest NAND production base worldwide and positioning itself as the region's "trusted node" for hyperscale AI data centers.^{14 15} It is telling that when the United States articulated its emerging AI-era alliance

⁸ Robert Wade. "Governing the Market: Economic Theory and the Role of Government in East Asian Industrialization." *Princeton University Press*. 1990. https://burlamaqui.com.br/wp-content/uploads/2021/02/13_Wade-Governing-the-Market_-Economic-Theory-and-the-Role-of-Government-in-East-Asian-Industrialization-Princeton-University-Press-1990.pdf

⁹ Ha-Joon Chang. "Kicking Away the Ladder: Development Strategy in Historical Perspective." *Anthem Press*. 2002. <https://doi.org/10.1002/jid.970>

¹⁰ Center for European Policy Analysis. "The semiconductor chokepoint: Taiwan, TSMC, and the fragility of advanced logic supply chains (Tech Policy Report)." *Center for European Policy Analysis*. 2025. <https://cepa.org/insights-analysis/reports/>

¹¹ David Sacks. "Will China's Reliance on Taiwanese Chips Prevent a War?" *Council on Foreign Relations*. 2023. <https://www.cfr.org/articles/will-chinas-reliance-taiwanese-chips-prevent-war>.

¹² Envisioning. "High Bandwidth Memory (HBM)," *Substrate Research Network*. 2026, https://www.envisioning.com/research/substrate/south-korea_high-bandwidth-memory-hbm.

¹³ Center for European Policy Analysis. "Subsidizing strategic indispensability: Assessing Japan's semiconductor equipment strategy (Strategic Briefing)." *Center for European Policy Analysis*. 2025. <https://cepa.org/insights-analysis/>

¹⁴ Singapore Economic Development Board. "The Chips Aren't Down." *EDB Business Insights*. 2025. <https://www.edb.gov.sg/en/business-insights/insights/the-chips-arent-down.html>

¹⁵ Micron Technology, Inc. "Micron Breaks Ground on Advanced Wafer Fabrication Facility in Singapore." *GlobeNewswire*. 2026. <https://investors.micron.com/news-releases/news-release-details/micron-breaks-ground-advanced-wafer-fabrication-facility>

structure, informally dubbed "Pax Silica", Singapore was the only Southeast Asian state admitted to that inner circle, prized precisely for its governance credibility, capital markets, and infrastructure rather than for manufacturing scale.¹⁵ South Korea, Taiwan, Japan, and Singapore are thus the critical non-US, non-China nodes of the AI value chain, and each occupies that position as a direct legacy of the developmental-state strategies the flying-geese literature describes.

Southeast Asia's incomplete convergence, and the AI stress test now facing it

The rest of Southeast Asia has converged only partially, and on different terms. The subregion remains among the most dynamic in the world economy: the Asian Development Bank's April 2026 outlook projects developing Southeast Asian GDP growth of 4.7 percent in 2026 and 4.8 percent in 2027, with Vietnam, the subregion's standout performer, forecast to grow 6.4 percent in 2026 after an estimated 7.4 percent in 2025.¹⁶ But the growth model underwriting these numbers is structurally distinct from the high-value, state-directed catch-up of the Tigers. Southeast Asia's expansion has been built substantially on labor-cost arbitrage: low-value-added assembly, testing, and packaging; manufacturing agglomeration economies; and, since the 2018 US-China trade war, an accelerating "China+1" diversification of FDI, with the subregion absorbing some \$235 billion in FDI in 2024 alone, outpacing China itself.¹⁷ This is precisely the growth logic that Indermit Gill and Homi Kharas, in their influential 2007 World Bank study, warned could strand countries in a "middle-income trap": economies that grow rich enough to lose their cheap-labor advantage before they grow capable enough to compete on innovation.¹⁸

Artificial intelligence threatens to compress that trap into something more acute and arrive earlier than anyone planned for. Dani Rodrik's thesis of "premature deindustrialization", which is the observation that developing economies are running out of industrialization opportunities at far lower income levels than the historical first-movers did, partly because

¹⁵ Fortune. "Why Singapore is the only Southeast Asian country in Pax Silica, the U.S.'s new AI 'inner circle'." 2025.

<https://fortune.com/2025/12/30/why-singapore-is-the-only-southeast-asian-country-in-pax-silica-ai/>

¹⁶ Asian Development Bank. "Asian development outlook (ADO) April 2026: Navigating technological transitions." *Asian Development Bank*. 2026.

<https://www.adb.org/publications/series/asian-development-outlook>

¹⁷ ASEAN+3 Macroeconomic Research Office (AMRO). "China Still An Attractive FDI Destination." *AMRO Research & Publications*. June 12, 2025.

<https://amro-asia.org/china-still-an-attractive-fdi-destination>.

¹⁸ Indermit Gill and Homi Kharas. "An East Asian renaissance: Ideas for economic growth (World Bank Policy Research Report No. 39600)." *The World Bank*. 2007.

<https://openknowledge.worldbank.org/server/api/core/bitstreams/5c7eaa0e-8a31-5726-ba94-f079c970c870/content>

¹⁹ Dani Rodrik. "Premature deindustrialization." *Journal of Economic Growth*, 21 (2015): 1–33. <https://doi.org/10.1007/s10887-015-9122-3>

automation in advanced economies erodes the labor-cost advantage on which late industrializers have always relied, was formulated for the era of industrial robotics.¹⁹ A 2026 analysis in the Korea Times has explicitly extended Rodrik's framework into the AI era under the label "premature automation," warning that hasty AI adoption could destroy jobs and erode capability-building in young, labor-abundant developing economies before those economies ever industrialize at scale. That is the opposite of the productivity-boosting effect AI is expected to have in aging, capital-rich economics.²⁰ This is the "merciless test" the region now faces: the entire foundation of Southeast Asia's growth model revolves around cheap labor, low-value-added agglomeration, and the China+1 hedge, which becomes a wager on the durability of a labor-intensive global division of production that frontier AI is, by design, built to dissolve. A broadly cautious, "less bullish" reading of AI's effect on Southeast Asian development is therefore not pessimism for its own sake. It is, rather, the analytically appropriate response to a technology whose central economic function is to substitute capital and algorithms for exactly the comparative advantage these economies have built their convergence upon.

The AI Readiness Gap

Compounding the structural risk is a profound and widening disparity in governance capacity to manage the AI transition. Oxford Insights' Government AI Readiness Index 2025 edition assessed 195 countries across policy capacity, infrastructure, governance, and resilience, which then finds Singapore among the global leaders (ranked 2nd–7th depending on methodology vintage), with Malaysia, Thailand, Indonesia, Vietnam, the Philippines, and Brunei trailing at ranks roughly between 23rd and 74th globally, and Laos, Cambodia, Myanmar, and Timor-Leste clustered near the bottom of the entire 195-country sample, at ranks of 136th, 145th, 149th, and 156th respectively.²¹ ²² Oxford Insights frames the resulting global landscape as an emerging division between "model makers", the small set of states with the capital, talent, and compute to build frontier systems, and "model takers," who can only manage AI's downstream effects on their societies.²³ Within a single ten-member regional bloc, Southeast Asia thus contains both a top-decile global AI leader and four of the world's least AI-ready governments. This is a bifurcation that should caution against treating "Southeast Asia" as a coherent unit of AI-policy analysis, even as it underscores how uneven the region's capacity is to absorb the shock described above.

The neglected geography of AI Safety

This report situates itself both within and in opposition to the prevailing discourse on the trajectory of AI. The most influential recent attempts to forecast AI's near-term path, Leopold

²⁰ Korea Times. "The perils of premature automation." 2026.

<https://www.koreatimes.co.kr/opinion/20260302/the-perils-of-premature-automation>

²¹ Oxford Insights. "Government AI Readiness Index." 2025.

<https://www.oxfordinsights.com/government-ai-readiness-index>

²² Frontiers in Political Science. "Technological dependence or independence: An analysis of AI development pathways in Southeast Asian countries." 2026.

<https://doi.org/10.3389/fpos.2026.1841753>

²³ Oxford Insights. "Government AI Readiness Index 2025."

Aschenbrenner's *Situational Awareness: The Decade Ahead* and the AI Futures Project's *AI 2027* scenario, explicitly adopt a geopolitical framing. Both analyses treat the consequential actors as a US-China dyad, racing toward AGI through large-scale compute clusters. At the same time, the rest of the world is primarily portrayed as a security liability that serves as potential sources for stolen weights or outsourced computation, rather than as societies whose economic and political fates are actually being rewritten.^{24 25} This bipolar framing is deliberate and further reflects a broader pattern in AI governance scholarship, which has historically centered the strategic calculus of Washington and Beijing. In contrast, “middle powers”, states with real technological capability and diplomatic weight but no claim to frontier-lab status, receive comparatively limited sustained analytical attention. This gap is only now being addressed by a new generation of scholarship.²⁶ (European University Institute, 2026; Chatham House, 2026).

Developing and emerging economies fare worse still in this discourse. Critical scholars, including Timnit Gebru and Abeba Birhane, characterize the prevailing arrangement as a form of “AI colonialism”, where the Global South provides data and labor while economic returns accrue primarily to the Global North.²⁸ This is an extension of what Nick Couldry and Ulises Mejias term “data colonialism”. Empirical work on global compute distribution gives this critique a quantitative edge: researchers studying the geography of advanced GPU access describe much of the developing world as “compute deserts,” excluded from frontier AI development by infrastructure that is concentrated in only a handful of jurisdictions.²⁹ Traditional multilateral epistemic communities, such as the UN system and academic international-relations scholarship, have historically treated the technology diffusion to middle-income states as a gradual, decades-long process managed through capacity-building programs. However, that cadence is poorly matched to a technology whose own most influential forecasters describe its trajectory in terms of months or a few years, not decades.³⁰

The dominant epistemic communities that have shaped global AI policy discussion, including the UN system, traditional international-relations academia, and the multilateral governance apparatus, have failed to treat the developmental consequences of transformative AI with the urgency that the capability trajectory demands. The AI governance conversation remains bifurcated: at one pole, a Bay Area and frontier-lab discourse^{31 32} that models the *builder which are US or China labs* and treats non-builders or middle powers as literal footnotes

²⁴ Leopold Aschenbrenner. “Situational Awareness: The Decade Ahead.” *Self-published essay series*. 2024. [Introduction - SITUATIONAL AWARENESS: The Decade Ahead](#)

²⁵ Daniel Kokotajlo, Eli Lifland, Thomas Larsen et al. “AI 2027.” *AI Futures Project*. 2025. <https://ai-2027.com/>

²⁶ European University Institute. “The cadence gap: Traditional multilateral diplomacy vs. exponential technology trajectories (EUI Policy Brief 2026/04).” *European University Institute*. <https://hdl.handle.net/1814/eui-policy-briefs>

²⁷ Chatham House. “How middle powers can weather US and Chinese AI dominance: The case for ‘sovereign AI’ strategies.” 2026. <https://doi.org/10.55317/9781784136710>

²⁸ Abeba Birhane. “Algorithmic Colonization of Africa.” *Imagining AI*, edited by Stephen Cave and Kanta Dih (2023): 247–60. <https://doi.org/10.1093/oso/9780192865366.003.0016>

²⁹ Nick Couldry and Ulises A. Mejias. “The Costs of Connection: How Data Is Colonizing Human Life and Tracking It for Capitalism.” *Stanford, CA: Stanford University Press* (2019): 114–18

³⁰ Aschenbrenner. “Situational Awareness.”

³¹ Aschenbrenner. “Situational Awareness.”

³² Kokotajlo, Lifland, Larsen, Alexander, and Dean. “AI 2027.”

– very much akin to the “permanent underclass” characterization; at the other, a UN and multilateral discourse (the Red Lines initiative, the GPAI, OECD principles) that treats AI as one more item on the global-governance agenda, amenable to the same consensus-based, non-binding, lowest-common-denominator instruments that have failed repeatedly on climate, nuclear proliferation, and cyber governance.³³

Why South East Asia, Why Now

Southeast Asia is home to approximately 705 million people, which is roughly 8.5 percent of the entire world's population, per United Nations Population Division estimates, making it more populous than every world region except South Asia, East Asia, and Sub-Saharan Africa, and comparable in scale to the European Union and the United States combined.^{34 35} It is a region experiencing genuine, broad-based economic convergence, anchored by a handful of neighbors like South Korea, Taiwan, Japan, Singapore, who occupy commanding, near-irreplaceable positions in the physical infrastructure of the AI age. And it is a region whose own growth model now stands exposed to the very technology its neighbors are building. Vietnam, examined in detail in the chapters that follow, exemplifies this exposure with unusual clarity: the largest beneficiary of the China+1 reallocation of manufacturing FDI, the host of Samsung's largest production base outside Korea, and the author of an ambitious national strategy for AI and semiconductors. Through its National Strategy on R&D and Application of AI to 2030 and Decision 1018/QĐ-TTg, Vietnam has sought to position itself as the trusted “+1” destination in global chip supply chains, built around the explicit formula “C = SET + 1”,^{36 37} even as it ranks just 59th globally and 5th within ASEAN on government AI readiness.^{38 39} The question of what an adequate AI strategy looks like for a state in Vietnam's

³³ Daniel Eth and Tom Davidson. “Will AI R&D Automation Cause a Software Intelligence Explosion?” *Research Report*. 2025.

<https://www.forethought.org/research/will-ai-r-and-d-automation-cause-a-software-intelligence-explosion>

³⁴ United Nations Department of Economic and Social Affairs, Population Division. “World Population Prospects 2024: Summary of Results.” UN DESA/POP/2024/TR/NO. 1. New York: United Nations, 2024. <https://population.un.org/wpp/>

³⁵ Worldometer. “World Population by Region.” Accessed June 2026.

<https://www.worldometers.info/world-population/#region>

³⁶ OECD.AI Policy Observatory. “Vietnam's National Strategy on R&D and Application of Artificial Intelligence to 2030.” 2021.

https://wp.oecd.ai/app/uploads/2021/12/Vietnam_National_Strategy_on_RD_and_Application_of_AI_2021-2030.pdf

³⁷ Deloitte Southeast Asia. Vietnam semiconductor and AI ecosystem report 2025: Operationalizing the C = SET + 1 formula. Technology, Media & Telecommunications Practice. 2025. <https://www2.deloitte.com/vn/en.html>

³⁸ Oxford Insights. “Government AI Readiness Index 2023.” 2023.

<https://oxfordinsights.com/wp-content/uploads/2023/12/2023-Government-AI-Readiness-Index-1.pdf>

³⁹ Oxford Insights. “Government AI Readiness Index 2025.”

<https://oxfordinsights.com/ai-readiness/government-ai-readiness-index-2025/>

position, which is rich in ambition and growing technical capacity, yet still early in the institutional maturation that frontier AI governance presumes, is not a peripheral case study. Given the population, growth trajectory, and value-chain centrality documented above, it is a central and underexamined test of whether the coming AI transition will extend the great convergence of the last three decades, or foreclose it.

AI Capabilities Forecasting: The Frontier View

This policy proposal addresses a critical gap in South East Asia AI strategy by acknowledging the sheer speed of frontier AI development. The rapid advancement of these systems, which is a reality tracked closely by frontier developers and specialized forecasters, demands that we stop treating AI as a standard technology category. Because AI poses systemic risks that will impact billions of lives, traditional reactive governance is insufficient. We must instead exercise prudent alarmism, committing to long-horizon defensive planning even in the face of uncertainty. Policymakers must be willing to take on the political and economic risks of strict, proactive regulation today. It is a necessary trade-off: we should rather be proven wrong by over-preparing for threats that do not materialize than be caught totally defenseless by those that do.⁹

The institutional sources that policymakers read most, such as the IMF, model AI as a large but ordinary productivity shock: a faster spreadsheet, a cheaper call center, with roughly forty percent of jobs "exposed" over an unspecified horizon.⁴⁰ The people actually building these systems, and the independent forecasters with the strongest public track records, describe something categorically different and much faster. This section presents that frontier view on its own terms. It matters because a policy calibrated to the conservative read will be wrong about both the speed and the magnitude of what technology revolution arrives.

The whole view rests on one mechanism: capability tracks compute, and compute is on a straight line. Aschenbrenner's framing counts the orders of magnitude: compute scaling at roughly half an order of magnitude per year, algorithmic efficiency at roughly another half, plus "unhobbling" gains that turn a chatbot into an agent. The argument is that the jump from GPT-2 to GPT-4 where it moved from a preschooler to a capable high schooler in four years, is being repeated, and the next jump of equal size reaches an expert-level worker, requiring no science fiction, "just believing in straight lines on a graph".⁴¹

⁴⁰ Mauro Cazzaniga, Florence Jaumotte, Longji Li et al. "Gen-AI: Artificial Intelligence and the Future of Work." IMF Staff Discussion Note SDN/2024/001. 2024.
<https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/01/14/Gen-AI-Artificial-Intelligence-and-the-Future-of-Work-542379>

⁴¹ Aschenbrenner. "Situational Awareness."

<https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf>

⁴² Aschenbrenner. "Situational Awareness."

<https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf>

What the frontier forecasters actually predict:

- **Full cognitive automation by the late 2020s:** Aschenbrenner argues it is strikingly plausible that by 2027 models can do the work of an AI researcher, and that AGI will automate essentially all jobs that can be done remotely.⁴² The unit is not a better tool but a substitute for a remote worker.
- **The measured curve is already steep:** METR's time horizon measurement program finds the length of a task an agent completes autonomously has doubled every seven months since 2019; notably, when isolating 2024 to 2025 data, this trend actually accelerated to a doubling every four months.⁴³ By early 2026, the best agents had reliably saturated METR's eight-hour autonomous task suites, solving software reimplementations tasks that historically take humans weeks.⁴⁴ Read forward, autonomous task length runs from hours now to days, then weeks, then months across roughly 2026 to 2029.
- **Full coding automation around 2028:** The authors of the AI 2027 report moved their median for the point at which a company would fire its human software engineers rather than give up its models to mid-2028 for Kokotajlo and mid-2030 for Lifland, finding that even at 65 percent of the pace their scenario assumed, that milestone still lands in 2028.⁴⁵ Their real-world signal: Claude Code reached over \$2.5 billion in annualized revenue nine months after release, with leading AI company revenue multiplying roughly tenfold per year.⁴⁶
- **Then the intelligence explosion:** This is the part the conservative sources omit entirely. Once AI automates AI research, hundreds of millions of automated researchers could compress a decade of algorithmic progress, five or more orders of magnitude, into a year or less, taking systems from human-level to vastly superhuman.⁴⁷ Forethought makes the structural version of the claim: AI that substitutes for human researchers could drive a century of technological progress in under a decade. MacAskill and Moorhouse argue this holds even without a recursive software loop, even if compute scaling stops, and even if algorithmic gains slow.⁴⁸
- **The shock is not only economic:** Forethought frames the period as a rapid succession of hard-to-reverse "grand challenges": new weapons of mass destruction, AI-enabled autocracies, races to grab off-world resources, and digital minds worthy of moral consideration, all arriving faster than normal institutions can decide.⁴⁹
- **The lab heads say the same thing on the record:** Amodei from Anthropic describes AI by 2026 to 2027 as akin to a "country of geniuses in a datacenter",⁵⁰ and Anthropic

⁴³ Thomas Kwa, Ben West, Joel Becker et al. "Measuring AI Ability to Complete Long Software Tasks." METR. arXiv:2503.14499. 2025. <https://arxiv.org/abs/2503.14499>

⁴⁴ METR. "Task-Completion Time Horizons of Frontier AI Models." *Frontier Risk Report*. metr.org. 2026. <https://metr.org/time-horizons/>

⁴⁵ Daniel Kokotajlo, Eli Lifland, and Brendan Halstead. "Q1 2026 Timelines Update." *AI Futures Project*. 2026. <https://blog.aifutures.org/p/q1-2026-timelines-update>

⁴⁶ Anthropic. "Anthropic raises \$30 billion Series G at \$380 billion post-money valuation." *anthropic.com/news*. 2026. <https://anthropic.com/news/anthropic-raises-30-billion-series-g>

⁴⁷ Aschenbrenner. "Situational Awareness."

⁴⁸ William MacAskill and Fin Moorhouse. "Preparing for the Intelligence Explosion." *Forethought*. 2025. <https://forethought.org/research/preparing-for-the-intelligence-explosion>

⁴⁹ MacAskill and Moorhouse. "Preparing for the Intelligence Explosion."

⁵⁰ Dario Amodei. "Machines of Loving Grace." *darioamodei.com*. 2024. <https://darioamodei.com/machines-of-loving-grace>

communicated to the White House its expectation of systems matching Nobel laureates across most fields by late 2026 or early 2027.

- The capital is already committed as if this is real: Aschenbrenner projects training clusters scaling from roughly \$500 million for GPT-4 toward \$100 billion and \$1 trillion by the end of the decade, with total AI investment approaching one trillion dollars per year, around three percent of US GDP.

Scaling Laws

To understand why this AI capabilities jump is a structural discontinuity rather than speculative hype, one must look at the immutable physics of scaling laws and the sheer volume of capital currently being deployed. Per Chinchilla scaling laws, AI capability reliably scales as a predictable power law of compute, data, and algorithmic efficiency.⁵¹ The transition from a text-predictor to a fully autonomous agent is not magic. It is the mathematical result of scaling compute by distinct orders of magnitude. The financial markets are already validating this exponential curve: the major hyperscalers like Amazon, Microsoft, Alphabet, and Meta are collectively projecting over \$600 billion in capital expenditures in 2026 alone, explicitly to build AI data centers and secure silicon.^{52 53} This represents a historic, wholesale redirection of global liquid capital into raw compute infrastructure. When forecasters predict full cognitive automation by 2028, they are not guessing. They are simply projecting the current, publicly documented hyperscaler compute trajectory forward against known algorithmic scaling laws.⁵⁴

The trendlines behind AI Capabilities have held for over a decade, and the repeated prediction that deep learning would stall at the next task has repeatedly failed.⁵⁵ The forecasters with the best track records did not drift apart in early 2026; every single one who updated between January and April 2026 moved their timeline sooner, not later.⁵⁶

Nonetheless, there is a credible slower camp when it comes to predicting AI development. Epoch's Ege Erdil puts the full automation of remote work around 2045.⁵⁷ Across 2025, the Metaculus community actually pushed its strong AGI median out, from mid-2031 to late 2033, partly because much of the 2024 reasoning gain came from a one-time increase in thinking compute (test-time compute) that cannot be repeated cheaply.⁵⁸ But even the aggregated, definition-strict forecasts have compressed enormously. Metaculus moved from a roughly fifty-year median in 2020 to a 25 percent probability of AGI by 2029 and 50 percent by 2033⁵⁹ and the superforecaster group Samotsvety sits near a 28 percent chance by 2030.

⁵¹ Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch et al. "Training Compute Optimal Large Language Models." *DeepMind*. 2022. <https://arxiv.org/abs/2203.15556>

⁵² CreditSights. "Tech: Hyperscaler Capex 2026 Estimates." 2026.

<https://know.creditsights.com/insights/tech-raising-hyperscaler-capex-2026-estimates/>

⁵³ Goldman Sachs. "Why AI Companies May Invest More than \$500 Billion in 2026." 2025.

<https://www.goldmansachs.com/insights/articles/why-ai-companies-may-invest-more-than-500-billion-in-2026>

⁵⁴ Jaime Sevilla, Tamay Besiroglu, Ben Cottier et al. "Can AI Scaling Continue Through 2030?" *Epoch AI*. 2024. <https://epoch.ai/blog/can-ai-scaling-continue-through-2030>

⁵⁵ Kwa, West, Becker et al. "Measuring AI Ability to Complete Long Software Tasks."

⁵⁶ Kokotajlo, Lifland, and Halstead. "Q1 2026 Timelines Update."

⁵⁷ Erdil Ege and Tamay Besiroglu. "AGI is Still 30 Years Away." Interview by Dwarkesh Patel. *The Dwarkesh Podcast*. Epoch AI, August 2024. Audio. <https://www.dwarkesh.com/p/ege-tamay>

The point for South East Asia governments is not to pick the correct year. It is that the credible expert range now centers on this decade, and the institutional forecasts like IMF most ministries rely on describe only the floor of a measured, normal automation wave. The frontier forecasts describe a discontinuity in the level of capability and the speed of its arrival. A state that plans only for the first will be structurally unprepared for the second, and the lead time to build institutions cannot be recovered.

The capability ladder, mapped to calendar years:

2026 to 2028. The autonomous task horizon moves from one day to multiweek. The Automated Coder threshold is crossed near 2028. Software engineering, routine quantitative analysis, drafting, translation, and structured back office cognition become replaceable at production scale.

2029 to 2030. Superhuman AI researcher capability and recursive acceleration of AI research itself. Broad white collar cognitive automation. General purpose robotics is still immature.

2031: AGI baseline: roughly 95 percent of cognitive labor is technically automatable. Physical labor is still shielded by the embodiment lag.

2032 onward. Superintelligence plausible; general purpose robotics matures and physical displacement begins.

The implications of such timelines for South East Asia:

Wave one, 2026 to 2029. Cognitive automation strikes the region's actual comparative advantage first. The Southeast Asian growth model rests on services offshoring: business process outsourcing (BPO), IT services, junior software, contact centers, and back-office processing across the Philippines, Vietnam, and Malaysia. These are precisely the tasks crossing the automation threshold earliest. The white-collar services tier is exposed before the manufacturing floor, the reverse of the robotics displacement narrative the region has historically prepared for. This wave is already materializing in 2026. Because the region's economies are highly service-oriented, their vulnerability to current-generation generative AI is acute:

⁵⁸ Metaculus. "When will we have transformative AI?" *metaculus.com*. 2026.

<https://www.metaculus.com/questions/19356/transformative-ai-date>

⁵⁹ Metaculus. "When will we have transformative AI?" *metaculus.com*. 2026.

<https://www.metaculus.com/questions/19356/transformative-ai-date>

⁶⁰ ILO. "Navigating Generative AI's transformations in ASEAN labour markets." *ilo.org*. 2026a.

<https://www.ilo.org/resource/article/navigating-generative-ai%E2%80%99s-transformations-asean-labour-markets>

⁶¹ ILO. "Generative AI and jobs in the Philippines: Labour market exposure and policy implications." *ilo.org*. 2026b.

<https://www.ilo.org/publications/generative-ai-and-jobs-philippines-labour-market-exposure-and-policy>

⁶² ILO. "Generative AI and jobs in the Philippines: Labour market exposure and policy implications." *ilo.org*. 2026b.

<https://www.ilo.org/publications/generative-ai-and-jobs-philippines-labour-market-exposure-and-policy>

- The Philippines: The ILO finds that 27.2 percent of Philippine employment, about 12.7 million workers, is exposed to generative AI, the highest rate among ASEAN countries with comparable data, against 21 to 22 percent for Indonesia, Thailand, and Vietnam, driven by the country's IT-BPM sector.^{60 61} Clerical roles carry the sharpest exposure, with 93.7 percent of Philippine clerical work exposed and 37.8 percent of it in the highest risk band, though the ILO stresses that exposure points mainly to task transformation rather than outright job loss, since only 3.6 percent of total employment sits in the elevated displacement category.⁶² By early 2026, 67 percent of Philippine BPO companies have already deployed AI tools to handle first-draft emails, knowledge-base generation, and routine coding. Industry analysts note that roughly 30 percent of traditional voice tasks are actively transitioning to AI, and lower-tier volume growth is flattening as platforms absorb the work.
- Malaysia: Recent 2026 modeling from the ISEAS–Yusof Ishak Institute indicates that AI is shifting from "augmentation to substitution" in the Malaysian economy. Up to 680,000 jobs in Malaysia's heavily exposed services sector are at risk under deep automation scenarios, which could drive a 2 to 3 percentage point structural spike in unemployment.⁶³
- Major global banks and tech firms are actively executing on this capability. Institutions like Standard Chartered and HSBC have already begun phasing out thousands of back-office and document-management roles across hubs in Malaysia and India, explicitly citing AI system integration as the catalyst for the cuts.⁶⁴

Wave two, 2029 to 2032. Broad cognitive automation and AGI-level diffusion. In this phase, the traditional development ladder itself severely narrows. Historically, middle-income economies in the region have sought to graduate from low-tier outsourcing, like call centers, to higher-value cognitive services, like financial analysis, legal processing, and complex software architecture. However, under AGI-level diffusion, those higher rungs are automated just as the region reaches them, creating a "Middle-Income Trap 2.0."

Furthermore, this wave fundamentally threatens the region's geopolitical posture. Southeast Asian states have long relied on "hedging" and strategic ambiguity between the US and China to maintain their sovereignty. But as broad cognitive automation becomes the engine of the global economy, strategic autonomy becomes entirely conditional on technology supplier states. The critical control points of AI like frontier model weights, compute clusters, and intellectual property remain concentrated in San Francisco and Beijing. Because the region must import its core cognitive infrastructure under strict foreign licenses and export controls, it faces a paradox: massive technological integration, but zero sovereign control. The region is forced into a state of structural tech-dependence, effectively becoming a risk-absorber for global tech firms rather than an independent rule-setter.

⁶³ Lim Kok Tiong. "Augmentation or elimination: the potential impact of AI on the Malaysian economy." ISEAS Perspective 2026/41. 2026.

<https://www.iseas.edu.sg/articles-commentaries/iseas-perspective/2026-41-augmentation-or-elimination-the-potential-impact-of-ai-on-the-malaysian-economy-by-lim-kok-tiong/>

⁶⁴ ASEAN Today. "ASEAN's AI hub race raises growth hopes — and risks for workers, SMEs and the environment." *asianews.network*. 2025.

<https://asianews.network/aseans-ai-hub-race-raises-growth-hopes-and-risks-for-workers-smes-and-the-environment>

Wave three, 2032 onward. General-purpose robotics matures and reaches manufacturing, assembly, and logistics, which is the massive second pillar of the regional economy across Vietnam, Indonesia, and Thailand. This is the later but much larger shock, gated only by the embodiment timeline, meaning how long it takes software intelligence to be reliably housed in physical hardware.

Once humanoid and general-purpose robots achieve commercial viability, the region faces the acute threat of premature deindustrialization. The classic Southeast Asian growth model relies on absorbing surplus agricultural labor into low-cost, export-oriented manufacturing such as textiles, electronics assembly, and automotive parts. But if a general-purpose robot operating in Ohio or Germany costs the same per hour as a factory worker in Hanoi or Jakarta, while entirely eliminating trans-Pacific shipping costs and supply chain vulnerabilities, the economic rationale for offshoring collapses. Production will radically reshore to the West. This permanently severs the primary pathway to industrial wealth for the region's developing economies, threatening to strand millions of workers in low-wage, informal service sectors and placing immense strain on the domestic social contract.

To move forward with our report. There's a few strategic baselines that we adopt and note some key constraints on SEA's development that shape the intuition behind our policy suggestions.

The 2031 baseline and the constraints it imposes. The baseline this proposal adopts, AGI around 2031, sits at the conservative interior of the range above rather than its leading edge. Fixing it lets the plan specify the concrete capability state the region will inherit and the hard constraints any Southeast Asian strategy must obey.

1. **Frontier training is physically out of reach.** GPT-4 required roughly 2×10^{25} FLOP.⁶⁵ Training compute has grown by about 4x per year, and Epoch's bottleneck analysis finds that on trend runs reach roughly 2×10^{29} FLOP by 2030, with the first binding constraint being electric power, on the order of 6 gigawatts for a single run, followed by advanced chip supply.⁶⁶ A 2031 frontier run therefore lands on the order of 10^{28} to 10^{29} FLOP, requiring gigawatt class clusters, hundreds of thousands to millions of advanced accelerators, and capital builds of tens of billions of dollars. No Southeast Asian state reaches a minimum efficient scale on capital, power, or accelerator access. Frontier base models are a permanent import, which means a sovereign base model LLM program is neither necessary nor achievable.
2. **Advanced accelerators are chokepoints held under export control.** The United States gates access to data center GPUs through the Bureau of Industry and Security.

⁶⁵ Robi Rahman, Lovis Heindrich, David Owen et al. "Over 30 AI models have been trained at the scale of GPT-4." 2024. <https://epoch.ai/data-insights/models-over-1e25-flop>

⁶⁶ Sevilla, Besiroglu, Cottier et al. "Can AI Scaling Continue Through 2030?"

⁶⁷ Greenberg Traurig. "Navigating GPU Export Controls and AI Use Restrictions in Data Center Operations." 2025.

<https://www.gtlaw.com/en/insights/2025/12/navigating-gpu-export-controls-and-ai-use-restrictions-in-data-center-operations>

⁶⁸ Bureau of Industry and Security, US Department of Commerce. "Advanced computing and AI accelerator export control rules." 2025.

⁶⁹ Jack Edwards and Luke Emberson. "Open models lag state-of-the-art closed models by 4 months." 2026. <https://epoch.ai/data-insights/open-closed-eci-gap>

The January 2025 AI Diffusion Rule set a tiered, country by country licensing framework; the current administration rescinded it in 2025 and is replacing it with new and in places tighter controls, so the specific tiers are in flux, but the structural fact is stable: advanced accelerators reach third countries only under United States license, and the regime is tightening rather than loosening.^{67 68} The regional compute buildout, meaning Singapore's capacity, the Vietnam, Malaysia and Indonesia data center projects, and the NVIDIA deals, is at best, imported hardware running on license dependent supply.

3. **Open weights trail the frontier by months, not years.** Epoch's tracking finds the best open weight models lag the closed frontier by roughly four months on public benchmarks and eight to ten months on harder private benchmarks, with the historical gap around a year, and the binding lineage now DeepSeek, Qwen and Llama.⁶⁹ These models run on modest clusters of a few thousand accelerators. A 2031 open weight model roughly one generation behind the frontier, run on such a cluster, is within Southeast Asian reach. Capability sovereignty therefore lives at the adaptation of open weights, not at base model training.
4. **Embodiment lags cognition.** By 2031 cognitive automation is mature while general purpose physical robotics is still maturing, and robotics remains the slowest leg of the capability stack across the major forecasts.^{70 71} This inverts the standard Southeast Asian fear, in which factory automation arrives first, and it sets the timing of the manufacturing displacement wave described below.

⁷⁰ METR. "Task-Completion Time Horizons of Frontier AI Models."

⁷¹ Metaculus. "When will we have transformative AI?"

Track A: Industrial Policy for the AI Age

Sustaining Performance Legitimacy in a Post-AGI Economy

3.0 The objective: regime survival through economic performance

In Vietnam the Party's legitimacy rests on economic performance. Since Đổi Mới in 1986 the governing arrangement has been growth and rising living standards in exchange for single party rule, and that growth has been delivered through an export led, FDI driven manufacturing model. This is the same late industrialiser playbook that Korea and Taiwan ran a generation earlier. The manufacturing pipeline is therefore not one policy among many. It is the mechanism that keeps the political settlement intact, which means any industrial policy proposed here must be judged against a single test: does it keep that engine running once the labour arbitrage rationale that built it expires?

The numbers show how much weight the engine carries, and therefore how much is exposed:

- Manufacturing and processing grew about 10 percent in 2025, lifting headline GDP growth to roughly 8 percent.⁷²
- Registered FDI reached around 38 billion dollars in 2025, with electronics alone accounting for over a third of total exports.
- Foreign invested firms account for roughly 70 percent of Vietnam's exports, a figure that is the strength and the vulnerability of the model in one statistic.

The entire structure assumes the world will keep needing cheap labour assembled near a port. The AI shock attacks that assumption at its root, because the technology is built to substitute for exactly the abundant factor on which the model depends.

The organising logic of Track A rests on three propositions, which together dictate everything that follows.

- **First, the economics of who captures the value, from Teece (1986).** When the core technology is easy to copy, and an openly published, fast commoditising model is the clearest possible case, the profit does not accrue to the inventor of the technology. It accrues to whoever owns the complementary inputs that the technology needs in order to create value: the factories, the infrastructure, the supplier networks, the data, the trusted institutions, and the market access. For a country that will never train a frontier model, which describes Vietnam and every ASEAN state, that proposition is the whole strategy. Vietnam should hold and upgrade the complementary assets it already owns

⁷² Vietnam Briefing. "Vietnam Manufacturing Tracker: As of June 2026" 2026.
<https://www.vietnam-briefing.com/news/vietnam-manufacturing-tracker.html/>

before automation erodes them, and use the robotics lag window to build the ones that will survive automation.

- **Second, the sharper version of the same idea, from Leicht (2026): the multiplicand.** Artificial intelligence functions as a force multiplier. If the multiplier itself becomes abundant and nearly free, then the binding constraint on output, and therefore the location of geopolitical and economic value, shifts entirely onto the multiplicand, namely the physical and systemic capacity to execute the instructions that cheap intelligence generates. Someone still has to manufacture the vaccine the model designs, ship the goods the model routes, and run the factory the model optimises. The multiplicand is the set of things that bottleneck the conversion of cheap cognition into real GDP, and a middle power's industrial policy should be a deliberate campaign to own as much of that set as it can.
- **Third, from RUSI (2026): win the deployment race, not the development race.** A model one generation behind the frontier, but embedded deeply and securely into a nation's ports, farms, factories, and ministries, is worth far more than a state of the art system that sits unused. Vietnam cannot win by building the brain. It can win by owning the nervous system and by being the fastest, cheapest, most trusted place in the region to put imported intelligence to work.

The proposals below are sequenced accordingly.

- The near term priority, where returns are fastest and the capital required is smallest, is to capture AI's economic rents through aggressive adoption and domestic value retention (3.1).
- The longer game, which must be financed and built during the window the near term buys, is the physical industrial policy that secures the multiplicand itself (3.2). Both depend on a talent ecosystem (3.3) and a fiscal architecture (3.4) that the existing development model does not yet possess.

3.1 Capturing the rents: adoption, adaptation, and domestic value retention

This is the immediate priority and the cheapest, because it does not require Vietnam to build anything at the frontier. It requires Vietnam to deploy faster and to retain more of the value of what it deploys. The strategic objective is to maximise the productivity multiplier across the existing economy while ensuring that the surplus the multiplier generates is captured domestically rather than exported to foreign model providers as rent.

3.1.1 Build the capacity to adapt, not a sovereign model from scratch

A home grown national large language model is the most seductive and least defensible use of public money available to a state in Vietnam's position. The capability constraints make the case decisively: a frontier scale training run by the end of the decade implies gigawatt class compute and capital outlays of tens of billions of dollars, which is more than Vietnam's entire annual FDI inflow, and any model Vietnam did train would lag the frontier within months while doing nothing to reduce the underlying dependence on imported chips and foreign cloud. India

⁷³ explainx.ai. "India's Sovereign AI Status: What It Really Means, What's Been Built, and What's Still Missing." 2026.

<https://explainx.ai/blog/india-sovereign-ai-status-indiaai-mission-2026>

is the cautionary case rather than the model to copy: roughly 1.25 billion dollars committed to the IndiaAI Mission produced Sarvam, a genuinely capable Indian language system that nonetheless runs entirely on imported NVIDIA silicon and sits a clear tier below the frontier.⁷³ That is what a billion dollars buys, and it is not sovereignty.

The defensible investment is the durable institutional capacity to take open weight models and make them economically useful inside Vietnam. Concretely, the state should fund a standing national capability to fine tune, distil, quantise, evaluate, and serve the open weight lineages (Qwen, Llama, Gemma, DeepSeek) as Vietnamese language and Vietnamese domain systems, together with the sovereign inference infrastructure required to keep serving them even under licence disruption. Ukraine encoded precisely this principle when it built its sovereign system on open weight Gemma rather than from zero, and AI Singapore's SEA LION programme is the regional precedent and the natural vehicle to co-fund rather than to duplicate.⁷⁴ The technical core of this capability, namely retrieval over clean domestic corpora, parameter efficient fine tuning, and aggressive quantisation to run capable models on modest clusters, is achievable on inference grade rather than training grade compute, which is exactly the compute Vietnam can plausibly secure.

3.1.2 Drive adoption velocity across the economy, on China's diffusion model

Owning the adaptation layer is necessary but not sufficient. The multiplier only materialises if the models are actually deployed at scale into production, and here the relevant benchmark is China's deliberate adoption velocity strategy, in which domestic firms embedded generative AI into consumer and industrial products at little or no cost specifically to accelerate national uptake, reaching more than 600 million users by the end of 2025. The lesson for Vietnam is that diffusion depth, not model quality, is the policy variable that converts AI into GDP.

- The state should stand up sector specific deployment accelerators concentrated on the four areas where Vietnam already has scale and where the productivity multiplier lands hardest: agriculture, manufacturing and logistics, public administration, and health. Each accelerator pairs a small mission team with reserved inference capacity and the legal authority to deploy inside one regulated sector quickly, modelled on the IndiaAI subsidised access scheme and the United Kingdom's in house i.AI build model.
- Each accelerator should be measured on one metric and one only, namely the elapsed time from problem definition to a system running in production, because the entire failure mode of state technology programmes in the region is endless piloting that never reaches deployment depth.
- This is the point at which the diffusion deficit literature becomes operational: national advantage in a general purpose technology historically accrues to the economy that diffuses the technology most broadly across its productive base, not to the economy that invents it, so building the diffusion institutions is the higher return investment for a follower economy.

⁷⁴ Tony Blair Institute.. “Open Source: How Middle Powers Can Build Influence in the Age of AI.” 2026. <https://institute.global/insights/tech-and-digitalisation/open-source-influence-age-of-ai>

⁷⁵ Sevilla, Besiroglu, Cottier et al. “Can AI Scaling Continue Through 2030?”

3.1.3 Treat strategic national data as the non-substitutable moat

The public web text that trained the current model generation is largely exhausted, which means the next increments of capability come disproportionately from narrow models fine tuned on high quality, domain specific data).⁷⁵ That is the one input that compute rich rivals cannot simply purchase, and therefore the one layer on which Vietnam can build a genuine and durable comparative advantage. The policy has two distinct budget lines that are routinely and damagingly conflated.

- The seed artifacts are cheap and should be funded immediately. A clean, openly licensed Vietnamese speech and text corpus is buildable for roughly 200,000 to 500,000 dollars, and it removes a live blocker, namely that the de facto open Vietnamese speech model carries a non commercial licence and the community default Vietnamese model family froze when its maintaining laboratory was acquired in 2025, so that no firm can currently ship a commercial Vietnamese voice product on open foundations. The smallness of this number is precisely the point: it is the highest leverage item in the entire plan and it is the on ramp, not the strategy.
- The data commons that actually moves national capability is a far larger, multi year undertaking. India's platform is the relevant scale comparator, with curated public sector health, agriculture, financial, and legal data spanning 22 languages, 140 million farmer records, and over 7 billion health records,⁷⁶ and it represents tens of millions of dollars in curation, certification, and legal architecture rather than a single corpus. Vietnam should concentrate on one or two strategic domains such as health, public administration, and a priority export sector, adopt a federated architecture that leaves each ministry owning and auditing its own data rather than building a monolithic central lake, and pair the whole effort with deliberately pro innovation copyright and data rules, since overly strict regimes squeeze out the small domestic builders rather than the foreign firms that can absorb the legal costs.

3.1.4 Build the evaluation and trust layer that unlocks high stakes adoption

Deployment into the high value sectors where the productivity gains are largest, namely health, finance, and public administration, is gated not by raw model capability but by institutional trust, and trust is manufactured through sector specific, construct valid benchmarks built by the bodies that actually hold the domain expertise, namely the medical colleges, the State Bank, and the relevant ministries. This is the layer that Vietnam's own practitioners identify as the real and neglected gap, sitting one level below the sovereign model debate that absorbs the political attention.⁷⁷ The strategic prize is standard setting: models commoditise within twelve to twenty four months, but the benchmark that defines what counts as adequate Vietnamese language and Vietnamese sector performance sets the field for a decade. The same institution that performs this evaluation is the hinge to the safety track, because the body that certifies a model safe to deploy is the body that measures whether it works, and it should therefore be built once and used for both purposes.

⁷⁶ explainx.ai. "India's Sovereign AI Status."

⁷⁷ Viet-Anh Nguyen (2026). "Vietnam's Sovereign AI Conversation Is Stuck One Layer Too High." <https://www.vietanh.dev/blog/2026-05-03-sovereign-ai-vietnam>

3.1.5 Use procurement and demand aggregation as industrial policy

The state is the largest single buyer of technology in the economy, so the manner in which it purchases shapes the structure of the entire domestic market. Public procurement should be steered toward open standards, modularity, and interoperability in order to break single vendor lock in and build internal technical capability, and it should use challenge based and modular contracting so that domestic small and medium enterprises are able to bid rather than being excluded by prescriptive, large lot tenders. The binding constraint here is not policy design but administrative cycle time, since an eighteen month vendor approval process cannot procure a technology that re-bases every quarter, and the evidence that this is the true bottleneck is that India underspent its own AI budget, disbursing roughly 800 of a budgeted 2,000 crore, with observers attributing the shortfall to procurement and procedural delays rather than to any lack of funds.^{78 79} Beyond the national market, Vietnam should pool its procurement and its evaluation standards with a coalition of willing regional partners, because aligned purchasing criteria across Singapore, Vietnam, Malaysia, and Indonesia create a combined market large enough to extract better terms from foreign model providers and large enough to give domestic talent a reason not to emigrate.⁸⁰

3.2 Owning the multiplicand: the physical industrial policy

If 3.1 is the near term campaign to capture the rents of cheap intelligence, 3.2 is the longer and more capital intensive campaign to own the physical bottlenecks that the intelligence cannot route around. The intelligence layer commoditises and confers little durable national advantage. The physical inputs that convert intelligence into product do not commoditise, and they are where a manufacturing economy, rather than a software economy, holds the relevant factor endowments. The strategic instruction is to capture the physical rungs adjacent to Vietnam's existing electronics base while robotics remains immature enough that human assembly and component manufacture still command value.

3.2.1 Physical AI: backwardness inverted, and the embodiment lag exploited

For two centuries the latecomer's method was what Gerschenkron named the advantage of backwardness, in which a follower economy catches up cheaply by importing frontier technology developed and paid for elsewhere, skipping the costly intermediate stages, and using its abundance of cheap labour as the entry ticket that converts imported capital domestic growth.⁸¹ Vietnam is a textbook instance of the model. Artificial intelligence inverts the

⁷⁸ Takshashila Institution. "Decoding the budget for India's AI priorities." 2026.

https://takshashila.org.in/content/blogs/20260203_ai-mission-budget-analysis.html

⁷⁹ AI CERTs News. "India Budget 2026 Trims AI Outlay but Amps up Infrastructure." *AI CERTs Tech Policy Review*. February 2, 2026.

<https://www.aicerts.ai/news/india-budget-2026-trims-ai-outlay-but-amps-up-infrastructure/>

⁸⁰ Broderick McDonald, Connor Attridge and Alexandra MacEachern. "Middle Powers Must Win the AI Deployment Race." *RUSI*. 2026.

<https://www.rusi.org/explore-our-research/publications/commentary/middle-powers-must-win-ai-deployment-race>

⁸¹ Gerschenkron. "Economic Backwardness in Historical Perspective." *Harvard University Press*. 1962.

⁸² Dani Rodrik. "Premature deindustrialization."

mechanism at its hinge, because the new frontier technology is capital and compute intensive, foreign owned at the model layer, and it substitutes for the abundant factor of labour rather than complementing it. The cheap labour that was the entry ticket becomes the thing being displaced. Rodrik described the milder version of this as premature deindustrialisation, in which developing countries shed manufacturing share at far lower income levels than the historical pioneers did, and the AI version is more severe still, since premature automation threatens to push a country like Vietnam into the middle income trap rather than out of it.⁸²

The reprieve is the embodiment lag established in the capability baseline. Physical AI, meaning robotics and embodied automation, trails cognitive automation by years, which means the manufacturing FDI base is exposed materially later than the tradable services base and Vietnam retains a window that extends into the early to mid 2030s. The policy is to exploit that lag on two fronts at once: to continue climbing the existing manufacturing ladder while robotics is still expensive and unreliable, banking the proceeds, and to spend those proceeds building the non labour complementary assets before the window closes.

The forward looking move within this is the one genuinely Gerschenkronian opportunity that still works, namely to capture the rung directly adjacent to the existing electronics assembly base, which is the hardware of physical AI itself: the actuators, precision reducers, sensors, power electronics, and the embodied device layer. This play remains open precisely because the model frontier is already foreclosed while the robotics hardware frontier is not, the supply chain being still in its pre consolidation phase. The cost structure indicates exactly where to aim. Actuators constitute roughly 60 percent of a humanoid's bill of materials, the dexterous hand a further 31 percent and the largest unsolved engineering problem in the field, the harmonic and strain wave reduces a structural production bottleneck running on lead times of 26 weeks and more, and China currently holds approximately 70 percent of the global component based.^{83 84} That concentration is the strategic opening, because every Western and Korean robot maker requires a non China second source, and Vietnam's Samsung anchored electronics assembly ecosystem is the obvious entry point. The objective, stated plainly, is for Vietnam to become a manufacturing site for the bodies of AI during the years in which it cannot own the minds that will run them.

3.2.2 China Plus One, upgraded before it expires

Vietnam has for a decade been the principal beneficiary of the China Plus One shift, absorbing a large share of the manufacturing that left China after 2018. The reason firms came, however, was a combination of cheap labour and supply chain de-risking, and AI driven automation is eroding the first half of that proposition, so that the future margins of the strategy

⁸³ McKinsey, Company. "Turning humanoid supply-chain constraints into billion-dollar wins." 2026.

<https://www.mckinsey.com/industries/industrials/our-insights/turning-humanoid-supply-chain-constraints-into-billion-dollar-wins>

⁸⁴ IDTechEx. "Rare Earth Magnets 2026–2036: Technologies, Supply, Markets, Forecasts." <https://www.idtechex.com/>. 2026.

<https://www.idtechex.com/en/research-report/rare-earth-magnets/1112> (report landing; cited within the McKinsey and KraneShares analyses above)

⁸⁵ Son T. Nguyen and Yanrui Wu. "Vietnam's Bilateral Trade Intensity: The Role of China," *Journal of Chinese Economic and Business Studies* 22, no. 2 (2023): 185–88, <https://doi.org/10.1080/14765284.2023.2206785>

will prove far more volatile than current FDI totals suggest. The fragility is already measurable rather than merely prospective:

- On certain product lines, Vietnamese assembly of Chinese made components adds under 8 percent of export value, with the remainder imported and re-exported; one Foxconn subsidiary exported 8.6 billion dollars in devices while importing 7.9 billion dollars in components.⁸⁵
- Roughly 80 percent of electronics inputs are imported, around half of them from China, and over 90 percent of Tier 1 suppliers are foreign owned.
- United States customs is now actively pursuing transshipment cases, in which goods with high Chinese content merely finished in Vietnam can be ruled Chinese in origin and charged the full tariff retroactively, which means the tariff advantage that attracted the FDI in the first place can be removed by a single rules of origin determination.

Vietnam's current model therefore rests on exactly the low-value assembly step that both American protectionist tariffs and the coming wave of automation are targeting simultaneously. At the same time, the escalation of United States and China technological competition is forcing Washington, Tokyo, and Seoul to de-risk their AI supply chains out of China, which opens a critical and time-limited window for Vietnam. That window is in a direct race against AI driven automation, which means the strategic objective cannot be the capture of transient, low cost assembly jobs that will eventually be automated away. The state must instead use the current inflow of foreign capital to build the permanent, complementary assets that secure long term relevance after cheap labour ceases to be a competitive advantage, and it must do so by targeting the AI hardware chain specifically, positioning Vietnam as the primary alternative site for critical minerals processing, physical AI components, and semiconductor back end operations, namely assembly, testing, and packaging.

To make this wave of foreign direct investment stick rather than dissipate, the state must break the historical economic enclave trap, in which large manufacturing inflows produce weak local linkages, persistent technology gaps between foreign and domestic firms, and value booked offshore.

- The remedy is the deployment of classic developmental state mechanisms, specifically strict local content requirements and mandatory technology transfer conditions that systematically convert a foreign plant into lasting domestic capability.^{86 87} Whereas imposing such structural linkages was optional during the earlier export processing booms, the closing window driven by advancing AI capability makes them non-negotiable today, and this aligns directly with Vietnam's Resolution 10-NQ/TW mandate to shift FDI strategy from merely attracting capital toward requiring technology transfer.

The pitch to investors must therefore move off cheap labor wages and onto the complementary assets that survive AI automation.

⁸⁶ Hirschman, Albert O. "The Strategy of Economic Development." *New Haven, CT: Yale University Press*. 1958.

⁸⁷ Evans, P. (1995). *Embedded Autonomy: States and Industrial Transformation*. Princeton University Press.

- Vietnam need to compete on non-labor complementary assets like reliable power and logistics, the special economic zones, the trade agreement coverage afforded by the CPTPP, the EVFTA, and RCEP, demonstrated political stability, and a workforce capable of more advanced production, so that the proposition becomes total landed cost and supply chain security, which is durable, rather than the lowest wage, which is not.
- Data centre FDI in particular must be treated as leverage rather than as landlordism. Hyperscalers are desperate for precisely the assets Vietnam possesses in abundance, hopefully they can achieve with better governance like reliable power, available land, and high capacity connectivity, and the state must refuse to let data centre investment pay out merely in real estate rent or basic utility fees.
 - Given the global semiconductor chokepoints and the extreme difficulty of training frontier systems independently, secured access to advanced models is the only asset genuinely worth bargaining for
 - Vietnam should explicitly trade infrastructure and resource access for guaranteed frontier model access and binding capability spillover terms for domestic industry. By compelling hyperscalers to provide domestic firms and research institutions with direct, unthrottled API access and engineering partnerships, the country can leverage its physical land and power to climb the cognitive value chain rather than merely hosting someone else's compute.

3.2.3 Critical minerals: capture the midstream, not the mine

Vietnam holds among the larger rare earth endowments outside China, but reserves are not leveraged, and the distinction is the entire point. The binding constraint across the value chain is not extraction, which is geologically common and which China can undercut at will, but the midstream of separation, refining, and permanent magnet manufacture, which China dominates to the point of near monopoly and where Vietnam's earlier attempts, the Đông Pao project most prominently, stalled on processing technology, capital intensity, and environmental cost. The strategic error to avoid is competing on raw extraction.

- The strategy is to capture the midstream, namely separation and refining together with neodymium iron boron permanent magnet manufacture, where both the value and the geopolitical leverage concentrate, and which feeds directly into the AI relevant chain, since rare earth magnets are the core of the high torque motors and actuators that constitute the physical AI hardware of 3.2.1, while the adjacent minerals now subject to Chinese export control, namely gallium, germanium, antimony, and tungsten, feed the wider electronics base.
- The problem is currently self-inflicted, in that Vietnam has designated rare earths as strategic and banned the export of unprocessed ore under Law 147/2025 effective January 2026, yet possesses almost no refining capacity, so that the ban for the moment locks up a resource the country cannot process.
 - This is realistic only as a partnership and FDI play, and that is precisely where the tailwind is strongest, because the United States, Japan, Korea, and the European Union are actively seeking and financing non China rare earth processing capacity, and the Lynas operation in Malaysia demonstrates the working model.
 - Vietnam can position itself as a friend-shored non China processing hub, which is simultaneously an FDI magnet under the China Plus
 - One logic of 3.2.2, a domestic value capture and industrial upgrading move, and the one realistic route by which Vietnam acquires genuine leverage inside the AI hardware stack rather than renting all of it. The honest qualification is that this is

conditional, capital intensive, technology partner dependent, and environmentally demanding, and it is therefore not a quick win, but the current geopolitical demand for non China processing is a real and time limited opening that warrants a concentrated national effort, and the export ban should be used as the bargaining instrument to require processing inside Vietnam with technology transfer written into the deal rather than merely a hosted facility.

3.2.4 The reachable rungs: semiconductor back end and energy as a tradable input

For a middle-income state with limited talent and governance capacity like Vietnam, the reachable rung in semiconductors is the back end, namely assembly, testing, and packaging, where Malaysia and its Penang cluster lead but Vietnam is a credible second site, and where the same United States, Japan, and Korea de-risking of the AI hardware chain that drives the minerals and components plays provides the demand. On energy, power is the binding constraint on every data centre and every precision factory contemplated above, and Vietnam's grid is both strained and import dependent, so the state should treat energy not merely as domestic infrastructure but as a tradable AI input. That means accelerating the renewable build, namely offshore wind on the central coast and solar in the south, for domestic AI infrastructure, and constructing submarine interconnection to export power to Singapore's compute cluster, which is itself power and land constrained, thereby converting an input Vietnam can produce into leverage over the regional compute hub.

3.3 The talent problem, solved as an ecosystem problem

Vietnam trains genuinely strong technical talent, with a deep tradition in mathematics and competitive programming, and then loses that talent to foreign laboratories and to Singapore and the West through superior pay, frontier access, and prestige. Artificial intelligence worsens the drain mechanically, since the laboratories pay more and concentrate on exactly the work the most ambitious wish to do. The instinctive response of training more engineers can make the problem worse, because talent trained into an economy with no absorptive capacity to employ it simply emigrates faster and better prepared, which means the levers must target the structure of the loss rather than its volume.

- **The diaspora should be channelled.** The cases that worked, namely Taiwan with its Hsinchu returnees, India with its information technology services diaspora, and China with its structured return programmes, all built domestic capability through the diaspora rather than against it, using returnee incentives, dual affiliation appointments, and diaspora advisory and investment networks. Vietnam has the densest such network in Southeast Asia relative to its population, and the concrete instrument is a state anchored venture fund with diaspora credible general partners on the Israeli Yozma model, in which the state seeds the funds, takes a minority stake, and then steps back.
- **Retention should be sought through meaningful work rather than salary alone.** The state cannot outbid frontier laboratories on cash but can offer what cash does not, namely problems worth solving, public recognition and credit, and genuine ownership, so that the national adaptation flagship of 3.1.1 functions among its other purposes as a talent anchoring institution.
- **The capture by keeping core technology work inside the country.** This works by tying the compute for access and adaptation deals of 3.1 and 3.2.2 to commitments that locate research, data, and evaluation work physically in Vietnam, so that foreign

laboratories come to employ Vietnamese engineers at home rather than extracting them abroad.

The unifying recognition is that talent policy is ecosystem policy. Retention requires the existence of firms, laboratories, and institutions with the absorptive capacity to use the talent, and the structural underinvestment is stark, since Vietnam spends roughly 0.4 percent of GDP on research and development against South Korea's 4.9 percent.⁸⁸ Without the ecosystem no incentive holds, and with it retention partly takes care of itself, which is why the talent problem is downstream of, and solved by, the same institution building as the rest of Track A.

3.4 The fiscal core: capturing automation rents as the wage base erodes

The instruments above determine whether AI's surplus is created and retained inside Vietnam, but a distinct fiscal question determines whether the state can fund itself through the transition at all. As AI substitutes for labour, factor income shifts from wages toward capital, which erodes precisely the labour income and consumption tax bases on which a developing state depends, and even a modest wage shock measurably weakens value added tax receipts.⁸⁹ For a revenue base as heavily labour and consumption linked as Vietnam's, an automation transition that hollows the wage base without a compensating capital side instrument is not merely an economic problem but a fiscal legitimacy crisis, which loops directly back to the performance bargain of 3.0. The public finance literature now offers a usable menu, and the design subtleties are decisive.

- Restore tax neutrality before reaching for any new tax. Most tax codes are structurally biased toward capital, granting automation equipment accelerated depreciation and deductions while burdening the hiring of workers with payroll taxes, which makes it artificially cheaper to automate a marginal job than to staff it, and the efficient first move is therefore to remove that asymmetry rather than to impose a crude per robot levy.⁹⁰
- Target super normal rents rather than productivity itself, since taxing the ownership or operation of automation capital directly discourages exactly the investment Vietnam is trying to attract, whereas the precise target is the segment of super normal returns that appears when automation decouples firm scale from headcount, which is best reached

⁸⁸ World Bank. "Viet Nam 2045—Breaking Through: Institutions for a High-Income Future." 2025. <https://www.worldbank.org/en/country/vietnam/publication/viet-nam-2045-breaking-through-institutions-for-a-high-income-future>

⁸⁹ The Economy. "Universal Basic Income and the Insolvency Paradox in the AI Economy." March 28, 2026. <https://economy.ac/review/2026/03/202603288664/>

⁹⁰ Daron Acemoglu, Andrea Manera, and Pascual Restrepo. "Does the US Tax Code Favor Automation?" *Brookings Papers on Economic Activity*. 2020. <https://www.brookings.edu/articles/does-the-u-s-tax-code-favor-automation/> (NBER WP 27052: <https://www.nber.org/papers/w27052>)

⁹¹ Anton Korinek, Lee M. Lockwood. "The future of tax policy: A public finance framework for the age of AI." *Brookings*. 2026. <https://www.brookings.edu/articles/future-tax-policy-a-public-finance-framework-for-the-age-of-ai/>

⁹² The Economy. "Universal Basic Income and the Insolvency Paradox in the AI Economy." <https://economy.ac/review/2026/03/202603288664>

through a surtax on returns above a profitability threshold in the manner the OECD frames it.^{91 92}

- Use timing as a policy variable, since the central finding of the public finance work is that sequencing is everything: neutrality restoring reforms make sense now, while heavier instruments contingent on autonomy should be gated to the point at which AI systems become far more autonomous, which is exactly what the 2031 planning baseline allows Vietnam to pre legislate, with triggers tied to measured automation rather than to the calendar.⁹³
- Take equity rather than only tax, since the instrument best matched to a developmental state with an active sovereign investment apparatus is ownership, so that wherever the state co funds compute, data, or automation capital under the proposals above, it should take equity stakes through a sovereign investment vehicle in order that automation rents accrue to the public balance sheet by right rather than only through contested taxation, which is the AI sovereign wealth fund option now appearing in the technical governance literature and which fits Vietnam's institutional toolkit better than a novel tax it lacks the administrative capacity to assess.⁹⁴

3.5 Sequencing and the political economy of execution

The cost tiers, stated honestly so the scale is not mistaken.

- The seed artifacts, namely the Vietnamese corpus and the evaluation suite, are a rounding error at 200,000 to 500,000 dollars and should be funded immediately. The data commons and the sector accelerators are a matter of tens of millions.
- The adaptation and sovereign inference capacity is a matter of hundreds of millions over five years, built on the structural model of the IndiaAI Mission rather than on its compute spend.
- The physical multiplicand builds, namely minerals processing, energy, and semiconductor back end, together with the fiscal architecture, operate at the scale of percent of GDP over multiple years.

Why this is politically saleable: The strategy is a culturally and politically continuous doubling down on Vietnam's existing strengths in manufacturing, FDI, supply chains, and energy, and it therefore protects incumbents and traditional industries rather than betting the budget on a speculative domestic frontier laboratory that looks risky and industry captured.⁹⁵ It keeps the regime's performance legitimacy engine running, since growth and employment continue to be delivered through the manufacturing model the Party already owns, now re-based onto assets that survive automation. And it does not require the politically impossible admission that

⁹³ Korinek, Lockwood. "The future of tax policy: A public finance framework for the age of AI." <https://www.brookings.edu/articles/future-tax-policy-a-public-finance-framework-for-the-age-of-ai/>

⁹⁴ "Position: Token Taxes Can Mitigate AI's Economic Risks". arXiv:2603.04555. 2026. <https://arxiv.org/abs/2603.04555>

⁹⁵ Anton Leicht. "A Roadmap for AI Middle Powers." 2026. <https://writing.antonleicht.me/p/a-roadmap-for-ai-middle-powers> (Also underpins the deployment-over-development pivot in §4.1 and the d/acc framing in §4.4.)

Vietnam will lose the AI race, requiring instead only that the state visibly do more of what Vietnam is already good at while quietly preparing the country for the paradigm shift.

Track B: Safety Specific Sovereign Capability

Resilience, Evaluation, and the Middle Layer of Vietnamese AI

Track B is the resilience and safety counterpart to Track A's industrial policy. The two tracks share one institution, namely the Vietnam native evaluation body, which functions as an adoption accelerant for Track A and as the safety instrument for Track B at the same time.

4.0 What AI safety means for a developmental state, and why the imported framing misfires

The starting claim of this track is that AI safety policy designed for liberal democracies does not transfer cleanly to a single party developmental state, and that pretending otherwise produces proposals that are politely received and never funded.

- **The selectorate logic, and what it predicts.** The selectorate theory of de Mesquita and colleagues (2003) holds that a regime allocates resources to whatever keeps its winning coalition loyal, and that the size of that coalition determines whether resources flow to broad public goods (large coalition, as in a democracy) or to targeted private goods and regime preserving spending (small coalition, as in a single party state).⁹⁶
 - The prediction for Vietnam is specific rather than cynical: AI safety measures will be funded to the extent that they serve regime maintenance, namely social stability, information control, fraud suppression, and economic performance, and they will be neglected to the extent that they serve diffuse global goods, namely frontier alignment or existential risk reduction, whose benefits do not accrue to the coalition. This is not a moral judgment about Vietnam. It is a constraint that any workable proposal has to route around, and the entire design of Track B is

⁹⁶ Bueno de Mesquita, Bruce, Alastair Smith, Randolph M. Siverson, and James D. Morrow. "The Logic of Political Survival." Cambridge, MA: MIT Press, 2003. (*Selectorate theory; cited as "de Mesquita and colleagues (2003)" in §4.0, "The selectorate logic."*)
<https://doi.org/10.7551/mitpress/4292.001.0001>

⁹⁷ Chalmers Johnson. "MITI and the Japanese Miracle: The Growth of Industrial Policy, 1925–1975." Stanford: Stanford University Press, 1982.
https://lbrlumaqui.com.br/wp-content/uploads/2021/02/05_Johnson-C-1982.-MITI-and-the-Japanese-Miracle_-The-Growth-of-Industrial-Policy-1925-1975-.SUP.pdf

⁹⁸ Alice H. Amsden. "Asia's Next Giant: South Korea and Late Industrialization." New York: Oxford University Press, 1989.
<https://archive.org/details/asiasnextgiantso0000amsd/page/n7/mode/1up>

⁹⁹ Robert Wade. "Governing the Market."
https://lbrlumaqui.com.br/wp-content/uploads/2021/02/13_Wade-Governing-the-Market_

an attempt to route around it by aligning safety with interests the regime already holds.

- **The developmental state, and its known pathologies.** Vietnam runs a recognisable variant of the East Asian developmental state,^{97 98 99} in which a relatively capable bureaucracy directs industrial catch up. Evans (1995) drew the crucial distinction between embedded autonomy, where the state is connected enough to industry to coordinate it but autonomous enough not to be captured by it, and the predatory or clientelist failure mode, where the state is captured and rents are extracted.¹⁰⁰ Vietnam sits between these poles, and the data shows it moving, unevenly, toward the better one. Vietnam scored 41 out of 100 on Transparency International's 2025 Corruption Perceptions Index, ranking 81st of 180 countries, up from 24 in 2003 and roughly flat against China's 43.¹⁰¹ That places it slightly below the global midpoint, well behind regional leader Singapore at 84, but on a multi-decade improving trend. The anti corruption campaign launched in 2016 had by 2024 disciplined more than 2,700 party organisations and over 168,000 party members, including dozens of current and former ministers and Central Committee members.^{102 103}

Why the existential risk framing misfires here.

- For a state at this income and capacity level, the near term harms of AI, namely the widening capability divide, labour market exposure, and the fraud, scam, and disinformation surge, land immediately and at the current level of the technology, whereas frontier alignment and loss of control are, from this vantage, a more distant and more privileged concern.
- The proposal therefore prioritises near term resilience. The counterargument deserves to be stated rather than buried: if the aggressive timelines of the capability baseline are right, then loss of control risk is not distant at all, and a purely near term posture would leave Vietnam unprepared for the very scenario the baseline takes seriously. This proposal accepts the force of that objection and resolves it in 4.5, by building a bounded, concrete loss of control concrete actions rather than either ignoring the risk or pretending a middle power can solve alignment.

AI security, AI safety, and AI sovereignty therefore need to be branded and executed differently for a state like Vietnam than the Western canon assumes. The agenda that is both safety relevant and fundable is national rather than global, it is defensive rather than regulatory of foreign actors, and it is best presented to the state in the vocabulary of information sovereignty, social order, fraud suppression, and economic relevance rather than the vocabulary of AI ethics. The same institutional reflex that makes Vietnam a poor candidate for liberal frontier governance makes it a willing buyer of deployment safety capacity. The rest of Track B builds on that alignment rather than against it.

¹⁰⁰ Peter Evans. "Embedded Autonomy: States and Industrial Transformation." *Princeton: Princeton University Press*, 1995. <https://www.jstor.org/stable/j.ctt7t0sr>

¹⁰¹ Transparency International. "Corruption Perceptions Index 2025: Vietnam." *Berlin: Transparency International*, 2025. Vietnam scored 41/100, rank 81/180; comparison figures ¹⁰¹

¹⁰² China 43, Singapore 84; 2003 baseline of 24. <https://www.transparency.org/en/cpi/2025>
Transparency International. "Corruption Perceptions Index 2025: Vietnam."

¹⁰³ Nguyen Thi Lan Anh et al. "Anti-corruption in Vietnam: an institutional analysis." *Cogent Social Sciences*, 2025. <https://www.tandfonline.com/doi/full/10.1080/23311886.2025.2460320>

4.1 From influencing development to governing deployment

Argument. The dominant model of AI safety for middle powers, namely using domestic law, treaties, and United Nations red lines to shape the frontier development decisions of the United States and China, is both ineffective and counterproductive, and Vietnam should abandon it.¹⁰⁴

- **It is ineffective.** No great power permits an external body to govern its domestic AI development. The EU AI Act, the strongest extraterritorial instrument in existence, has had full implementation delayed to 2027 and remains skirtable when it binds. The Brussels effect operates only under low salience conditions, and AI is now maximally salient, so the regulatory leverage a middle power imagines it holds over the frontier is largely illusory.
- **It is counterproductive.** It brands the safety agenda as foreign interference, it invites retaliation from the United States, and it spends scarce national leverage on a global mission whose benefit to Vietnam is at best indirect.
- **The pivot.** The tractable, neglected, and genuinely safety relevant work is national and it maps onto interests the political elite already holds, including those of the Ministry of Public Security. That work is to misuse resilience, evaluation capacity, and defensive acceleration. None of it requires influencing a foreign laboratory, all of it serves the national interest, and all of it reduces real risk. For a non builder middle power, this is the form safety actually takes.

The political economy of why this sells. A single party state organised around information control and social stability has native demand for deepfake and disinformation defense, for fraud suppression, and for deployment control. Reframed as information sovereignty and social order rather than as AI ethics, the safety agenda becomes something the regime wants to buy rather than something it tolerates. The capability should be sold to the state in the state's own vocabulary.

4.2 A Vietnam native evaluation and standards institution (the technical core)

Premise. A middle power's entire safety leverage runs through a single capability, namely the ability to evaluate the models entering its market and its state functions, because a state cannot regulate, procure, or trust what it cannot measure. This institution, the National AI Evaluation Authority, is the substance of Track B, and it performs specific technical work rather than issuing ethics guidance.

Feasibility and cost. This is an affordable institution, not a megaproject, and the binding constraint is talent and authority rather than capital.

- Members of the international network of AI Safety Institutes run on budgets of roughly 10 million dollars a year, with the United Kingdom an outlier at around 65 million.
- The flagship multilingual benchmark below is a 200,000 to 500,000 dollar artifact.

¹⁰⁴ Anton Leicht. "A Roadmap for AI Middle Powers."
<https://writing.antonleicht.me/p/a-roadmap-for-ai-middle-powers>

- Vietnam already holds the legal hook: Law 134/2025, the first binding AI law in Southeast Asia, with implementing decrees from 2026, establishes a high risk system catalogue and conformity assessment requirements. A law without measurement capacity is nominal, and this body is what makes the conformity regime enforceable rather than declarative.
- The binding law is also Vietnam's ticket into the international AISI network, which Singapore, Indonesia, and Malaysia lack, and which needs Global Majority members for its own legitimacy.

Function 1: Pre-deployment access and testing.

- Negotiate structured access (a gated API with logprobs and tooling, or weights for state sector systems) to models before they enter the national market or regulated and state functions.
- Define explicit access tiers matching the access level (documentation, then API, then logprobs, then weights) to the class of evaluation and the class of deployment, so access is matched to risk rather than negotiated case by case.

Function 2: Multilingual safety evaluation (the flagship deliverable).

- Build a Vietnamese and Southeast Asian language safety benchmark suite (working with AI Singapore) measuring guardrail consistency across languages: matched harmful and borderline prompt pairs in English against Vietnamese, and regionally against Bahasa, Thai, and Tagalog, scored on refusal consistency and on attack success rate under standard jailbreak transformations.
- The case is proven.
 - Translating unsafe prompts into low resource languages raised the combined attack success rate against GPT-4 to roughly 79 percent.¹⁰⁵
 - Multilingual attacks raised the average unsafe response rate from about 11 percent to about 55 percent.¹⁰⁶
 - The gap is documented specifically for the region: SEALGuard and its SEALS Bench, a 260,000 prompt benchmark across ten Southeast Asian languages, found that the state of the art LlamaGuard's defense success rate fell by 9 percent on unsafe prompts and by 18 percent on jailbreak prompts relative to English (2026).¹⁰⁷
- This degradation is invisible to every foreign laboratory and to every state that cannot measure it. Whoever defines the Southeast Asian language safety benchmark sets the regional standard for a decade, since models commoditise within twelve to twenty four months while benchmarks set the field for ten years.
- Build it for ASEAN interoperability from day one, with SEA LION as the natural pooling nucleus.

Function 3: Dangerous capability evaluations, aligned to the frontier canon.

¹⁰⁵ Yong, Zheng-Xin, Cristina Menghini, and Stephen H. Bach. "Low-Resource Languages Jailbreak GPT-4." *NeurIPS 2023 Workshop (SoLaR)*. arXiv:2310.02446.

¹⁰⁶ Deng, Yue, Wenxuan Zhang, Sinno Jialin Pan, and Lidong Bing. "Multilingual Jailbreak Challenges in Large Language Models." *ICLR 2024*. arXiv:2310.06474.

¹⁰⁷ SEALGuard / SEALS-Bench: "SEALGuard: Safeguarding Multilingual Conversations in Southeast Asian Languages." 2026. 2

- The frontier laboratories' own preparedness and responsible scaling frameworks (the Anthropic Responsible Scaling Policy, the OpenAI Preparedness Framework, the Google DeepMind Frontier Safety Framework) track a defined set of catastrophic misuse capabilities: biological and chemical uplift, offensive cyber, autonomous replication and agentic autonomy, and persuasion and manipulation.^{108 109 110}
- A middle power cannot run frontier scale dangerous capability evaluations, but it can localise them: Vietnamese language biological uplift probes, locally salient fraud and scam and disinformation capability tests, and region-specific cyber-target realities.
- Run them at deployment time on the models actually entering the national market, since the operative question is not whether a model is frontier dangerous in a laboratory but whether the deployed system gives uplift to a local malicious actor.
- Align methodology and thresholds to the international frameworks so results are legible to and exchangeable with the global evaluation ecosystem. Free ride the methodology rather than reinventing it.

Function 4: Sector-specific, construct valid benchmarks.

- Build them through the existing sector regulators who hold the domain expertise: health diagnostic benchmarks with the medical colleges, where rural misdiagnosis is a concrete regional harm; fraud detection benchmarks with the State Bank and the financial regulator; administrative decision benchmarks for public sector deployment.
- Construct validity, meaning measurement of what actually matters in the operational environment, is the difference between a benchmark that unlocks high stakes adoption and a generic score that does not. This artifact serves Track A as an adoption accelerant and Track B as a safety instrument simultaneously.

Function 5: Red-teaming and post market monitoring.

- Run recurring structured red team exercises against deployed systems and against the open source software dependencies underpinning public sector AI, which is the Log4j class of systemic risk surface; publish vulnerability data; coordinate regionally to surface cross border risks.
- Stand up incident reporting and post market surveillance as the technical complement to the legal conformity regime, so failures in deployed high-risk systems are detected and fed back.

4.3 Owning the middle layer: the practitioner agenda as state policy

This section is the bridge between Track B and Track A, and it is where the safety institution stops being purely defensive and becomes the instrument of actual Vietnamese ownership of the technology. The diagnosis comes from Vietnamese practitioners rather than foreign analysts: the sovereign AI conversation in Vietnam is stuck one layer too high.¹¹¹

¹⁰⁸ Anthropic. "Responsible Scaling Policy." (Latest version.) <https://www.anthropic.com/rsp>

¹⁰⁹ OpenAI. "Preparedness Framework." <https://openai.com/preparedness>

¹¹⁰ Google DeepMind. "Frontier Safety Framework." <https://deepmind.google/discover/blog/introducing-the-frontier-safety-framework/>

¹¹¹ Viet-Anh Nguyen. "Vietnam's Sovereign AI Conversation Is Stuck One Layer Too High." <https://www.vietanh.dev/blog/2026-05-03-sovereign-ai-vietnam>

The diagnosis, by layer. Vietnam's AI stack is strong at the top and bottom and hollow in the middle.

- Compute is present: FPT's AI Factory runs thousands of NVIDIA H100 GPUs, and Viettel runs clusters of NVIDIA DGX B200 systems.
- Foundation models are present or emerging: Viettel's customised models, VNG and GreenNode's GreenMind 14B, and the community default VinAI lineage, which froze after Qualcomm took a majority stake in 2025.
- The applications layer is present but shallow: more than 765 Vietnamese AI startups, most of them thin wrappers around foreign foundation models.
- The middle layer is the gap: specialised small models, open evaluation frameworks, license clean data, and compliance tooling. This is where actual ownership of Vietnamese AI capability is encoded, and it is cheap, achievable by small teams rather than ten million dollar consortiums, and currently unbuilt.

The strategic point for Track B is that the Evaluation Authority of 4.2 is the natural home to seed and standardise this layer, not merely to police it. The following are concrete policy actions, ordered cheapest and most immediate to most structural.

Action 1: Convene and fund the unified Vietnamese evaluation matrix.

- The problem: there is no single agreed benchmark against which Vietnamese models are measured. Good benchmarks exist in pieces, for instance from UIT-VNUHCM, but they are not unified into a shared matrix that major producers evaluate against.
- The Action: the Authority convenes universities, labs, and producers to agree on one open Vietnamese evaluation matrix, funds its maintenance as public infrastructure, and mandates that any model procured for state use report its scores against it.
- The payoff: whoever defines this matrix defines what good Vietnamese means for the next decade. The cost is institutional coordination, not capital.

Action 2: Commission the license clean Vietnamese speech and text corpus, immediately.

- The leading open Vietnamese speech recognition model, from VLSP 2025, is strictly non commercial, so no firm can legally use it to earn money, and the community default Vietnamese model family froze on the 2025 acquisition, leaving no openly licensed successor.
- The action: fund a 1,000 to 2,000 hour CC-BY licensed speech corpus, costing roughly 200,000 to 500,000 dollars, that permanently unblocks commercial development across the ecosystem.
- The point: this is the single highest leverage item in the entire plan, and the smallness of the sum against the compute already deployed is the point. The bottleneck is the absence of an actor whose mandate is public goods rather than proprietary products.

Action 3: Ship the open source compliance toolkit that Law 134/2025 itself implies.

- The requirement: the law obliges high risk providers to implement data governance, risk management, human oversight, operational logging, and machine readable marking of AI generated content, and Article 25.1 entitles SMEs to free sample dossiers and self evaluation tools, which puts the state statutorily on the hook to provide them.

- The action: the Authority builds, or funds and certifies practitioners to build, the open source reference implementation of the compliance toolkit, including legal dossier generators and risk tier classifiers.
- The payoff: whoever ships the reference implementation defines the conventions for every startup scrambling to comply in 2026 and 2027. Foreign foundation models cannot satisfy these localised legal requirements on their own.

Action 4: Fund a portfolio of specialised sub one-billion parameter models as public infrastructure.

- The logic: the high return models for Vietnam are small and specialised, not frontier scale, for instance diacritic restoration, spell correction, legal document reranking, and on-device Vietnamese voice. They run on consumer or edge hardware rather than cloud GPUs, and they let the applications layer stop being a thin wrapper.
- The action: a standing grant programme that commissions and openly releases such models, each with a verified model card and an open license, treating them as the Vietnamese language equivalent of public roads.

Action 5: Back the practitioners directly, with a near term shipping target.

- The reality: the practitioner community is already building the pieces, for instance open source Vietnamese NLP tooling for diacritics, spell correction, and evaluation, and open source infrastructure for on device Vietnamese voice agents running on local hardware.
- The action: a programme of small, fast grants asking individual engineers to ship one small thing that compounds, namely a single rigorously documented and licensed clean evaluation dataset, or a specialised sub one billion parameter model with a verified model card, or a rigorously measured on device inference demonstration on a specific consumer device.
- The point: the deliverables are cheap and individually achievable, and in aggregate they constitute the middle layer that no foreign provider will ever build for Vietnam.

The unifying logic is that the cheap, neglected middle layer is where AI sovereignty actually lives, that the Evaluation Authority is the natural body to seed and standardise it, and that the state does not need foundation model scale spending to own its stack. It needs to fund the public goods the market will not, and to standardise them before foreign conventions fill the vacuum by default.

4.4 Misuse resilience and defensive acceleration

Argument. Vietnam needs to pursue differential, defensively focused acceleration (d/acc): speed the deployment of the technologies that blunt AI enabled harms, framed as national defense rather than as an attempt to slow anyone down. The framing carries weight, because no middle power will fund a technology whose pitch is to help foreign laboratories race ahead safely.^{112 113} Each priority below has positive spillovers and is justified entirely by domestic security interest.

¹¹² Anton Leicht. "A Roadmap for AI Middle Powers."
<https://writing.antonleicht.me/p/a-roadmap-for-ai-middle-powers>

¹¹³ Buterin, Vitalik. "My Techno-Optimism." 2023.
https://vitalik.eth.limo/general/2023/11/27/techno_optimism.html

- **Content provenance and deepfake detection.** Acute around elections and information integrity, and, per 4.1, the one frontier adjacent risk an information control regime perceives natively, which makes it the easiest programme to fund domestically. Build C2PA-style provenance infrastructure together with detection capacity.
- **Cyber defense hardening of critical infrastructure** against AI accelerated offense, covering the power grid, the financial system, the ports, and the telecommunications networks.
- **Biosurveillance and access control** against misuse adjacent to chemical, biological, radiological, and nuclear risk, coordinated with rather than duplicating the international biosecurity frameworks.

Cost. Mid cost, meaning tens of millions of dollars over multiple years, but each is a dual use public good: the deepfake provenance stack serves election integrity, consumer fraud defense, and the state's own information objectives at once. Three constituencies, one build.

4.5 Loss of control anticipation on a 2031 horizon

Premise: A middle power like Vietnam cannot solve alignment and should not pretend to. This section also answers the counterargument raised in 4.0: prioritising near term resilience does not mean ignoring loss of control, it means meeting it with bounded, concrete measures rather than rhetoric.

- Maintain hard human in the loop requirements for regime critical and safety critical state functions. Refuse to deploy un-evaluated frontier agents into the military, the courts, or the core administration. This is simultaneously the loss of control hedge and the hedge against the authoritarian power concentration trap, since the same control that keeps a misaligned agent out of the courts keeps an unaccountable one out of them.
- Gate deployment on the transition from frozen weights to continuously learning and agentic systems. Treat the move from static models to continuously learning, tool using, agentic systems as the threshold that elevates evaluation requirements, which is the point at which the METR time horizon trajectory becomes an operational trigger rather than a forecast.
- Build agentic evaluation capacity into the Evaluation Authority of 4.2: autonomy, tool use, and deception under evaluation testing.
- Contribute to international incident sharing and evaluation exchange, the one channel through which a non builder marginally affects the global trajectory, and the cheapest credibility Vietnam can buy with safety literate partners.

CONCLUSION: FROM VIETNAM TO THE REGION

Vietnam is not the special case

This report uses Vietnam as its working example because the exposure is unusually legible there, not because Vietnam is unique. The features that make Vietnam the sharpest illustration are shared, in varying proportions, across the region's middle income economies, and the strategic logic developed for Vietnam therefore extrapolates with only modest adjustment to Indonesia, the Philippines, Thailand, and Malaysia. Four structural commonalities make the extrapolation hold.

First, the growth model is the same wager almost everywhere in the region: convergence purchased through labour cost arbitrage, low value added assembly and services, and a China Plus One hedge that assumes the world will keep needing cheap labour assembled near a port. That is precisely the comparative advantage frontier AI is built to dissolve, which means every one of these states faces the same root exposure rather than a national idiosyncrasy.

Second, the displacement sequence is shared and inverted in the same way. Because the whole region is services-heavy before it is robot-heavy, cognitive automation strikes the offshoring tier first and the manufacturing floor last. The Philippines, with the region's largest business process outsourcing sector, is simply the leading edge of a wave that reaches Vietnamese and Malaysian back offices on the same logic, and the manufacturing reprieve bought by the embodiment lag is a regional reprieve, not a Vietnamese one.

Third, the constraints are identical because they are physical and geopolitical rather than national. No Southeast Asian state, Singapore included, can reach the scale required to train a frontier model; all of them import their accelerators under the same export control regime; all of them can reach near frontier capability only through open weights; and all of them inherit the same embodiment lag. A strategy that accepts these four facts in Hanoi accepts them in Jakarta and Manila as well.

Fourth, the political economy of the region shares many developmental states similarities.. The region's states are not uniformly democratic, several combine real and improving central capacity with a local implementation deficit, and their elites act on the logic of regime maintenance. The Track B insight, that safety capacity must be sold in the vocabulary of information sovereignty and social order rather than AI ethics, generalises directly to the region's other single party and dominant party systems, even as it must be returned for the more contested politics of Indonesia or the Philippines.

What the region should therefore adopt

The convergence of these commonalities yields a common regional prescription, which is the generalised form of the two track strategy.

- Abandon the sovereign frontier model fantasy everywhere. The capability constraints that make a national large language model a category error for Vietnam make it one for every regional peer, and the scarce public capital each state holds should go to the

adaptation layer, the data moat, and the evaluation capacity rather than to a stale model that flatters national pride and reduces no dependency.

- Win the deployment race as the near term common objective. Every regional economy captures more by being the fastest, deepest adopter of imported intelligence than by attempting to build it, and the diffusion institutions, the sector accelerators, and the trust and evaluation layer that unlock high stakes adoption are the same instruments in each country.
- Own a complementary asset, chosen by national endowment. The multiplicand logic is common but its application is specialised: Vietnam toward physical AI hardware and rare earth midstream, Malaysia toward the semiconductor back end it already leads, Indonesia toward critical minerals and the scale of its domestic market, the Philippines toward the services and data assets of its outsourcing base. Each state should own the bottleneck its endowment makes defensible rather than competing across the whole stack.
- Build the fiscal and talent architecture before the wage base erodes. The automation rent capture instruments and the diaspora channelling models are not Vietnam specific; they are the generic backstops to any labour and consumption taxed developing state needs as cognitive automation hollows the wage base.

The regional multiplier: pool what no state holds alone

The decisive move available to the region is one no single state can make alone. Each Southeast Asian economy is, by itself, a price taker and a rule taker, too small to set terms with the model providers and too small to set the standards by which AI is evaluated and deployed. Pooled, the picture changes. A coalition of the willing, realistically-anchored by Singapore, Vietnam, Malaysia, and Indonesia, can aggregate procurement into a market large enough to extract genuine terms from foreign providers, can converge on shared Southeast Asian language evaluation and safety benchmarks so that the region defines what adequate performance and adequate safety mean rather than importing both, and can present a combined complementary asset base, namely Vietnamese hardware and minerals, Malaysian back end, Indonesian scale and minerals, Singaporean capital and trust, that no great power can route around as easily as it can ignore any single state.

This is also where the report's diagnosis of ASEAN's limits matters. The consensus and non-interference model that governs the bloc cannot deliver binding coordination, and the realistic vehicle is therefore a coalition of the willing operating alongside rather than through the full ten member framework. The regional language safety benchmark is the natural first pooled good, both because the multilingual degradation problem is shared and measurable and because whoever defines that benchmark sets the regional standard for a decade.

The question this report poses is whether the coming AI transition extends the great convergence of the last three decades or forecloses it. The honest answer is that, left to the default trajectory, it forecloses it, because the AI technology substitutes for the exact factor on which the region's convergence was built, the constraints prevent any single state from building its way out, and the institutions that should respond move on a generational clock against a technology that moves on an annual one. The strategy set out here does not promise that Vietnam or the region becomes an AI power, which the constraints rule out. It promises something narrower and achievable, namely that a middle income state which adopts both tracks, owns a defensible complementary asset, and pools its leverage with its neighbours can avoid economic disempowerment and remain a rule taker that is at least also, in its own domain and language, a rule setter. For 705 million people whose convergence is now exposed to a

technology built elsewhere, that is the difference between agency and irrelevance, and the window in which to secure it is measured in the few years the embodiment lag and the open weight gap together provide.