PEREIRA INTERNATIONAL

ARTIFICIAL INTELLIGENCE AND SOUTHEAST ASIA: AN ECONOMIC REVIEW

MAY 2026

I. AI IS CHANGING THE GLOBAL & REGIONAL ECONOMIES

The economic significance of artificial intelligence lies not in any single application but in how it alters the cost structure of knowledge work, affecting a broad range of sectors, from medical research to consumer electronics. AI systems function as scalable engines for search, synthesis, and analytical reasoning. Tasks that previously required costly, skilled labour and faced prohibitive scaling barriers can now be executed at minimal marginal cost. This is the core economic shift from which everything else follows.

AI is best understood as a general-purpose technology — one applicable across virtually every sector of the economy, improving continuously, and generating further innovations rather than being confined to a fixed set of uses. Electricity and computing were the clearest prior examples. Their economic effects unfolded over decades and were determined as much by complementary investments in infrastructure, skills and institutional change as by the technologies themselves. AI will follow the same pattern. The implication is that near-term productivity gains will likely be modest while long-run effects may be very large. Hence, the countries and firms that invest early in the surrounding ecosystem tend to capture a disproportionate share of the eventual returns.

Three Channels of Economic Change

The first and most immediate channel is labour substitution and augmentation. AI substitutes for a range of cognitive tasks that previously defined the work of junior and mid-level professionals: research synthesis, first-draft writing, code generation, or query resolution. The substitution effect is sharpest at the bottom of the professional skill distribution — the roles that typically serve as entry points into white-collar careers. The augmentation effect runs in the opposite direction: a senior professional with AI tools can now produce output that previously required a supporting team, compressing the ratio of support staff to experienced workers across financial services, consulting, and legal sectors. Both effects are already evident in hiring patterns in advanced economies, where demand for entry-level analytical roles is softening even as demand for senior professionals holds firm.

The second channel is innovation acceleration. AI compresses research and development (R&D) cycles by narrowing the search space in disciplines where the number of possible combinations — molecular compounds, material configurations, software architectures — is far too large for human researchers to explore efficiently. The canonical example is protein folding: a problem that had occupied structural biologists for fifty years was effectively resolved by AI within a short timeframe, removing a bottleneck that had constrained the entire pharmaceutical discovery pipeline. Similar dynamics are playing out in materials science, climate modelling, and semiconductor design. The economic consequence is a reduction in the cost of scientific discovery, which, all else equal, should raise the rate at which new commercially viable products enter the market. The qualifier matters: the gap between scientific discovery and commercial deployment remains wide, slow, and expensive, and AI does not obviously shorten that part of the journey.

The third channel is the creation of new products and services. AI enables product categories that previously did not exist at a viable cost: personalised learning systems, diagnostic tools that extend specialist medical expertise to low-resource settings, autonomous logistics, and AI-native software that adapts continuously to user behaviour. These expand what the economy can produce rather than simply producing existing things more efficiently. The contribution of genuinely new product categories is notoriously difficult to capture in standard GDP measures, but the demand effects and the investment flows chasing them are real.

Capital Formation Multipliers

AI is substantially more investment-intensive than prior software generations, spurring capital formation activity across the full value chain. Earlier software could be deployed on general-purpose hardware at modest upfront cost; the economics of frontier AI are fundamentally different. Training and running large models requires purpose-built GPU clusters, specialised cooling infrastructure, and reliable power at scale — capital commitments that are large, long-dated, and practically irreversible.

At the upstream end, the numbers are striking. NVIDIA's revenue grew from roughly $27 billion in fiscal year 2023 to over $130 billion in fiscal year 2025. Microsoft, Google, Amazon, and Meta have each announced annual capital expenditure programmes in the $60–80 billion range for 2025, with AI infrastructure dominating. Global data centre investment is projected to exceed $1 trillion cumulatively through 2030. Upfront commitments on this scale have already locked in a substantial volume of expenditure to be materialised over the coming years.

The investment cycle does not stop at the data centre. AI capability is increasingly a hardware requirement rather than a software feature: end-users — enterprises and consumers — cannot access it without replacing their physical devices. On the enterprise side, AI-capable PCs equipped with dedicated neural processing units are becoming the baseline for corporate procurement; analysts estimate upwards of 1.7 billion devices globally require replacement over the coming years. On the consumer side, generative AI smartphones are projected to grow from a nascent category to over 900 million annual shipments by 2028. Both cycles are being pulled forward simultaneously, compressing what would ordinarily be a gradual refresh into a concentrated demand event.

Taken together, this is a capex supercycle running across the full stack: infrastructure, enterprise devices, and consumer hardware. Hardware orders drive semiconductor production, which in turn drives demand for materials, packaging, testing, and logistics across the global supply chain. The geographic distribution of that supply chain is where Southeast Asia enters the picture.

The Productivity Debate

How large AI's aggregate productivity effect will prove, and on what timescale, remains genuinely contested among economists. The range of serious estimates is wide enough to make confident projection in either direction implausible. Both the optimistic and pessimistic positions have merit. The honest answer is that the outcome range remains wide, and precise-figure projections — whether bearish or bullish — carry more confidence than the evidence supports.

The sceptical case rests on a structural argument about task coverage. AI currently substitutes for tasks that account for a limited share of total economic value added; fewer than a quarter of AI-exposed tasks can be profitably automated in the near term, translating into aggregate productivity gains that are modest by historical standards. The deeper sceptical argument is about lags. Computing power expanded by orders of magnitude through the 1970s and 1980s without showing up in measured productivity. This gap persisted for around two decades before being followed by a productivity surge in the 1990s. There are reasonable grounds to expect something similar here: integration of AI into existing workflows takes time, requires organisational adaptation, and generates verification costs that partially offset headline efficiency gains.

The optimistic counter is that the adoption curve for AI is compressing relative to prior general-purpose technologies. Early diffusion data suggest generative AI is spreading through the workforce faster than personal computers did at a comparable stage — reaching a significant share of knowledge workers within two years of mainstream availability, compared with a decade or more for PCs. If that pace is sustained, the standard argument for long adoption lags loses some of its force. The optimistic view also holds that task coverage will expand as models improve and that the macro effect could ultimately be large.

II. HOW SOUTHEAST ASIA BENEFITS FROM AI

The Export Channel: Embedded in the Hardware Supply Chain

Southeast Asia's role in the semiconductor industry lies in critical downstream capacity — back-end assembly, testing, packaging, and electronics components — rather than frontier chip design, which remains concentrated in the United States, Taiwan, and South Korea. As AI capex drives semiconductor throughput higher, demand for these services rises in tandem.

Singapore and Malaysia are the clearest beneficiaries. Malaysia accounts for roughly 13% of global outsourced semiconductor assembly and test capacity, with operations by major international chipmakers and a well-developed ecosystem of local suppliers. Meanwhile, Singapore has a relatively small but critical base of semiconductor foundries, in addition to hosting sizeable production of AI-adjacent electronics components. For both nations, semiconductors are the single largest component of domestic exports, meaning that rising global chip volumes translate fairly directly into export revenue. Vietnam's electronics export sector, already the country's largest goods export category valued at over $100 billion annually, benefits from rising demand for servers, networking equipment, and consumer devices incorporating AI capabilities. Other countries like Thailand and the Philippines benefit through modest electronics assembly and component manufacturing bases, though with less direct exposure to AI-specific hardware.

The mechanism is straightforward: the AI capex boom is, among other things, a large, positive demand shock for the kinds of manufactured goods that Southeast Asian countries produce and export. For as long as the capex cycle continues, the region's electronics exporters will benefit from a meaningful demand tailwind.

Driving Investments in Data Centres and Semiconductor Capacity

Beyond exports, AI is driving a more structural wave of inbound FDI into data centre infrastructure across the region. These investments are partly to serve growing Asian cloud markets, but also to reduce supply chain concentration risk, and respond to data localisation requirements in countries such as Indonesia.

The infrastructure investment pipeline is substantial. Microsoft has committed approximately $2.2 billion to Malaysian data centre and AI infrastructure and $1.7 billion to Indonesia; Google has committed $2 billion to Malaysia and $1 billion to Indonesia; Amazon, ByteDance, and others have announced comparably significant figures. Singapore remains the established regional hub, but land scarcity and power constraints have pushed new capacity into Johor (adjacent to Singapore), the Klang Valley, and the Jakarta-Batam corridor.

Semiconductor manufacturing FDI is more selective. Established chip assembly and test players have increased capacity in Malaysia, Vietnam, and Thailand, while Singapore continues to extend its lead in the foundry sector. The US-China technology decoupling has been an accelerant: firms seeking supply chain redundancy outside China have found Southeast Asia a viable alternative, given its existing electronics manufacturing base and generally stable investment environment.

AI as a Productivity Tool for Domestic Economies

AI's potential to raise productivity within Southeast Asia's domestic sectors — independent of the region's position in global supply chains — is the least examined of the three channels. Structural productivity gaps run deep across agriculture, retail, financial services, and healthcare. These sectors are SME-dominated, employ the majority of the workforce, and sit far below the technological frontier — not because the technology is now unavailable, but because the firms operating within them face compounding constraints that trap them at sub-scale. Informal bookkeeping, weak management practices, and thin margins make it impossible to afford the specialist services that larger firms take for granted: logistics consulting, marketing expertise, human resource management, and agronomic advice.

AI-powered tools, accessible from frontier providers at rapidly falling cost, can begin to close that gap at a price point that previously had no equivalent. A smallholder farmer who never had access to agronomic advice can now receive localised yield guidance at near-zero cost. A small trader who could never justify a logistics consultant gains access to route optimisation. A small retail business can use AI for inventory management and record-keeping.

The binding limits remain digital infrastructure, usable data, and human capacity to integrate these tools into existing workflows. But Southeast Asia has demonstrated it can execute technology leapfrogs when those constraints are addressed. AI adoption in high-value domestic applications is a plausible continuation of that pattern — provided it is met with deliberate policy effort rather than left to organic diffusion, which risks inertia and friction suppressing take-up.

III. DOWNSIDES AND STRUCTURAL RISKS

Labour Market Displacement

AI's key labour-market risk is not headline unemployment but the erosion of the entry-level career pathways through which workers acquire skills and move up the income ladder. Apocalyptic predictions about AI 'replacing all jobs' are, in their simple form, reductionist. The more corrosive threat is structural: what type of jobs disappear, at what stage of a career, and whether traditional routes to upward mobility remain intact.

AI's advantage lies in routine, codifiable, and language-based cognitive tasks — precisely the work concentrated in the entry and lower-middle tiers of the white-collar labour market. These roles are economically crucial because they serve as the rungs through which workers in developing economies accumulate experience, build professional skills, and move up the income ladder. When AI compresses demand for such roles, it does not simply displace labour in the narrow sense. It risks narrowing the mechanisms by which workers gain skills and know-how over their careers.

Southeast Asia is exposed to this dynamic. The Philippines' business process outsourcing industry — employing roughly 1.3 million workers and generating around US$30 billion in annual revenue — is structurally vulnerable. Many core BPO functions, including customer service, data processing, back-office administration, and basic software support, fall squarely within the range of tasks that generative AI performs increasingly well. Other countries as Malaysia and Thailand also have sizeable BPO sectors, even if smaller as a proportion of their respective economies.

The broader regional concern extends well beyond BPO. Across Southeast Asia's urbanising economies, junior white-collar jobs have become the primary gateway into the professional middle class. Yet these are precisely the roles AI is most likely to compress before workers have accumulated the experience necessary to move into higher-value positions. The result may not be mass unemployment in the conventional sense, but something potentially more corrosive: a weakening of the career ladder itself. If sustained over time and across millions of workers, the loss of these early-career pathways could produce persistent labour-market hysteresis — a structural deterioration in lifetime earnings potential and social mobility that no short-run employment statistic will fully capture.

Automation and the Manufacturing Development Model

AI-augmented automation threatens to close the export-manufacturing development path before much of Southeast Asia has fully exploited it. The standard model is well established: attract labour-intensive manufacturing investment on the basis of wage cost advantage, generate employment and productivity spillovers, and gradually upgrade the industrial base into higher-value segments. The original "Four Asian Tigers" ascended this way; China sustained it for three decades; Vietnam, Cambodia, and Bangladesh are at earlier stages of the same trajectory.

Automation and AI-assisted manufacturing are eroding the wage advantage on which this model was premised. China's ratio of industrial robots to manufacturing workers has risen from approximately 2.5 per 1,000 in 2015 to over 10 per 1,000 in 2023, approaching South Korean levels. Automated Chinese factories, and increasingly, facilities in Mexico or Eastern Europe, can close the cost gap with low-wage Southeast Asian production more rapidly than prior technology transitions allowed, because the productivity gains from automation are not offset by the distance to final markets or the limits of local labour supply.

This does not mean the manufacturing development window has closed entirely. Product complexity, supply chain depth, and institutional factors still give labour-abundant countries advantages in certain segments. But the window is narrowing with each incremental advance in manufacturing automation. Countries like Cambodia and Myanmar, which have barely begun their manufacturing ascent, face a more difficult trajectory than those that have established themselves in manufacturing value chains. Those that have, such as Vietnam, Thailand, and Malaysia, face a different risk: they have moved beyond pure labour-intensive manufacturing but have not yet achieved the technological intensity that would insulate them from competition on automation at the bottom of their current product mix.

Geopolitical Dependency and the Squeeze Between Providers

Southeast Asia faces an additional positioning risk amid the US-China technology bifurcation. The region's access to frontier AI capabilities runs primarily through US providers — cloud platforms, model APIs, and hardware. Meanwhile, its manufacturing supply chains remain deeply integrated with China. Neither dependency is easily unwound, even over the medium term.

The economic risk extends beyond geopolitics. If AI capabilities become a critical productivity input across sectors, access to them on commercially neutral terms becomes a strategic economic asset. US export controls on advanced semiconductors, imposed from 2022 and progressively tightened, have already constrained access to high-end compute for some regional firms. The controls primarily target China, but their extraterritorial scope has limited regional data centres and research institutions' ability to acquire the hardware required to run frontier models domestically.

There is also a value-capture problem. When a Vietnamese logistics firm uses an AI routing tool built on a US frontier model, the economic rent accrues upstream to the model developer and the cloud provider, not to a Vietnamese technology company. At the economy-wide level, this represents a structural terms-of-trade issue: the region pays for AI capabilities as an imported input analogous to capital equipment imports. The parallel to manufacturing is instructive: countries that moved from importing finished goods to producing components to designing products captured progressively more of the value chain. This trajectory appears more difficult for the region to replicate in the AI value chain, but opportunities remain with appropriate goals and policies.

The Risk of Disruptive Capital Expenditure Correction

The core vulnerability in the current AI investment boom is the gap between committed capital and validated returns. Valuations of AI-adjacent companies are elevated by historical standards, and capex commitments already in place imply a trajectory of enterprise adoption and monetisation that is aggressive and has yet to be demonstrated. If the investment thesis is tested against actual earnings, and monetisation disappoints, the equity rerating that follows will tighten financing conditions and trigger a pullback in capital spending.

The consequences for Southeast Asia would be direct. FDI commitments would be deferred or cancelled while semiconductor and electronics demand would soften, eroding the export tailwinds that the region is currently enjoying. The 1999–2001 technology bust offers an instructive, if imperfect, precedent: electronics demand fell sharply, and Southeast Asian exports were adversely affected. The current buildout is more anchored in near-term commercial applications and real revenue, which moderates the analogy, but does not eliminate the risk. Data centre overbuilding in some markets could further amplify the correction.

Policymakers should thus treat the current investment environment as favourable but contingent. The structural case for AI-linked FDI into the region is durable; the cyclical case is not, and the two should not be conflated when making infrastructure and incentive commitments that carry long time horizons.

IV. POLICY IMPLICATIONS

Getting the Basics Right

The foundational conditions for AI adoption — basic schooling, connectivity, and policy certainty — are not uniformly in place across most of the region, a prior deficiency the policy debate tends to skip over. Most are not adequately across the majority of the region. The common thread across these domains is chronic underinvestment relative to known economic returns, and a political economy that has consistently prioritised visible short-term expenditure, such as visible physical infrastructure, over the slower-payoff foundational investments. What AI changes is the cost of continued underinvestment, since they are now the determinants of whether the region captures any meaningful share of the AI productivity dividend.

Energy and digital infrastructure. Data centres require reliable, affordable power at scale. But the problem runs deeper than the data centre buildout. AI tools deployed in areas such as logistics, agriculture, and services require connectivity that reaches beyond the urban core. Fixed broadband penetration remains low outside major cities across most of the region, and mobile internet quality is uneven. Even before AI, energy reliability and digital connectivity were key to countries' participation in complex global value chains. AI raises those returns sharply: without the infrastructure to enable the population to access AI-enabled tools and services, countries risk failing to fully reap the productivity dividend that mass AI adoption could deliver.

Education and skills. Two foundational gaps matter most and are multi-decade investments. The first is numeracy. The region's PISA performance in mathematics is weak relative to its income level, which constrains the pool of workers who can use, adapt, and build on AI tools in any meaningful way. The second gap is English proficiency. Frontier AI models are predominantly trained on English-language text, and performance degrades for locally specific domains in non-English languages. The practical consequence is that English-fluent users interact with AI more effectively than non-fluent users due to better prompt construction, richer outputs, and fuller access to the global knowledge base that models encode. General English proficiency, which is already economically valuable for trade and services integration, now carries an additional AI productivity premium that compounds over time.

Government digitalisation. Digital identity systems, interoperable payment infrastructure, and open government data frameworks are the platforms on which AI applications in public services and domestic sectors get built. The region has genuine strengths here — GovTech Singapore, the Philippines' PhilSys, Indonesia and Malaysia's digital payment ecosystem — but coverage is uneven within countries, with rural and lower-income populations systematically excluded. Completing this infrastructure is justified by the existing digital economy; AI substantially increases the return on investment in expanding the ecosystem and further increasing user access.

Rule of law and contract enforcement. AI-adjacent industries — software, data services, AI applications — are more IP-intensive than the manufacturing sectors that currently anchor most of the region's economies. Weak IP protection and unreliable contract enforcement suppress the private investment in AI application development that public R&D co-funding is meant to complement. This is not a new problem, but one whose economic cost rises as the region tries to move up the value chain from hardware assembly to software and services.

Developing an AI-Friendly Human Capital Policy

The central human capital challenge is to adapt education and training systems to continuously generate the capabilities an AI-integrated economy requires, not merely to manage a temporary disruption. The two tasks are related but distinct, and conflating them risks leading to policy confusion that siphons resources without delivering results.

In the short term, the challenge is to manage the dislocation. Workers displaced from BPO and other vulnerable roles face a genuine skills mismatch: the jobs being created in AI-adjacent fields require capabilities that take time and structured investment to acquire, and the institutional machinery to support that transition is underdeveloped across most of the region. Effective responses share several characteristics. They are sector-specific rather than generic: a reskilling programme for displaced BPO workers must be designed around the actual gap between existing competencies and the roles into which workers can credibly transition, not around a standardised AI literacy curriculum. Programmes should be outcomes-linked, with public funding at least partly contingent on employment results, to create accountability for training providers. Lastly, such programmes should be accompanied by income support adequate to allow workers to complete meaningful retraining rather than accepting the first available lower-wage alternative out of financial necessity.

In the longer term, the task is more fundamental: building a human capital system tailored to an economy in which AI capabilities will continue to advance. This requires clarity, first, on what to build toward. “AI skills” is the wrong policy target — it implies training workers to use tools that will be superseded before curricula are finalised. The durable objective is to develop capabilities that AI cannot systematically replicate: domain judgment, contextual reasoning in complex institutional environments, relationship-intensive service delivery, and the tacit knowledge that accumulates through structured professional experience. These are not acquired through short courses and programmes, but require learning architectures that enable individuals to access reskilling and upskilling at various points in their careers.

The fundamental constraint is institutional. Firms underinvest in training because worker mobility allows competitors to free-ride on that investment; public systems are the conventional corrective, but most of the region's TVET and tertiary infrastructure was designed around an industrial economy and has not been substantively restructured. A fit-for-purpose architecture would involve employers as co-investors with genuine accountability for outcomes; tie public funding to employment results rather than enrolment; and offer modular, stackable credentialling that allows workers to accumulate qualifications incrementally. That last feature matters considerably in Southeast Asia, where a large share of the prime-age workforce cannot sustain multi-year retraining commitments without income support adequate to make completion viable.

Singapore's SkillsFuture programme is the regional benchmark, but it operates under conditions — near-full employment, high institutional capacity, and a small, administratively legible workforce — that do not generalise to larger, middle-income economies. Sectoral training funds built on pooled financing and employer co-ownership offer a more transferable model; several countries have existing training and skill development programmes that should be re-equipped to deliver AI-related skills rather than building new infrastructure from scratch.

Building Domestic AI Capacity Sustainably

No Southeast Asian economy can compete in frontier foundation model development, and public resources directed toward that ambition will be wasted. Training models at the capability frontier require massive capital expenditure in the billions, the market structure tends naturally toward oligopoly, and the returns to being a third or fourth mover are low. In essence, the opportunity cost of chasing this position is high, while the expected economic return is low.

Domain-specific model development is a more defensible use of public investment, with a clear goal of addressing a current underprovision. US and China-based frontier developers have weak commercial incentives to invest in fine-tuning for low-resource languages or niche sectoral datasets, because the addressable market is small relative to the cost. Fine-tuning is currently several orders of magnitude cheaper than pre-training, shifting the feasibility calculus — though this cost differential is narrowing as techniques become more accessible, meaning frontier developers may eventually find these niches commercially viable. Hence, the window for regional public investment to establish a meaningful position is real but not indefinite.

Application development built on genuine regional advantages offers the highest near-term returns. The region holds decades of operational data in sectors such as palm oil, rice, aquaculture, and tropical disease surveillance — but this volume is currently hindered by usability. Much of this data sits fragmented across private operators or poorly digitised public systems. Realising the comparative advantage requires prior investment in data curation, labelling, and interoperability infrastructure before model development can begin in earnest. Where that groundwork exists or is tractable, the case for publicly supported applications in precision agriculture, tropical disease surveillance, and small business credit scoring is strong — each addresses an information failure with significant productivity or welfare consequences, and each involves domain complexity that makes locally adapted tools superior to off-the-shelf alternatives.

The right policy instruments are co-investment frameworks pairing public R&D funding with private deployment commitments, shared compute infrastructure accessible to universities and startups, and—selectively — procurement policy as a demand-side tool. The procurement recommendation carries an implicit quality-cost trade-off: committing to domestically developed tools is appropriate for lower-stakes applications such as agricultural extension or logistics, where performance gaps are manageable. In more sensitive domains such as credit decisioning or healthcare diagnostics, the same logic could impose real costs on end users and should not apply without demonstrated capability thresholds.

The long-run objective is absorptive capacity, with a focus on execution. Korea and Taiwan built world-class semiconductor positions through licensed technology and sustained engineering capability development — not through protecting domestic champions indefinitely. The AI analogue is a cumulative process that arguably takes place over decades: building the workforce, data infrastructure, and institutional knowledge required to deploy, adapt, and incrementally improve tools developed elsewhere. That is a more achievable, and more valuable, ambition than misguided attempts at frontier model development.

Promoting Wider AI Access and Adoption

Without deliberate intervention, AI is likely to widen the

productivity gap between large firms and the SMEs that account for most employment across Southeast Asia — the long tail that government policy must actively reach. Large multinationals will adopt AI regardless of policy; the challenge is ensuring that adoption diffuses downward to smaller businesses, where AI uptake may face frictions, leaving a large degree of the economic benefits untapped.

The problem is not simply access to AI tools. Many SMEs in the region still operate on a "pen and paper" basis with fragmented records, paper-based processes, inconsistent accounting practices, and limited technical capacity. Firms cannot deploy AI systems effectively if inventories, invoices, customer records, and operational data are incomplete or non-standardised. In practice, AI adoption depends on prior digitisation. Before firms can automate processes or use AI-driven analytics, they need basic digital systems that generate reliable data.

AI adoption will likely require similar forms of shared digital infrastructure: interoperable e-invoicing systems, standardised accounting platforms, cloud-based business management tools, and trusted frameworks for data sharing. This is where public policy makes a difference. One useful example is India's approach to digital public infrastructure. Systems such as the Unified Payments Interface (UPI) dramatically lowered the cost of digital participation for smaller firms by creating interoperable payment rails used across the economy. The significance of UPI lay not only in digital payments themselves but also in how it brought millions of businesses into a broader digital ecosystem. These investments are less visible than national AI strategies, but they are far more important for broad-based adoption.

Governments also tend to over-rely on generic training programmes.

Awareness is rarely the main constraint. Most firms already know AI exists; what they often lack is a clear understanding of how it generates commercial value in their specific operating environment. Sector-specific demonstration is usually more effective than broad “AI literacy” initiatives. A logistics company is more likely to adopt route optimisation tools after seeing measurable reductions in fuel costs and delivery times among peers than after attending a general seminar on digital transformation.

At the same time, governments should avoid mistaking uptake for success. Subsidy programmes frequently measure easy metrics such as software licences issued or firms trained, even though these say little about productivity gains. A firm that purchases an AI tool but never meaningfully integrates it into operations has not become more competitive.

Ultimately, AI adoption is an organisational challenge as much as a technological one. Firms often need to redesign workflows, retrain staff, and change management practices before productivity gains materialise. Industry associations and trade bodies are therefore likely to be more effective channels for adoption than top-down national campaigns, because they can translate technological change into commercially relevant terms that firms can actually use.

Creating an Enabling Policy and Regulatory Environment

The immediate regulatory challenge facing Southeast Asia is not excessive AI regulation but fragmented, uncertain rules that force firms to apply existing rules never designed for the technology. The result is legal ambiguity that slows adoption among legitimate businesses while doing relatively little to deter bad actors. In the near term, governments do not necessarily need sweeping AI-specific legislation. More important is clear guidance explaining how existing laws apply, where gaps remain, and what firms are expected to do while those gaps are addressed.

One area where policy incoherence carries high economic costs is data localisation. Several Southeast Asian governments have moved toward localisation requirements in response to concerns over data sovereignty cybersecurity. While these concerns have merit, blanket localisation rules are often a costly and blunt instrument. Restricting cross-border data flows fragments regional markets, raises infrastructure costs, and limits the development of AI systems that depend on scale and diverse datasets. For a region where individual national markets are relatively small, this can become a constraint on competitiveness. In many cases, targeted rules governing data access, accountability, and usage are likely to be more effective than strict geographic restrictions.

Efforts should also be made to smooth cross-border data movements. Regulatory fragmentation makes it harder for domestic firms to scale regionally. A company operating across ASEAN may face multiple sets of AI-related standards, compliance processes, and liability expectations. Large multinational platforms can absorb these costs; mid-sized regional firms often cannot. This weakens Southeast Asia's broader digital ecosystem and limits the region's collective influence relative to larger regulatory powers. ASEAN already possesses institutional mechanisms for regulatory coordination, but political commitment to deeper harmonisation remains limited.

A differentiated approach that is proportionate to the risks should also be developed. Not all AI applications warrant the same degree of oversight. Systems used in credit scoring, hiring, healthcare, or law enforcement clearly require stricter scrutiny because errors in these areas can produce lasting harm. Applying similarly burdensome rules to lower-risk applications, however, imposes compliance costs that fall disproportionately on legitimate firms while generating limited public benefit.

Ultimately, effective AI governance depends less on ambitious legislation than on regulatory capability. Regulators need sufficient technical expertise to strike the right balance between harm prevention and enabling development. Without that capacity, AI regulation risks becoming largely symbolic: complex rules on paper that sophisticated firms can navigate around with ease.

V. WHERE SOUTH-EAST ASIA STANDS

	Singapore	Malaysia	Indonesia	Vietnam	Philippines	Thailand
IMF AI Preparedness Index 2023[1]	**0.80 (#1)**	**0.63 (#34)**	**0.52 (#60)**	**0.48 (#77)**	**0.50 (#67)**	**0.54 (#52)**
WIPO Global Innovation Index 2025[2]	**59.9 (#5)**	**40.6 (#34)**	**31.3 (#55)**	**37.1 (#44)**	**33.6 (#50)**	**36.7 (#45)**
Network Readiness Index 2025[3]	**75.46 (#3)**	**57.37 (#38)**	**53.75 (#49)**	**56.00 (#40)**	**48.89 (#66)**	**54.54 (#44)**

Source: IMF, WIPO, Portulans Institute

Across varied indices of preparedness and capacity to benefit from AI, a consistent pattern has emerged. Singapore's combination of leading infrastructure, policy and innovation focus makes it the top regional performer by a significant margin and globally competitive. Malaysia is the regional runner-up with strong infrastructure and innovation capacity, while policy lags behind. The rest of the region lags not too far behind, with infrastructure being a persistent concern despite the significant improvement in digital connectivity over the past decade.

Governance

Singapore leads decisively, with its Model AI Governance Framework and National AI Strategy 2.0 providing a coherent regulatory architecture, complemented by active labour market policies including SkillsFuture credits and sector-specific reskilling pathways. Malaysia

1. The IMF AI Preparedness Index covers four categories of digital infrastructure, human capital and labour market policies, innovation and economic integration and regulation and ethics. A total of 174 countries were assessed.

2. The WIPO Global Innovation Index scores economies across four categories of science and innovation investment, technological progress, technological adoption and socioeconomic impact. A total of 139 countries were assessed.

3. The Portulans Institute Network Readiness Index measures the application and impact of ICT across four categories of technology, people, governance and impact. A total of 127 countries were assessed.

has moved more deliberately, releasing its National AI Roadmap and a governance code, though enforcement mechanisms and workforce transition policies remain underdeveloped. Thailand and Indonesia have published national AI strategies but implementation is uneven and labour adjustment frameworks are nascent. The Philippines and Vietnam trail furthest, with policy frameworks still largely aspirational and limited institutional capacity to manage AI-driven labour displacement in their large, informal-sector-heavy workforces.

Infrastructure

Singapore's near-universal fibre coverage, dense data centre capacity and the region's most mature cloud ecosystem places it in a category of its own. Malaysia is well-positioned as a second-tier hub, benefiting from significant recent data centre investment in Johor and steady network upgrades. Thailand and Vietnam occupy the middle ground—both have made meaningful strides in mobile and fixed broadband penetration, though enterprise-grade digital infrastructure remains uneven outside major urban centres. Indonesia's vast archipelagic geography continues to constrain connectivity despite improving headline metrics, while the Philippines lags most visibly, with network quality and reliability persistent bottlenecks to broader AI adoption across the economy.

Economic Structure

Singapore's high-value services orientation in industries like finance, professional services and wholesale trade positions it to capture AI productivity gains most directly, supported by strong innovation capacity and deep R&D linkages. Malaysia's diversified base, including an established electronics manufacturing ecosystem, offers meaningful upside, though innovation intensity remains modest. Vietnam and Thailand are export-manufacturing-heavy economies where AI's impact will likely be channelled through process optimisation rather than frontier innovation. Indonesia and the Philippines, with larger shares of agriculture and informal services and weaker innovation systems, face the steepest path to translating AI capabilities into broad-based productivity gains, despite sizeable domestic markets that could anchor adoption.

VI. COUNTRY SCORECARD: WINNERS, LOSERS, AND WHAT EACH NEEDS TO GET RIGHT

Economy	Key Advantages	Principal Risks	Net Position
Singapore	Established data centre and financial hub; strong regulatory capacity; high-skill workforce; gateway for regional AI deployment and financial flows	Small domestic market; power and land constraints on data centre growth; high dependence on global investment cycle; limited manufacturing base to capture hardware demand	Winner. Singapore's position as the region's AI hub is entrenched, and its institutional quality gives it a structural advantage in attracting AI-related financial and professional services. The binding constraint is physical: it cannot absorb the data centre investment that its competitive position would otherwise attract. The macro risk is also significant: a global capex correction hits Singapore disproportionately hard given its dependence on foreign investment flows.

Economy	Key Advantages	Principal Risks	Net Position
Malaysia	Largest semiconductor assembly and test base in the region; major recipient of hyperscaler investment (Johor corridor); established electronics manufacturing ecosystem; English-speaking professional workforce	Middle-income trap risk if it fails to move up the AI value chain beyond assembly and hosting; brain drain of technical talent; political governance uncertainty; exposed to the capex correction cycle	Significant winner, with caveats. Malaysia is the most direct ASEAN beneficiary of the AI hardware supply chain, and the Johor data centre corridor represents a genuine structural opportunity. The critical question is whether it captures higher-value activity — chip design, cloud operations, AI R&D — or remains in back-end manufacturing and real estate hosting. The latter is profitable but does not solve the middle-income trap.
Indonesia	Largest domestic market in the region (270 million), creating genuine scale for domestic AI deployment; growing data centre investment; fintech and digital economy base; natural resource wealth	Very large informal and semi-formal labour force highly exposed to automation displacement; manufacturing development model at risk; weak institutional capacity to execute complex policy interventions; archipelagic geography complicates digital infrastructure roll-out	Mixed. Indonesia's domestic market scale is its most distinctive asset — large enough to sustain meaningful AI deployment in fintech, agriculture, and logistics, and potentially large enough to support domain-specific model development. But it carries the region's most acute labour displacement risk, given its large low-skilled workforce and dependence on manufacturing-led growth. Whether the opportunity outweighs the risk turns almost entirely on policy execution capacity, which has historically been the country's binding constraint.

Economy	Key Advantages	Principal Risks	Net Position
Vietnam	Strong electronics export base; successful FDI attraction record; young and literate workforce; beneficiary of US-China decoupling in supply chain diversification	Electronics manufacturing concentration creates cyclical exposure; lower-end manufacturing (garments, footwear) faces automation competition from China; limited institutional capacity for rapid workforce upskilling	Near-term winner, medium-term structural risk. Vietnam benefits more than almost any other regional economy from the AI hardware capex cycle through its electronics export exposure. The structural risk is that the manufacturing trajectory it is ascending — from garments into electronics assembly — faces competition from automated Chinese manufacturing faster than Vietnam can climb into higher-value segments. The window is open but narrowing.
Thailand	Established electronics and automotive manufacturing base; growing data centre investment; relatively high per-capita income reduces acute labour displacement pressure; strong logistics infrastructure	Automotive-heavy manufacturing has different AI exposure to semiconductor supply chain; ageing population constrains adjustment capacity; political instability creates investment uncertainty	Moderate gainer. Thailand is a secondary beneficiary of the AI capex cycle — its electronics exposure is real but less direct, and its automotive manufacturing faces a separate structural challenge from EV transition that is independent of AI. Data centre investment is growing from a low base. The ageing population, which creates labour market pressure in most contexts, actually reduces the acute displacement risk because workforce supply is already constrained.

Economy	Key Advantages	Principal Risks	Net Position
Philippines	Large, English-speaking, educated workforce; established services export base; diaspora remittances provide a macro buffer; growing domestic digital economy	BPO sector — the country's largest source of formal private employment — faces the most acute AI displacement risk in the region; limited fiscal space for large-scale adjustment programmes; infrastructure gaps constrain data centre investment; political governance risk	Structural loser without policy intervention. The Philippines' economic model has been built on traded services — BPO and overseas remittances — precisely the activities most exposed to AI substitution. The BPO sector will not collapse overnight, and firms that pivot successfully to AI-augmented higher-value services can partially offset the structural decline. But the scale of adjustment required, affecting over a million direct BPO employees plus a large indirect workforce, is large relative to the country's institutional capacity to manage it. Without serious investment in reskilling and income support, the Philippines risks absorbing the downside of AI adoption without capturing much of the upside.

This overview reflects conditions as of early 2026. Figures cited for corporate investment commitments and semiconductor trade data draw on publicly available company disclosures, trade statistics, and industry estimates. Country assessments are the author’s judgement based on available economic evidence; they are not predictions.

Pereira International Pte Ltd
1 Fullerton Road,
#02-01, One Fullerton,
Singapore 049213
Tel : +65 6248 4790
Cell : +65 9622 3361

www.pereiraintl.com

email :
info@pereiraintl.com
derwin.pereira@pereiraintl.com