


THE ECONOMICS OF SOVEREIGN AI: BALANCING AUTONOMY, INNOVATION, AND GROWTH IN THE ASIA-PACIFIC

A Report for the AI Adoption Initiative




ABOUT OXFORD ECONOMICS

Oxford Economics was founded in 1981 as a commercial venture with Oxford University's business college to provide economic forecasting and modelling to UK companies and financial institutions expanding abroad. Since then, we have become one of the world's foremost independent global advisory firms, providing reports, forecasts and analytical tools on more than 200 countries, 100 industrial sectors, and 8,000 cities and regions. Our best-in-class global economic and industry models and analytical tools give us an unparalleled ability to forecast external market trends and assess their economic, social and business impact.

Headquartered in Oxford, England, with regional centres in New York, London, Frankfurt, and Singapore, Oxford Economics has offices across the globe in Belfast, Boston, Cape Town, Chicago, Dubai, Dublin, Hong Kong, Los Angeles, Mexico City, Milan, Paris, Philadelphia, Stockholm, Sydney, Tokyo, and Toronto.

We employ 700 staff, including more than 450 professional economists, industry experts, and business editors—one of the largest teams of macroeconomists and thought leadership specialists. Our global team is highly skilled in a full range of research techniques and thought leadership capabilities from econometric modelling, scenario framing, and economic impact analysis to market surveys, case studies, expert panels, and web analytics.

Oxford Economics is a key adviser to corporate, financial and government decision-makers and thought leaders. Our worldwide client base now comprises over 3,000 international organisations, including leading multinational companies and financial institutions; key government bodies and trade associations; and top universities, consultancies, and think tanks.



CONTENTS

Executive summary	4
Section 1. AI and the path to prosperity	16
Section 2. Government’s role in diffusing AI and safeguarding trust	20
Section 3. AI sovereignty and economic outcomes	32
Section 4. Quantifying the economic impacts	41
Section 5. Direct costs of domestic capacity build-out	42
Section 6. Opportunity cost from AI sovereignty restrictions ..	52
Section 7. Environmental costs	60
Section 8. The role of industry and global cloud providers	68
Section 9. Way forward and recommendations	75
Appendix 1: AI sovereignty policies in APJ countries	82

All data presented in tables and charts in this report are proprietary to Oxford Economics, except where otherwise stated and cited in footnotes, and are copyright © Oxford Economics Ltd.


The modelling and results are based on information provided by third parties, upon which Oxford Economics has relied in good faith in the preparation of its analysis and forecasts. Any subsequent revision or update to these data may affect the assessments and projections presented.

To discuss the report further please contact:

Bali Kaur Sodhi: bsodhi@oxfordeconomics.com

Oxford Economics
6 Battery Road, #38-05, Singapore 049909
Tel: +65 6850 0110


MAY 2026



EXECUTIVE SUMMARY

Artificial intelligence (AI) is increasingly recognised as a general-purpose technology with the potential to lift productivity, strengthen competitiveness, and support long-term economic growth across Asia-Pacific and Japan (APJ). At the same time, governments are placing greater emphasis on sovereignty—seeking to ensure that data, infrastructure, and AI systems remain subject to domestic control, oversight, and values. This shift reflects legitimate economic security and societal concerns, but it also raises important questions about cost, scalability, and the pace of AI adoption.

AI sovereignty can be pursued through a range of policy designs that vary in how far they require the AI stack to be domestically owned and localised. These choices can strengthen control, but can also introduce material trade-offs—higher costs, slower innovation cycles, constrained access to talent, and reduced interoperability. Where these costs translate into higher prices for AI services, compute, and compliance, they are ultimately borne by businesses and filter through the wider economy, reducing efficiency and, over time, constraining overall prosperity.



Against this backdrop, the report examines the economic implications of approaches that seek to develop the full AI stack domestically, relative to pathways that maintain access to global technology supply chains through strategic partnerships and risk-based assurance frameworks. In particular, the research quantifies both the direct costs of building domestic AI capacity and the opportunity costs arising from slower diffusion and reduced productivity gains.

It also highlights the additional risks and challenges, such as cybersecurity, that can emerge under different sovereignty approaches. The analysis is supplemented with in-depth interviews with technical, legal, and sector experts to develop a research report that supports policymakers as they navigate a growing dilemma: how to strengthen sovereignty without undermining economic efficiency, innovation, or sustainability.

A POLICY DILEMMA AT THE HEART OF AI SOVEREIGNTY

Governments face a fundamental trade-off. On the one hand, AI sovereignty is seen as a means to reduce strategic dependence on foreign providers, safeguard sensitive data, ensure that AI systems align with national laws and values, and create opportunities to promote local firms. On the other hand, AI adoption and productivity gains not only depend on but are accelerated by access to global cloud infrastructure, advanced models, skills, and innovation ecosystems. AI sovereignty and adoption can complement each other, but when sovereignty is pursued through protectionist measures that limit access to global technologies and their security resilience, it can raise costs, increase environmental pressures, and impede adoption goals.

The pursuit of AI sovereignty has risen rapidly across APJ. Governments are seeking greater control over the data, infrastructure, and digital systems that underpin their economies. As AI becomes embedded in critical services and commercial activity, policymakers are placing more emphasis on ensuring that sensitive workloads can be governed within trusted frameworks and remain subject to domestic oversight.

The motivations behind this shift are economic, security-driven, and societal. Economically, governments aim to strengthen domestic capability, capture more value along the AI stack,

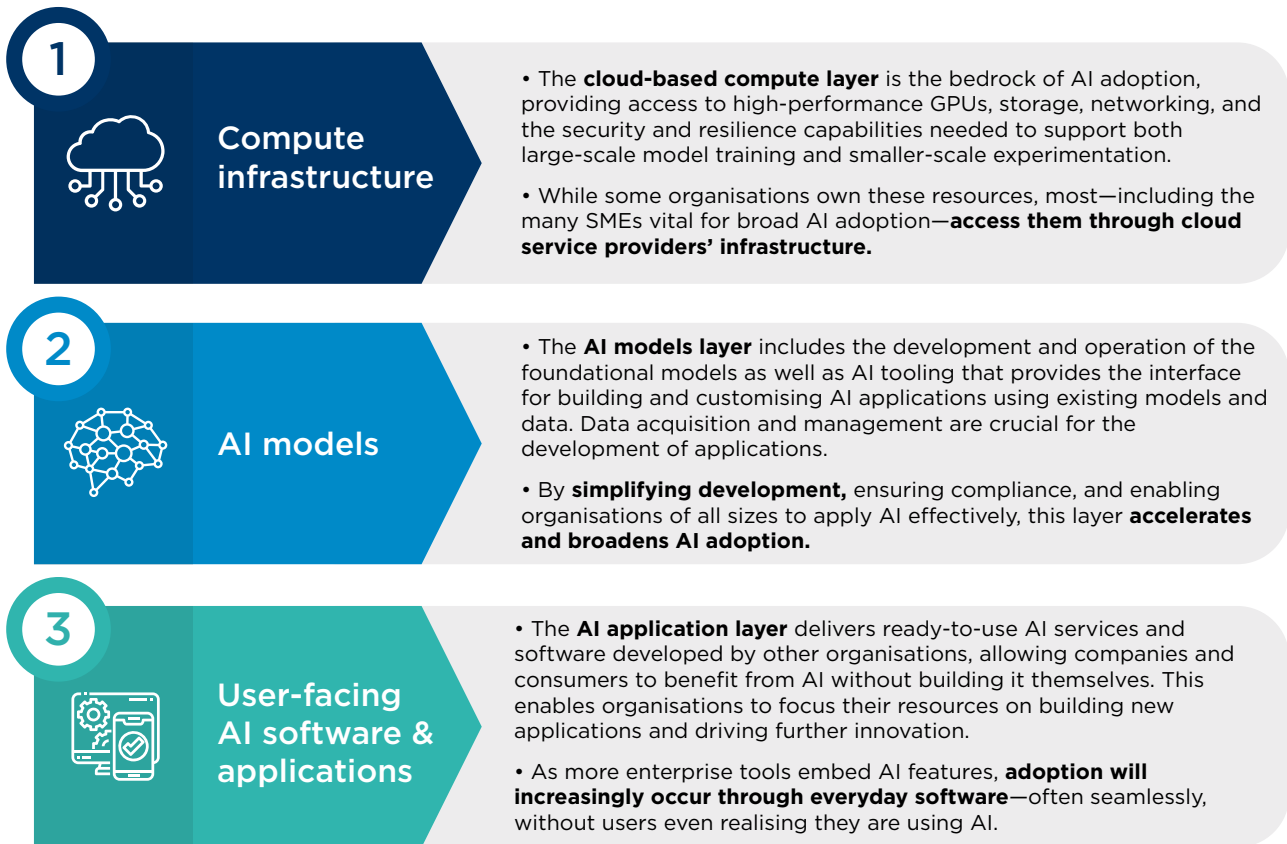
and retain a greater share of economic value within the domestic economy. Policymakers also cite security considerations, notably concentrated dependence on a small number of global cloud and AI suppliers, supply-chain fragility, and geopolitical tensions, as a principal justification for sovereign AI policies. Societal considerations include ensuring that AI systems reflect domestic languages, cultural norms, and ethical expectations rather than importing assumptions embedded in foreign-trained models. However, the empirical basis for several of these motivations remains contested: cross-country studies on cloud market structure and on cross-border data flows suggest that restrictive policies can raise costs and slow adoption, while any security or strategic-autonomy benefits depend on policy design rather than data localisation alone.^{1,2}

How these motivations translate into policy becomes clearer when viewed through the three layers of the AI stack: compute infrastructure, AI models, and user-facing applications. Viewing sovereign AI measures through this three-layer structure provides a clear understanding of how organisations access and deploy AI, and where restrictions have the greatest impact on economic outcomes. Figure 1 summarises these layers and their role in AI adoption and the section after sets out how policy measures map to each layer.

1 OECD, “[Competition in the Provision of Cloud Computing Services](#)”, *OECD Competition Policy Roundtable Background Note*, 2025, OECD Publishing, Paris.

2 IMF, “[Goeconomic Fragmentation and the Future of Multilateralism](#)”, 2023, *Staff Discussion Note SDN/2023/001*, International Monetary Fund, Washington, DC.

Figure 1: Three key layers in the AI Stack



Source: Oxford Economics

FIVE APPROACHES TO AI SOVEREIGNTY

Approaches to sovereign AI vary across the region. Some economies adopt assurance-led frameworks that maintain access to global cloud providers, while others introduce stricter localisation or ownership requirements. These choices shape how quickly organisations can access compute, develop, deploy or fine-tune models, and adopt AI-enabled applications.

Each approach reflects different policy motivations—from assurance and resilience to strategic autonomy—but carries distinct economic consequences. When mapped to the stack, the scenarios highlight three main channels through which restrictions affect economic outcomes:

- **Compute infrastructure:** AI sovereignty restrictions raise costs because domestic data centres lack the scale and efficiency of global

hyperscale providers, making workloads more expensive to run. Data-localisation rules and limits on eligible providers further reduce utilisation and increase operating costs. Cybersecurity capabilities and operational resilience are also affected: smaller domestic providers typically operate fewer availability zones and invest less in pooled cyber defence resources than global hyperscalers. They can also create fiscal risks when governments procure GPUs or build facilities significantly before demand materialises, leaving hardware underused or outdated by the time systems are operational.

- **AI models:** Stricter controls on model development, fine-tuning, and data movement fragment datasets and limit access to global tooling. This lowers model quality, raises

training costs, and slows the pace at which organisations can deploy or adapt advanced AI systems.

- **Application layer:** Restrictions on foreign software providers narrow competition and reduce access to global APIs and AI-enabled enterprise tools. This leads to higher prices, slower innovation cycles, and delays in bringing mature AI capabilities into everyday business operations. Together, these effects translate into slower diffusion and weaker productivity gains compared with more open, assurance-led approaches.






Our analysis considers five broad approaches to policy design, ranging from open, assurance-led frameworks to highly restrictive, ownership-centric models. **These approaches are defined by the differing levels of restriction they introduce across the AI stack**, allowing us to compare their economic implications consistently.

In practice, many governments pursue hybrid sovereignty models, combining measures from multiple points along this continuum—tightening controls where they deem risks are highest while continuing to rely on global providers. The combination of these policies means that **countries rarely sit neatly at a single level**. Yet, the continuum provides a clear framework for understanding the economic and operational implications of different design choices.

The report assesses five levels of restrictiveness:

- **Level 1: Control-and-choice (assurance-led).** Sovereignty is pursued primarily through governance, transparency, and risk-based controls rather than localisation and ownership mandates. Both public and private sectors retain broad access to global cloud service providers and AI systems. Data-residency requirements apply only to a narrow set of highly sensitive public-sector workloads, enabling governments to maintain regulatory control while preserving access to global AI scale and innovation.
- **Level 2: Limited public-sector restrictions.** Moderate localisation and data-residency requirements are introduced, or implicitly required for government workloads - primarily applicable to a wider set of sensitive government workloads. Meanwhile private-sector access to global AI and cloud services remains largely unrestricted. These measures strengthen control over public-sector data and infrastructure but increase costs and complexity for government AI deployment, slowing adoption relative to Level 1.
- **Level 3: Expanded public-sector controls with domestic capability build-out.** Public-sector restrictions extend to AI tooling, software, and applications, alongside increased investment in developing domestic AI capabilities. While foreign providers may still supply underlying infrastructure or models, deployment increasingly occurs in domestic environments. The private sector remains largely unrestricted, but public-sector AI adoption slows further due to higher costs and narrower supplier choice.
- **Level 4: Economy-wide restrictions.** Sovereignty requirements extend beyond government into the private sector. Data residency and local-preference rules significantly increase domestic hosting and compute requirements across the economy, requiring substantial infrastructure expansion. Higher costs and reduced scale efficiencies slow AI diffusion, particularly among smaller firms and cost-sensitive sectors.
- **Level 5: Ownership-centric (domestically owned full AI stack).** The most restrictive model, under which both public and private sectors are required to predominantly run on domestically owned and operated infrastructure, models, software, and applications. This approach maximises formal control and strategic autonomy but entails very high investment requirements, reduced scalability, and slower innovation cycles compared with less restrictive regimes.

Figure 2: AI sovereignty-related restriction levels

		Restriction Level 1	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Sector	Policy dimension	Control-and-choice (Assurance-led)	Government sector-led restrictions	Government sector-led restrictions, AI tooling, software & applications	Economy-wide restrictions	Domestically owned full AI stack
 Government sector	Local data hosting	Limited	Moderate	Moderate	Moderate	High
	Cloud provision	Limited	Moderate	Moderate	Moderate	High
	AI tooling, software & applications	Limited	Limited	Moderate	Moderate	High
 Private sector	Local data hosting	Limited	Limited	Limited	Moderate	High
	Cloud provision	Limited	Limited	Limited	Moderate	High
	AI tooling, software & applications	Limited	Limited	Limited	Moderate	High
Implications for AI diffusion and the broader economy						
 Likely economic implications	<ul style="list-style-type: none"> • Lowest cost • Rapid diffusion and strong productivity gains 	<ul style="list-style-type: none"> • Moderate cost increase and time to deployment • Limited drag on government adoption and productivity gains 	<ul style="list-style-type: none"> • Moderate cost increase, but longer deployment time • Larger drag on government adoption and productivity gains 	<ul style="list-style-type: none"> • Higher costs and time to deploy • Slower overall adoption and material impact on productivity gains 	<ul style="list-style-type: none"> • Highest cost and deployment time • Slowest overall adoption and significant impact on productivity gains 	
 Trade-off	<ul style="list-style-type: none"> • Assurance-led protection allows flexibility, scalability, and global cooperation 				<ul style="list-style-type: none"> • Maximises autonomy but reduces scalability and global collaboration 	

Source: Oxford Economics

HEADLINE ECONOMIC FINDINGS

The quantitative results point to a clear and consistent pattern across the economies assessed. While targeted domestic investment in AI infrastructure can support sovereignty, resilience, industrial-policy, and capability-building objectives, costs rise sharply as restrictiveness is applied more broadly across sectors and use cases, leading to materially lower AI adoption and rising opportunity costs to the wider economy beyond fiscal costs. These **trade-offs should therefore be carefully assessed** when designing sovereignty policies, particularly where measures extend beyond select highly sensitive public-sector workloads into economy-wide application.

DIRECT COSTS OF BUILDING OUT DOMESTIC CAPACITY

Direct costs increase non-linearly as sovereign AI requirements tighten across the AI stack and will depend on their starting level of digital and AI infrastructure development. More restrictive sovereignty approaches can also create fiscal risk. Where governments procure large volumes of advanced chips or invest heavily in domestic compute infrastructure ahead of demand, they may lock in high upfront costs while facing the risk of underutilisation, delayed deployment, and hardware obsolescence. This is particularly relevant in AI, where underlying technologies evolve quickly and the value of early capital outlays can erode before systems become fully operational.

Under moderate restrictiveness (Levels 2-3), additional costs are relatively contained across most economies. However, extending restrictions economy-wide (Levels 4-5) produces a step change in required investment, reflecting the need to build domestic capacity across infrastructure, software, skills and—at the highest level of restrictions—model development.

Large and AI-intensive economies face the highest absolute costs. By 2035, Japan and India each incur US\$149.7 billion and US\$102.5 billion, respectively, in additional direct costs under Restriction Level 5. That is equivalent to around 0.3% and 0.2% of GDP, respectively, over 2025–2035. South Korea also faces



substantial costs of around US\$88.4 billion (0.4% of 2025–2035 GDP), reflecting its high baseline adoption and the scale of workloads that must be supported domestically under full onshoring. These outcomes are driven primarily by economic size, expected AI uptake, and the volume of AI activity affected by restrictions.

Smaller and more open economies incur lower absolute costs, but the relative burden can still be significant. Singapore’s direct costs rise from around US\$1 billion at Restriction Level 2 to over US\$29.1 billion (0.4% of 2025–2035 GDP) at Level 5, highlighting the sensitivity of highly digitalised economies to limits on access to global

Figure 3: Total additional direct costs (compared to Level 1), 2025–2035

	Total additional direct costs, \$ bn (relative to Restriction Level 1)			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0	0.0	1.1	1.5
India	3.2	5.2	61.1	102.5
Indonesia	0.8	1.8	20.6	33.4
Japan	4.9	7.2	86.2	149.7
Lao PDR	0.0	0.0	1.0	1.3
Malaysia	0.4	0.6	7.5	13.0
Myanmar	0.0	0.1	1.6	2.3
Nepal	0.0	0.1	1.4	2.0
Philippines	0.4	0.6	8.1	13.2
Singapore	1.0	1.4	17.4	29.1
South Korea	2.8	4.4	52.3	88.4
Taiwan	0.9	1.4	17.1	30.0
Thailand	0.3	0.6	7.3	11.5
Vietnam	0.3	0.6	8.0	11.8

	Total additional direct costs, % of 2025-2035 GDP			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.00%	0.00%	0.2%	0.2%
India	0.01%	0.01%	0.1%	0.2%
Indonesia	0.00%	0.01%	0.1%	0.2%
Japan	0.01%	0.02%	0.2%	0.3%
Lao PDR	0.00%	0.01%	0.4%	0.5%
Malaysia	0.01%	0.01%	0.1%	0.2%
Myanmar	0.00%	0.01%	0.2%	0.2%
Nepal	0.00%	0.01%	0.2%	0.3%
Philippines	0.01%	0.01%	0.1%	0.2%
Singapore	0.01%	0.02%	0.2%	0.4%
South Korea	0.01%	0.02%	0.2%	0.4%
Taiwan	0.01%	0.01%	0.2%	0.3%
Thailand	0.00%	0.01%	0.1%	0.2%
Vietnam	0.00%	0.01%	0.1%	0.2%

Source: Oxford Economics
 Note: All values in 2025 prices.

infrastructure. Mid-sized economies see costs increase several-fold as restrictions tighten, even though totals remain below those of the largest economies. Malaysia and Nepal incur the lowest absolute costs, but these still represent a material share of economic output, ranging from 0.2% to 0.3% of 2025–2035 GDP.

DELAYS IN ADOPTION

Adoption effects amplify these costs. Under higher restrictiveness levels, AI adoption by firms is delayed by three to five years, reflecting the time required to build domestic infrastructure, tools, and skills. Once adoption resumes, it follows a permanently lower trajectory due to higher costs, reduced choice, and weaker innovation incentives. Under Restriction Levels 4 and 5, firm-level adoption rates fall sharply to single-digit levels compared with adoption rates of 15%–44% under less restrictive approaches.

There are substantial spillover risks to private-sector adoption under Restriction Levels 2 and 3, even if the overall decline in adoption is moderated by the relatively small size

of the public sector. Evidence shows that public-sector procurement often sets the pace for wider adoption, with firms far more likely to adopt when government leads. Restrictions on public-sector AI and cloud services therefore weaken those demonstration effects, reduce confidence, and dampen private sector investment. They can also elevate national-security risks if critical public-sector workloads are shifted to less mature domestic environments with lower cyber-defence investment, increasing exposure to service disruptions or successful attacks. Although our modelling captures this spillover, the impact may be conservative, and real world effects could be larger.

Sovereign AI measures are expected to place a disproportionate burden on SMEs. Smaller firms depend more on low-cost cloud services and embedded AI features in enterprise software, and have less capacity to absorb higher compliance and infrastructure costs. Restrictions that limit access to global tooling or require domestic-only hosting raise barriers to adoption and slow diffusion of AI-enabled applications that SMEs rely on most. A further risk under Restriction

Figure 4: Adoption rates by levels of restriction, 2035 (OECD definition, see note below)

	Adoption rates (OECD definition), % of firms				
	Restriction Level 1	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	6.3%	6.0%	6.0%	2.1%	0.9%
India	18.8%	18.0%	17.9%	6.9%	3.2%
Indonesia	14.5%	13.9%	13.8%	5.0%	2.3%
Japan	20.3%	19.4%	19.3%	7.4%	3.4%
Lao PDR	8.7%	8.3%	8.3%	2.9%	1.3%
Malaysia	18.9%	18.0%	17.9%	6.8%	3.1%
Myanmar	10.0%	9.6%	9.5%	3.5%	1.6%
Nepal	15.0%	14.4%	14.3%	5.5%	2.5%
Philippines	15.5%	14.8%	14.7%	5.4%	2.4%
Singapore	43.8%	42.1%	41.8%	20.3%	9.6%
South Korea	28.6%	27.4%	27.1%	11.2%	5.2%
Taiwan	28.0%	26.8%	26.6%	10.9%	5.0%
Thailand	15.4%	14.8%	14.6%	5.4%	2.4%
Vietnam	16.4%	15.7%	15.5%	5.8%	2.6%

Source: Oxford Economics.

Note: Adoption rates (OECD definition) refer to firms that have integrated AI in the production of goods and services. It excludes firms that are experimenting, piloting, or scaling AI in their operations.

Levels 2 and 3 is that slower government adoption weakens the demonstration effects that normally support early SME uptake, reducing confidence and slowing diffusion across the wider economy. As a result, SMEs experience larger delays in adoption under more restrictive approaches, widening productivity gaps across the economy.

OPPORTUNITY COSTS

These adoption impacts translate into substantial opportunity costs. Opportunity costs—defined as foregone GDP gains relative to an unrestricted adoption path—rise steeply once restrictions extend beyond the public sector. Under Restriction Level 5, Japan faces cumulative losses exceeding US\$58.2 billion (around 1.4% of its 2035 GDP),

Figure 5: Opportunity costs by levels of restriction, 2035

	Opportunity cost, \$ bn			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0	0.0	0.2	0.2
India	2.9	3.3	41.5	54.7
Indonesia	0.8	1.0	12.2	15.7
Japan	3.1	3.5	44.4	58.2
Lao PDR	0.0	0.0	0.1	0.1
Malaysia	0.4	0.5	6.4	8.4
Myanmar	0.0	0.0	0.3	0.4
Nepal	0.0	0.0	0.3	0.4
Philippines	0.3	0.4	5.0	6.5
Singapore	1.2	1.4	16.2	23.5
South Korea	2.1	2.5	30.7	41.4
Taiwan	1.0	1.2	14.7	19.8
Thailand	0.3	0.3	4.3	5.5
Vietnam	0.3	0.4	4.7	6.1

	Opportunity cost, % of 2035 GDP			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0%	0.0%	0.2%	0.3%
India	0.0%	0.0%	0.6%	0.8%
Indonesia	0.0%	0.0%	0.5%	0.7%
Japan	0.1%	0.1%	1.0%	1.4%
Lao PDR	0.0%	0.0%	0.3%	0.3%
Malaysia	0.1%	0.1%	0.9%	1.2%
Myanmar	0.0%	0.0%	0.3%	0.4%
Nepal	0.0%	0.0%	0.5%	0.7%
Philippines	0.0%	0.0%	0.6%	0.8%
Singapore	0.2%	0.2%	2.2%	3.2%
South Korea	0.1%	0.1%	1.4%	1.8%
Taiwan	0.1%	0.1%	1.4%	1.8%
Thailand	0.0%	0.0%	0.6%	0.8%
Vietnam	0.0%	0.0%	0.5%	0.7%

Source: Oxford Economics

reflecting both its economic scale and high baseline adoption potential. India incurs opportunity costs approaching US\$55 billion (around 0.8% of 2035 GDP), driven by rapid projected adoption under less restrictive settings. Lower restrictiveness levels generate smaller, but still material, losses across all economies.

Taken together, the results show that **ownership-centric and highly restrictive approaches impose large upfront investment requirements while simultaneously eroding the productivity gains AI is expected to deliver.** The combination of higher direct costs, delayed adoption, substantial opportunity costs underscores the economic trade-offs inherent in more restrictive sovereign AI policy designs.

ENVIRONMENTAL IMPACT

The analysis shows that expanding domestic AI capacity under more restrictive policy settings can increase environmental pressures. The scale of these impacts depends on the efficiency and utilisation of data centre infrastructure. Hyperscale platforms typically benefit from advanced technical capabilities, higher utilisation rates, and more efficient cooling systems, reducing electricity and water use per unit of AI output.³ By contrast, smaller or sub-scale facilities may require more energy and water to deliver the same level of compute. Where more restrictive AI governance strategies potentially limit access to the more advanced, energy-efficient technologies and operating models, environmental outcomes can worsen. This reinforces the need to align AI governance objectives with energy and water constraints.

As restrictions increase and a greater share of workloads must be hosted domestically, **energy demand rises sharply.** This can lead to materially higher carbon emissions, especially in economies with carbon-intensive power systems. Large and AI-intensive economies account for the largest absolute increases reflecting both the scale of projected AI adoption and the emissions intensity of local electricity generation. These effects are particularly pronounced in Asia, where many economies continue to rely on carbon-intensive energy systems.

³ IEA, “[Data centres and data transmission networks](#)”, accessed December 2025.



Water consumption represents an additional and often underappreciated environmental pressure, particularly where AI capacity is delivered through less advanced or less efficient data-centre technologies. Facilities with older cooling systems

or lower utilisation can require substantially more water per unit of compute, both directly for cooling and indirectly through electricity generation, **amplifying risks in economies already facing high water stress.**

IMPLICATIONS FOR POLICY DESIGN

Our findings point to a clear implication for policy design. Governments are seeking to balance sovereignty, economic efficiency, and environmental sustainability as AI drives higher energy and water use. Achieving all three simultaneously is challenging, and outcomes depend critically on how restrictions are designed. Our research quantifies the economic costs associated with policy restrictions linked to AI sovereignty across the stack. The findings contribute to the broader evidence base that policymakers can draw on when assessing the trade-offs involved in designing sovereignty strategies.

The realisation of the projected economic benefits associated with AI depend on openness. Overly restrictive localisation or preferential treatment for domestic providers can weaken competition and limit access to leading technologies, while portability and open standards help reduce lock-in. Sovereignty supported through assurance, transparency, and auditability, using verifiable safeguards—such as data residency controls, key management, and operator accountability—rather than restrictions is associated with higher economic benefits. These benefits come from lower direct costs, faster adoption, and higher AI-enabled productivity gains.

Furthermore, where regulation is introduced, clear and consistent regulatory guidance is essential to avoid unintended behavioural spillovers. When requirements are ambiguous, firms may overinterpret rules—localising data or avoiding global partnerships unnecessarily—raising costs, slowing adoption, and leading to inefficient infrastructure choices.

Therefore, models that blend global capability with local control recognise that sovereignty does not require self-sufficiency. Instead, they acknowledge the ability to configure AI systems on domestic terms while retaining access to global R&D, talent, and infrastructure. Importantly, this approach enables countries to focus investment on the parts of the AI stack where domestic value creation is strongest—skills, high-quality datasets, sector-specific applications, and responsible governance—without incurring the high fixed costs and rapid hardware refresh cycles of sovereign compute infrastructure. Recent developments in European cloud policy illustrate this shift towards risk-based technical assurance.

Overall, collaborative and innovation-friendly sovereignty models offer a pragmatic path to meeting national objectives without unduly constraining AI adoption or its productivity benefits.



SECTION 1. AI AND THE PATH TO PROSPERITY

Across the world, organisations are racing to harness the power of AI to drive innovation and efficiency—from accelerating scientific research and improving patient outcomes to enhancing customer experience and modernising public services. Enterprises are now embedding AI models into telecommunications, finance, manufacturing, and retail operations, while governments are using them to strengthen service delivery and public resilience.

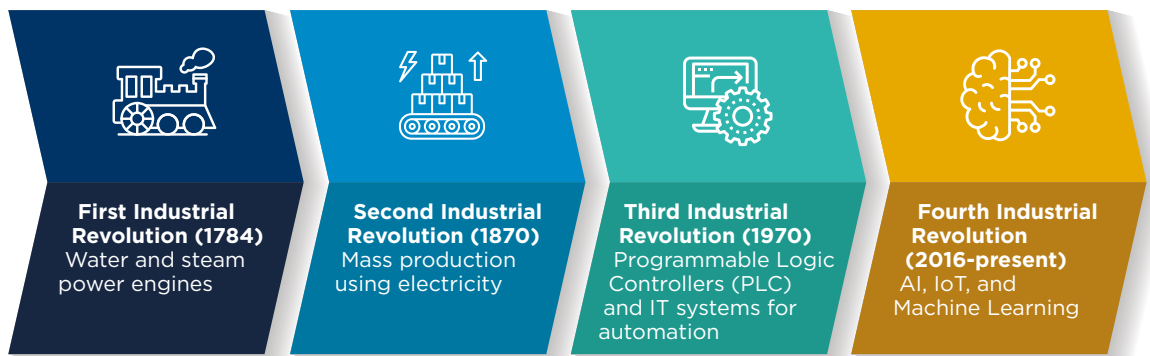


AI IS AT THE CORE OF THE FOURTH INDUSTRIAL REVOLUTION

AI will sit at the heart of what will be described as the fourth industrial revolution—a transformation defined by the fusion of physical and digital technologies and expected to extend far beyond automation. Smart manufacturing lines will be designed to use AI-driven systems to monitor

quality and predict maintenance; logistics firms will rely on machine-learning algorithms to optimise routing; financial institutions will deploy generative models for fraud detection and customer engagement.

Figure 6: Industrial revolutions and their foundations⁴



Source: Rashid and Kausik (2024).

AI USAGE AND ADOPTION

Adoption is accelerating rapidly. According to McKinsey (2025), 88% of organisations reported using AI in at least one business function—up 10 percentage points from 2024. The greatest progress has come in areas such as marketing, product development, and customer operations.⁵ Among developers, Stack Overflow’s 2024 survey found that 84% were using or planning to use AI tools in their development processes, up from 76% the previous year.⁶

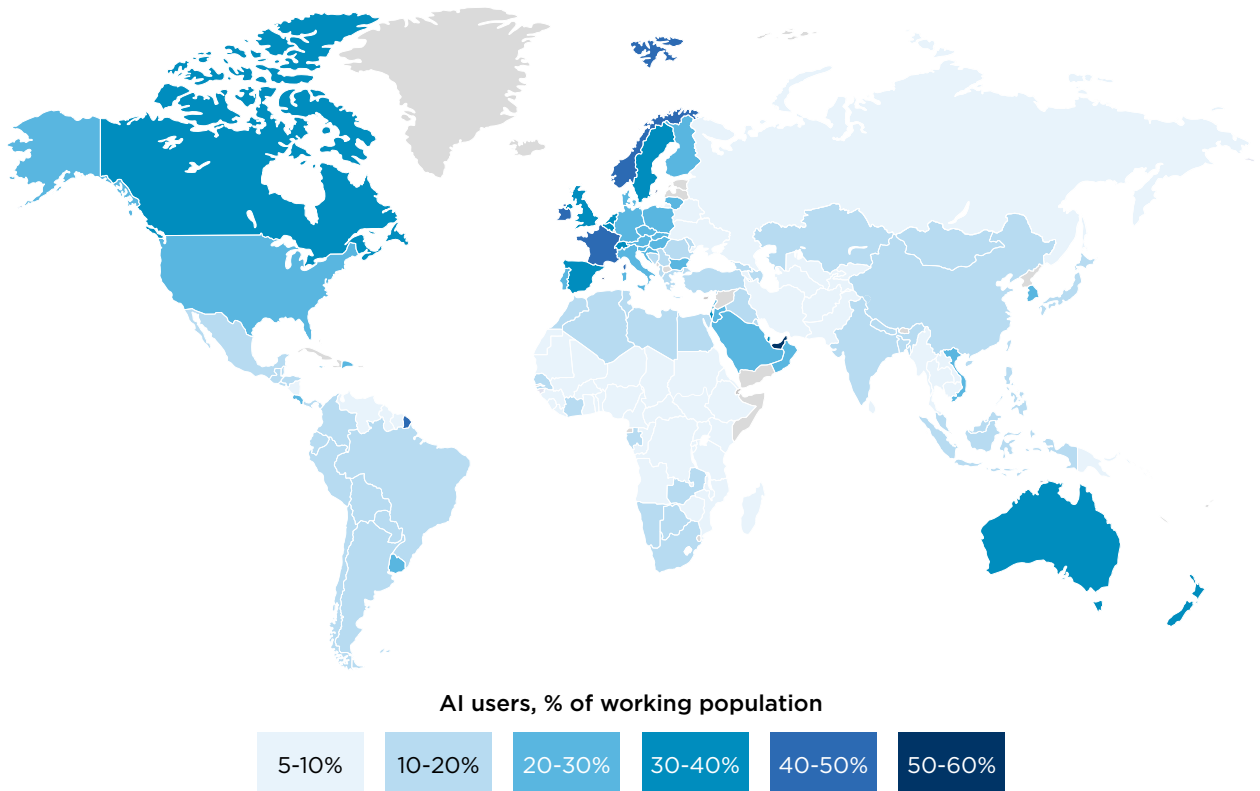
This rapid diffusion underscores AI’s transition from an experimental technology to a mainstream driver of productivity and competitiveness. Microsoft estimates that around 15% of the global

working-age population now use AI tools in their daily work. Adoption is highest in the United Arab Emirates and Singapore, where more than half of workers report using AI. Most other high-adoption economies are concentrated in Europe and North America, as shown in Figure 7, reflecting greater access to digital infrastructure and enterprise AI systems.⁷

Comparing across regions, Asia Pacific and Japan (APJ)⁸—the regions of focus of this study—lag behind, with AI user share at 13.5% of the working population, significantly below North America (33.5%) and Europe & Central Asia (21.7%).

4 Bin Rashid, A., Kausik, A., “AI revolutionizing industries worldwide: A comprehensive overview of its diverse applications”, Hybrid Advances 7, 2024, accessed December 2025.
 5 Singla, A., et al., “The State of AI in 2025: Agents, Innovation, and Transformation”, McKinsey & Company, 2025, accessed November 2025.
 6 Stack Overflow, “2024 Stack Overflow Developer Survey”, 2024, accessed December 2025.
 7 AI Economy Institute, “AI Diffusion Report: Where AI Is Most Used, Developed, and Built”, Microsoft, 2025, accessed December 2025.
 8 This study defines Asia Pacific and Japan (APJ) as comprising the following economies: Bangladesh, Cambodia, India, Indonesia, Japan, South Korea, Lao PDR, Malaysia, Myanmar, Nepal, Pakistan, Philippines, Singapore, Sri Lanka, Republic of China (Taiwan)—henceforth referred to as Taiwan—Thailand, and Vietnam.


Figure 7: Global distribution of AI adoption rates



Source: Oxford Economics.

Note: Adoption rates in the figures on this page and the associated discussion only refer to firms that are experimenting, piloting, scaling or have fully integrated AI in their operations. Adoption rates in the rest of this report refer to firms that have integrated AI in the production of goods and services. It excludes firms that are experimenting, piloting or scaling AI in their operations.

Figure 8: AI usage among the working-age population by region



Source: Microsoft, Oxford Economics

THE SIGNIFICANT ECONOMIC POTENTIAL OF AI

The economic potential of AI is reflected in macroeconomic assessments. International institutions now view AI as a general-purpose technology with the capacity to reverse the post-global financial crisis-slowdown in productivity growth. The IMF argues that, if properly deployed, AI is “our best chance” to relax structural supply-side constraints and generate a sustained resurgence in productivity and GDP growth after more than a decade of stagnation. The OECD similarly highlights AI’s potential to act as a new general-purpose technology reshaping production, innovation, and living standards across a wide range of sectors.^{9,10} Simulations from IMF and OECD indicate that, under plausible adoption paths, AI-driven productivity gains could materially raise long-run output in advanced and emerging economies alike, even after accounting for transition costs and labour-market adjustment.^{11,12}

Firm-level experiments and early deployments reinforce this view. Controlled trials across customer service, professional writing, software development, and management consulting typically report efficiency gains in the range of 10%–55%, with average improvements of around 25% in task completion speed and quality—evidence that AI is already augmenting workers’ cognitive capabilities rather than simply automating routine tasks.¹³

In the APJ region, AI is increasingly recognised as a pillar of long-term economic strategy. The World Economic Forum describes the region

as having entered a “golden era of the digital economy”, with AI and computing power emerging as the core productivity engine of that transformation.¹⁴ An ADB study estimates that AI-related spending and its spillover effects generated around US\$247.1 billion in economic activity across Asia-Pacific in 2024 alone, with indirect productivity spillovers accounting for roughly 87% of this total.¹⁵

Country-level assessments point to sizeable upside potential. For example, in Japan generative AI is estimated to unlock roughly JPY 148.7 trillion (around US\$1.1 trillion) in productive capacity, contingent on effective adoption and workforce transitions.¹⁶ More broadly, the World Bank and Asian Development Bank emphasise that digital technologies, including AI, have already enhanced productive capacities, expanded opportunities, and improved government services. These technologies will be central to sustaining growth as populations age and traditional growth engines mature.¹⁷

While APJ currently lags the global frontier in the diffusion of AI, there is significant benefit from closing this gap. Not only would it narrow regional disparities in productivity and wages but also position APJ to capture a larger share of the global AI dividend. Faster AI adoption—underpinned by investment in digital infrastructure, skills, and balanced AI policies—could deliver substantial increments to regional GDP and labour productivity over the coming decade.

9 Spence, M., “AI’s promise for the Global Economy”, IMF, 2024, accessed December 2025.

10 OECD, “The impact of Artificial Intelligence on productivity, distribution and growth”, 2024, accessed December 2025.

11 Spence, M., “AI’s promise for the Global Economy”, IMF, 2024, accessed December 2025.

12 OECD, “Macroeconomic productivity gains from Artificial Intelligence in G7 economies”, 2025, accessed December 2025.

13 Alex, A., “The Projected Impact of Generative AI on Future Productivity Growth”, Penn Wharton Budget Model, 2025, accessed December 2025.

14 Lin, S., “How accelerating AI is the foundation for industry intelligence in Asia-Pacific”, World Economic Forum, 2023, accessed December 2025.

15 Asian Development Bank, “Leveraging Artificial Intelligence and Cloud Computing to Accelerate Growth in Asia and the Pacific”, 2025, accessed December 2025.

16 Grigorian, S., “5 ways Asia-Pacific economies can operationalise AI to unlock economic opportunity”, World Economic Forum, 2023, accessed December 2025.

17 Asian Development Bank, “Digital Technology”, accessed December 2025.

SECTION 2. GOVERNMENT'S ROLE IN DIFFUSING AI AND SAFEGUARDING TRUST

AI policy is shifting from an emphasis on frontier breakthroughs toward a broader agenda that prioritises widespread deployment and adoption across the economy. Policymakers increasingly recognise that the most immediate and substantial economic gains will come from integrating practical, sector-ready AI tools—often less complex than frontier systems—into business operations, public services, and critical infrastructure. Evidence from multiple studies shows that generative and assistive AI can deliver double-digit improvements in task speed and quality, underscoring the importance of accelerating adoption—a theme discussed in the previous chapter.



As AI's economic potential becomes clearer, policy frameworks are rapidly evolving. The OECD Principles on Artificial Intelligence offer governments a shared foundation for advancing diffusion responsibly, emphasising inclusive growth, respect for human rights, transparency,

robustness, and accountability.¹⁸ These standards now guide national AI strategies, regulatory measures, and public-sector adoption across advanced and emerging economies, as shown in Figure 9.

Figure 9: Priorities to consider when shaping AI policies



Source: OECD Principles on Artificial Intelligence

THE RISE OF DIGITAL SOVEREIGNTY

Data, compute, and semiconductors have become strategic assets. AI has shifted from a purely technological domain to a central feature of economic statecraft. Export controls on advanced chips and manufacturing equipment—particularly between the United States and China—illustrate how governments are leveraging technology dependencies to advance their geopolitical objectives.^{19,20} Several governments have moved to secure large allocations of advanced AI chips and associated compute infrastructure. Recent

examples include Saudi Arabia's PIF-backed HUMAIN agreement with Nvidia for an initial 18,000-GPU deployment, and the UAE's state-backed push—via G42 and US-approved export arrangements—to procure advanced Nvidia processors and build national-scale AI capacity.^{21,22}

AI sovereignty builds on digital sovereignty, which seeks control over critical digital infrastructure, data, and services. While digital sovereignty addresses broad technology ecosystems, AI

18 OECD, "OECD Updates AI Principles to Stay Abreast of Rapid Technological Developments", 2024, accessed December 2025.

19 Allen, G., Goldston, I., "Understanding US Allies' Current Legal Authority to Implement AI and Semiconductor Export Controls", Center for Strategic and International Studies, 2025, accessed December 2025.

20 Barczentewicz, M., "US Export Controls on AI and Semiconductors", International Centre for Law and Economics, 2025, accessed December 2025.

21 AP News, "Nvidia to send 18,000 AI chips to Saudi Arabia", May 2025, accessed January 2026.

22 Albergotti, R., Warner, K., "How the UAE got the US to bless its AI ambitions", Semafor, September 2024, accessed January 2026.

sovereignty goes further by focusing on the unique requirements of AI systems such as access to high-performance compute, foundational models, specialised tooling (including MLOps platforms), and AI-specific skills. It aims to secure greater discretion across the three core layers of AI: compute infrastructure, AI models, and user-facing applications, while balancing openness where it accelerates adoption and resilience.

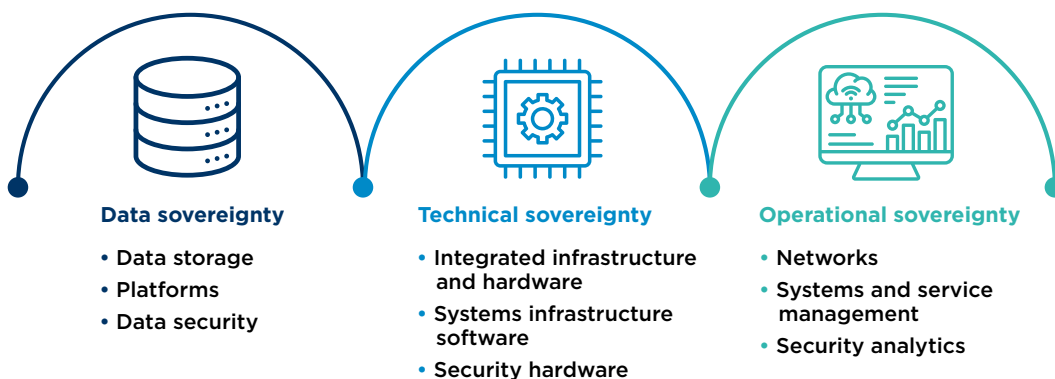
As AI adoption accelerates, policymakers face a dilemma. On one hand, governments recognise that AI is becoming a core driver of productivity, innovation, and long-term competitiveness. On the other, they must manage legitimate concerns about dependence on foreign technology providers for critical AI infrastructure and the geopolitical leverage such dependence creates, and societal cohesion.

In addition, a related consideration is how sovereignty is interpreted in practice. Nicklas Lundblad, Senior Fellow with the Tech Policy Program at the Center for European Policy Analysis, observes that many governments still default to a territorial understanding of

sovereignty centred on physical control of data centres and compute infrastructure. This approach, motivated by a desire to be secure and resilient, can be misleading, as it risks directing attention and resources toward the most capital-intensive and rapidly depreciating parts of the AI stack rather than the areas where countries are best positioned to build durable capability—such as skills, high-quality datasets, assurance, and sector-specific applications [see page 25].

This has brought questions of digital sovereignty to the forefront of the policy agenda. Digital sovereignty can be understood as a state’s ability to exercise effective control, oversight, and strategic choice over the digital infrastructure, data, and technologies that underpin its economy and security. It encompasses not only where data are stored or processed, but also who designs and operates core systems, which legal regimes apply, and how far national authorities can exercise effective oversight of these systems—a framing that increasingly extends to AI-specific concerns such as ensuring that models reflect domestic values, languages, and regulatory priorities.²³

Figure 10: Aspects of digital sovereignty²⁴



Source: IDC (2024)

23 Roberts, S., “‘Digital Sovereignty key’ amid growing AI investment, geopolitical tensions”, Verdict, 2025, accessed December 2025.

24 Francis, L., Sharma, R., “State of Sovereign and Industry Cloud Investment by Asia/Pacific Governments”, IDC, 2024, accessed December 2025.

MOTIVATIONS UNDERPINNING THE RISE OF DIGITAL SOVEREIGNTY

The policy objectives underpinning AI sovereignty are multifaceted. For many governments, including those in APJ, sovereignty policies are motivated by a blend of technological self-reliance, national security, economic competitiveness, resilience, political motivations, and cultural autonomy—aimed at ensuring that critical digital systems remain under trusted and accountable control. While national priorities differ, the underlying motivations are broadly similar and are summarised in Figure 11.

Economic motives are prominent. Many governments see AI sovereignty as a way to capture economic value via domestic AI development. This can help economies move up the value chain and strengthen long-run productivity. By investing in domestic data assets, model development, and cloud or on-premises infrastructure, they aim to capture more of the value created along the AI stack, support local talent and startups, and retain a greater share of economic value within the domestic economy.

For APJ economies—which vary widely in digital maturity—these strategies serve different aims: for more advanced economies, positioning as regional or global AI leaders; for others, securing a foothold in fast-growing AI-related sectors and narrowing, rather than widening, the competitiveness gap with global leaders.

This framing often overlooks the scale of foreign investment that currently underpins AI ecosystems in many APJ economies. International cloud and technology firms have made significant commitments to local data centre infrastructure, workforce skilling initiatives, and partnerships with domestic institutions. These investments

contribute directly to innovation capacity and economic development, meaning that an exclusive focus on sovereignty may understate the benefits of continued external participation.

Security motives are equally important. As AI becomes embedded in critical services, policymakers are placing greater emphasis on resilience, continuity of access, legal oversight, and exposure to external dependencies across the digital supply chain.²⁵ Geopolitical tensions and export controls on advanced chips and AI technologies have highlighted the fragility of cross-border supply chains. In this context, AI sovereignty policies are designed to protect sensitive government, defence, financial, and critical-infrastructure systems. They aim to ensure that data and models that are central to national security remain under domestic jurisdiction and are therefore less prone to potential disruptions in technology supply.

Societal considerations add a third dimension. Governments are increasingly concerned that AI systems should reflect domestic ethical standards, cultural norms, and linguistic diversity, rather than importing assumptions embedded in models used for training. A recent study by Song et al. (2025) revealed that while the response of LLMs differed based on the language they were prompted in—English or Chinese, the cultural values aligned towards the norms of linguistically dominant cultures (i.e., in English the responses represented more individualistic tendencies).²⁶ Sovereign AI is therefore equally about ensuring that AI systems reflect local languages and cultural contexts, and remain accessible to relevant local populations, as it is about control and ownership of the underlying technology.

25 Gambacorta L., Shreeti V., "The AI Supply Chain", Bank of International Settlement, 2025, accessed December 2025.

26 MIT Sloan Office of Communications, "Generative AI's hidden cultural tendencies", MIT Management Sloan School, 2025, accessed December 2025.

Figure 11: Motivations for pursuing AI sovereignty^{27,28,29,30,31}

Drivers	Key motivations
Economic considerations: Developing domestic AI and tech competitiveness	<ul style="list-style-type: none"> • Capturing value from the AI stack (by building national data, models, compute, and applications). • Develop domestic competitiveness through local talent, startups, and infrastructure to drive productivity and innovation.
Economic and security considerations: Increasing technological self-reliance and building domestic capability	<ul style="list-style-type: none"> • Reduce the country’s dependence on leading foreign technology powers for core digital infrastructure, compute, chips, and cloud services. • Safeguard against potential geopolitical leverage exerted through a nation’s dominance in global AI value chains (e.g., export controls, sanctions). • Reduce exposure to foreign legal and jurisdictional requirements.
Security considerations: Building secure and resilient national AI systems	<ul style="list-style-type: none"> • Ensure that national data, infrastructure, and AI systems in critical sectors are protected from and/or can withstand foreign cyberattacks and sabotage. This is especially important for sensitive government, military, and other highly confidential or sensitive data.
Societal considerations: Cultural and ethical sovereignty	<ul style="list-style-type: none"> • Ensuring AI reflects national values, languages, regulatory standards, and strategic priorities.

Source: Bain & Company, Kearney, Atlantic Council, Accenture.

27 Gerosa, M., et al., “The State of Sovereign AI”, The Linux Foundation, 2025, accessed December 2025.

28 Hoecker, A., et al., “Sovereign Tech, Fragmented World”, Bain & Company, 2025, accessed December 2025.

29 Dobberstein, N., et al., “Will sovereign AI be a game changer?”, Kearney, 2024, accessed December 2025.

30 Ray, T., “Sovereign remedies: Between AI autonomy and control”, Atlantic Council, 2025, accessed December 2025.

31 Wood, D., et al., “Sovereign AI: Own your AI future”, Accenture, 2025, accessed December 2025.

NICKLAS LUNDBLAD: REFRAMING NATIONAL APPROACHES TO AI SOVEREIGNTY

“At what point do we start—now or wait for the industry and applications to mature? That’s a key trade off countries grapple with when engaging with advanced technologies,” highlighted Nicklas Lundblad, thought leader, writer, researcher, and public policy expert with decades of experience in some of the world’s most prestigious technology organisations.

According to Nicklas, the increasingly popular policy pivot towards onshoring AI infrastructure and supply chains in pursuit of AI sovereignty is somewhat misguided. Many equate sovereignty with territorial control over data centres, compute, and cloud infrastructure. Yet that approach, which Nicklas calls “territorial sovereignty” approach, diverts valuable resources away from the areas where most countries can genuinely build comparative advantage; productivity-enhancing AI applications, services, and solutions.

Attempting to replicate the full AI stack domestically is not only prohibitively expensive but also strategically inefficient, given the rapid pace at which underlying technologies become obsolete. The idea that each country should develop their own domestic cloud and computing market is unrealistic.

To move beyond this narrow framing, Nicklas proposes a more nuanced understanding of AI sovereignty built around two alternative models: functional and reciprocal sovereignty.

Functional sovereignty relies on guarantees, safeguards, and controls (assurance-led) rather than physical ownership of infrastructure. For example, international cloud providers have for years been offering “sovereignty-as-a-service”,

through innovative solutions such as customer-controlled encryption keys, jurisdictional controls, auditable security assurances, and many other innovative and creative solutions. Governments can in turn reinforce these arrangements through regulatory tools, such as sanctions, to ensure compliance and mitigate risks. This, Nicklas argues, allows countries to capitalise on their competitive advantage, enhance their sovereignty, while maintaining access to global cutting-edge technology.

The other is reciprocal sovereignty, which is based on interdependence and trade integration. Nicklas shared that in this scenario, a country specialises in something that other countries desire and rely on. This form of sovereignty emerges when a nation develops expertise, technologies, or services that are indispensable to global partners. In this case, specialised countries are de facto sovereign as other countries depend on them. Achieving this, however, requires sustained investment in skills, innovation, and research and development.

Pursuing territorial sovereignty, Nicklas suggests, “takes money away from what could produce your reciprocal sovereignty”. While the three types of sovereignty are not entirely mutually exclusive, governments need to weigh their competitive advantage and other ways through which they can increase their sovereignty, resilience, security. Similar to the approach in other industries such as manufacturing, energy, or agricultural production, a more sophisticated approach to sovereign AI, Nicklas concludes, lies not in replicating global supply chains but in leveraging them while building unique strengths at home.

AI SOVEREIGNTY POLICIES IN APJ THROUGH THE LENS OF THE AI STACK

Several governments in Asia Pacific are positioning their economies as global or regional leaders in AI. Concerns about technology concentration risk and geopolitical tensions are leading some governments to advocate for “AI sovereignty”. While governments’ definitions and approaches to “AI sovereignty” differ, common elements include measures to localise various elements of AI infrastructure to promote domestic industry and national economic resilience.

Viewing AI sovereignty policies through the lens of a three-layer cloud stack—compute infrastructure, AI models, and AI applications—offers a clear framework for understanding how organisations access and deploy AI. As shown in Figure 12 and detailed below, this perspective explains adoption dynamics by highlighting how

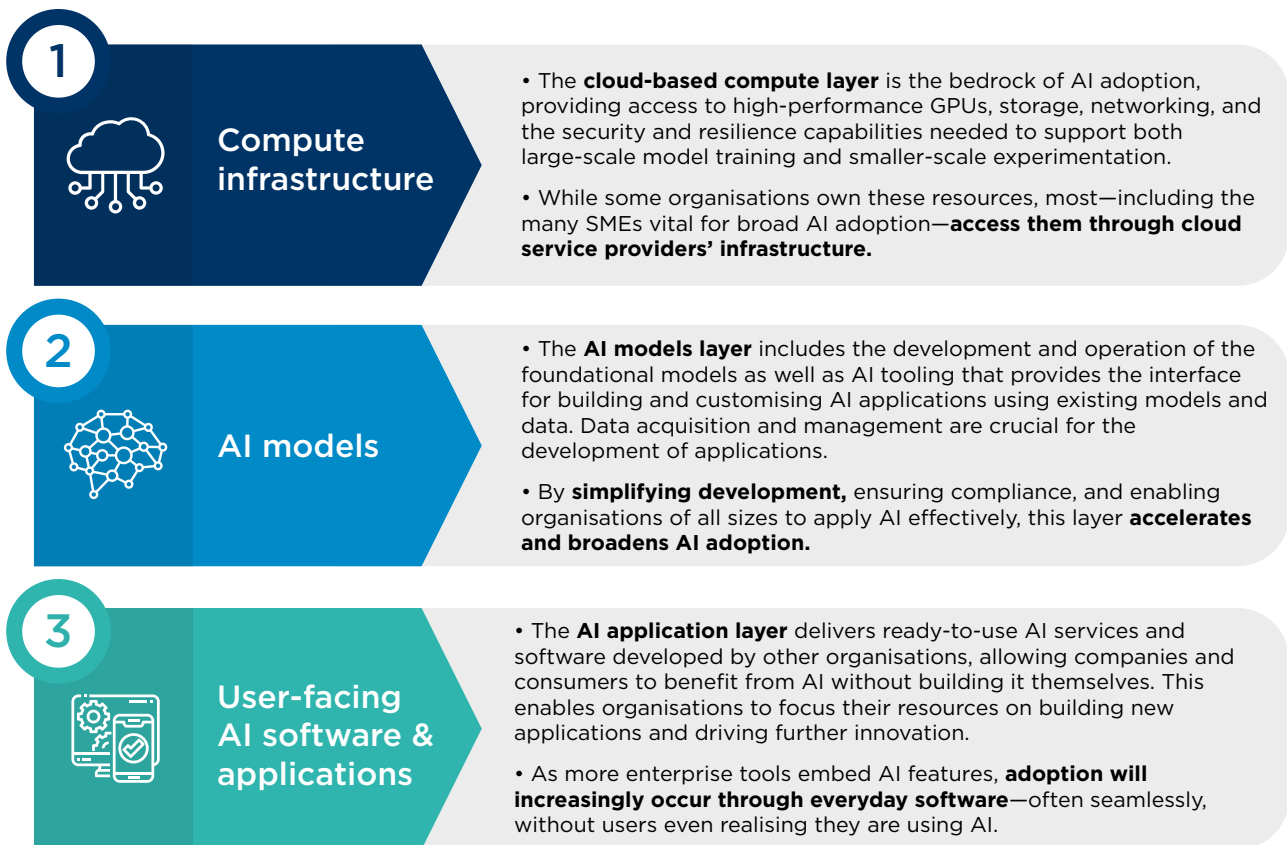
sovereignty-driven regulations influence each layer.

COMPUTE INFRASTRUCTURE

Compute infrastructure is the foundation of the AI stack. Compute infrastructure underpins AI development and deployment, providing access to GPUs, storage and networking—primarily via cloud platforms. As most firms, particularly SMEs, rely on cloud providers rather than owning infrastructure, sovereignty measures at this layer focus on the conditions under which cloud services can be used, especially for public-sector workloads.

Some economies adopt assurance-based approaches. For example, Japan allows government use of commercial clouds certified under its ISMAP programme and promotes

Figure 12: Three key layers in the AI stack



Source: Oxford Economics

trusted cross-border data flows through its “Data Free Flow with Trust” framework (also see Box 1).^{32,33} Singapore follows a similar model through its Government on Commercial Cloud (GCC 2.0), which enables agencies to use major cloud platforms within a government-secured environment.^{34,35}

Other economies apply hybrid or more restrictive models. India follows a hybrid approach that leverages global cloud service providers through its MeghRaj 2.0 initiative—mandating common standards and India-based data centres—while simultaneously expanding national AI compute capacity under the IndiaAI Mission.³⁶ The recent budget proposes long-term tax incentives for global cloud and data center players with a view to attract global know-how while deepening domestic capabilities.³⁷ Malaysia operates a hybrid-plus-sovereign model, with the public sector using MyGovCloud, which pairs a government-run environment with approved global CSPs, while the Strategic AI Infrastructure launched in May 2025 provides a sovereign environment for sensitive AI workloads (also see Box 2). Budget 2026 also announced a ‘Sovereign

AI Cloud’ initiative to expand domestically governed AI compute capacity.^{38,39,40,41} Indonesia permits offshore processing for much of the private sector but retains localisation requirements for certain strategic data.⁴² Taiwan represents a unique case as it does not impose broad, economy-wide restrictions on the use of global cloud service providers, but benefits from a near end-to-end semiconductor and AI-hardware supply chain that enables sovereign-grade compute capacity—reflected in Nvidia’s expanding AI infrastructure projects and supercomputing investments in Taiwan.⁴³ South Korea applies stricter controls through its Cloud Security Assurance Program (CSAP), which requires providers to meet detailed security standards for public-sector use.^{44,45,46} Recent procurement moves suggest a more pragmatic opening: in 2026, the National AI Computing Center call for proposals was opened to domestic and international cloud service providers and related consortiums, even as CSAP certification requirements remained in force for sensitive workloads.⁴⁷

32 Digital Agency of Japan, “[Updated ISMAP cloud services list](#)”, 2024, accessed December 2025.

33 Digital Agency of Japan, “[Data Free Flow with Trust \(DFFT\)](#)”, 2024, accessed December 2025.

34 Pradana, D., “[Singapore’s Government on Commercial Cloud \(GCC\): 2025 Guide to a Secure, Scalable Public-Sector Cloud](#)”, Accrets International, July 2025, accessed December 2025.

35 Ministry of Digital Development and Information (MDDI), “[Factsheet - Government Cyber Security Operations Centre \(GCSOC\)](#)”, February 2023, accessed December 2025.

36 Kumar, S., et al., “[India’s AI Revolution: A Roadmap to Viksit Bharat](#)”, Ministry of Electronics and Information Technology, 2025, accessed December 2025.

37 The Times of India, “[Union budget 2026: Tax break aims to clear skies for cloud players](#)”, February 2026, accessed February 2026.

38 Raj, A., “[Malaysia has a new government hybrid cloud service](#)”, Techwire, May 2022, accessed December 2025.

39 MyGovCloud, “[MyGovCloud](#)”, 2025, accessed December 2025.

40 Bernama, “[Malaysia Launches Region’s First Sovereign Full-Stack AI Infrastructure](#)”, May 2025, accessed December 2025.

41 MyDIGITAL, “[Budget 2026: Accelerating Malaysia’s Digital Transformation for All](#)”, 2025, accessed March 2026.

42 Information Technology and Innovation Foundation, “[Indonesia’s Data Localisation Regulation](#)”, 2025, accessed December 2025.

43 Channel News Asia, “[Taiwan opens new cloud centre to bolster ‘sovereign AI’ effort](#)”, December 2025, accessed February 2026.

44 Microsoft Learn, “[Korea CSAP](#)”, August 2025, accessed December 2025.

45 Wong, S, “[Explainer: Korea’s AI Basic Act](#)”, Asian Legal Business, 2025, accessed December 2025.

46 Kim, J., et al., “[2025 AI Governance in Korea: Strategic Investments and Regulatory Reforms](#)”, Jipyong LLC, 2025, accessed December 2025.

47 Recent Korean government AI tenders illustrate this shift. In 2026, two major national AI projects removed long-standing domestic-company requirements from their request for proposals (RFPs)—rules that had previously excluded foreign cloud providers. Earlier versions of these tenders (through 2025) required either a Korean main business location or restricted bidding to domestic firms, but the 2026 releases replaced these with more open eligibility criteria. CSAP obligations, however, continue to apply to all public-sector cloud services.

AI MODELS

The model layer supports the development, fine-tuning, and deployment of AI systems. Policy attention here reflects concerns around safety, ethics, and alignment with domestic norms. Singapore has operationalised this approach through tools such as AI Verify and its Model AI Governance Framework, which provide practical guidance on testing, transparency, and risk management.^{48,49} Japan similarly relies on non-binding guidelines and voluntary compliance, emphasising international alignment and flexibility.^{50,51,52} Taiwan also follows a voluntary governance approach, issuing non-binding guidelines and supporting local-language model development through state-backed supercomputing initiatives rather than imposing binding model-level regulation.⁵³ India has recently published AI Governance Guidelines that set out principles for responsible and ethical AI development without imposing binding rules, complemented by voluntary technical standards such as the TEC's fairness-assessment framework.^{54,55} South Korea has moved further toward formal regulation, introducing a tiered framework under its AI Basic Act that imposes additional obligations on high-impact AI systems⁵⁶, Malaysia⁵⁷ and Indonesia^{58,59} are converging on baseline safeguards—issuing ethical guidelines

and transparency requirements while national strategies continue to evolve. Malaysia, in particular, is moving from its 2024 National Guidelines on AI Governance and Ethics towards binding legislation: the AI Governance Bill that is expected to be tabled in late 2026, with a focus on reducing risks in sensitive sectors such as healthcare and biometrics.⁶⁰

Broadly, across APJ, regulators are converging on governance that supports innovation while mitigating risk, rather than mandating local-only development.

USER-FACING AI SOFTWARE AND APPLICATIONS

The AI application layer delivers ready-to-use AI services and software created by other organisations, often embedded in everyday enterprise software, enabling widespread adoption without in-house development.

Governments across APJ are increasingly setting expectations for responsible use, with differing degrees of formality and enforcement.

Singapore's AI Governance Framework for Generative AI sets clear expectations for developers and deployers around explainability, human oversight, and safety, complemented by

48 Infocomm Media Development Authority (IMDA), “Singapore launches world’s first AI testing framework and toolkit to promote transparency; Invites companies to pilot and contribute to international standards development”, 2022, accessed December 2025.

49 IMDA, “Singapore proposes framework to foster trusted Generative AI development”, 2024, accessed December 2025.

50 Inoue, K., Kamata, C., “Japan’s emerging framework for responsible AI: legislation, guidelines and guidance”, International Bar Association, 2025, accessed December 2025.

51 Ministry of Economy, Trade and Industry (METI), “AI Guidelines for Business Ver 1.0 Compiled”, 2024, accessed December 2025.

52 Habuka, H., “Japan’s Approach to AI Regulation and Its Impact on the 2023 G7 Presidency”, Center for Strategic and International Studies, 2023, accessed December 2025.

53 White & Case, “AI watch global regulatory tracker Taiwan”, February 2026, accessed February 2026.

54 IndiaAI, “India AI Governance Guidelines”, November 2025, accessed December 2025.

55 TEC, “Fairness Assessment and Rating of Artificial Intelligence Systems”, 2023, accessed December 2025.

56 Lee, S., “South Korea’s Evolving AI Regulations”, Stimson, 2025, accessed December 2025.

57 Buza, M., van Mutius, S., “DPA Digital Digest: Malaysia”, Digital Policy Alert, 2024, accessed December 2025.

58 Rolindrawan, W. Y., Ismayudha, Q. P., “Artificial Intelligence Comparative Guide”, Mondaq, 2025, accessed December 2025.

59 Virgiany, M. Amatullah, N., “Ethical guidelines on use of artificial intelligence (AI) in Indonesia”, Hiswara Bunjamin and Tandjung, 2024, accessed December 2025.

60 Isamudin, D., “Malaysia’s First AI Bill to Be Tabled by Mid-2026”, New Straits Times, August 2025, accessed April 2026.

data protection rules under the PDPA.^{61,62} Japan encourages responsible deployment through voluntary AI Business Guidelines. Taiwan adopts a similar guidance-led approach, relying on sector-specific rules and voluntary administrative guidelines—developed under its AI Action Plan 2.0 and reinforced by the AI Basic Act—to support responsible AI use without restricting access to foreign AI applications.⁶³ South Korea has adopted a more prescriptive stance, with specific

legal restrictions on harmful uses of AI—such as deepfakes and election-related disinformation—alongside broader obligations under the AI Basic Act.⁶⁴ Other economies, including Thailand and the Philippines⁶⁵, rely on combinations of AI-specific guidance and data protection laws to govern applications, focusing on accountability and privacy rather than restricting access to foreign-developed AI services.

KEY INSIGHTS FROM AI SOVEREIGNTY-RELATED POLICIES

Taken together, these examples highlight a spectrum of AI sovereignty measures across APJ—from assurance-led and open models to more restrictive or state-led approaches—applied differently across the AI stack. Three key points emerge:

- First, there is **no single model for AI sovereignty**. Approaches across APJ range from open, assurance-led frameworks to more restrictive, state-led models. Most economies impose some constraints—such as data residency rules or certification requirements for foreign cloud providers—which influence who can supply AI services and how quickly adoption can scale.
- Second, **restrictions on cloud provision can slow access to AI**. While localisation and ownership preferences can support domestic capability, they may delay access to scalable compute and AI services. Evidence shows that AI adoption is strongest when enabled by cloud services, highlighting the need to balance sovereignty objectives against the risk of constraining AI-driven growth.⁶⁶
- Third, **public-sector procurement could play a pivotal role**. Government use of AI often sets the pace for wider adoption: survey evidence shows that firms are significantly more likely to adopt AI when the public-sector leads.⁶⁷ Restrictions applied to public-sector AI and cloud services therefore have spillover effects on the private sector.

61 IMDA, “[Singapore proposes framework to foster trusted Generative AI development](#)”, 2024, accessed December 2025.

62 Chia, O., “[Temporary deepfake ban discussed as way to tackle AI falsehoods during Singapore election](#)”, The Straits Times, 2024, accessed December 2025.

63 Shao, G., Shih, S., “[Taiwan: AI basic Act](#)”, Baker McKenzie, 2026, accessed February 2026.

64 Ministry of Science and ICT, “[Technology and Innovation, A New Chapter in the Age of AI](#)”, 2024, accessed February 2026.

65 National Privacy Commission, “[NPC’s new initiative on ASEAN cross-border tools to boost PH digital competitiveness](#)”, 2022, accessed December 2025.

66 Asian Development Bank (ADB), “[Leveraging artificial intelligence and cloud computing to boost economic impact in Asia and the Pacific \(ADB Brief No. 352\)](#)”, 2025, accessed December 2025.

67 AWS, “[Unlocking AI potential country reports \(Japan and Australia\)](#)”, 2025, accessed December 2025.

BOX 1: JAPAN'S SOVEREIGN AI STRATEGY: BUILDING SELF-RELIANCE THROUGH GLOBAL COLLABORATION

Japan is seeking to strengthen domestic AI capability while maintaining a highly collaborative and internationally aligned policy framework. The government's overarching policy is to make Japan "the world's most AI-friendly country for development and utilisation," balancing aggressive innovation with prudent risk management. Former Prime Minister Fumio Kishida, for example, has championed global AI governance: under Japan's G7 presidency he launched the Hiroshima AI Process to craft common guidelines for generative AI, a framework now supported by more than 50 countries and regions.⁶⁸

Balancing autonomy with international engagement

One pillar of Japan's approach is boosting computing power on home soil. For this the government has pledged roughly JPY 115 billion (~US\$740 million) in subsidies to expand Japan's AI cloud and supercomputing infrastructure.⁶⁹ This includes partnerships with foreign tech leaders like Nvidia and AWS.⁷⁰ By working with international chip and cloud experts, Japan ensures its researchers and companies have access to world-class computing within national borders. The approach secures critical infrastructure at home while tapping global know-how gives Japan a robust, independent foundation for AI development.

Home-grown AI models and innovation

Japan is also funding the creation of its own AI models, including large language models (LLMs)

tuned to Japanese language and culture. A flagship effort is the Fugaku LLM project, which uses Japan's Fugaku supercomputer to train an AI model on Japanese and English datasets.⁷¹ Crucially, this model is intended to be open and widely available, reflecting Japan's strategy of fostering innovation through transparency. Developing indigenous AI models provides Japan with greater strategic control, while maintaining an open approach ensures continued integration with the global research community. This dual strategy enables both domestic and international researchers to build on Japanese innovations, fostering knowledge exchange. It represents a balanced model of autonomy and openness—allowing Japan to tailor AI solutions to national priorities while contributing to and benefiting from global advancements.

AI adoption & governance—leading by example

Japan's strategy promotes widespread AI use domestically, guided by light-touch regulation aligned with international norms. The government's AI plan calls for public institutions to lead by example in adopting AI for public service.⁷² To regulate these advances, Japan favours flexible guidelines over strict laws.^{73,74} At the same time, Japan actively engages in global rulemaking for AI. For Japanese companies and developers, this policy environment means they can innovate rapidly at home (due to fewer regulatory hurdles) and easily operate abroad (as Japan's AI governance is in harmony with global norms).

68 The Government of Japan, "The Hiroshima AI Process: Leading the Global Challenge to Shape Inclusive Governance for Generative AI", February 2024, accessed December 2025.

69 Nvidia, "NVIDIA to help elevate Japan's Sovereign AI efforts through Generative AI infrastructure build-out", May 2024, accessed December 2025.

70 AWS, "Japan's digital agency accelerates government cloud migration with AWS generative AI-powered architecture reviews", September 2024, accessed December 2025.

71 Fujitsu, "Release of 'Fugaku-LLM' - a large language model trained on the supercomputer 'Fugaku'", May 2024, accessed December 2025.

72 CSIS, "Japan's Agile AI Governance in Action: Fostering a Global Nexus Through Pluralistic Interoperability", October 2025, accessed December 2025.

73 Inoue, K., Kamata, C., "Japan's emerging framework for responsible AI: legislation, guidelines and guidance", July 2025, accessed December 2025.

74 METI, "AI Guidelines for Business Ver 1.0 Compiled", 2024, accessed December 2025.

BOX 2: MALAYSIA'S NATIONAL DRIVE TOWARDS BEING AN AI NATION BY 2030

Malaysia's government is pursuing an AI sovereignty strategy aimed at transforming the country into an AI-driven nation by 2030.⁷⁵ This strategy balances national interests such as data security and technological autonomy with the need to foster innovation and economic growth. Malaysia is investing heavily across the AI stack, from hardware to applications, to build domestic capabilities and position the country as a regional leader in AI.

Bolstering domestic compute infrastructure

A cornerstone of Malaysia's AI sovereignty policy is developing local AI computing infrastructure. In May 2025, Malaysia launched a Strategic AI Infrastructure, the region's first sovereign full-stack AI ecosystem, which brings data storage, high-performance AI servers, and cloud services under national control.⁷⁶ By localising data centres and large-scale AI computing, sensitive data can be processed and stored within Malaysia's borders, safeguarding user privacy and national data security. This push builds on Malaysia's broader data-centre boom where the country approved over RM 114.7 billion in data centre and cloud investments from 2021 to 2023, leveraging ample land, affordable power, and strategic location to become a regional digital hub.⁷⁷

This infrastructure is expected to sit alongside Malaysia's hybrid public-sector cloud, MyGovCloud, which combines a government-operated cloud with vetted national and global cloud providers. Budget 2026 also announced that the Malaysian Communications and Multimedia Commission (MCMC) is investing RM 2 billion to develop a "Sovereign AI Cloud" for Malaysia, a government-backed platform to provide domestic AI computing capacity for

agencies and businesses.⁷⁸ This model allows the government to retain strong assurance, oversight, and data-residency controls while continuing to benefit from the scale, resilience, and innovation of international cloud platforms.

Tailoring solutions to local needs

Malaysia is also championing the development of its own AI models to tailor solutions to local needs. A landmark example is Inteltek Luhur Malaysia Untukmu (ILMU), launched in 2025 as Malaysia's first large-scale multimodal AI model that understands Bahasa Malaysia and local cultural context.⁷⁹ By investing in indigenous AI algorithms, Malaysia aims to ensure AI systems reflect national values and languages. At the application layer, Malaysia's policy encourages widespread AI adoption and responsible use. To spur uptake, the government offers incentives. For example, Budget 2026 proposes a 50% tax deduction for SMEs that invest in AI and cybersecurity training via a programme led by TalentCorp and MyDigital.⁸⁰

Government-wide strategy for mobilising AI capabilities

Authorities have introduced guidelines to keep AI development in check. In late 2024, the government issued National AI Governance and Ethics Guidelines, which outline principles for fairness, transparency, and accountability in AI.⁸¹ Malaysia is also proactively addressing new AI-related opportunities through its forthcoming National AI action plan 2030.⁸² By coupling pro-AI growth policies with ethical guardrails, Malaysia aims to ensure citizens get the benefits of AI-enhanced services without compromising security or public trust.

75 MyDIGITAL Corporation, "Malaysia and WEF Drive ASEAN's Next Leap in AI Governance and Industrial Innovation", 2025, accessed December 2025.

76 Bernama, "Malaysia Launches Region's First Sovereign Full-Stack AI Infrastructure", May 2025, accessed December 2025.

77 Malaysian Investment Development Authority (MIDA), "Malaysia approved RM114.7 billion investments in data centres and cloud services from 2021 to 2023", 2024, accessed December 2025.

78 Bernama, "RM 5.9 billion cross-ministry allocation to keep Malaysia at forefront of AI development-PM Anwar", 2025, accessed December 2025.

79 Institute of Strategic & International Studies Malaysia, "Owning the future: Malaysia's sovereign AI cloud", 2025, accessed December 2025.

80 MyDIGITAL Corporation, "Budget 2026: Accelerating Malaysia's Digital Transformation for All", 2025, accessed December 2025.

81 Ministry of Science, Technology and Innovation, "The national guidelines on AI governance and Ethics", 2024, accessed December 2025.

82 The National AI Office, "Draft of National AI Action Plan 2030", 2025, accessed December 2025.

SECTION 3. AI SOVEREIGNTY AND ECONOMIC OUTCOMES

BALANCING GROWTH, SECURITY, AND DIGITAL SOVEREIGNTY

Governments aiming to balance economic growth with national capacity building, security, and resilience face difficult trade-offs. Many administrations see a need to assert greater control over the digital ecosystems that their economies and defence capabilities. Sovereignty measures, however, carry material costs that extend beyond the upfront investment required to build and operate advanced technology systems. They also impose environmental as well as significant opportunity costs that may not be immediately visible.



A further challenge lies in the way sovereignty ambitions are operationalised. Dr Jaijit Bhattacharya, President of the Centre for Digital Economy Policy in India who we interviewed for this study, notes that effective sovereignty is achieved not through attempting to fully replicate AI infrastructure domestically, but by ensuring resilience and strategic optionality, allowing countries to maintain freedom to operate even as technologies evolve and global dependencies shift. This perspective underscores the need to weigh both the direct and indirect costs of different sovereignty pathways [see page 36].

Our research contributes to the assessment of trade-offs by providing an evidence base on the economic costs of sovereign-AI measures. These costs must be considered when evaluating alternative policy tools to achieve sovereignty objectives.

The remainder of this section outlines these economic costs, and the following section describes how we are quantifying them.

ECONOMIC IMPACTS OF AI SOVEREIGNTY RESTRICTIONS

AI sovereignty measures affect the economy through several policy channels:

- **Data controls** impose localisation, residency, or other data-management obligations. These policies determine where data can be stored, how they are accessed, and under what conditions they can be transferred (often across borders) or processed. These most commonly impact government data, but a wider set of sectors are being covered in policy updates.
- **Provider eligibility and procurement restrictions** determine which firms can supply cloud, AI infrastructure, models, tools, or applications—particularly to the public sector. These may include local-presence requirements, certification thresholds, “buy local” preferences, joint-venture requirements, or restrictions on foreign participation in government tenders. Such measures can support domestic capability-building, but may also narrow supplier choice, reduce competition, and slow access to frontier technologies.
- **Ownership and operational control requirements** govern who can build, own, operate, or manage key components of the AI stack. These restrictions can apply to data centres, cloud platforms, model-hosting environments, and AI service provision. Where requirements favour domestic ownership or

operation, they can limit participation by global providers and increase the need for duplicative local investment.

- **Model development, training, and technology-access requirements** affect how foundation models are developed, fine-tuned, hosted, and updated. These may include expectations to train models domestically, use local datasets, develop local-language models, or rely on locally hosted tooling and inference environments. In some cases, external constraints—such as export controls on advanced chips or AI infrastructure—can further shape the cost and feasibility of building domestic capability.

Although motivated by legitimate security and resilience objectives, these measures can introduce frictions in the rapidly evolving AI ecosystem and can place a drag on the economic benefits from AI.

HIGHER COSTS AND REDUCED EFFICIENCY

Local data-handling requirements and limits on which providers may offer cloud services both significantly raise the cost of delivering AI. When data must remain within national borders, infrastructure and/or applications must have some degree of local ownership, and/or only certain providers are permitted to host or operate cloud services, organisations are pushed toward domestic or restricted-provider data-

centre capacity rather than tapping on the global hyperscale infrastructure.

These constraints can also create fiscal pressures for governments. These constraints can also create fiscal and execution risks for governments. Public investment in GPUs and local AI infrastructure can play an important role in expanding national capability, particularly where it improves access for researchers, startups, SMEs, and public agencies. However, the timing and sequencing of that investment matters. Where large capital commitments are made ahead of clear demand, supporting skills, operating capability, or downstream adoption, there is a risk that capacity is underutilised or delivers lower economic returns than expected. This risk is amplified by long procurement and construction timelines, as AI hardware evolves quickly and equipment may no longer be at the frontier by the time systems are fully operational. The policy challenge is therefore not whether governments should invest, but how to align public investment with ecosystem readiness, expected demand, and the complementary capabilities needed to generate public value.

These pressures ultimately fall on firms, appearing across the entire AI delivery chain. Firms incur higher capital and operating expenditure because restricted-provider or domestic-only facilities have lower utilisation rates and require parallel security, monitoring, and compliance systems. Model development becomes more expensive, as smaller, siloed datasets and less flexible compute resources increase the cost of training and fine-tuning. Software and analytics teams also face higher delivery costs, since restrictions on data movement and cloud-provider participation limit access to cross-border analytics services, APIs, and global model updates, slowing development and reducing efficiency.

Environmental footprint concerns also rise.

The IEA projects global data-centre electricity-use to more than double by 2030⁸³, and local infrastructure buildup accelerates this trend by preventing optimal allocation of resources and reducing utilisation efficiency.

LOWER QUALITY AND CAPABILITY

AI systems rely on scale, both in data and tooling. Restrictions that fragment datasets or limit access to external providers narrow this base and introduce additional cybersecurity trade-offs:

- **AI systems benefit from diverse, large-scale datasets**, so forcing all training data to be local, which can limit access to the most suitable sources and weaken model accuracy and robustness. Local data can still play an important role, but their value comes from being applied selectively—for example, to fine-tune models for domestic languages, regulatory requirements, or sector-specific use cases—rather than replacing broader datasets entirely.
- **Application capability is constrained** when software teams cannot draw on global APIs or cloud-native services, slowing down feature development and limiting integration options.
- **Cybersecurity resources, physical security, and operational resilience may be more limited** when shifting away from global hyperscale providers that invest heavily in cyber defence. Given the smaller size of domestic providers, they may not have the pooled resources to invest appropriately in advanced cybersecurity and resilience measures, which could increase exposure to sophisticated threats.

Over time, these capability gaps widen as global AI ecosystems evolve faster than locally confined alternatives.

SLOWER DIFFUSION AND REDUCED READINESS

Sovereign infrastructure could take years to build, often requiring new facilities, dedicated energy supply, specialised cybersecurity controls and skilled personnel. These long lead times delay the availability of compliant AI services. Once built, sovereign infrastructure also requires sustained reinvestment to remain current—covering hardware refresh cycles, networking upgrades, and ongoing cybersecurity maintenance—adding to the long-run cost of self-supplied capability.

83 IEA, “Energy Demand from AI”, 2025, accessed December 2025.

In parallel, public and private budgets are diverted towards compliance and infrastructure at the expense of workforce training and organisational change, reducing the capacity of firms and agencies to adopt AI. The combined effect is lower readiness, with fewer organisations able to adopt or scale AI solutions.

The policy challenge is therefore not whether governments should invest, but how to align public investment with ecosystem readiness, expected demand, and the complementary capabilities needed to generate public value.

WEAKER COMPETITIVE PRESSURES AND SLOWER SPILLOVERS

Rules that narrow the pool of eligible suppliers (such as those through certification requirements, local-presence conditions, or participation constraints) reduce competitive pressures in early deployments. This often leads to higher prices, slower innovation, and lower-quality solutions in the public sector, which sets many of the standards and reference points that private firms follow.


International evidence shows that when organisations cannot access the full range of global tools, partnerships, and service providers, the effect of spillovers diminishes and diffusion slows across the wider economy.^{84,85,86,87}

84 Katz, R., et al., “[Economic impact of cloud adoption in Asia-Pacific](#)”, Telecom Advisory Services LLC, 2023, accessed December 2025.

85 Wedekind, C., Böhning, C., “[Sovereign cloud usage in Europe. Is there a world without U.S hyperscalers?](#)”, Amaranth Advisory, n.d., accessed December 2025.

86 Yan M., Liu, H., “[The Impact of Digital Trade Barriers on Technological Innovation Efficiency and Sustainable Development](#)”, MDPI, 2024, accessed December 2025.

87 Skare M., Soriano, D., “[How globalization is changing digital technology adoption: An international perspective](#)”, *Journal of Innovation & Knowledge* 6, no. 4 (2021), pp. 222-233, accessed December 2025.



DR JAIJIT BHATTACHARYA: BALANCING OPENNESS AND CAPABILITY—INDIA'S PRAGMATIC PATH TO AI SOVEREIGNTY

“AI sovereignty is not about cutting yourself off from the world. It's about ensuring that your freedom to govern and to grow is never compromised,” says Dr Jaijit Bhattacharya, President of the Centre for Digital Economy Policy in India.

An engineer by training and long-time voice on digital sovereignty, Dr Bhattacharya frames the issue as one of strategic dependence. The core question, he argues, is whether a nation relies predominantly on digital infrastructure governed by a single foreign jurisdiction, or whether it maintains diversified and reliable sources of AI capability. Over-reliance on providers governed by a single foreign legal and political system can expose countries to systemic vulnerabilities—from supply cut-offs to discriminatory provisioning or pricing that indirectly constrains access. In this context, he describes India's approach to AI as resilience-driven: safeguarding continuity of access, preserving choice, and sustaining innovation even as global technological and political conditions evolve.

Crucially, this does not translate into having a closed market for AI. India remains one of the world's most open markets for global technology firms. Resilience requires credible alternatives, especially domestic options that preserve India's freedom to operate even under adverse external conditions. In Bhattacharya's formulation, the objective is not exclusion but optionality: global players should participate freely while India simultaneously develops the capabilities needed to withstand future shocks.

This emphasis on resilience is matched by a strong push for adoption. “Nobody can stop the march of technology. One can only accelerate it. That's what the government is trying to do, in a cautious manner,” Dr Bhattacharya observes. India has therefore refrained from imposing overarching AI regulations that could slow innovation or raise barriers for startups. Regulation, where needed, is intentionally light-touch and adaptive. It is layered onto existing sector rules and led by the relevant

ministries, such as health or transport, rather than through a single economy-wide AI law.

Dr Bhattacharya also draws a distinction between enterprise adoption and citizen-facing risk. In his view, a majority of enterprises make strategic choices about platforms and vendor ecosystems that best suit their needs and often have the governance capacity to manage associated risks. Government intervention should be focused on cases where AI directly affects individuals, where personal data are involved, or where market structures create lock-in and limit meaningful choice. Even then, he argues, interventions should be proportionate and limited.

Where government can be most catalytic is through its own adoption as well as facilitating access to data. Public-sector procurement and deployment of AI solutions can accelerate diffusion across the economy. Further, Dr Bhattacharya points to the value of open APIs—made possible through India's open public digital infrastructure—that allow startups and private firms to build on the same backbone. In cases where raw data cannot be shared, governments can still release models trained on internal datasets, enabling innovation while protecting confidentiality.

At the same time, India is conscious of AI's disruptive effects on jobs and skills. The pace of change, Dr Bhattacharya warns, is faster and more far-reaching than with previous technologies. This will require closer monitoring and targeted interventions to manage short-term displacements and enable long-term gains from productivity and new job creation.

Taken together, India's approach offers a nuanced model for AI sovereignty. By combining openness with domestic capability-building, and light-touch, sector-led governance with active public-sector adoption, India seeks to sustain innovation, resilience, and economic growth in the age of AI.

MAPPING TO THE AI STACK

The two channels—data controls and ownership restrictions—feed through the stack as shown in Figure 13 below:

Figure 13: Impact of data controls and ownership restrictions mapped to the AI stack

AI stack layer	Data controls	Ownership restrictions
Compute infrastructure	<ul style="list-style-type: none"> Higher capital and operating costs Duplicated security systems; lower utilisation efficiency 	<ul style="list-style-type: none"> Fewer qualified suppliers Delayed construction Loss of scale efficiencies Potential lag in using the latest technology Early public GPU procurement increases fiscal risk
AI models	<ul style="list-style-type: none"> Fragmented data reduces model quality On-shore training raises compute cost Delayed fine-tuning and smaller parameter scales 	<ul style="list-style-type: none"> Reduced access to global models/tooling Limited partnerships Higher cost per inference
User-facing AI software and applications	<ul style="list-style-type: none"> Limited access to global APIs and data ecosystems Reduced ability to integrate cross-border services 	<ul style="list-style-type: none"> Smaller app marketplace Lower competition and innovation Public-sector demand confined to domestic vendors

Source: Oxford Economics

The combined outcome is higher costs, lower quality, and longer delays across all layers, which are then transmitted into the wider economy through adoption and productivity channels.

MODELLING THE ESCALATION OF SOVEREIGNTY REQUIREMENTS

To assess the economic impacts of different sovereignty measures, we have defined a set of **stylised policy restrictiveness levels**. These do not represent any single country’s current approach. Rather, they provide a coherent way to examine how policies influence data and cloud infrastructure, and how AI services affect costs, adoption, sustainability, and long-run economic performance. In practice, many governments pursue hybrid sovereignty models, drawing on measures from multiple points along this continuum—tightening controls where they deem the risks are highest while continuing to rely on global providers.

Sovereignty requirements directly shape the pace and scale of AI uptake. Organising these measures into five restrictiveness levels, as

shown in Figure 14, allows us to compare their economic implications in a structured way, while acknowledging that real world configurations may not fit neatly within a single level.

Restrictiveness level 1: Control and choice (assurance led)

In Restrictiveness level 1, policy measures are limited across both the government and private sectors. Public bodies and businesses retain full access to global cloud service providers (CSPs) and a broad range of AI tooling, software, and applications. Data residency requirements apply only to a small subset of highly confidential public-sector workloads, with assurance and oversight concentrated where risks are most acute.

Policies at this level aim to safeguard sensitive data through targeted controls while maintaining openness to global cloud and AI services to maximise innovation and productivity. This balanced approach enables governments to uphold data protection and regulatory oversight without constraining broader economic participation in global AI development. It strengthens trust in high-risk domains while allowing the wider public and private sectors to benefit from state-of-the-art global capabilities. Complementing this, governments focus on enabling adoption through development programmes and industry partnerships rather than through restrictive mandates.

Restrictiveness level 2: Government sector-led restrictions

Restrictiveness level 2 introduces moderate residency and cloud requirements for government workloads, while restrictions on the private sector remain limited. Data residency requirements are prevalent or implicit for public data, and local preferences are applied to public-sector cloud services. CSPs may be required to form local partnerships, operate domestic facilities, integrate government purchased GPUs into sovereign environments, or meet classifications that are more difficult for providers running global or regional cloud systems.

Policies at this level aim to strengthen governmental control and resilience over public-sector data and infrastructure, while preserving open access to global providers for private-sector users. These measures, however, raise the cost and complexity of cloud and AI deployment in government, slowing adoption as public agencies invest time and resources in building local capabilities. Once compliant infrastructure is established, government entities can still use foreign AI tools, software, and applications, moderating the impact on AI functionality relative to more restrictive levels. The private sector continues to benefit from broad access to global CSPs and AI services.

The economic effects at this level materialise through three main channels:

- higher public-sector infrastructure investment requirements;
- slower government AI uptake due to increased costs and reduced supplier choice; and
- weaker diffusion of innovation to the private sector, where public-sector procurement often acts as an anchor.

Restrictiveness level 3: Government sector-led restrictions, which includes added restriction on AI tooling, software, and applications

Restrictiveness level 3 builds on level 2 by extending restrictions in the government sector to AI tooling, software, and applications. Local preferences are applied to AI tooling, software, and applications for the public sector, limiting the services provided by foreign firms. With these restrictions, local firms can build and develop tooling solutions and fine-tune models based on foundation models provided by foreign firms, but store these within the country. The private sector is unrestricted.

Policies at this level aim to build domestic cloud and AI capabilities in the government sector across the entire AI stack. The policy aims to reduce reliance on foreign suppliers, while keeping private-sector use of global services relatively unconstrained. The effect of these restrictions is a further increase in the cost and time required for government to deploy AI solutions, as investment is now needed not only in infrastructure but also in domestic AI tooling, software, and application development. Skills and capacity constraints become more binding, and supplier choice narrows along both the infrastructure and application layers.

As in Restrictiveness level 2, private-sector access to global cloud and AI providers limits the direct negative impact on business-sector productivity.

However, weaker public-sector diffusion channels and fewer public-private partnerships dampen AI-enabled benefits to the wider economy.

Restrictiveness level 4: Economy-wide restrictions

Restrictiveness level 4 extends moderate data residency and local-preference measures to both the government and private sectors. Restrictions are expanded, imposing restrictions on data, CSPs, AI tooling, software, and applications for specific private-sector workloads and sectors.

This level of restrictions is typically in place to strengthen domestic AI capabilities. However, the economic trade-off is more pronounced. A significant build-out of domestic capabilities is required to serve both sectors, and higher costs coupled with more limited scale push up the price of AI services. As a result, AI diffusion slows across the economy, particularly in smaller firms and sectors where margins are thinner and the ability to absorb higher costs is constrained. Innovation becomes more reliant on a relatively small, domestically focused ecosystem, which can limit exposure to global best practice and reduce competitive pressure.

Restrictiveness level 5: Ownership-centric (Domestically owned full AI stack)






Restrictiveness level 5 represents the most restrictive regime with high levels of restriction across all dimensions and sectors. Both businesses and government agencies rely primarily on domestically operated infrastructure, with strict controls on data transfer, local development of AI models, and predominantly local procurement for

cloud and AI services. International partnerships are limited, and global CSPs are largely excluded from public-sector procurement unless they operate in fully sovereign environments. Policies within this scenario aim to maximise strategic autonomy and national control but would be associated with significant capital and operational costs, lower scalability, and typically slower innovation cycles.

Policies at this level aim to maximise strategic autonomy and national control over the AI stack. Economically, this entails very high capital and operational costs to build and sustain a full domestic stack, including, where pursued, large-scale compute and national foundation models. Both public and private sectors face significant delays in AI deployment as domestic capacity is built, and innovation cycles tend to be slower due to reduced scale, limited external collaboration, and narrower technology choice. The trade-off is stark. As strategic autonomy and control are maximised, scalability, global integration, and the speed at which AI-driven productivity gains are realised are materially reduced.

Figure 14 summarises the sovereign-AI measures applied in each restriction level across policy dimensions and sectors. Levels 2–5 introduce increasing limitations on the participation of global CSPs along with AI tooling, software, and applications. Comparing these with Level 1 illustrates how different degrees of restriction affect the costs of delivering AI, the pace of adoption, and the productivity gains that follow.

Figure 14: AI sovereignty-related restriction levels

		Restriction Level 1	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Sector	Policy dimension	Control-and-choice (Assurance-led)	Government sector-led restrictions	Government sector-led restrictions, AI tooling, software & applications	Economy-wide restrictions	Domestically owned full AI stack
 Government sector	Local data hosting	Limited	Moderate	Moderate	Moderate	High
	Cloud provision	Limited	Moderate	Moderate	Moderate	High
	AI tooling, software & applications	Limited	Limited	Moderate	Moderate	High
 Private sector	Local data hosting	Limited	Limited	Limited	Moderate	High
	Cloud provision	Limited	Limited	Limited	Moderate	High
	AI tooling, software & applications	Limited	Limited	Limited	Moderate	High
Implications for AI diffusion and the broader economy						
 Likely economic implications		<ul style="list-style-type: none"> • Lowest cost • Rapid diffusion and strong productivity gains 	<ul style="list-style-type: none"> • Moderate cost increase and time to deployment • Limited drag on government adoption and productivity gains 	<ul style="list-style-type: none"> • Moderate cost increase, but longer deployment time • Larger drag on government adoption and productivity gains 	<ul style="list-style-type: none"> • Higher costs and time to deploy • Slower overall adoption and material impact on productivity gains 	<ul style="list-style-type: none"> • Highest cost and deployment time • Slowest overall adoption and significant impact on productivity gains
	 Trade-off		<ul style="list-style-type: none"> • Assurance-led protection allows flexibility, scalability, and global cooperation 			<ul style="list-style-type: none"> • Maximises autonomy but reduces scalability and global collaboration

Source: Oxford Economics

SECTION 4. QUANTIFYING THE ECONOMIC IMPACTS

AI sovereignty measures can affect economic performance through three linked pathways: higher costs of supplying AI domestically; slower diffusion across firms; and weaker or delayed productivity gains. Our analysis focuses on direct infrastructure and opportunity costs, and environmental impacts, based on cross-country evidence. It does not quantify broader risks that may arise under more restrictive sovereignty models. Cybersecurity is one example: shifting from hyperscale cloud to domestic-only provision may reduce the level of cyber defence investment and increase exposure to attacks.

We organise the assessment around three components, each developed in the sections that follow:

- **Direct cost estimation across the AI stack**, described in SECTION 5, covers compute infrastructure, foundation model development, and user-facing applications. This includes capital and operating expenditures, alongside complementary investments in software, skills, and talent.

- » **Opportunity cost estimation**, presented in SECTION 6, captures the macroeconomic effects of constrained adoption:
- » **Adoption impacts**, where sovereignty-related cost and compliance differences are translated into alternative AI diffusion paths.
- » **Macroeconomic impacts**, which convert differences in adoption trajectories into foregone productivity and GDP gains using sector-level evidence and Oxford Economics' macroeconomic modelling framework.
- **Environmental impacts**, described in SECTION 7, presents the implications of different restriction levels for energy and water use.

Throughout, we distinguish outcomes by **restrictiveness level**—reflecting how far AI sovereignty requirements extend beyond the public sector into the wider economy—as set out in the previous section. This provides a consistent basis for comparing policy configurations and for highlighting the trade-offs between sovereignty objectives and economic and environmental outcomes.



SECTION 5. DIRECT COSTS OF DOMESTIC CAPACITY BUILD-OUT

OVERVIEW

In this section, we present the estimated economic costs associated with building domestic AI capability under alternative AI sovereignty-related restriction levels. It focuses on how restrictions on data, infrastructure location, and openness affect the scale of required investment across the AI value chain.



The analysis proceeds in four steps:

- Estimating **AI adoption** and projecting how demand for AI evolves over time
- Translating adoption into domestic **infrastructure requirements**

- Estimating the **costs of infrastructure**
- Estimating the **costs for software, skills, and model development**

Together, these steps quantify how increasingly restrictive sovereignty policies raise both the scale and complexity of investment required to support AI use domestically.

STEP 1: AI ADOPTION AND PROJECTED DEMAND

AI adoption is defined, in line with OECD definitions, as firms using AI to improve the production of goods and services, as this is where productivity gains are concentrated.⁸⁸

As no single globally consistent dataset exists, we triangulate estimates across multiple datasets to calculate adoption in 2024. To project adoption over the next decade, we apply a diffusion model using an S-shaped logistic curve anchored to each country's current adoption level. The speed of diffusion is calibrated to historical adoption patterns observed for general-purpose technologies such as personal computers and the internet, reflecting a policy environment broadly supportive of AI. Projections are constrained by Oxford Economics' country-specific forecasts of internet penetration.⁸⁹

Under this baseline that corresponds to Restrictiveness Level 1, AI adoption is projected to rise sharply across all APJ economies between 2024 and 2035, though starting points and rates of convergence differ. Across the APJ economies shown, AI adoption among firms is currently very low in 2024 (small, single-digit shares in all economies). By 2035, we project a step-change: the adoption rate rises to 44% and 29% in Singapore and **South Korea respectively, having been** enabled by world-class broadband penetration and AI infrastructure. AI adoption among firms is projected to reach between 15% and 20% in the other APJ countries (Figure 15) through improvements in digital infrastructure, AI readiness, and improvements in skills.

STEP 2: DOMESTIC INFRASTRUCTURE REQUIREMENTS

LINKING ADOPTION TO INFRASTRUCTURE DEMAND

We translate projected adoption into infrastructure demand by allocating the global stock of AI-related compute infrastructure in 2024 across countries based on two factors: AI adoption rates; and economic size (proxied by real GDP). This provides country-specific estimates of current

infrastructure needs.⁹⁰ Future infrastructure demand is then projected in line with projected changes in adoption, as described above.

AI-related infrastructure demand in APJ countries is expected to increase from around 0.5 gigawatts (GW) in 2024 to 5 GW by 2035, a tenfold increase.⁹¹ Most of this growth will come from larger economies, driven by higher

88 OECD, "Macroeconomic productivity gains from Artificial Intelligence in G7 economies", 2025, accessed December 2025.


89 Forecasts published by Oxford Economics in the global forecasting database.

90 We cannot rely on country-specific AI data centre availability because most infrastructure demand is met through facilities located in other countries.

91 AI demand in this analysis refers exclusively to inference. We assume that inference accounts for 80% of total AI compute demand, in line with estimates reported in other studies. See appendix for more details.

Figure 15: Business AI adoption (OECD), 2024 and 2035

Firm AI adoption (OECD definition), % of firms




Source: Oxford Economics.

Note: Adoption rates refer to firms that have integrated AI in the production of goods and services. It excludes firms that are experimenting, piloting, or scaling AI in their operations.

Figure 16: AI demand is anticipated to increase between 2024 and 2035

Estimated AI demand, GW

Growth multiple, 2035 vs 2024



Source: Oxford Economics

adoption rates and deeper integration of AI into production. Japan (approx. 1.7 GW) and India (approx. 1.2 GW) lead the forecast, reflecting their economic scale and expanding AI capacity. South Korea, Singapore and Taiwan show high intensity relative to size, supported by strong adoption

(29%, 44% and 28% in 2035, respectively), though their incremental growth is smaller given already elevated adoption levels. Demand in smaller economies is projected to remain below 0.2 GW, constrained by limited infrastructure and investment.

Figure 17: AI sovereignty-related restriction levels for local data hosting

		Restriction Level 1	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Sector	Policy dimension	Control-and-choice (Assurance-led)	Government sector-led restrictions	Government sector-led restrictions, AI tooling, software & applications	Economy-wide restrictions	Domestically owned full AI stack
Government sector	Local data hosting	Limited	Moderate	Moderate	Moderate	High
Private sector	Local data hosting	Limited	Limited	Limited	Moderate	High

Source: Oxford Economics

SOVEREIGNTY-RELATED RESTRICTION LEVELS AND ONSHOREING REQUIREMENTS

AI sovereignty policies can impose varying degrees of restrictions in their requirements for domestic data and compute retention, as outlined earlier in this report, and shown in Figure 17.

To estimate supply needs under different restriction levels, we draw on expenditure data showing that the government accounts for roughly 10% of total AI spending in the Asia-Pacific region.⁹² Evidence from the UK suggests only a small share (10%) of government data is highly sensitive⁹³, while other studies indicate that sensitive private-sector data already represents more than half (54%) of cloud assets in 2025, up from 47% in 2024.⁹⁴ Drawing on the range of these estimates, we categorise the proportion of data retained within the economy into three levels:

- *Limited*: Sensitive public-sector workloads (assumed 10%) and no data restricted in private sector workloads (i.e. 0%)

- *Moderate*: Restrictions to 50% of the data in both the public and private sector workloads
- *High*: All workloads in the private and public sector are completely restricted (i.e., 100%)

While practices differ by country, this framework provides a consistent basis for comparison—particularly against the most restrictive case, where all AI workloads must be hosted domestically.

As restrictions tighten, domestic infrastructure requirements are projected to rise sharply (as shown in Figure 18). The largest increase is expected to occur when restrictions extend from public-sector workloads to the private sector, with the high-restriction level requiring a step-change in capacity. For example, in Figure 18, Japan’s domestic AI supply requirement is projected to increase to 1.7 GW, India to 1.2 GW, and South Korea to 1 GW under full onshoring (level 5). Comparatively, countries face minimal requirements under limited restrictions (level 1).

92 Katz, R., et al., “Economic impact of artificial intelligence and cloud computing developing in Asia and the Pacific”, Telecom Advisory Services LLC, 2025, accessed December 2025.

93 AWS, “Data Classification”, 2025, accessed December 2025.

94 Moore, T., Thales Group, “Thales 2025 Cloud Security Study Insights”, 2025, accessed December 2025.

Figure 18: Additional domestic AI supply required (compared to Level 1), 2035

	Additional domestic AI supply required, GW 2035 (relative to Restriction Level 1)		
	Restriction Levels 2 and 3	Restriction Level 4	Restriction Level 5
Cambodia	0.00	0.00	0.00
India	0.05	0.58	1.17
Indonesia	0.01	0.15	0.31
Japan	0.07	0.83	1.68
Lao PDR	0.00	0.00	0.00
Malaysia	0.01	0.06	0.14
Myanmar	0.00	0.01	0.01
Nepal	0.00	0.00	0.01
Philippines	0.01	0.06	0.13
Singapore	0.01	0.16	0.32
South Korea	0.04	0.49	0.99
Taiwan	0.01	0.16	0.35
Thailand	0.00	0.06	0.13
Vietnam	0.00	0.06	0.12

Source: Oxford Economics

STEP 3: INFRASTRUCTURE COST

The additional costs of developing AI compute infrastructure under different levels of restrictiveness are based on estimates of capital expenditure (CapEx) and operating expenditure (OpEx) from a range of sources.

CapEx covers upfront investment in construction and equipment. For hyperscale AI facilities, capital costs are estimated at around US\$44 billion per gigawatt of server power, including approximately US\$30 billion for IT hardware (servers and GPUs) and US\$14 billion for buildings and supporting infrastructure.⁹⁵ Additional costs include connectivity and grid expansion, with grid upgrades alone estimated at US\$0.75 billion per gigawatt.⁹⁶ These figures are averages and actual costs differ by country. Some economies can deliver capacity more cheaply due to stronger existing infrastructure, while others face higher costs because of weaker energy, construction, or

connectivity networks. These large, fixed costs remain a major barrier for countries pursuing AI sovereignty strategies.

Operating expenditure (OpEx) comprises ongoing costs such as electricity, cooling, water, staffing, maintenance, and hardware depreciation. These costs persist throughout the facility’s life and vary by location, energy prices, and labour costs. Hardware replacement is a major component, as GPUs typically require renewal every five to six years.⁹⁷

Infrastructure costs rise steeply as sovereignty requirements become more restrictive. With the highest levels of restrictions modelled, expenditure in Japan is expected to be highest at US\$122 billion, driven by large-scale infrastructure needs, followed by India at US\$79 billion and South Korea at US\$69 billion. In Figure 19, smaller economies

95 Epoch AI, “Frontier Data Centers Documentation”, 2025, accessed December 2025.

96 Global Energy Monitor, “Asia drives global gas expansion despite price volatility and plummeting cost of renewables”, 2023, accessed December 2025.

97 To keep up with the latest frontier GPUs, this cycle may be less.

such as Nepal, Thailand, and Vietnam are expected to incur much lower absolute costs, even under the most restrictive levels modelled, though the

relative burden can still be substantial given their economic size.

Figure 19: Additional compute infrastructure costs (compared to Level 1), 2025–2035

	Additional compute infrastructure costs, \$ bn (relative to Restriction Level 1)		
	Restriction Levels 2 and 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0	0.2	0.3
India	3.2	39.0	78.8
Indonesia	0.8	9.9	21.0
Japan	4.9	60.3	121.8
Lao PDR	0.0	0.1	0.2
Malaysia	0.4	4.3	9.1
Myanmar	0.0	0.4	0.9
Nepal	0.0	0.3	0.6
Philippines	0.4	4.5	9.1
Singapore	1.0	11.8	23.8
South Korea	2.8	34.4	69.4
Taiwan	0.9	10.9	23.4
Thailand	0.3	4.1	8.7
Vietnam	0.3	4.1	8.3

Source: Oxford Economics
 Note: All values in 2025 prices.

STEP 4: COMPLEMENTARY INVESTMENT—SOFTWARE, SKILLS, AND MODELS

AI TOOLING, SOFTWARE, AND APPLICATIONS

Infrastructure alone does not deliver AI capability. Compute infrastructure must be paired with software, development tools, and applications to enable real-world use. Evidence from the 2024 Stanford HAI Index shows that start-ups invested around **US\$80 billion** in AI software and applications—nearly twice the amount spent on AI infrastructure, research, and governance. Large technology firms show a similar pattern: in 2024, major firms spent **US\$155 billion on R&D**.^{98,99}

Combining start-up and corporate investment suggests that **non-infrastructure investment**

can reach up to 70% of compute infrastructure costs, underscoring the scale of complementary spending required under sovereignty regimes.

DEVELOPERS AND SKILLS

A skilled workforce is essential to translate investment into economic value, and most Asian economies lag the United States in AI-related developer intensity. With over half of developers already using, learning about, or showing interest in AI developments—an important foundation for future capability growth—these economies still face a significant gap.¹⁰⁰ Developers account for around **2% of the US workforce**, compared with

98 Stanford University HAI, “The 2025 AI Index Report”, 2025, accessed December 2025.

99 Trendline HQ, “Big techs big R and D bill”, 2025, accessed December 2025.

100 Muir, R., “Who Are the Developers Working on Generative AI Projects?”, Heavybit, 2023, accessed December 2025.

closer to **0.1%** in countries such as **Indonesia**. Only Japan and Singapore approach US levels.^{101,102}

Where AI sovereignty policies impose wide-ranging restrictions (restriction levels 3–5), we assume countries partially close this gap. This has been scaled to reflect the anticipated proportion of developers working on AI activities and the degree of data restriction.¹⁰³ Training costs are estimated using World Bank data on public expenditure per tertiary student. Indonesia faces the largest training costs due to the scale of additional demand. If Indonesia aimed to employ a similar proportion to the US, it would need to train about 1.1 million developers. At roughly US\$880 per student per year, this would translate to around US\$4.1 billion in extra spending.

RETENTION AND MIGRATION

Training alone is insufficient if skilled workers migrate abroad. Evidence from LinkedIn data shows that several Asian countries see net outflows of AI-skilled workers.¹⁰⁴ Specifically, the results show that India ranks highest, recording a net outflow of AI engineering talent of about 1.5 per 10,000 LinkedIn members between 2019 and 2024.¹⁰⁵ In addition, Indonesia and South Korea also exhibit net outflows, though to a lesser extent. Countries experiencing significant outmigration may therefore need to raise wages to retain talent.

To account for additional expenditure required to retain talent, we examined wage differentials

between countries with net outflows and those with net inflows of AI talent. Our analysis indicates that countries facing outflows offer, on average, wages around 20% lower (in PPP terms) than those with inflows.^{106,107} Therefore, we include a 20% increase in developer wages to capture the stronger requirements needed to retain AI talent.

AI MODEL TRAINING DEMAND

Where sovereign AI-related restrictions are widespread (under levels 4 and 5 only), countries are assumed to develop domestic foundation model capabilities. This requires:

- **Supercomputer capacity:** Countries without existing capability will need to build or acquire a supercomputer to train foundation models. Estimates suggest that constructing a suitable system can cost around US\$1 billion, depending on specifications.¹⁰⁸
- **Model training:** Training frontier-scale models can require substantial compute resources, with costs running into the tens of millions of US dollars. For example, the capital expenditure required to develop a foundation at the scale of ChatGPT is estimated at approximately US\$130 million.¹⁰⁹ Beyond financial costs, training incurs environmental impacts. For instance, training ChatGPT-4 was estimated to emit 5,184 tons of CO₂¹¹⁰, and we estimate it consumed approximately 1,400,000 litres of water.¹¹¹

101 JetBrain, “[Developer Ecosystem Data Playground: discover the latest software development trends and insights](#)”, 2025, accessed December 2025.

102 Wachs, J., et al., “[The geography of open source software: evidence from Github](#)”, Technological Forecasting and Social Change, Volume 176, (2022), accessed December 2025.

103 The estimated number of developers involved in AI is calculated by combining those currently working on AI, those learning about it, and half of those who have expressed interest.

104 [OECD.AI](#) (2025), data from [LinkedIn Economic Graph](#), last updated 2025-04-07, accessed on 2025-12-15

105 As detailed in the appendix, our estimates suggest that Cambodia experiences a higher net outflow than India. However, this result is derived from a model based on total migration and income variables rather than LinkedIn data.

106 Salary Expert “[Software engineer](#)”, n.d., accessed December 2025.

107 Software engineer salary converted to PPP terms using [World Bank PPP conversion factor](#).

108 Japan’s Fugaku supercomputer cost approximately [US\\$1 billion to build](#), while the UK government has committed £750 million to develop a new national AI supercomputer as part of its [AI Research Resource initiative](#).

109 Cottier, B., et al., “[The Rising Costs of Training Frontier AI Models](#)”, ArXIV, 2024, accessed December 2025.

110 Stanford University HAI, “[The 2025 AI Index Report](#)”, 2025, accessed December 2025.

111 Li, P., et al., “[Making AI Less “Thirsty”: Uncovering and Addressing the Secret Water Footprint of AI Models](#)”, ArXIV, 2023, accessed December 2025.

Models also need continuous updates, with new versions released regularly. We assume four major updates over the next 10 years, reflecting recent trends—such as the release of three ChatGPT versions between 2022 and 2025. Based on these assumptions, we increase the projected costs in these scenarios to account for higher compute requirements, as well as the associated carbon and water impacts.

TOTAL COMPLEMENTARY COST ESTIMATES

Taken together, AI sovereignty strategies require substantial investment across infrastructure, software, skills, and models. These costs rise non-linearly with the degree of restriction:

- Level 3 restrictions would involve moderate additional investment focused on public-sector needs;

- Level 4 restrictions would involve significant investment and costs, as private-sector capabilities must also be developed domestically
- Level 5 restrictions add further complexity and cost, including frontier model development and continuous upgrades

Complementary investments—software, applications, and developer ecosystems—represent roughly **70% of infrastructure costs**, highlighting that sovereignty strategies extend far beyond data centres alone. Under Level 5 restrictions, relative to Level 1, Japan is estimated to incur combined complementary costs of around **US\$28 billion**, followed by India at around **US\$24 billion**, while smaller economies face lower absolute but still material burdens relative to their size.

Figure 20: Complementary costs assumptions relative to Restriction Level 1

	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Policy dimension	Government sector-led restrictions	Government sector-led restrictions, AI tooling, software & applications	Economy-wide restrictions	Domestically owned full AI stack
Tooling, software, and applications	None	Moderate Public sector needs only	Moderate Economy-wide	High Economy-wide
Developers and skills	None	Moderate Public sector needs only	Moderate Economy-wide	High Economy-wide
Increased wages related to retention	None	Moderate Public sector needs only	Moderate Economy-wide	High Economy-wide
AI model training demand	None	None	Frontier model development only	Frontier model and supercomputer

Source: Oxford Economics

Figure 21: Combined additional complementary costs (compared to Level 1), 2025–2035

	Additional complementary costs (relative to Restriction Level 1), \$ bn		
	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0	0.9	1.1
India	2.0	22.1	23.6
Indonesia	1.0	10.6	12.4
Japan	2.2	25.9	27.9
Lao PDR	0.0	0.9	1.1
Malaysia	0.2	3.2	3.9
Myanmar	0.0	1.2	1.4
Nepal	0.0	1.1	1.3
Philippines	0.3	3.6	4.1
Singapore	0.4	5.6	5.3
South Korea	1.6	17.9	19.0
Taiwan	0.5	6.1	6.5
Thailand	0.2	3.2	2.8
Vietnam	0.3	3.9	3.5

Source: Oxford Economics
 Note: All values in 2025 prices.

TOTAL ADDITIONAL DIRECT COST ESTIMATES

The estimated **total additional direct costs in 2035**, relative to Restriction Level 1, associated with progressively more restrictive AI sovereignty frameworks across the APJ economies considered, are shown in Figure 22. These costs capture the additional expenditure required to build and operate domestic AI capability across infrastructure, software, skills and—at higher restrictiveness levels—model development.

Larger and more AI-intensive economies face the highest absolute costs. Japan and India incur the largest increases in direct costs under the most restrictive settings, reaching around US\$150 billion and US\$102 billion, respectively, under Restriction Level 5, equivalent to 0.3% and 0.2% of their 2025–2035 GDP. South Korea also faces substantial costs (around US\$88 billion, or 0.4% of 2025–2035 GDP), reflecting its high baseline adoption and the scale of domestic capacity required under full onshoring. These outcomes are driven by economic size, expected AI uptake, and the volume of workloads that must be supported domestically.

Smaller and more open economies show lower absolute costs but are nevertheless substantial. Singapore’s costs rise from US\$1 billion (0.01% of 2025–2035 GDP) at Restriction Level 2 to over US\$29 billion (0.4% of 2025–2035 GDP) at Level 5, highlighting the sensitivity of highly digitalised economies to restrictions that limit access to global infrastructure and providers. Similarly, mid-sized economies such as Indonesia, the Philippines, Thailand, and Vietnam see costs increase several-fold as restrictions tighten, even though absolute totals remain below those of the largest economies. Smaller economies such as Malaysia and Nepal incur lower absolute costs, but the relative burden can still be material, ranging from 0.2% to 0.3% of their 2025–2035 GDP.

While Restriction Levels 2 and 3 show the lowest additional direct costs, these estimates do not capture additional risks that could amplify real-world impacts. In particular, if public-sector AI and cloud workloads shift to less mature domestic environments with lower cyber-defence investment, exposure to service disruptions or successful attacks on critical systems may increase.

Figure 22: Total direct costs (compared to Level 1), 2025–2035

	Total additional direct costs, \$ bn (relative to Restriction Level 1)			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0	0.0	1.1	1.5
India	3.2	5.2	61.1	102.5
Indonesia	0.8	1.8	20.6	33.4
Japan	4.9	7.2	86.2	149.7
Lao PDR	0.0	0.0	1.0	1.3
Malaysia	0.4	0.6	7.5	13.0
Myanmar	0.0	0.1	1.6	2.3
Nepal	0.0	0.1	1.4	2.0
Philippines	0.4	0.6	8.1	13.2
Singapore	1.0	1.4	17.4	29.1
South Korea	2.8	4.4	52.3	88.4
Taiwan	0.9	1.4	17.1	30.0
Thailand	0.3	0.6	7.3	11.5
Vietnam	0.3	0.6	8.0	11.8

	Total additional direct costs, % of 2025-2035 GDP			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.00%	0.00%	0.2%	0.2%
India	0.01%	0.01%	0.1%	0.2%
Indonesia	0.00%	0.01%	0.1%	0.2%
Japan	0.01%	0.02%	0.2%	0.3%
Lao PDR	0.00%	0.01%	0.4%	0.5%
Malaysia	0.01%	0.01%	0.1%	0.2%
Myanmar	0.00%	0.01%	0.2%	0.2%
Nepal	0.00%	0.01%	0.2%	0.3%
Philippines	0.01%	0.01%	0.1%	0.2%
Singapore	0.01%	0.02%	0.2%	0.4%
South Korea	0.01%	0.02%	0.2%	0.4%
Taiwan	0.01%	0.01%	0.2%	0.3%
Thailand	0.00%	0.01%	0.1%	0.2%
Vietnam	0.00%	0.01%	0.1%	0.2%

Source: Oxford Economics
 Note: All values in 2025 prices.

SECTION 6. OPPORTUNITY COST FROM AI SOVEREIGNTY RESTRICTIONS

OVERVIEW

Restrictions introduced to achieve AI sovereignty can generate significant opportunity costs by slowing the adoption of AI and reducing the productivity gains that would otherwise accrue. These opportunity costs arise through two main channels:

- **delays in adoption** while domestic infrastructure, tools, and skills are developed; and
- **persistently lower adoption levels** once infrastructure is in place, reflecting higher costs, limited access to global technologies, and weaker innovation incentives.

This section quantifies these effects and translates them into impacts on total factor productivity (TFP) and GDP. The analysis distinguishes between **different levels of restrictiveness**, as set out in SECTION 4, reflecting the extent to which sovereignty measures apply to the public sector alone or extend economy-wide to private-sector activity.



The analysis proceeds in three steps.

1. We assess the time required to build domestic AI capabilities and the resulting delays to adoption.
2. We estimate how regulatory and market

frictions affect adoption once infrastructure is established.

3. We translate changes in adoption trajectories into foregone productivity and GDP gains using Oxford Economics' macroeconomic modelling framework.

STEP 1: DELAYS FROM BUILDING DOMESTIC AI CAPABILITIES

SOURCES OF ADOPTION DELAYS

Delays in AI adoption arise from the time needed to build the physical, technological, and human capital required to support domestic AI deployment. All restrictiveness levels assume the need for AI-ready infrastructure, while higher levels (3-5) also require domestic development of AI tools, applications, and skilled labour.

AI-ready data centres

Constructing an AI-capable data centre typically takes two to three years. However, connecting new facilities to the power grid is often a binding constraint. In the United States, grid interconnection delays have lengthened sharply since 2022, with average wait times now approaching six years—well beyond construction timelines.¹¹² As countries scale AI infrastructure simultaneously, similar bottlenecks are likely to emerge elsewhere. To reflect this, lower and mid-level restrictiveness assumptions (levels 2-4) incorporate a three-year delay for data centre construction and grid connection, while the highest restrictiveness level includes an additional year to reflect the complexity of scaling multiple facilities in parallel.

AI tooling, applications, and skills

Beyond infrastructure, AI adoption depends on the availability of applications, enterprise software, and skilled workers. Development timelines vary widely: simple applications may take several months, while large-scale systems—such as those

used in finance or healthcare—often require one to two years. Building foundation models, which underpin advanced AI capabilities, can take many months of compute-intensive training and engineering effort.^{113,114} Training skilled workers through tertiary education typically takes three to four years. To capture these constraints, higher restrictiveness levels (Levels 3-5) include an additional one-year delay to reflect application development and workforce readiness.

Adoption delays by restrictiveness level

Under these assumptions, restrictiveness levels limited to the public sector imply shorter delays, while broader economy-wide restrictions imply longer ones. At the highest level of restrictiveness, adoption is assumed to be delayed by up to five years, reflecting the scale and complexity of building domestic AI capability across infrastructure, software, and skills. These delays may lengthen if difficulties arise in building out infrastructure, securing specialist talent, or operationalising new AI systems. In summary,

- Level 2 reflects a three-year delay for public-sector adoption;
- Level 3 adds time for AI tooling and application development;
- Level 4 applies similar delays but across both public and private sectors; and
- Level 5 assumes a five-year delay, reflecting the significant challenge of ramping up domestic capabilities at scale.

¹¹² Brookfield, "Building the backbone of AI", August 2025, accessed December 2025.

¹¹³ Seifeur, "Unveiling the Training Timeline of GPT-4: A Detailed Exploration", n.d., accessed December 2025.

¹¹⁴ Meta, "The llama 3 herd of models", 2024, accessed December 2025.

STEP 2: IMPACT OF SOVEREIGNTY RESTRICTIONS ON ADOPTION LEVELS

CHANNELS AFFECTING ADOPTION ONCE INFRASTRUCTURE IS IN PLACE

Even after domestic capabilities are established, sovereignty-related policies can continue to suppress AI adoption. Evidence from the literature identifies three main channels through which restrictions—particularly data localisation and ownership limits—affect diffusion:

- **Higher costs.** Research indicates sovereign cloud regions typically cost 15%–30% more than public cloud, while data localisation measures can raise data management costs by 15%–55% as added compliance and duplication requirements further raise operating costs. Further, the purchase of GPUs from the US also entails complex and specific reporting requirements, borne by the users of these chips.¹¹⁵ Higher prices reduce firms’ willingness and ability to adopt AI, particularly among smaller businesses.^{116,117,118,119,120}
- **Constrained choice.** Open public-sector procurement often accelerates diffusion by showcasing successful use cases. When procurement is restricted to domestic providers, competition is reduced, and access to advanced tools narrows, slowing the spread of AI across sectors.^{121,122,123,124}

- **Quality and innovation.** Limits on cross-border data flows reduce access to diverse, high-quality datasets, while restrictions on procuring cutting-edge GPUs slow their integration into AI infrastructure. These constraints weaken model performance and hinder innovation, delaying the availability of effective AI tools and slowing uptake. In addition, access to global networks is important for scaling AI workloads efficiently—a capability that could become harder to achieve if such access is restricted.^{125,126}

Taken together, these frictions help explain why stricter sovereignty policies are associated with slower and more limited AI adoption.

MEASURING RESTRICTIVENESS

To quantify the effects of regulatory barriers on adoption, we draw on three established indices capturing different dimensions of restrictiveness:

- OECD Digital STRI, capturing regulatory barriers to digital trade, such as cross-border data restrictions and licensing requirements;¹²⁷
- OECD INDIGO (Digital Trade Integration and Openness), measuring broader integration factors like interoperability standards and e-commerce facilitation.¹²⁸

115 Greenberg Traurig LLC, “Navigating GPU Export Controls and AI Use Restrictions in Data Center Operations”, Lexology, 2025, accessed December 2025.

116 Scognamiglio, F., “Cloud Cover: Price Swings, Sovereignty Demands, and Wasted Resources”, BCG, 2025, accessed December 2025.

117 Wölbart, C., “Delos: ‘Sovereign’ Cloud 10 to 20 % More Expensive than Microsoft’s Public Cloud”, Heise Online, 2024, accessed December 2025.

118 OECD, “Cross-Border Data Flows”, 2024, accessed December 2025.

119 AWS, “Unlocking AI potential country reports”, 2025, accessed December 2025.

120 Capgemini, “The on-demand tech paradox”, 2025, accessed December 2025.

121 AWS, “Unlocking AI potential country reports”, 2025, accessed December 2025.

122 Wedekind, C., Böhning, C., “Sovereign cloud usage in Europe. Is there a world without U.S hyperscalers?”, Amaranth Advisory, n.d., accessed December 2025.

123 Yan, M., Liu, H., “The Impact of Digital Trade Barriers on Technological Innovation Efficiency and Sustainable Development”, *Sustainability* 16, 5169, (2024), accessed December 2025.

124 Skare, M., Soriano, D., “How globalization is changing digital technology adoption: An international perspective”, *Journal of Innovation & Knowledge* 6, no. 4 (2021): pp. 222-233, accessed December 2025.

125 Linux Foundation, “The State of Sovereign AI”, August 2025, accessed December 2025.

126 IDC, “Four Considerations for AI-Ready Infrastructure Buildout”, 2024, accessed December 2025.

127 OECD, “Digital Services Trade Restrictiveness Index”, n.d., accessed December 2025.

128 OECD, “Index of Digital Trade Integration and Openness”, n.d., accessed December 2025.

- OECD FDI Regulatory Restrictiveness Index, capturing limits on foreign ownership and participation in domestic markets.¹²⁹

We construct two composite measures by pairing each digital index with the FDI index. The two digital indices—OECD Digital STRI and INDIGO—capture related but distinct aspects of digital openness: STRI focuses on regulatory barriers such as licensing and cross-border data restrictions, while INDIGO emphasises integration factors like interoperability and e-commerce facilitation. Because they differ in indicators and country coverage, we evaluate them separately and then average the results to improve robustness.

ESTIMATED ADOPTION EFFECTS

As shown in Figure 23, the combined Digital STRI and FDI Restrictiveness Index and INDIGO and FDI Restrictiveness Index exhibit a clear negative relationship with AI adoption rates: countries with

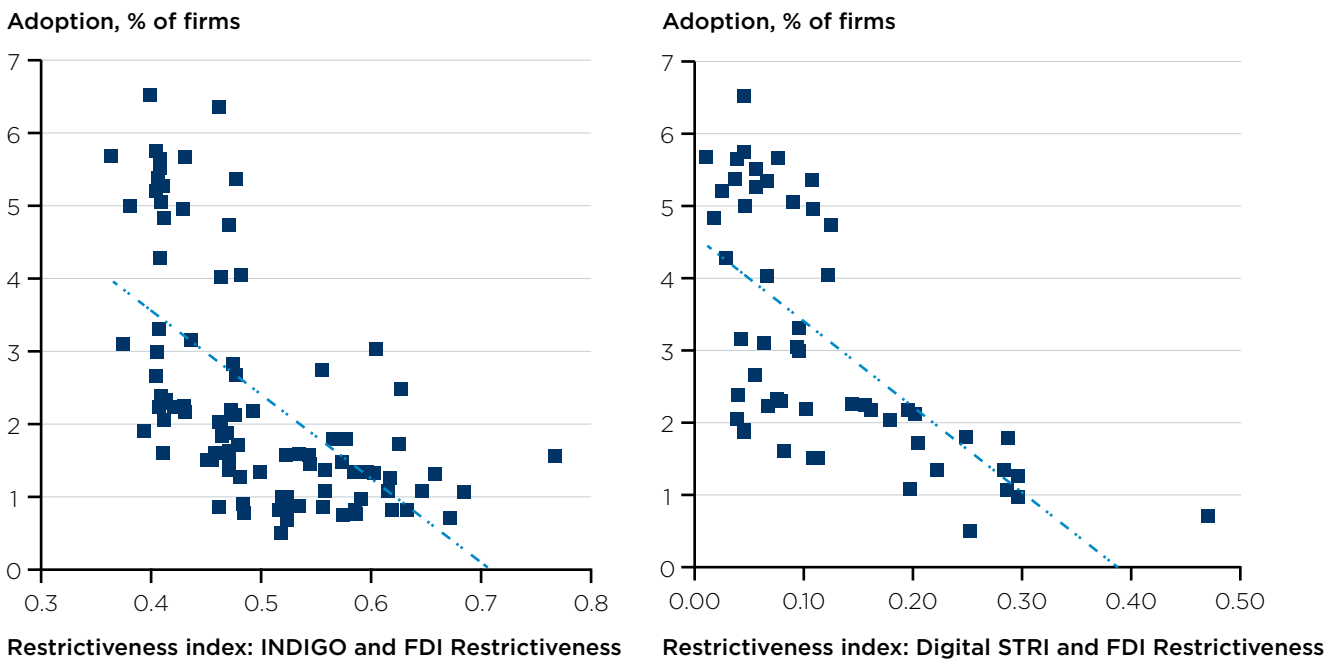
higher levels of digital and ownership restrictions tend to have lower AI adoption rates, underscoring the importance of openness and regulatory flexibility.

The link between restrictiveness, as measured by the two composite indices, and adoption levels is estimated using partial-correlation models that incorporate AI readiness, proxied by the Oxford Insights Government AI Readiness Index. Including readiness ensures policy effects are not overstated in economies with lower baseline capabilities.

Model results indicate that with the highest level of restriction (Level 5), adoption is projected to be 55% lower, whereas under moderate scenarios, it is projected to be 20% lower. Under Levels 2 and 3, restrictions apply mainly to the public sector, whereas Levels 4 and 5 extend across the entire economy.

Evidence suggests delays in public-sector adoption also slow private-sector uptake. Survey

Figure 23: Restrictiveness and adoption



Source: Oxford Economics

129 OECD, “FDI Regulatory Restrictiveness Index”, n.d., accessed December 2025.



data indicate 70% of firms are more likely to adopt AI when the public sector leads, highlighting the importance of government leadership.¹³⁰ While direct evidence on AI adoption is limited, research shows private-sector cloud adoption is influenced by public-sector restrictions. Given the complementarity between cloud and AI, we assume a 2% reduction in private-sector adoption under Restriction Levels 2 and 3, where restrictions apply only to the public sector, while noting that survey evidence suggests the impact could be greater.

COMBINED ADOPTION TRAJECTORIES

The final adoption paths combine two effects:

- an initial period of stagnation reflecting infrastructure and capability-building delays; and
- a downward adjustment to the long-run adoption curve reflecting persistent regulatory frictions.

Delays range from three years at lower restrictiveness levels to five years at the highest level. After this period, adoption resumes but at a lower trajectory, scaled by estimated reductions linked to digital and ownership barriers. This combined approach ensures that projections account for two distinct but reinforcing constraints: the lag in readiness due to construction and capability development, and the persistent drag on uptake from higher costs, limited choice, and reduced innovation.

¹³⁰ AWS, “Unlocking AI Potential Country reports”, 2025, accessed December 2025.

Figure 24: Adoption rates (OECD definition) by levels of restriction, 2035

	Adoption rates (OECD definition), % of firms				
	Restriction Level 1	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	6.3%	6.0%	6.0%	2.1%	0.9%
India	18.8%	18.0%	17.9%	6.9%	3.2%
Indonesia	14.5%	13.9%	13.8%	5.0%	2.3%
Japan	20.3%	19.4%	19.3%	7.4%	3.4%
Lao PDR	8.7%	8.3%	8.3%	2.9%	1.3%
Malaysia	18.9%	18.0%	17.9%	6.8%	3.1%
Myanmar	10.0%	9.6%	9.5%	3.5%	1.6%
Nepal	15.0%	14.4%	14.3%	5.5%	2.5%
Philippines	15.5%	14.8%	14.7%	5.4%	2.4%
Singapore	43.8%	42.1%	41.8%	20.3%	9.6%
South Korea	28.6%	27.4%	27.1%	11.2%	5.2%
Taiwan	28.0%	26.8%	26.6%	10.9%	5.0%
Thailand	15.4%	14.8%	14.6%	5.4%	2.4%
Vietnam	16.4%	15.7%	15.5%	5.8%	2.6%

Source: Oxford Economics. Note: Adoption rates refer to firms that have integrated AI in the production of goods and services. It excludes firms that are experimenting with, piloting, or scaling AI in their operations.

STEP 3: TRANSLATING ADOPTION INTO OPPORTUNITY COSTS

Changes in AI adoption are translated into opportunity costs using a phased productivity and macroeconomic framework consistent with OECD (2025).¹³¹ The approach links micro-level evidence on AI’s impact to economy-wide outcomes through four stages:

- **Occupational exposure to AI**, based on task-level analysis of around 18,000 tasks
- **Micro-level productivity gains**, drawing on experimental and firm-level studies
- **Adoption and learning dynamics**, scaling gains by restriction level-specific adoption paths
- **Macroeconomic impacts**, translating sectoral productivity changes into TFP and GDP outcomes

This framework ensures a consistent and evidence-based translation from task-level effects to

national economic performance. Further technical detail is provided in the appendix.

OPPORTUNITY COST ESTIMATES

The opportunity cost is defined as the difference between GDP gains under an **unrestricted adoption path** and those achieved under higher restrictiveness levels. As shown in Figure 25, these costs rise sharply as restrictions become more stringent and extend into the private sector.

Under the highest restrictiveness levels, cumulative opportunity costs are substantial. Japan faces the largest potential loss—exceeding **US\$58 billion (1.4% of 2035 GDP)**—reflecting both its economic size and high baseline adoption potential. India also incurs significant costs, approaching **US\$55 billion (0.8% of 2035 GDP)**, driven by its large economy and rapid adoption under an unrestricted path.

¹³¹ OECD, “Macroeconomic productivity gains from Artificial Intelligence in G7 economies”, 2025, accessed December 2025.

Lower restrictiveness levels generate smaller but still material losses. However, while the modelling includes a spillover impact of lower public-sector adoption on the private sector, the risk that this effect is larger in practice remains significant. Constraints on public-sector AI and cloud services can weaken demonstration effects,


reduce confidence, and dampen private sector investment, with the risks particularly acute for SMEs. If the true impact is larger than modelled, SME uptake may fall more sharply than captured in our results, delaying diffusion, widening productivity gaps, and increasing the overall economic impact of more restrictive approaches.

Figure 25: Opportunity costs by level of restriction, 2035

	Opportunity cost, \$ bn			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0	0.0	0.2	0.2
India	2.9	3.3	41.5	54.7
Indonesia	0.8	1.0	12.2	15.7
Japan	3.1	3.5	44.4	58.2
Lao PDR	0.0	0.0	0.1	0.1
Malaysia	0.4	0.5	6.4	8.4
Myanmar	0.0	0.0	0.3	0.4
Nepal	0.0	0.0	0.3	0.4
Philippines	0.3	0.4	5.0	6.5
Singapore	1.2	1.4	16.2	23.5
South Korea	2.1	2.5	30.7	41.4
Taiwan	1.0	1.2	14.7	19.8
Thailand	0.3	0.3	4.3	5.5
Vietnam	0.3	0.4	4.7	6.1

	Opportunity cost, % of 2035 GDP			
	Restriction Level 2	Restriction Level 3	Restriction Level 4	Restriction Level 5
Cambodia	0.0%	0.0%	0.2%	0.3%
India	0.0%	0.0%	0.6%	0.8%
Indonesia	0.0%	0.0%	0.5%	0.7%
Japan	0.1%	0.1%	1.0%	1.4%
Lao PDR	0.0%	0.0%	0.3%	0.3%
Malaysia	0.1%	0.1%	0.9%	1.2%
Myanmar	0.0%	0.0%	0.3%	0.4%
Nepal	0.0%	0.0%	0.5%	0.7%
Philippines	0.0%	0.0%	0.6%	0.8%
Singapore	0.2%	0.2%	2.2%	3.2%
South Korea	0.1%	0.1%	1.4%	1.8%
Taiwan	0.1%	0.1%	1.4%	1.8%
Thailand	0.0%	0.0%	0.6%	0.8%
Vietnam	0.0%	0.0%	0.5%	0.7%

Source: Oxford Economics



SECTION 7. ENVIRONMENTAL COSTS

Scaling AI capability will require a rapid expansion of data-centre capacity and associated electricity use under any policy model. What differs across approaches is where that capacity is built, how efficiently it is used, and whether workloads can access the most efficient available infrastructure.

This section therefore assesses the environmental implications of stricter localisation requirements which require domestic build-out. The marginal environmental cost of such stricter localisation depends on the difference between domestic build-out and the most efficient available alternative. Larger, newer hyperscale facilities typically achieve higher efficiency through purpose-built infrastructure, advanced cooling architectures, and better access to clean-energy procurement. By contrast, smaller or more fragmented domestic facilities often operate at higher effective energy and water intensities, reflecting lower utilisation rates, duplicated systems, and limited access to cutting-edge technologies.



Broadly, higher domestic build-out of data centres under restrictive policies creates environmental costs through two principal channels:

- **Carbon emissions**, driven by higher power demand and the carbon intensity of electricity generation; and
- **Water consumption**, reflecting cooling requirements in data centres and the water used in electricity generation.

These differences are particularly consequential in Asia. Many economies in the region retain

relatively carbon-intensive power systems and face elevated levels of water stress. As a result, additional AI-related energy demand can generate larger emissions and water impacts than in regions with cleaner generation mixes or lower water constraints.

In this section, we estimate the environmental implications of alternative restrictiveness levels for AI sovereignty. Higher restrictiveness levels require a greater share of AI workloads to be hosted domestically, increasing the volume of local infrastructure and electricity demand—and, in turn, the associated environmental footprint.

CARBON EMISSIONS

WHY ENERGY EFFICIENCY IS IMPORTANT FOR AI GROWTH

The energy demand associated with AI is increasingly well documented.^{132,133} AI model training and inference rely on data centres that house servers, storage, and networking equipment—systems that require continuous, electricity-intensive operation. As AI adoption expands, demand for data centre capacity rises, increasing electricity consumption.¹³⁴

Much of the sector's efficiency gains have come from the shift to hyperscale cloud architectures, whose advanced technologies have helped contain electricity growth despite surging compute demand. However, excluding these providers under strict AI sovereignty could mean forfeiting these efficiency benefits, resulting in higher local power demand and associated emissions.

It is important to note that the resulting carbon footprint is not driven by compute demand alone, but varies significantly with data-centre

design choices, utilisation rates, cooling architecture, climate of hosting country and the carbon intensity of the electricity supply. In this context, much of the sector's historical efficiency improvement has come from the shift toward hyperscale cloud architectures.¹³⁵ Purpose-built hyperscale facilities typically operate at materially lower Power Usage Effectiveness (PUE) levels—often around 1.1—compared with 1.6–2.0+ for traditional on-premise or 1.3–1.5 for smaller colocation data centres.¹³⁶ In addition, case-study evidence suggests that large hyperscale cloud environments can be up to around four times more energy-efficient than typical on-premise infrastructure, reflecting combined PUE and utilisation effects.¹³⁷ These differences reflect higher utilisation, advanced cooling and power systems, and integrated facility design, and have helped moderate electricity growth despite rapidly rising compute demand.

Excluding hyperscale providers under stricter AI sovereignty approaches, therefore, risks forfeiting these efficiency advantages, meaning that

132 The Financial Times, "[The state of AI: here comes the energy crunch](#)", November 2025, accessed December 2025.

133 The Financial Times, "[The power crunch threatening America's AI ambitions](#)", December 2025, accessed December 2025.

134 International Energy Agency, "[AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works](#)", April 2025, accessed December 2025.

135 International Energy Agency, "[Data centres and data transmission networks](#)", July 2023, accessed December 2025.

136 Sustainability Atlas, "[On-premise vs colocation vs hyperscale cloud: data center sustainability performance compared](#)", February 2026, accessed April 2026.

137 AWS, "[AWS Cloud sustainability](#)", 2024, accessed April 2026.

delivering the same AI workloads can require more electricity overall and result in higher associated carbon emissions at the local level.

While uncertainty remains around the pace of AI growth, future data centre efficiency, and national electricity mixes, there is broad consensus that AI will significantly contribute to rising power demand. For countries pursuing AI sovereignty strategies—where a larger share of workloads must run domestically—this amplifies the challenge, as limited access to hyperscale efficiencies could compound electricity and emissions impacts.

WHY ASIA FACES AN OUTSIZED EMISSIONS BURDEN

The carbon footprint of AI expansion is expected to be comparatively large in Asia due to:

- **High carbon intensity** in electricity generation across many economies; and
- **Large and growing demand** for AI tools and services given population scale and expected adoption.
- While many economies in Europe and the Americas have increased the share of renewables in power generation, several Asian economies remain relatively carbon intensive. As a result, they emit more CO₂-equivalent per unit of electricity generated.¹³⁸ For example, the emissions intensity of electricity generation in

Malaysia and the Philippines is around **2.8 times higher** than the EU average. This implies that the same increase in electricity demand can generate substantially higher emissions in these settings.

EMISSIONS OUTCOMES BY RESTRICTIVENESS LEVEL

To quantify these impacts, we combine projected domestic AI infrastructure requirements with country-specific electricity emissions factors. Given the uncertainty around future power-sector decarbonisation pathways, we adopt a conservative simplifying assumption that carbon intensity remains constant over the next decade.¹³⁹ This provides a transparent benchmark for comparing environmental implications across restrictiveness levels.

Under all restrictiveness levels, AI-driven emissions rise sharply, reflecting the underlying increase in compute demand. The increase is largest where policy settings require a greater share of AI workloads to be hosted domestically. Across the economies assessed, **India and Japan** emerge as the largest emitters at every restrictiveness level (Figure 27):

- In **India**, emissions growth is driven by scale—servicing AI demand across a very large economy and population.
- In **Japan**, emissions rise because expanding domestic AI capability builds on an already high baseline of AI uptake and investment, requiring substantial additional compute capacity.

WATER CONSUMPTION

WHY AI EXPANSION INCREASES WATER DEMAND

The environmental footprint of AI is not limited to carbon emissions. AI systems can also be water-intensive, both:

- **Directly**, through water usage in data-centre cooling systems to prevent overheating; and
- **Indirectly**, through water use in electricity generation, particularly in coal, gas, and nuclear systems where water is used for steam generation and cooling.^{140,141}

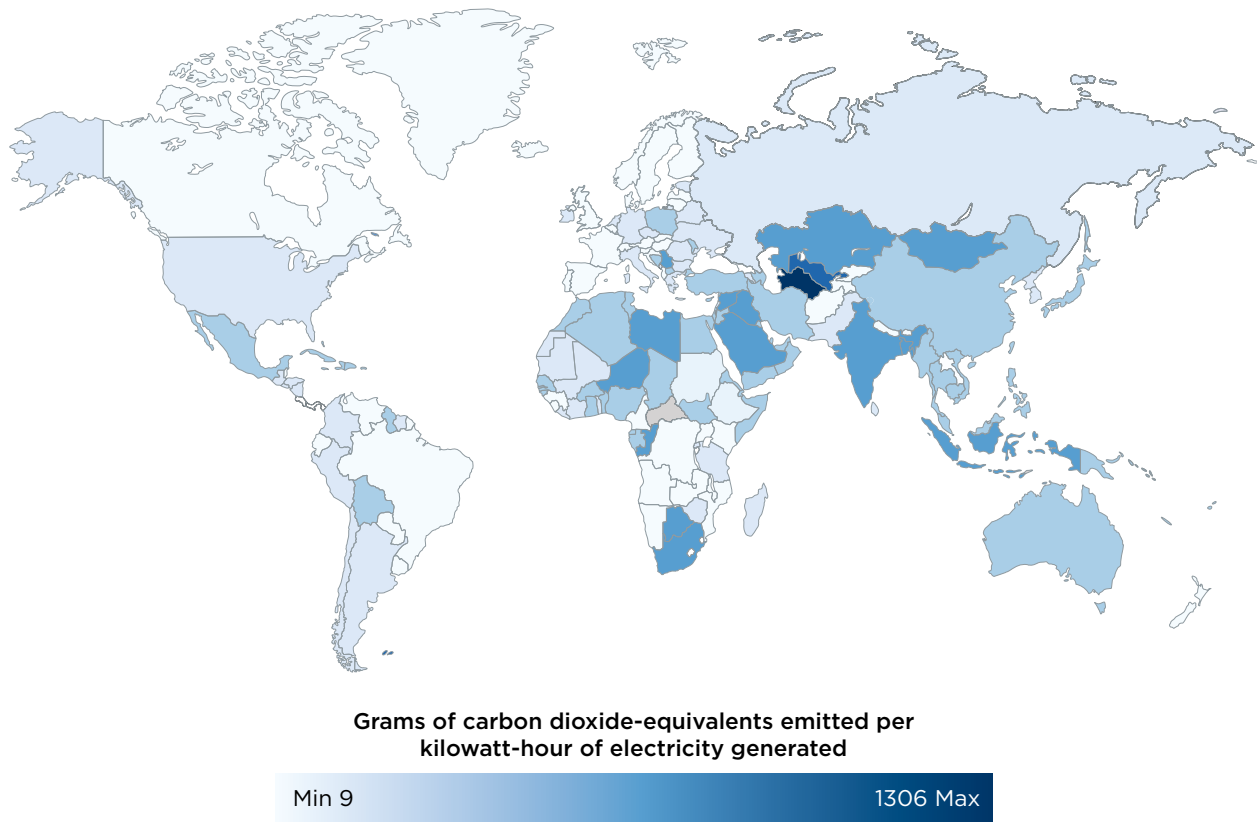
138 Ember Energy Institute – Statistical Review of World Energy and Our World in Data 2024, “Carbon intensity of electricity generation”, June 2025, accessed December 2025.

139 S&P Global, “Energy Asia: Asia’s fossil fuels addiction poses challenges for energy transition roadmap”, June 2025, accessed December 2025.

140 Environmental and Energy Study Institute (EESI). “Data centres and water consumption”. June 2025. Accessed December 2025.

141 Department for Environment, Food and Rural Affairs, UK Government. “AI’s thirst for water”. September 2025. Accessed December 2025.

Figure 26: Carbon emissions of electricity



Source: Ember, Energy Institute by our World in Data

Figure 27: Carbon emissions by level of restriction, 2035

	Additional CO ₂ -eq emissions, metric megatonnes (relative to Restriction Level 1)		
	Restriction levels 2 and 3	Restriction level 4	Restriction level 5
Cambodia	0.0	0.0	0.1
India	0.4	5.1	10.2
Indonesia	0.1	1.3	2.6
Japan	0.4	4.9	9.9
Lao PDR	0.0	0.0	0.0
Malaysia	0.0	0.5	1.0
Myanmar	0.0	0.1	0.1
Nepal	0.0	0.0	0.0
Philippines	0.0	0.5	1.0
Singapore	0.1	1.0	2.0
South Korea	0.2	2.5	5.0
Taiwan	0.1	1.3	2.7
Thailand	0.0	0.4	0.9
Vietnam	0.0	0.4	0.7

Source: Oxford Economics

Water outcomes, however, depend strongly on data-centre technology and cooling design. Traditional facilities rely heavily on evaporative cooling, which can consume large volumes of freshwater. By contrast, newer hyperscale AI facilities increasingly deploy advanced cooling architectures such as closed-loop liquid cooling, direct-to-chip cooling, and immersion cooling. These technologies can materially reduce, and in some cases largely eliminate, operational freshwater use.

Smaller domestic facilities may lack access to these technologies or the expertise needed to deploy them effectively, increasing water use per unit of AI output. As a result, water consumption and the technologies used to manage it become critical considerations for economies expanding AI capacity, especially where water resources are already under stress.

EXPOSURE: WATER STRESS ACROSS ASIA

Many economies across Central, South, and East Asia already face high water stress—defined as annual freshwater withdrawals relative to renewable freshwater resources. In some cases, withdrawal rates exceed thresholds associated with severe or extreme stress. For example, **South Korea and Sri Lanka** face **extreme** water stress, withdrawing more than **80%** of renewable freshwater resources annually whereas **India** faces **severe** water stress, withdrawing around **66%** annually (Figure 28).¹⁴²

Against this backdrop, additional water demand from AI-driven infrastructure can materially increase environmental pressure.

METHOD FOR ESTIMATING WATER CONSUMPTION

We estimate water requirements by linking projected domestic AI infrastructure needs to benchmark water consumption of large data centres. Evidence suggests a large data centre can consume up to **5 million gallons of water per day**, or roughly **1.8 billion gallons annually**, with the majority—around **67%**—attributable to indirect water consumption through electricity generation.¹⁴³

In the modelling, we:

1. estimate the gap between projected domestic demand and supply of AI and data services (expressed in energy terms);
2. translate this into the implied number of large AI data centres required, based on benchmark capacity from major facilities¹⁴⁴; and
3. apply benchmark water use per facility to estimate total annual water consumption.

This approach provides an internally consistent estimate of how water demand scales with domestic hosting requirements under different restrictiveness levels.

WATER OUTCOMES BY RESTRICTIVENESS LEVEL

Under all restrictiveness levels, water demand rises as domestic AI infrastructure expands. The largest water requirements are concentrated in **South Korea, India, and Japan**, reflecting a combination of high projected domestic compute demand and the scale of infrastructure implied by domestic hosting requirements. Given the current level of water stress in countries, this underscores the need to consider water constraints alongside energy and emissions in AI sovereignty strategies.

142 World Bank Group, “Levels of water stress: freshwater withdrawal as a proportion of available freshwater resources”, Food and Agriculture Organization of the United Nations (FAO), 2022, accessed December 2025.

143 EESI, “Data centres and water consumption”, June 2025, accessed December 2025.

144 Epoch AI, “Frontier Data Centres”, 2025, accessed December 2025.

Figure 28: Water stress levels

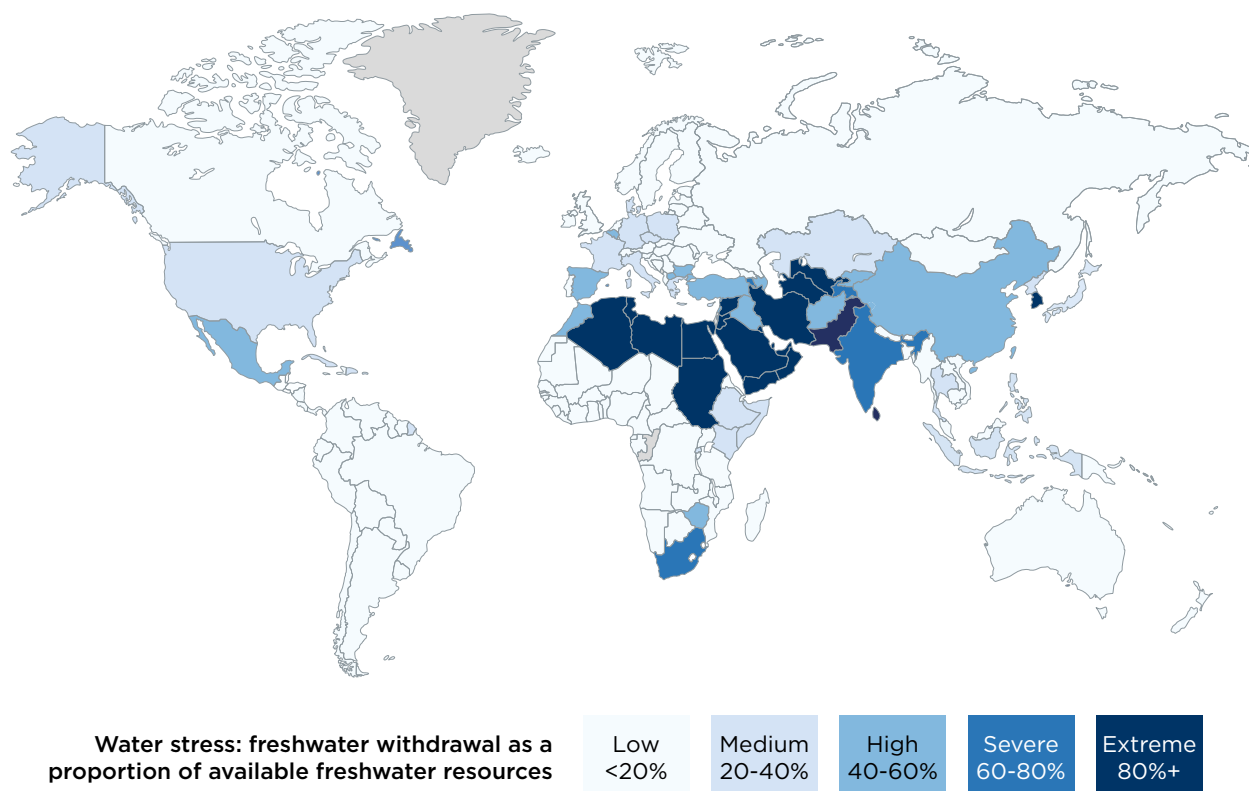


Figure 29: Water consumption by level of restriction, 2035

	Additional gallons of water required, bn (relative to Restriction level 1)		
	Restriction levels 2 and 3	Restriction level 4	Restriction level 5
Cambodia	0.0	0.0	0.0
India	0.2	2.5	5.1
Indonesia	0.1	0.6	1.4
Japan	0.3	3.6	7.3
Lao PDR	0.0	0.0	0.0
Malaysia	0.0	0.3	0.6
Myanmar	0.0	0.0	0.1
Nepal	0.0	0.0	0.0
Philippines	0.0	0.3	0.6
Singapore	0.1	0.7	1.4
South Korea	0.2	2.1	4.3
Taiwan	0.1	0.7	1.5
Thailand	0.0	0.3	0.6
Vietnam	0.0	0.3	0.5

Source: Oxford Economics

BOX 3: SINGAPORE'S GREEN DATA CENTRE ROADMAP: BALANCING AI-DRIVEN COMPUTE GROWTH WITH A NET ZERO TARGET

The challenge: a binding environmental constraint in Asia's data centre boom

The Asia Pacific region is poised to be one of the fastest growing regions, with data centre capacity projected to double to 37.6 GW by 2030.¹⁴⁵ However, Asia's power mix remains dominated by fossil fuels, turning data centres into a major "energy-climate dilemma".¹⁴⁶

Singapore is at the forefront of this climate dilemma. As a tropical city state, with no natural resources and limited potential to harness renewables¹⁴⁷, Singapore faces high cooling loads, with data centres accounting for a large share of sectoral emissions. The challenge is to keep scaling top tier compute for AI while reducing energy and water impacts.

The dual mandate: net zero by 2050 and an AI hub under National AI Strategy 2.0

Singapore's pledge to reach net-zero by 2050 sits alongside the National AI Strategy 2.0's (NAIS 2.0's) push to make AI a growth engine, with investments in research, talent, and compute.¹⁴⁸ This creates a trade-off between expanding high-density compute and reducing emissions. To achieve the vision set out in the roadmap, Singapore has adopted a coordinated, whole of government approach with industry and academia.

The Green Data Centre Roadmap and capacity "green-lane"

The Green Data Centre Roadmap was set out to balance these aims: sustain the ambitions of the NAIS 2.0 and the broader digital economy, while ensuring growth is environmentally responsible. It provides a clear pathway to add at least 300 MW of near-term capacity, and to unlock more by pairing best-in-class efficiency with credible low-carbon energy.¹⁴⁹

On 1 December 2025, Singapore announced its second Data Centre - Call For Application which emphasises the dual mandate.¹⁵⁰ Future capacity allocation requires operators to demonstrate credible plans for using viable low-carbon energy sources and deploying best-in-class efficiency solutions. For instance, power usage efficiency requirements for new builds were tightened to 1.25 in 2025, from 1.3 in 2022. To improve water efficiency, new builds must also achieve the BCA-IMDA Green Mark for Data Centres (2024) Platinum certification, which requires a water usage effectiveness (WUE) below 2.2, where highest points are allocated to builds with water usage efficiency of less than or equal to 2.0m³/MWh/year.¹⁵¹

145 KPMG, "The Asia Data Centre Landscape", 2025, accessed December 2025.

146 IEA, "Total energy supply, Asia Pacific, 2023", accessed December 2025.

147 Energy Market Authority, "Energy 2050 Committee Report", 2022, accessed December 2025.

148 Smart Nation, "Singapore National AI Strategy 2.0", 2023, accessed December 2025.

149 IMDA, "Green Data Centre Roadmap", 2024, accessed December 2025.

150 Economic Development Board, "Launch of Second Data Centre - Call For Application", 2025, accessed December 2025.

151 Building and Construction Authority, "BCA-IMDA Green Mark for Data Centres", 2024, accessed December 2025.

Decarbonising the whole stack

Institutionally, the roadmap adopts a **whole-of-stack model**. Facilities will be progressively upgraded and operated more efficiently, guided by standards tailored to tropical conditions¹⁵² and the refreshed Green Mark for Data Centres (2024), along with complementary IT energy efficiency and liquid cooling standards.¹⁵³ Water efficiency will also be strengthened in parallel. The water and energy usage of data centres are a global concern. With open procurement markets, Singapore can benefit from innovations and solutions being developed around the world.

On the compute infrastructure and software front, operators and end users are encouraged to upgrade to energy efficient equipment and optimise software. This is supported by global public-private partnerships, including the Sustainable Software Development Guidelines developed by IMDA and Microsoft¹⁵⁴, and dedicated funding to advance green computing.¹⁵⁵ Innovation is also being catalysed through partnerships with universities and industry, including a national testbed for green technologies.¹⁵⁶

Singapore has turned environmental limits into strategic design rules for green data centres. After the moratorium, new capacity is treated as a scarce national resource and prioritised for projects with top-tier PUE/WUE, energy-efficient IT, and low-carbon power. Trade-offs remain, as Singapore will depend on regional low-carbon imports and fuels (such as hydrogen) and on global advances and innovation in AI efficiency research. Overall, it shows how a land-scarce, tropical economy can decarbonise the stack, from facilities and hardware to software and energy, while growing its digital economy.


152 IMDA, “[Tropical Data Centre \(DC\) Standard](#)”, 2023, accessed December 2025.

153 IMDA, “[Singapore Standard SS 715:2025: Energy Efficiency of Data Centre IT Equipment](#)”, 2025, accessed December 2025.

154 IMDA & Microsoft, “[Sustainable Software Development Guidelines](#)”, 2023, accessed December 2025.

155 IMDA, “[Green Computing Funding Initiative](#)”, 2024, accessed December 2025.

156 Begum, S., “[Data centre test bed powered by green energy among new projects to fuel Jurong Island’s green push](#)”, JTC, November 2025, accessed December 2025.



SECTION 8. THE ROLE OF INDUSTRY AND GLOBAL CLOUD PROVIDERS

CONTROL-AND-CHOICE AS AN ENABLER OF SOVEREIGNTY

Our analysis indicates that highly restrictive approaches to AI sovereignty tend to force difficult trade-offs. Tight controls can raise costs, constrain access to skills, and slow progress towards decarbonisation. By contrast, control-and-choice models—which combine robust assurance over data and operations with continued access to best-in-class international systems—are associated with faster productivity growth, broader diffusion of innovation, and lower overall risk.



Emerging international evidence suggests that sovereignty does not require self-sufficiency.¹⁵⁷

Instead, for AI—particularly at the emerging stage of the lifecycle—priority lies in enabling innovation, accelerating adoption, and setting strong foundations for responsible AI. Sovereignty rests on the ability to select and configure AI systems on domestic terms, blending global and local components while meeting legal, ethical, and security requirements. In practice, governments seek to balance three objectives: protecting sovereignty; enabling economic efficiency through access to frontier AI at competitive cost; and advancing environmental sustainability as AI demand increases energy and water use. Achieving all three simultaneously is challenging, and outcomes depend critically on the extent and design of restrictions.

Control-and-choice approaches manage this tension by relying on verifiable safeguards—such as data residency requirements, encryption with key management, and operator accountability—without excluding global providers. This helps align sovereignty objectives with efficiency and sustainability. More restrictive, isolation-based models impose sharper trade-offs. Limiting access to global vendors and talent can divert **workloads to smaller, less efficient facilities, raising costs** and slowing progress towards net-zero goals. Evidence from the data centre sector shows that larger, more integrated facilities typically achieve higher energy efficiency.¹⁵⁸

Recent developments in Europe illustrate this dynamic. Early proposals for the EU Cloud Certification Scheme included stringent sovereignty requirements, such as immunity from non-EU laws. Subsequent revisions shifted towards risk-based technical assurances rather than categorical exclusions, reflecting efforts to balance security objectives with competitiveness and environmental performance.¹⁵⁹ In this framework, sovereignty is strengthened by managing choice rather than restricting it.

Control-and-choice models introduce clear, reliable safeguards while allowing global providers to compete and innovate. They align with emerging definitions of AI sovereignty in which assurance levels are calibrated to the sensitivity and risk of specific workloads, rather than applied uniformly. Core elements of an effective control-and-choice framework include standards-based governance using internationally recognised risk frameworks; clear cloud security principles covering identity and data governance; identity and data access management; and portability and reversibility to reduce lock-in. Together, these components protect sovereign interests while preserving access to global R&D, talent and infrastructure (see Box 4).¹⁶⁰

157 Accenture, “Sovereign AI: Own Your AI Future from Managing Risk to Accelerating Growth”, 2025, accessed December 2025.

158 Accenture, “Sovereign AI: Own Your AI Future from Managing Risk to Accelerating Growth”, 2025, accessed December 2025.

159 Foo, C., “EU drops sovereignty requirements in cybersecurity certification scheme, document shows”, Reuters, April 2024, accessed December 2025.

160 Accenture, “Sovereign AI: Own Your AI Future from Managing Risk to Accelerating Growth”, 2025, accessed December 2025.



BOX 4: CORE COMPONENTS OF A CONTROL-AND-CHOICE FRAMEWORK

1. Standards-based governance: Using internationally recognised risk frameworks, such as the **NIST AI Risk Management Framework** and the **ISO/IEC 42001 AI management standard**, improves transparency and interoperability.^{161,162}

2. Clear cloud security principles: The UK National Cyber Security Centre's (NCSC) **14 Cloud Security Principles** provide a benchmark for identity management, data governance, supply-chain risk, and operational security.^{163,164}

3. Transparency and auditability: Mechanisms such as Google's **Access Transparency** and **Access Justifications** allow governments to monitor when and why cloud providers access customer data.^{165,166}

4. Portability and reversibility: The ability to move workloads, models, or data between environments reduces vendor lock-in. Open standards and registries, such as the **CSA STAR** registry, support choice.¹⁶⁷

This model protects sovereign interests while preserving access to global R&D, talent, and infrastructure.

161 Tabassi, E., "Artificial Intelligence Risk Management Framework (AI RMF 1.0)", NIST, 2023, accessed December 2025.

162 ISO, "ISO/IEC 42001: Information technology—Artificial Intelligence—Management System", 2023, accessed December 2025.

163 UK National Cyber Security Centre, "Cloud security guidance", 2023, accessed December 2025.

164 UK Government Digital Service, "Cloud guide for the public sector", 2023, accessed December 2025.

165 Google Cloud, "Access Transparency and Access Approval", n.d., accessed December 2025.

166 Google Cloud, "Overview of Access Transparency", n.d., accessed December 2025.

167 Cloud Security Alliance, "Security, Trust, Assurance and Risk (STAR)", n.d., accessed December 2025.

WHAT BALANCED SOVEREIGNTY COULD LOOK LIKE

Sovereignty, when designed well, can enable organisations across a shared ecosystem to collaborate securely. In regulated environments this can unlock joint model development, data sharing, and coordinated insights without exposing sensitive assets. However, most businesses perceive AI sovereignty with a defensive posture. This posture is motivated by compliance requirements (46%), control over critical data (28%), and cybersecurity concerns (27%), according to Accenture (2025). In the same survey, fewer than 13% cite monetisation or cultural alignment as strategic motivators (Figure 30).

Effective sovereign platforms can support data marketplaces, attract talent, create models tuned to local contexts, and nurture domestic innovation. Examples are emerging across sectors:

- In Saudi Arabia, **AWS and HUMAIN** are investing over US\$5 billion to build a sovereign AI Zone that provides dedicated AI infrastructure, supports locally governed model development, and expands national AI skills through large-scale training programmes.¹⁶⁸
- In Indonesia, **Indosat Ooredoo Hutchison**, working with Accenture and Nvidia, is building a sovereign AI cloud to support local start-ups and government clients while ensuring national data remain onshore.¹⁶⁹
- In Europe, **Oracle's EU Sovereign Cloud**, as well as **Microsoft's Bleu** and **Delos** partnerships, enable organisations to run sensitive workloads entirely under EU law.¹⁷⁰
- In healthcare, the **Sovereign AI Factory Frankfurt** provides a compliant environment for hospitals and research institutes to develop and deploy models under GDPR and the AI Act.¹⁷¹

In addition to these examples, major cloud providers are now offering sovereign-aligned AI infrastructure—such as AWS's AI Factories—which deliver dedicated, on-premises environments for secure data processing, locally governed model development, and workforce training.

If done well, sovereign assurance becomes an enabler of innovation and sustainability rather than a barrier and helps governments navigate the policy trilemma with greater flexibility.


168 AWS, "AWS and HUMAIN announce a more than \$5B investment to accelerate AI adoption in Saudi Arabia and globally", n.d., accessed March 2026.

169 Accenture, "Indosat Ooredoo Hutchison Group and Accenture Accelerate Sovereign AI Cloud to Power Indonesia's Digital Future", 2024, accessed December 2025.

170 Oracle, "Oracle EU Sovereign Cloud Helps Organizations Across Germany Securely Manage Critical Business Data", 2025, accessed December 2025.

171 Microsoft, "Announcing comprehensive sovereign solutions empowering European organizations", 2025, accessed December 2025.

Figure 30: Key motivations for Sovereign AI



Source: Accenture (2025), recreated by Oxford Economics

HOW GLOBAL PROVIDERS CAN SUPPORT SOVEREIGN OBJECTIVES

At the foundation of sovereignty is security, which means ensuring the confidentiality, integrity, and availability of sensitive workloads while retaining meaningful control over who can access data and under what conditions. Major AI and cloud providers increasingly offer features that help governments meet **sovereignty-aligned requirements** without compromising capability.

COMPLIANCE AND DATA-RESIDENCY CONTROLS

Major cloud providers are strengthening data storage practices within Europe. For example,

Microsoft’s EU Data Boundary ensures that core customer data, logs, and support information are stored and processed entirely within the EU or EFTA.^{172,173} Google’s EU Data Boundary, supported by its Access Justifications feature, keeps processing within the region and requires a specific, documented reason for any administrative access.¹⁷⁴ AWS provides comparable assurances through Control Tower data-residency controls and a broad set of independently validated certifications, as well as the AWS Nitro System, which offers hardware-based safeguards designed to prevent unauthorised access to data during processing.¹⁷⁵

172 Microsoft, “Microsoft EU Data Boundary”, accessed December 2025.

173 AP News, “Microsoft lets cloud users keep personal data within Europe to ease privacy fears”, 2024, accessed December 2025.

174 Google Cloud, “EU Data Boundary with Access Justifications”, n.d., accessed December 2025.

175 AWS, “Control tower controls data residency requirements”, 2021, accessed March 2026.

TRANSPARENCY MECHANISMS

Greater transparency is emerging as a core safeguard. Google's Access Transparency tools give organisations a clear audit trail of when, how, and why a provider accesses their environment. This improves accountability and supports compliance with regulatory expectations.¹⁷⁶ AWS provides equivalent transparency through the Sovereign Reference Framework for the AWS European Sovereign Cloud, which sets out independently validated sovereignty controls and makes auditable evidence available to customers via AWS Artifact. These reports give organisations clear insight into how sovereignty expectations are met across governance, operations, and data residency.¹⁷⁷

ENVIRONMENTAL LEADERSHIP

Global cloud providers can also accelerate progress towards national sustainability goals. Their commitments such as 24/7 carbon-free energy¹⁷⁸, water-positive operations, and integration with district-heating systems often outperform what smaller, standalone facilities can achieve. However, these benefits need to be balanced against the rising energy demands associated with AI.

LOCAL PARTNERSHIPS AND SOVEREIGN OPERATIONS

Several countries are adopting "trusted cloud" models that blend global technology with local oversight. Across Europe, providers such as Microsoft and Google are partnering with national operators to deliver cloud services with in-region governance, local personnel, and enhanced

control over data access and operations. AWS's European Sovereign Cloud will similarly operate with EU-based staff and in-region controls to reinforce legal and operational independence.¹⁷⁹ In addition, AWS and the German Federal Office for Information Security (BSI) are jointly developing sovereignty and cybersecurity standards to support national requirements.¹⁸⁰ These arrangements aim to strengthen national innovation ecosystems and help domestic firms build new capabilities on top of global platforms.

Recent cross-country evidence underscores how widespread these blended models have become. The Center for a New American Security's (CNAS) Sovereign AI Index which tracks over 130 national AI infrastructure and model projects across countries, records that foreign companies are involved in roughly 70% of projects, and that around four-fifths of those foreign-involved projects feature a U.S. company.¹⁸¹

Professor Yoon, Professor of Administrative Law at Hanyang University School of Law who we interviewed for this report, emphasises the importance of recognising the opportunity costs involved in pursuing sovereignty across the AI stack. She notes that in the private sector, global sovereign cloud solutions can meaningfully mitigate some of these trade offs. To enable effective partnerships between global cloud providers and local firms, she explained that governments can establish structured frameworks in which providers commit to locally relevant standards, oversight mechanisms, and partnership obligations as a condition of access or certification [see page 74].

176 Google Cloud, "Access Transparency and Access Approval", n.d., accessed December 2025.

177 AWS, "Exploring the new AWS European Sovereign Cloud", 2025, accessed March 2026.

178 Google Cloud, "A policy roadmap for achieving 24/7 carbon-free energy", 2022, accessed December 2025.

179 AWS, "In the Works - AWS European Sovereign Cloud", 2023, accessed December 2025.

180 AWS, "AWS and BSI cooperation agreement", 2025, accessed March 2026.

181 CNAS, "Sovereign AI Index: Tracking the Global Push for AI Self-Reliance", 2025, accessed April 2025.

SOUTH KOREA'S PATH TO AI SOVEREIGNTY

Among APJ countries, South Korea stands out as a highly digitalised economy with world-leading semiconductor capabilities and one of the most explicit ambitions in the region to become a global AI power. South Korea's AI strategy is publicly framed around becoming a top three global AI powerhouse, supported by expanded compute capacity, AI adoption, and domestic capability-building.¹⁸²

Professor Melissa Hyesun Yoon—a Professor of Administrative Law at Hanyang University School of Law, affiliated with its Graduate School of Artificial Intelligence, and an adviser to several Korean government bodies on AI regulation and data governance—describes that Korea's governance framework as a layered system as opposed to a single, coherent doctrine. At the compute layer, the Cloud Security Assurance Program (CSAP) creates tiered requirements that shape which providers can serve sensitive public-sector workloads. At the model layer, the government encourages domestic foundation model development, while at the application layer the AI Basic Act imposes transparency, risk-management, and impact-assessment obligations on developers and deployers regardless of where the underlying model or infrastructure originates. Korea's Personal Information Protection Act (PIPA) cuts across all three layers as a baseline sovereignty mechanism. In Professor Yoon's assessment, the dominant motivations are national security, economic competitiveness, and industrial policy, with resilience an important but secondary rationale.

Professor Yoon characterises Korea's model as a hybrid, drawing from both the United States and Europe. From the EU, Korea has adopted elements of risk-based regulation and public-interest oversight; from the US, it has taken cues from innovation-first policy design and the central role of private-sector dynamism. This does not, in her view, amount to a requirement for full domestic control across the board. Rather, sovereignty in Korea operates as a spectrum: controls are tightest for sensitive public-sector and national-security workloads, while commercial and lower-sensitivity workloads remain more open to global providers.

At the model layer, domestic development is encouraged but foreign models are not prohibited; at the application layer, sovereignty is exercised primarily through Korean regulatory jurisdiction rather than ownership.

Professor Yoon's assessment is that Korea recognises the limits of pursuing a fully home-grown AI stack at the frontier. She argues that, even with substantial public investment, the resource gap between Korean firms and major US AI companies remains very large at the compute layer of the stack. Domestic model development still has strategic value—particularly for Korean-language capability and high-sensitivity domains—but she does not see full frontier self-sufficiency as a realistic objective. The more important question, in her view, is one of opportunity cost: whether public and private resources are best directed to domestic compute and foundation models, or to skills, data infrastructure and applications, where Korea may derive stronger returns. For that reason, she sees access to global cloud infrastructure as economically important, especially for the private sector, where it can mitigate performance and cost trade-offs by providing frontier compute without equivalent domestic capital expenditure.

This framing also helps explain Korea's more recent, albeit measured, shift towards greater participation by foreign cloud providers. Professor Yoon noted that public-sector procurement is particularly influential in Korea because government acts as both the largest consumer and a leader for diffusion. In 2026, the Korean government opened procurement for various AI projects to domestic and international cloud service providers, even as the broader CSAP framework remained in place for sensitive workloads.¹⁸³

These developments suggest that South Korea, with its advanced digital infrastructure and already high levels of AI readiness and adoption, is pursuing sovereignty through a balanced approach that combines domestic capability-building and regulatory control with strategic openness to global providers, in order to sustain innovation, scale, and economic competitiveness.

182 South Korea Ministry of Science and ICT, "Blueprint for Korea's Leap to become one of the top three Global AI powerhouse (AI G3)", 2024, accessed April 2026.

183 Notably, the AI Computing Resource Utilization Infrastructure Enhancement Project and the High-Performance Computing Support Project.

SECTION 9. WAY FORWARD AND RECOMMENDATIONS

This paper quantifies the economic costs associated with different degrees of restrictiveness, providing one of several inputs to support governments' assessment of sovereign AI policy choices. Our results underline a central implication for policy design: highly restrictive approaches to AI sovereignty can raise costs, constrain access to skills, and slow progress towards decarbonisation. Conversely, control-and-choice approaches are more likely to preserve productivity gains by keeping access to best-in-class global systems under robust assurance, enabling governments to configure AI on domestic terms while still drawing on global R&D, talent, and infrastructure. This allows countries to focus investment where domestic value creation is strongest—skills, datasets, and applications—rather than diverting resources into replicating capital intensive compute.





As the Tony Blair Institute—whom we consulted as part of this research—emphasises, sovereignty should be understood not as technological isolation but as strategic agency: the ability of governments to make autonomous decisions over critical digital systems while remaining connected to global innovation and supply chains. This perspective reinforces the evidence presented in this report, showing that balanced, assurance led approaches can deliver stronger long term autonomy than models focused on full self sufficiency [see page 77].

In practice, governments are seeking to balance three objectives—sovereignty, economic efficiency, and environmental sustainability—as AI demand increases energy and water use. Achieving all three simultaneously is challenging, and outcomes depend critically on the extent and design of restrictions.

This aligns with the view of Dr Supheakmungkol Sarin, cofounder of AI Safety Asia and one of our interviewees for this report. He argues that most emerging economies lack the scale to build and operate a full domestic AI stack. As a result, they can achieve greater resilience through regional cooperation, shared infrastructure, and interoperable governance frameworks, rather than through costly localisation efforts [see page 81].

TONY BLAIR INSTITUTE FOR GLOBAL CHANGE: COMPETING AT THE FRONTIER - AI SOVEREIGNTY THROUGH STRATEGIC NATIONAL INVESTMENTS

AI sovereignty is not a binary condition that countries either find themselves in or not, nor is it synonymous with full technological self-sufficiency. Instead, sovereignty in the age of AI should be understood as the ability of governments to exercise strategic agency. That is, to make targeted, future-oriented choices about how AI is accessed, governed, and deployed in line with national priorities.


The Tony Blair Institute's research finds that governments are increasingly moving away from a narrow conception of AI sovereignty centred on owning and onshoring the entire AI stack (from semiconductors and data centres to models and applications). They suggest that, while politically understandable, this approach is often economically inefficient, technologically unrealistic and ultimately counterproductive.

Sovereignty should instead be defined by the ability of governments to shape outcomes and utilise available technology in pursuit of their national goals and objectives within an interdependent global ecosystem underpinned by supply chains that cross borders and continents. Against this background, strategic choices on AI development and diffusion must navigate a persistent trilemma of pursuing domestic control, accessing frontier capability through global systems, and maintaining coherence across regulatory, industrial and diplomatic strategies, with each choice calibrated to national comparative advantages, resource constraints, and strategic objectives.

A central tension within that trilemma concerns the trade-off, for example, of speed versus control. Importing and deploying frontier AI systems from global providers can significantly accelerate technological diffusion and economic impact, but may increase dependence on external actors. Conversely, building domestic systems may enhance control but at considerable cost and with significant time delays, and is realistically feasible only for countries with advanced digital infrastructure and strong talent ecosystems.

As more countries realise that they can't compete at frontier model development, they are increasingly focusing on tuning foreign cutting-edge models to their needs, domestic languages, and sectors. This is achieved through international partnerships that secure access to frontier models and compute infrastructure, while also leveraging national comparative advantages, such as in energy production, to assert influence in different parts of the AI stack.

Critically, the Tony Blair Institute highlights that the economic benefits of AI will only materialise through fast adoption and diffusion of AI technologies. An isolationist approach to AI sovereignty can significantly jeopardise those economic opportunities. At the national level, importantly, the advantage of embracing AI lies in improving services for citizens as well. This requires more than investment in infrastructure. It depends on strong leadership, institutional reform, and a rethinking of how government operates.



PRINCIPLES FOR BALANCED AI SOVEREIGNTY POLICIES

1) Risk-based data classification: Only around 10% of government data are highly sensitive, yet many sovereignty measures already extend beyond this to cover specific sectors or so-called “strategic” datasets. Targeting localisation narrowly at genuinely high-risk workloads allows governments to meet security objectives while avoiding the higher costs that arise when restrictions are applied more broadly than necessary.

2) Assurance through transparency and auditability: Control-and-choice models rely on verifiable safeguards—such as data residency requirements, encryption with key management, zero-operator access, and operator accountability—rather than categorical exclusions of global providers. This can strengthen sovereign oversight while maintaining capability and scale. Importantly, governments can anchor assurance in internationally recognised governance and risk-management standards as well as AI-specific standards (e.g., ISO 42001 on responsible AI development), which is particularly valuable for jurisdictions where deep technical expertise may be limited. Leveraging global standards provides a consistent, trusted baseline for oversight while avoiding the need to build bespoke technical requirements in every country.

3) Economic competitiveness through openness to global innovation: Policy should safeguard the open competition that underpins AI-driven productivity gains. Overly restrictive localisation or preferential treatment for domestic providers can weaken competition and limit access to leading technologies, including the advanced cybersecurity capabilities offered by hyperscale providers. Portability and reversibility—enabled by open standards, assurance registries, and

governance frameworks—help ensure that providers compete on performance and innovation rather than lock-in, reinforcing openness.

4) Partnership models that blend global capability with local control: A recurring theme in the evidence base is that sovereignty does not require self-sufficiency and isolationism, but depends on the ability to configure AI systems on domestic terms while combining global and local components; a configuration sometimes described as “managed interdependency”.¹⁸⁴ Trusted cloud approaches can operationalise this by pairing global technology with local oversight and compliance controls; examples include partnership and sovereign-operations models that preserve access to global R&D, talent, and infrastructure. When partnership models are used by government organisations, they also strengthen governments’ role as an innovation catalyst: early public-sector adoption demonstrates value, reduces uncertainty for firms, and supports diffusion across the wider economy. By contrast, restrictive sovereignty measures that delay government uptake weaken this effect.

As Arifah Sharifuddin (Institute Director at the Tech for Good Institute (TFGI) and an interviewee for our research) notes, these models work best when global capability is paired with contextualisation, trust, and strong public-sector capacity. Localising solutions to linguistic and cultural contexts, particularly for smaller enterprises, needs to be supported by practical governance processes and ongoing engagement with technology partners. Mechanisms such as regulatory sandboxes can help build confidence, clarify responsibilities, and create the conditions for adoption at scale. [see page 79].

184 Tanner, B., et al., “Is AI sovereignty possible? Balancing autonomy and interdependence”, Brookings, February 2026, accessed April 2026.

ARIFAH SHARIFUDDIN: BALANCING SOVEREIGNTY AND GLOBAL PARTNERSHIPS IN SOUTHEAST ASIA'S AI PATHWAY


“In Southeast Asia, the AI policy conversation is moving from ‘protecting the now’ to ‘building the next’ – from safeguarding national interests to shaping competitive digital industries,” notes Arifah Sharifuddin, Institute Director at the Tech for Good Institute (TFGI), a Singapore-based policy institute focused on inclusion and sustainability in the digital economy.

Drawing on TFGI’s work tracking technology governance across Southeast Asia’s six largest countries, Arifah describes governments as actively engaging across the full AI stack. This includes infrastructure readiness, data governance, talent development, and efforts to translate public investment into economic value, while also signaling the advancement of national interests.

“You can only enable adoption by micro and small enterprises when you localise your context,” she explains, highlighting the region’s linguistic and cultural diversity. Customised solutions, however, must be accompanied by practical steps to improve market access: how governments create space for local producers to participate in public contracts, even as they remain open to global providers.

Arifah expects Southeast Asia’s AI pathway forward will be defined less by a single regional “rulebook” and more by a collection of national approaches anchored in a shared objective: accelerating adoption while maintaining trust. Regional initiatives such as the ASEAN AI Safety Network and the ASEAN Digital Economy Framework Agreement are important. Yet, implementation remains difficult given differing domestic legal systems and institutional capacity. With AI use evolving rapidly, from generative to increasingly autonomous systems, governments often cannot legislate fast enough and instead rely on revising existing acts, a process that can still take years.

In this environment, Arifah sees a crucial role for global cloud and technology players. Policymakers are not the technical builders of AI systems, so sustained engagement and concrete use cases are essential to avoid rules that look good on paper but fail in practice. She highlights regulatory sandboxes as a pragmatic partnership mechanism, from controlled pilots in Singapore to emerging cross-border concepts like the Singapore-Johor-Batam Special Economic Zone’s trusted data corridors. Done well, these partnerships can help governments build trust, clarify responsibilities, and create the conditions for adoption to scale.





IMPLICATIONS AND WAY FORWARD

Governments across APJ are increasingly motivated to invest in national digital infrastructure and AI-enabled public services. The policy choice is not simply “open vs closed”; it is about calibrating assurance to the sensitivity and risk of workloads, while protecting the adoption and scale dynamics needed for economy-wide diffusion. India’s MANAV Vision—presented at the India AI Impact Summit 2026—illustrates this direction, framing sovereignty around ethical design, accountable governance, inclusivity, and legitimacy rather than strict localisation or self-sufficiency.¹⁸⁵

A control-and-choice approach is best placed to manage the sovereignty-efficiency-sustainability trade-offs. This enables robust assurance without forcing workloads onto smaller, less efficient facilities—a shift that can raise costs and slow progress towards net zero, given evidence that larger, more integrated data-centre facilities typically achieve higher energy efficiency.

Europe’s evolving cloud policy debate illustrates the same tension, with movement towards risk-based technical assurance rather than blanket exclusions as policymakers seek to balance security objectives with competitiveness.¹⁸⁶

In addition, where policy has been set, clear and consistent regulatory guidance is also important to avoid unintended behavioural spillovers. When rules are ambiguous, firms may over interpret requirements—localising data or avoiding global partnerships even when this is unnecessary—raising costs, slowing adoption, and increasing the risk of inefficient infrastructure choices. Providing clarity on what is permitted helps ensure policy objectives are met without creating avoidable burdens.

In practical terms, this points to collaborative, innovation-friendly sovereignty models—combining clear assurance requirements with interoperable market access and structured partnerships—so that sovereignty objectives are met without unnecessarily curtailing AI adoption and the associated productivity gains.

185 Prime Minister Shri Narendra Modi, “[Inauguration speech at AI Impact Summit 2026](#)”, 2026, accessed March 2026.

186 The European Commission, “[EU Cloud Certification Scheme](#)”, 2021, accessed December 2025.

DR SUPHEAKMUNGKOL SARIN: BALANCING OPENNESS AND DOMESTIC CAPABILITY IN SOUTHEAST ASIA'S AI PATHWAY

Dr Supheakmungkol Sarin, the co-founder of AI Safety Asia, argues that Southeast Asia's ability to capture the AI opportunity rests on a deceptively simple foundation: trust. This trust is earned through credible governance, practical capability, and institutional readiness.

After two decades in AI, from machine-learning research to nearly a decade at Google AI advancing inclusive technologies across emerging economies, Dr Sarin has seen first-hand how gaps in data, talent, and institutional capacity shape the region's trajectory. His experience in Cambodia, Myanmar, Bangladesh, Sri Lanka, and Nepal showed how scarce local data and technical talent can stall progress. "There's growing understanding around AI," he notes, "but gaps remain." In many smaller economies, even basic AI policy functions remain under-resourced. The result is a widening mismatch between fast-moving capability and the slower pace at which governments can respond.

This imbalance shapes how risk is identified and addressed. While concepts such as "AI sovereignty" are gaining attention, Dr Sarin notes that the term remains loosely defined in the region, with little consensus on whether it means building the full stack at home or retaining control over key dependencies and auditing systems used locally. At present, policymakers focus on visible, near-term harms, but he warns that risks are flying under the radar, including agentic systems that can plan and use tools with limited oversight, and wider resilience shocks to jobs, wages, and public trust. Openness to technology is relatively high in Southeast Asia, driven by a young, digitally-fluent population. But that trust is fragile. "If something bad happens," he cautions, "the trust could be really difficult to gain back."

For governments, the challenge is to enable rapid adoption without undermining long-term resilience. Dr Sarin is unequivocal: most Southeast

Asian economies should avoid building the full AI stack domestically. The demands make this unrealistic. Instead, he advocates three strategies. First, pooling resources regionally to improve bargaining power, evaluation capacity, and access to compute. Second, specialising in areas where each economy has established strengths, such as financial services in Singapore, semiconductors in Malaysia, and healthcare, wellness, and medical tourism in Thailand. Third, adopting a hybrid approach that keeps markets open while reducing overreliance on any single provider or model family.

Regulation should follow a similar logic. Rather than importing a regulatory framework, governments should build Asia-rooted playbooks by drawing on global approaches and adapting them to local labour markets, industrial structures, and state capacity. This requires region-specific evidence, continuous monitoring, shared foresight, scenario planning, and engagement across labour, digital, education, economic, and social-policy communities.

He highlights public-sector procurement as one of the most powerful levers for governments. In his view, procurement should embed responsible AI principles from the outset, including clear accountability for harms, redress mechanisms for citizens, basic audit requirements, and limits on what systems can access or do without approval. "The government must ensure who is responsible, and what is the mechanism to redress," he says. Building this governance layer transparently and inclusively is essential to sustaining societal trust.

For Southeast Asia, where adoption is already accelerating, earned trust may be the true competitive advantage. As Dr Sarin puts it: "With trust, adoption will follow."

APPENDIX 1: AI SOVEREIGNTY POLICIES IN APJ COUNTRIES

This appendix details the various AI sovereignty policies adopted or under consideration in the various APJ countries in this report. These are organised by the layers of the AI stack.

COMPUTE INFRASTRUCTURE

Compute infrastructure is the foundation of the AI stack. The cloud-based compute layer is the bedrock of AI adoption, providing access to high performance graphics processing units (GPUs), storage, and networking that enable both large-scale model training and smaller-scale experimentation. While some organisations own these resources, most—including the many small and medium enterprises (SMEs) vital for broad AI adoption— will access them through cloud service providers' infrastructure. At this layer, sovereignty measures vary widely especially in the context of provision of services to the public sector.

- **Japan** focuses on assurance rather than isolation, emphasising international operability of AI systems: its Information System Security Management and Assessment Program (ISMAP) certification allows government workloads on commercial clouds that meet security standards, while broader policy supports trusted cross-border data flows under the “Data Free Flow with Trust” framework.^{187,188} Japan's approach emphasises secure data sharing and adherence to voluntary industry guidelines.

- **Singapore** takes a similar approach with its Government on Commercial Cloud (GCC 2.0) platform, which lets agencies use major cloud platforms within a government-secured framework with enhanced monitoring and encryption.^{189,190} It is an industry-led approach that promotes corporate AI governance through non-binding guidelines rather than rigid enforcement. This illustrates Singapore's priority on agility and innovation via cloud, while ensuring sovereignty through rigorous cybersecurity.
- **Indonesia** and **Malaysia** currently adopt hybrid models—allowing offshore processing for the private sector under safeguards and operating government cloud programs to balance openness with security.
 - » In Indonesia, offshoring data storage for private sector services is allowed, provided that companies obtain user consent for cross-border data transfers. Certain sectors that involve “strategic data”, such as the finance sector, still face localisation or onshore processing requirements.¹⁹¹
 - » Malaysia has explicitly adopted a hybrid cloud policy for its public sector: the MyGovCloud initiative combines a government-operated private cloud with services from appointed public cloud providers—including AWS, Google, Microsoft, and Telekom Malaysia.^{192,193} The government also has strategic partnerships with those providers, which have all announced data

187 Digital Agency of Japan, “[Updated ISMAP cloud services list](#)”, 2025, accessed December 2025.

188 Digital Agency of Japan, “[Data Free Flow with Trust \(DFFT\)](#)”, 2024, accessed December 2025.

189 Pradana, D., “[Singapore's Government on Commercial Cloud \(GCC\): 2025 Guide to a Secure, Scalable Public-Sector Cloud](#)”, Accrets International, July 2025, accessed December 2025.

190 MDDI, “[Factsheet - Government Cyber Security Operations Centre \(GCSOC\)](#)”, 2023, accessed December 2025.

191 Information Technology and Innovation Foundation, “[Indonesia's Data Localisation Regulation](#)”, 2025, accessed December 2025.

192 Raj, A., “[Malaysia has a new government hybrid cloud service](#)”, Techwire, May 2022, accessed December 2025.

193 MyGovCloud, “[MyGovCloud](#)” 2025, accessed December 2025.

centre, cloud and AI projects in Malaysia.^{194,195,196} This allowed the nation to migrate most of its government systems to the cloud, while keeping oversight through vetted partners and in-country infrastructure. In May 2025, Malaysia expanded this model by launching the Strategic AI Infrastructure—a sovereign, full-stack AI environment designed to host sensitive AI workloads domestically—with a planned Sovereign AI Cloud set to further increase domestically governed AI capacity.¹⁹⁷

- **South Korea** applies stricter controls: under its Cloud Security Assurance Program (CSAP), cloud service providers must meet administrative, technical, and security measures required for public-sector use. Government agencies can only use cloud services that passed these audits and earned the CSAP certification.¹⁹⁸ Within this certification are three tiers, with the most stringent classification for systems containing sensitive or critical internal administrative information. In addition to these regulatory controls, South Korea previously imposed domestic-company requirements in tenders that effectively restricted participation by foreign providers. These barriers were relaxed in 2026—most notably in the request for proposals for major national AI compute and HPC projects—opening the door for foreign cloud service providers to bid, even though CSAP certification requirements continue to apply. South Korea’s AI Basic Act (effective from 2026) will also introduce further governance requirements, including risk assessments and possible registration for high-performance AI systems, to bolster trust in AI built on these clouds.^{199,200}

194 Prime Minister’s Office of Malaysia, “[PM Anwar Unveils National Cloud Policy Focus on Four Key Areas](#)”, 2024, accessed December 2025.

195 Azhar, D., Latiff, R., “[Malaysia plans national cloud policy, AI regulations](#)”, Reuters, October 2024, accessed December 2025.

196 Reuters, “[Malaysia launches national AI office for policy, regulation](#)”, December 2024, accessed December 2025.

197 Bernama, “[Malaysia Launches Region’s First Sovereign Full-Stack AI Infrastructure](#)”, 2025, accessed December 2025.

198 Microsoft Learn, “[Korea CSAP](#)”, August 2025., accessed December 2025.

199 Wong, S, “[Explainer: Korea’s AI Basic Act](#)”, Asian Legal Business, 2025, accessed December 2025.

200 Kim, J., et al., “[2025 AI Governance in Korea: Strategic Investments and Regulatory Reforms](#)”, Jipyong LLC, 2025, accessed December 2025.



- Elsewhere, **India** follows a hybrid model that uses global cloud service providers through its MeghRaj 2.0 framework, which sets common standards and requires India-based data centres, while the IndiaAI Mission expands national compute capacity through public-private partnerships that procure tens of thousands of GPUs and make them available to researchers and start-ups at up to 40% lower cost, supporting affordability, wider access, and the growth of domestic AI capability.²⁰¹
- Similarly, **Vietnam** enforces local storage for certain services under Decree 53 as part of the Cybersecurity Law. It mandates local storage of personal information and user-generated data in Vietnam, ensuring that Vietnamese authorities have jurisdictional control over critical data.²⁰² This reflects a focus on national security and data sovereignty at the infrastructure layer—even if general AI compute might still be cloud-based.
- In **Taiwan**, sector-specific regulations—not a blanket localisation rule—create tighter controls for certain categories of financial and healthcare data. While there are no localisation requirements for public-sector use of cloud services, government agencies that procure or utilise cloud providers must make sure those are not based in Mainland China (including Hong Kong and Macao). The physical locations for access and back-up of cloud data must also not be located in Mainland China.²⁰³
- In **Cambodia, Lao PDR, and Myanmar**, AI regulation is largely still under development. There are currently no restrictions or localisation rules, but it is likely that sectoral data protection clauses, e.g., in finance and health, will be introduced.^{204,205}

AI MODELS

The AI model layer provides the interface for developing, operating, and customising AI models. This allows customers to build their own models or to create and scale customised generative AI applications using models and datasets from a number of different providers. Easy access to AI tooling allows customers to fine-tune and augment the models they build with their own proprietary company data. It also makes capabilities such as retrieval augmented generation, which supplements user prompts with information from proprietary company data to improve reliability and performance. This layer faces growing governance attention in APJ countries, to ensure that as AI is built and deployed, it aligns with safety, ethics, culture, and other domestic regulatory requirements:

- **Singapore** leads with operational frameworks such as AI Verify and its Model AI Governance Framework for Generative AI. AI Verify—an open-source testing toolkit—helps developers self-assess model performance based on bias and explainability.²⁰⁶ The Model AI Governance Framework complements this with detailed guidance on risk assessment, transparency measures, and human oversight.²⁰⁷ In January 2026, Singapore also launched a Model AI Governance Framework for Agentic AI, extending its guidance-led approach to newer system types.²⁰⁸ These tools and guidelines exemplify Singapore’s emphasis on “governance-by-design”—providing practical means for AI developers to comply with safety, transparency, and accountability principles.
- **Japan, South Korea, and Taiwan** have issued guidelines and laws that introduce risk-based oversight without imposing blanket restrictions.

201 Ministry of Electronics and Information Technology (MeitY), “[India’s AI Revolution: A Roadmap to Viksit Bharat](#)”, 2025, accessed December 2025.

202 PwC Vietnam, “[Decree 53 guiding Cybersecurity Law](#)”, 2022, accessed December 2025.

203 Chambers and Partners, “[Cloud Computing](#)”, 2025, accessed January 2026.

204 Rouse Insights, “[Data Localisation and Transfer Issues in Southeast Asia](#)”, 2025, accessed January 2026.

205 DLA PIPER, “[Data Protection Laws of the World](#)”, 2025, accessed January 2026.

206 IMDA, “[Singapore launches world’s first AI testing framework and toolkit to promote transparency; Invites companies to pilot and contribute to international standards development](#)”, 2022, accessed December 2025.

207 IMDA, “[Singapore proposes framework to foster trusted Generative AI development](#)”, 2024, accessed December 2025.

208 IMDA, “[Singapore launches new model AI governance framework for Agentic AI](#)”, 2026, accessed February 2026.

- » Japan adopts a non-binding, “agile” approach to AI governance, favouring voluntary compliance over rigid rules. It steers industries through detailed guidance such as Ministry of Economy, Trade and Industry’s (METI’s) AI Governance Guidelines and the AI Guidelines for Business, and through the 2025 AI Promotion Act, which similarly encourages voluntary adherence to principles of transparency, safety, and international alignment. Implementation efforts focus on supporting research and development (R&D), computing infrastructure, talent development, and public literacy.^{209,210,211}
- » In South Korea, the AI Basic Act (effective from 2026) institutes a risk-tiered system: providers of “high-impact AI” will have to perform rigorous self-assessments and obtain government approval before deployment. Obligations will include risk-mitigation plans, transparency to users, human oversight, and documentation of AI systems. The aim here is to mandate responsible model development practices and disclosures for certain types of AI systems that pose risks to safety or fundamental human rights, instead of mandating AI models to be locally owned or developed.²¹²
- » While there is currently no specific law or regulation in Taiwan directly focusing on AI, the government has issued voluntary AI R&D guidelines, setting out principles and compliance requirements for model developers.²¹³
- **Malaysia, India, and Indonesia** are moving in the same direction, signalling requirements for transparency and safety while national strategies mature. These countries are converging towards an “innovation-friendly but safe” AI governance approach. While developing their national AI strategies and studying international models, these three countries currently prioritise having baseline safeguards in place.
 - » India’s regulatory focus on AI so far has been to lay the groundwork for safe AI innovation and is introducing targeted guidelines and principles. In 2026, India is moving towards a “techno-legal” approach that emphasises technical safeguards, oversight mechanisms, and lifecycle audits to ensure that AI systems are developed in a safe and people-centric way. India aims to mitigate harmful misinformation through transparency rather than broad ban.²¹⁴
 - » Indonesia has issued a non-binding AI Ethics Guidelines that encourage responsible use of AI in business. These guidelines spell out principles—inclusivity, transparency, accountability, data privacy etc.—for companies to follow when deploying AI.²¹⁵ Sectoral bodies are also taking action: the Financial Services Authority (OJK) for instance, released an ethical AI guideline for fintech companies to ensure risk controls.²¹⁶
 - » Malaysia has begun institutionalising AI governance through a phased approach. In 2024, the government launched the National Guidelines on AI Governance and Ethics, an initial voluntary framework setting out principles for responsible AI development and deployment, including fairness, reliability, transparency, privacy, risk assessment, and accountability. The government is now moving towards a more formal legislative framework: the Ministry of Digital is drafting Malaysia’s first AI Governance Bill, expected to be tabled in

209 Inoue, K., Kamata, C., “Japan’s emerging framework for responsible AI: legislation, guidelines and guidance”, International Bar Association, July 2025, accessed December 2025.

210 METI, “AI Guidelines for Business Ver 1.0 Compiled”, 2024, accessed December 2025.

211 Habuka, H., “Japan’s Approach to AI Regulation and Its Impact on the 2023 G7 Presidency”, Center for Strategic and International Studies, 2023, accessed December 2025.

212 Lee, S., “South Korea’s Evolving AI Regulations”, Stimson, 2025, accessed December 2025.

213 White & Case, “AI Watch: Global Regulatory tracker – Taiwan”, 2026, accessed January 2026.

214 Office of the Principal Scientific Adviser to the Government of India, “Strengthening AI governance through techno-legal framework”, 2026, accessed January 2026.

215 Rolindrawan, W. Y., Ismayudha, Q. P., “Artificial Intelligence Comparative Guide”, Mondaq, 2025, accessed December 2025.

216 Virgiany, M. Amatullah, N., “Ethical guidelines on use of artificial intelligence (AI) in Indonesia”, Hiswara Bunjamin and Tandjung, 2024, accessed December 2025.

2026, to address AI-related risks and strengthen governance across the AI lifecycle. This shift also extends to AI model development and use. Public statements by the Prime Minister and Ministry of Digital indicate that the proposed Bill is expected to address copyright and intellectual property issues linked to AI, including the use of copyrighted material in training data. This suggests a gradual move from voluntary guidance towards more binding oversight of AI systems, particularly where model development, content generation, and sensitive use cases create legal, ethical, or societal risks.^{217,218,219,220,221}

- **Lao PDR, Myanmar, and Cambodia** do not yet have fully formed AI strategies, regulations and laws, or guidelines regarding the development of AI models, apps, and services.

Broadly, across APJ, regulators are converging on governance that supports innovation while mitigating risk, rather than mandating local-only development.

USER FACING AI SOFTWARE AND APPLICATIONS

The AI application layer delivers ready-to-use AI services and software created by other organisations. This enables companies, especially smaller organisations, and consumers to benefit from AI without building it themselves. This allows organisations to focus their resources on developing new applications and driving further innovation. As enterprise tools increasingly embed AI features, adoption will often occur through everyday software. This occurs seamlessly, without users realising they are using AI. Across Asia Pacific and Japan, governments are introducing

frameworks to guide responsible AI use, with varying levels of formality and enforcement:

- **Singapore** sets clear expectations through its AI Governance Framework for Generative AI. This framework applies not just to developers but also to deployers and users, emphasising principles such as explainability, level of human involvement, accountability, and safety testing of AI in real-world use.²²² It also anticipates future sector-specific rules, such as those governing elections and deepfake content.²²³ In addition, the nation's Personal Data Protection Act (PDPA) already applies to AI-driven services that handle personal data, requiring consent and protection measures. Overall, Singapore's stance is to embrace AI applications, regardless of origin, but hold them to transparency and safety standards, potentially enforcing specific laws for sensitive domains.
- **Japan** favours risk-based governance that supports AI innovation while managing deployment risks, but their policy approaches differ. Japan continues to rely heavily on non-binding AI Business Guidelines, which encourage firms to identify AI risks and adopt voluntary countermeasures across the AI lifecycle.²²⁴
- **Taiwan's** framework is more closely anchored in its Artificial Intelligence Basic Act, the National Science and Technology Council's role as the central competent authority, and MODA's AI Risk Classification Framework, which is intended to support risk-based implementation by sectoral agencies. The emphasis is more on developing a common governance framework for public-sector and sector-specific AI deployment. The policy direction is to manage

217 Buza, M., van Mutius, S., "DPA Digital Digest: Malaysia", Digital Policy Alert, 2024, accessed December 2025.

218 Isamudin, D., "Malaysia's First AI Bill to Be Tabled by Mid-2026", New Straits Times, August 2025, accessed April 2026.

219 The National AI Office (NAIO), "Governance: aiming to promote the responsible and ethical use of AI", 2026, accessed April 2026.

220 Hakim, L., Mahari, H., "AI Governance Bill to hold developers, users liable for copyright breaches", New Straits Times, January 2026, accessed April 2026.

221 Jun, W., "Malaysia AI governance bill still taking shape", MalayMail, February 2026, accessed April 2026.

222 IMDA, "Singapore proposes framework to foster trusted Generative AI development", 2024, accessed December 2025.

223 Chia, O., "Temporary deepfake ban discussed as way to tackle AI falsehoods during Singapore election", The Straits Times, July 2024, accessed December 2025.

224 Ministry of Economy, Trade and Industry (METI) and Ministry of Internal Affairs and Communications (MIC), "AI Guidelines for Business Ver. 1.0", 2024, accessed December 2025.

AI risks through guidance, classification, and responsible-use frameworks, rather than mandating broad domestic ownership or excluding foreign AI services.^{225,226,227}

- **South Korea** takes a more prescriptive approach. The AI Basic Act introduces obligations for “high-risk” applications and calls for registries to enhance oversight. Outside of the AI Basic Act, South Korea’s other laws are increasingly focused on AI use. For instance, after surges in deepfake content and election disinformation, specific legal amendments banned certain malicious AI uses. A 2024 law criminalised the creation or viewing of deepfakes without consent, and the Public Official Election Act was amended to prohibit AI-generated election disinformation.²²⁸ The sovereignty goal here is to protect users from AI harms, and South Korea is willing to enforce more compliance more strictly to prevent misuses of AI applications.
- **Thailand** is drafting AI law principles that emphasise controls for high-risk applications and technical standards. High-risk applications refer to those that could significantly affect public health, safety, or rights, and will only be allowed subject to strict conformity assessments and oversight. Thailand’s Personal Data Protection Act (PDPA) also applies to all

AI-enabled applications, and mandates consent for personal data use, purpose limitations, and obligations on data controllers to protect user data. Together, these frameworks show Thailand’s multi-faceted approach of controlling AI application risks via dedicated AI laws (targeting the AI’s behaviour and technical robustness) and via data security laws (protecting information used by AI).

- **Philippines** provides guidance through the National Privacy Commission (NPC) Model Clauses, supporting lawful processing and cross-border data use in SaaS and AI applications, in line with the nation’s Data Privacy Act. The NPC also promoted a Data Management Framework (DMF) to guide organisations on proper data governance. This law cover areas like oversight, risk assessment, and continuous monitoring of data use.²²⁹ These measures support the Philippines’ *Digital Personal Data Protection* goals and explicitly encourage cloud and AI adoption by clarifying how to do so lawfully. In essence, the Philippines recognises the value of imported AI applications (e.g., enterprise AI software, cloud AI APIs) and seeks to make them accountable and privacy-compliant rather than excluding them.

225 Chambers and Partners, “[Data Protection & Privacy 2026](#)”, 2026, accessed April 2026.

226 Chambers and Partners, “[Cloud Computing 2025](#)”, 2025, accessed April 2026.

227 Taiwan Ministry of Digital Affairs, “[Public Sector AI Playbook](#)”, 2024, accessed April 2026.

228 Singh, N., “[South Korea criminalises explicit deepfake possession amid public outcry](#)”, Business Standard, 2024, accessed December 2025.

229 National Privacy Commission, “[NPC’s new initiative on ASEAN cross-border tools to boost PH digital competitiveness](#)”, 2022, accessed December 2025.

CONTACT

Global headquarters

Oxford Economics Ltd
60 St Aldates, Oxford,
OX1 1ST, UK

Tel: +44 (0)1865 268900

London

4 Millbank, London,
SW1P 3JA, UK

Tel: +44 (0)203 910 8000

Frankfurt

Marienstr. 15
60329 Frankfurt am Main
Germany

Tel: +49 69 96 758 658

New York

5 Hanover Square,
8th Floor, New York
NY 10004, USA

Tel: +1 (646) 786 1879

Singapore

6 Battery Road
#38-05

Singapore 049909

Tel: +65 6850 0110

Email:

mailbox@oxfordeconomics.com

Website:

www.oxfordeconomics.com

Further contact details:

[www.oxfordeconomics.com/
about-us/worldwide-offices](http://www.oxfordeconomics.com/about-us/worldwide-offices)

EUROPE, MIDDLE EAST AND AFRICA: OXFORD • LONDON • BELFAST • DUBLIN • FRANKFURT • PARIS • MILAN • STOCKHOLM • CAPE TOWN • DUBAI • **AMERICAS:** NEW YORK • PHILADELPHIA • BOSTON • CHICAGO • LOS ANGELES • TORONTO • MEXICO CITY • **ASIA PACIFIC:** SINGAPORE • HONG KONG • TOKYO • SYDNEY