

JANUARY 2026
HILDA BARASA
PEICHIN TAY
KEEGAN MCBRIDE
ALEXANDER IOSAD
JAKOB MÖKANDER



Sovereignty in the Age of AI: Strategic Choices, Structural Dependencies and the Long Game Ahead

Contents

4 Foreword

7 Executive Summary

11 Introduction: A New Architecture of Power

14 Understanding Sovereignty in the Age of AI

18 The Trade-Offs Shaping AI Sovereignty

27 A New Framework for Sovereign Posture:
Control, Steer, Depend

32 From Trade-Offs to Strategy: Navigating
Choices for Sovereignty in the Age of AI

40 Sovereignty in Practice: National Pathways
Through the AI Stack

51 Policy Levers for Expanding Agency in the Age of
AI

61 Conclusion: The Long Game of Strategic Resilience

63 Annex

64 Acknowledgements

Contributors: Helen Song, Koichi Tsukioka, Barbara Ubaldi, Rasmus Andersen, Johan Harvard, Guy Ward-Jackson, Tone Langengen, Devorah West, Marie Teo

Foreword

The creation of highly capable, general-purpose AI systems marks the beginning of a new era for our world. The pace of change is already extraordinary – and it is accelerating.

This is a revolution that leaders cannot opt out of. Those who move fast to deploy AI across their economies and institutions will gain a lasting advantage. Those who do not will see their ability to influence events – abroad and at home – progressively diminished.

This is why AI matters so much to the question of sovereignty today. The systems that will define the future require massive amounts of capital, talent and energy. They depend on global supply chains, and economies of scale mean that frontier AI development is likely to remain concentrated among a small number of actors.

Faced with this reality, many leaders are now asking what steps they should take to stay in control of their country’s future. The instinctive response is often to try to do everything at home – to build fully “sovereign” AI and treat reliance on partners as a threat.

This report argues that while this instinct is understandable, it is also wrong. Full self-sufficiency is too expensive, too slow and, for most countries, simply impossible. More importantly, it misrepresents what sovereignty really means in a digital, global and interconnected world.

Sovereignty should not mean independence from all others. Rather, it should be viewed as the ability to act strategically – with agency and choice – in a world that is irreversibly interdependent. In the context of AI, that means

being able to shape how systems are used in your country, having real options over infrastructure, models and partners, and retaining the flexibility to adapt as the technology evolves.

No state can dominate every layer of the AI stack. Leaders must make deliberate choices about where they want to build strength and influence. And by becoming indispensable in specific parts of the AI ecosystem – whether in data assets, specialised models, regulatory standards, energy capacity or talent pipelines – countries gain leverage across it, even if they do not control it all.

Crucially, sovereignty must not become an excuse for avoiding frontier AI. Failing to access and apply the best systems is itself one of the greatest threats to sovereignty today. Countries that cannot use these tools will become dependent on those that can – less able to defend themselves, drive economic growth or deliver public services effectively. Isolation and protectionism do not protect sovereignty; they weaken it.

Leaders should therefore resist the urge to own every model and data centre, and instead prioritise securing access to frontier capabilities, building domestic strengths where they matter most, and governing AI in line with national priorities and values.

This is also a question about the kind of state we want to build. When I meet with leaders, I often stress that governing in the age of AI requires reimagining the state: its capabilities, its partnership with industry and its ability to move at the speed of technological change. AI can make governments more effective and more strategic. But that will not happen by accident. It requires political focus, institutional capability and a strategy to integrate AI into the government machinery.

The value of this report lies in its practical guidance on how to make these choices. It provides a framework for assessing national posture across the AI stack, lessons from countries pursuing different paths and a set of concrete levers for expanding national agency over time.

Above all, it makes the case for confidence rather than fatalism. Countries that treat AI as a central pillar of their national purpose, deploying it widely and negotiating their place in the global ecosystem with clarity and ambition, will not see their sovereignty eroded. They will renew it for a new age.

Tony Blair

Executive Summary

Artificial intelligence is transforming the foundations of state power, reshaping economies, accelerating scientific discovery and recasting geopolitical relations. For political leaders willing to embrace it, AI is not simply another technological wave but a foundational governing infrastructure that will influence how states make decisions, deliver services and project strategic influence. Countries that fail to adopt and deploy AI at scale risk ceding their competitiveness and, ultimately, elements of their sovereignty to those who do.

In the age of AI, no country can claim complete self-sufficiency. Developing and deploying frontier AI requires enormous resources: billions of dollars in compute, data and engineering talent, alongside hyperscale data centres and cutting-edge semiconductors. These capabilities are overwhelmingly concentrated in the United States and China, which together control more than 90 per cent of global AI data-centre capacity. Most states will never be able to build or sustain frontier AI infrastructure on their own.

The dependencies this creates are increasingly seen as either strategic vulnerabilities or levers of geopolitical influence. Countries are left to confront a choice about how to remain competitive in the AI era: join the resource-intensive race to train the world's most advanced models or focus on deploying AI at scale, bolstering competitiveness through deployment rather than domestic frontier systems.

This landscape leaves governments confronting two defining questions: what is sovereign AI capability, and how can nations develop and exercise it? In response, calls for “AI sovereignty” are emerging, presented as an imperative to exercise exclusive control over frontier technology and its use, to build domestic capabilities and to enforce stricter regulations on foreign technology. Although motivated by legitimate economic and security concerns, this instinct reflects a narrow and ultimately counterproductive understanding of sovereignty: one that equates autonomy with full technological control and treats interdependence as a vulnerability to be eliminated.

A more realistic and grounded understanding of sovereignty that reflects the realities of interdependence rather than the illusion of isolation is therefore needed. Sovereignty in the age of AI is not a binary condition to be achieved or lost. It is fundamentally a question of agency and choice – the ability of a state to make deliberate, future-oriented decisions about how AI is integrated, governed and used in line with its national goals.

Sovereignty is shaped by how well countries configure and negotiate their position within an inherently interdependent technological system. This requires balancing a persistent trilemma: pursuing control by investing in domestic capability, accessing frontier capability through global systems, and ensuring coherence across regulatory, industrial, fiscal and diplomatic strategies. No state can maximise all three simultaneously. The task of modern statecraft is to manage these trade-offs within the layers of the AI stack in ways that preserve strategic autonomy and expand national agency over time.

Effective AI sovereignty therefore cannot be pursued through isolation. It must be deliberately negotiated. Governments will need to cultivate domestic strengths where they matter most, secure predictable access to frontier capabilities, design partnerships that preserve flexibility, and invest in the institutions and talent required to evaluate, govern and adapt AI systems. Above all, national competitiveness in the AI era will depend as much on deploying AI widely across the economy and public sector as on building frontier models themselves. For many countries, this path will deliver far greater returns than entering the resource-intensive race to train the world's most advanced systems.

The choices that decision-makers make today will have considerable, enduring consequences for their countries' futures.

This paper offers political leaders and key decision-makers a practical guide to shaping the strongest possible national position in this rapidly changing techno-geopolitical landscape. It demonstrates that AI sovereignty lies not in a futile pursuit of technological self-sufficiency but in cultivating informed,

strategic agency. It provides a new framework to help governments assess the trade-offs at each layer of the AI stack, illustrated with examples of how countries are tackling these choices today.

From this, we identify seven strategic levers through which governments can expand agency in the age of AI:

1. **Secure access to frontier AI models and compute.** Governments should prioritise accessing frontier AI capabilities. They need to formulate a clear view on what can be done using internationally derived models and sources of compute, and what must absolutely be done domestically and how much compute that would require.
2. **Accelerate AI adoption and diffusion across sectors.** Countries must ensure they are able to adopt and diffuse AI at scale. The value of AI will only be realised when it is widely adopted across the economy, which will bolster national prospects.
3. **Aggregate and signal national demand to shape the AI market.** Every country is competing for frontier access, and global providers and investors will prioritise countries that demonstrate strong, coordinated demand for AI.
4. **Treat interoperability as a core component of sovereignty.** Governments should prioritise building open, modular and interoperable systems to improve customisability, avoid lock-in and enhance resilience to better control their digital futures.
5. **Build and scale smaller, efficient and contextually relevant models.** Not every country needs to build frontier models, but the use of open-weight and small language models tailored to national needs can create significant value.
6. **Invest in talent and state capacity.** Leveraging the opportunities offered by AI requires a strong talent base, a workforce that can leverage AI, and high levels of state capacity to deploy and govern new technologies effectively.

7. **Align AI infrastructure with sustainable energy planning.** Countries must ensure that their energy systems are able to support domestic AI infrastructure now and in the future, as it will place increasing pressure on national power systems.

Together, these levers form a coherent agenda for strengthening national sovereignty in the AI era, one rooted in strategic positioning, deliberate interdependence and the effective governance of transformative technology.

01

Introduction: A New Architecture of Power

Technological shifts have always redefined the boundaries of state power. The ability to build, own and govern critical technologies has historically determined who created value, who set standards and who exercised sovereignty. Governments once directed much of the innovation that shaped national development, but today private firms increasingly operate at the technological frontier, moving faster than public institutions and legal frameworks can adapt. Public authority now often sits downstream from technical and commercial decisions taken elsewhere that can influence the functioning of economies, markets and even state institutions.

Many decision-makers are unprepared for the scale and speed of change now underway. Unlike previous technological waves, progress in AI unfolds in days and weeks, not years or budget cycles. AI has evolved from a specialised tool into a general-purpose capability embedded across sectors and the core operations of government, affecting not only how services are delivered but how decisions are made. An algorithm used in public administration, for example, can determine who receives services and when they receive them, how public resources are allocated and where enforcement is directed. This in turn shapes not only institutional behaviour but the everyday experiences and expectations citizens have of their government. The choices governments make today about their AI systems will determine whether they can expand capability in the future or find themselves locked into systems that are prohibitively and politically costly to change. In this sense, every AI-related decision becomes an act of statecraft, with implications that extend far beyond technology.

AI further magnifies interdependence across the digital economy. Advanced systems depend on scarce and highly concentrated resources including hyperscale compute, global cloud networks, semiconductors and highly skilled technical talent. These inputs are controlled by a small number of

firms and jurisdictions, and for many countries, especially those with limited fiscal space, energy supply or digital infrastructure, this concentration shapes the boundaries of what is possible.

The implications for sovereignty are profound. Economies of scale, high capital costs and strong network effects (where systems become more valuable and dominant as more users adopt them) mean that much of the emerging foundational AI infrastructure is likely to be developed and owned by a small number of countries and corporations. This centralisation offers clear benefits, such as enabling global access to cutting-edge technology at lower prices while ensuring strong technical security and reliability.

However, it also creates structural points of dependency and vulnerability. As AI becomes essential to the functioning of public administration and the wider economy, governments must ensure they have sufficient domestic capability and credible fallback options to maintain continuity of critical services, such as national security and health care, should external access to compute or data infrastructure be disrupted.

These technological dynamics sit within a shifting geopolitical landscape. At the AI frontier, the United States and China have emerged as the two poles of AI power, offering capability and partnership on contrasting terms. The US promotes a “trusted” AI ecosystem that is rooted in innovation, scale and private enterprise, but reinforced through an industrial strategy that combines subsidies with export controls to secure national security and strategic advantage.¹ As part of a new “promote” strategy, major US firms are expanding their global footprint through international investments, philanthropic initiatives and workforce initiatives.²

China advances a state-directed model that treats AI as a strategic public good while embedding ideological alignment into governance and promoting “openness” as a narrative to normalise its tightly controlled ecosystem.³ Through its Digital Silk Road initiative, China offers affordable, turnkey AI infrastructure that lowers adoption costs while deeply embedding Chinese technology and standards within national digital architectures.

This strategic alignment is already visible in practice. The UAE-US AI Acceleration Partnership, for example, provides the United Arab Emirates with secure access to US frontier compute by accepting export controls and limits on Chinese collaboration.⁴ Other states, such as Singapore, are experimenting with a more balanced strategy, engaging both Western and Chinese ecosystems to diversify risk and preserve market access.^{5,6} Across Africa, Asia and Latin America, governments face the same dilemma.

Decision-makers around the world now face two fundamental questions. First, what does sovereign AI capability truly mean in a world defined by interdependence and concentrated technical power, and how can it be built? Second, how can states harness AI in ways that bolster, rather than diminish, national sovereignty?

02

Understanding Sovereignty in the Age of AI

The earliest debates on digital sovereignty focused on data protection, privacy and economic dependence.⁷ In Europe, for example, concerns over growing reliance on US cloud companies led to new regulations designed to assert greater control over foreign cloud providers and strengthen domestic data governance.⁸ These debates, however, belonged to an era when sovereignty was primarily understood through the management of data flows and market power. In the age of AI, sovereignty debates have gained new urgency, but many governments still frame the issue through a narrow, capability-driven lens that misreads the structural realities of the AI ecosystem.⁹

AI now sits at the heart of national competitiveness, public-service delivery and economic transformation. It offers extraordinary potential to accelerate scientific discovery, boost productivity and widen access to essential services. At the same time, rapid advancements in AI systems, including the potential development of artificial general intelligence (AGI), have raised concerns about potential security risks, including in areas such as chemical and biological threats, and the possible displacement of significant parts of the workforce due to automation.¹⁰ These dual dynamics of risk and opportunity have made sovereignty a defining concern of the AI era.

As the race to develop and adopt AI intensifies, it has increasingly been interpreted as a contest for controlling both intelligence and the technology that will define the world's digital future. In response, many countries have begun pursuing "AI sovereignty" by purchasing vast amounts of new compute, investing in national energy capacity, supporting the creation of new localised AI models and launching programmes to help domestic AI firms.¹¹ The paradox is that even as demand for sovereign control grows louder, technology firms increasingly market "sovereignty as a service", offering localised clouds, compliance wrappers (regulatory, technical and institutional safeguards that adapt global AI systems to local laws and govern how they operate within a country) or hardware bundles that

simulate autonomy while deepening dependency.¹² These arrangements can create the appearance of control without delivering meaningful agency, reinforcing why sovereignty cannot be reduced to infrastructure localisation or the nominal ownership of digital assets.

As the Tony Blair Institute for Global Change (TBI) has previously argued in *Sovereignty, Security, Scale: A UK Strategy for AI Infrastructure*, countries will need some level of domestic compute capacity to ensure resilience for mission-critical areas, such as health care or national security. However, this requirement is often overstated in the pursuit of AI sovereignty. A baseline of sovereign capability matters, but replicating frontier-scale capability is neither feasible nor necessary for most countries.¹³ At the same time, states that fail to adopt AI and deploy it at scale will find themselves rapidly falling behind, not because they lack frontier models but because they fail to convert available capability into public value. Sovereignty in the AI era will depend as much on a country's ability to use AI effectively as on how much of the technology it builds itself.

Even for countries that could build their own isolated or sovereign AI ecosystems, the AI supply chain is inherently global: semiconductors are designed in one region and fabricated in another, cloud platforms operate across borders, models are trained on proprietary architectures, and technical talent circulates globally. Sovereignty must therefore be understood not as independence, but as the ability to act deliberately within an interdependent system.

Managing this interdependence, however, demands navigating trade-offs across three dimensions: the degree of control a state seeks over critical systems, the capability it needs to remain competitive and secure, and the degree of coherence it can achieve across its regulatory, industrial, diplomatic and fiscal strategies. In an environment defined by high levels of global dependencies across the entire supply chain, AI sovereignty should not be viewed as a binary choice i.e. sovereign or not. Instead, countries must work to shape domestic and international AI strategies that align with their strengths and weaknesses across the “AI stack” (shown in Figure 1).

Taken together, the layers of the AI stack illustrate the complexity of frontier AI ecosystems and why many countries will be unable to excel across the entire stack.

Policymakers must therefore evaluate their strategic configuration across the stack, determining where to build domestic strength, where partnerships can extend capacity and where managed dependencies are advantageous. For example, energy-rich countries may wish to leverage cheap electricity to attract the development of new data centres, and countries with strong talent pipelines and pro-innovation regulatory frameworks could become hubs for model development or AI services. By focusing on the application layer, many countries can also build contextually relevant use cases that deliver value while limiting the cost of supporting and sustaining domestic model training.

FIGURE 1

How layers in the AI stack influence sovereignty

AI stack layer	Why it matters for sovereignty
Compute infrastructure	Access to computational power determines which models a country can train, adapt or deploy, and whether it can maintain fallback capacity for critical systems.
Energy	Reliable, affordable and scalable energy is essential for running data centres and AI workloads. Energy constraints directly limit national AI capability and resilience.
Data	Training AI models requires significant amounts of data. High-quality, representative and machine-readable data shape how well AI reflects national contexts, languages and priorities, and who controls the associated value.
Models	Most of the leading frontier AI models are trained in a small number of countries, including both open- and closed-source models. Model choice influences performance, security and alignment with national norms.
Applications	AI creates value when it is deployed and diffused throughout society and the economy. Countries that scale deployment across sectors gain competitiveness and improve public services.
Talent and skills	The development and deployment of AI require the presence of a skilled and talented workforce to adapt, govern and effectively use AI.
Governance	Countries will have to make a choice on how they approach the governance of AI, both domestically and internationally. Regulatory frameworks, institutions and standards shape how AI is deployed and who sets the rules. Strong governance is essential for accountability, safety and international credibility.

Source: TBI analysis

Sovereignty in the age of AI is therefore a hybrid construct.¹⁴ It is a continuum of agency that is defined by a state's ability to make deliberate, future-oriented choices about how AI is integrated, governed and used in ways that protect public interests, create value, build domestic ecosystems, and preserve fallback capacity if external access is disrupted.

03

The Trade-Offs Shaping AI Sovereignty

Strengthening sovereignty in the age of AI will require governments to navigate a series of trade-offs across the layers of the AI stack. Decisions made at one layer of the stack may restrict or expand options available at another: compute capacity is constrained by energy supply, data access affects model relevance, and governance depends on the skills and institutions that uphold it. Collectively, these interactions define the decision space within which strategic choices about AI can practically be made.


Compute Infrastructure: Autonomy Versus Capability Versus Cost

Compute, the processing power provided by chips and measured in FLOPs (floating-point operations), underpins modern AI, yet it presents one of the most consequential trade-offs. The more autonomy a country seeks over its compute infrastructure, the more domestic investment it must make. The more it relies on external platforms, the more capability it can access but at the cost of control.

Frontier-scale compute is dominated by a handful of countries and firms. Taiwan manufactures nearly all advanced AI chips, while the US and its allies provide core funding, software and hardware.^{15,16} China is investing heavily to close this gap, but still lacks the scale and technological maturity of US-aligned supply chains.¹⁷ As a result, the US currently hosts around 75 per cent of the world's total AI compute capacity compared to 15 per cent in China and roughly 10 per cent distributed elsewhere, mostly in Europe.¹⁸ Only 32 countries worldwide host AI-specific data centres, leaving around 160 nations dependent on foreign infrastructure.¹⁹ TBI's *State of Compute Access* reports in [2023](#) and [2024](#) have found similar trends, highlighting the divergent availability of compute worldwide.

FIGURE 2

The US dominates global AI compute capacity, followed by China



Source: Epoch AI²⁰

This scale imbalance means training frontier models is not feasible for most countries. Yet domestic compute remains essential for resilience in critical domains such as health care, core public administration and national security. Decision-makers will have to decide how much of that compute should be held domestically to ensure autonomy and continuity, and how much can be reliably accessed via foreign cloud providers. The former will ensure higher levels of autonomy and security at a higher cost, while the latter offers scale and speed, but with less control and greater exposure to external conditions and shifting geopolitics.

Energy: Scale Versus Cost Versus Control

Training and hosting large AI models consumes vast amounts of electricity, and demand is skyrocketing. The International Energy Agency (IEA) projects that data-centre consumption will more than double by 2030, reaching around 3 per cent of the world's total electricity use. AI is driving a significant

share of this demand growth.²¹ Goldman Sachs similarly estimates that approximately 60 per cent of the energy needed for powering the growing data-centre demand will require new generation capacity.²²

As countries position themselves within the emerging AI economy, the cost of electricity, available generation capacity and national energy mix will increasingly determine their competitiveness. Crucially, countries enter this new landscape from profoundly different starting points. Some, such as Norway, the UAE, Saudi Arabia and France, possess deep structural advantages including abundant hydropower or solar potential, extensive nuclear capacity, strong grids and access to cheap capital.^{23,24,25} These conditions allow them to scale energy supply for AI rapidly and cost-effectively, making them attractive destinations for global data centres and compute clusters.

Other countries face harder constraints: limited domestic resources, weak grid infrastructure, high financing costs, or political and environmental barriers to expanding firm generation (a continuous, reliable supply of electricity). These inherited differences shape not only how much energy countries can produce, but also how quickly, affordably and securely they can expand supply in response to rapidly growing AI demand. Energy-constrained nations must often rely on regional grids, long-term imports, innovative financing models or foreign investment partnerships to secure the power base AI requires.

Decision-makers must therefore determine not only how much energy they possess, but how predictably and affordably they can expand energy supply, and on whose terms that expansion occurs.

Data: Representation Versus Openness Versus Sovereignty

Training frontier AI systems requires vast amounts of high-quality, machine-readable data, but the global corpus is heavily skewed towards English.²⁶ For example, in Common Crawl – an open repository of web-crawl data –

English makes up nearly half of all content, while major languages such as Arabic represent less than 1 per cent and many other languages have marginal representation.²⁷ In addition, there is a finite amount of available data for AI training, making access to high-quality data a major bottleneck to AI progress.²⁸ These imbalances mean that many national contexts, cultures and languages are structurally underrepresented in the data sets from which global models learn.

As the global supply of high-quality data tightens, countries are increasingly treating their data sets as sovereign assets. Rich, representative data can be used to fine-tune foreign models, build specialised applications, and accelerate scientific and economic innovation. The trade-off, however, lies between control and utility: tightly closed data sets may enhance sovereignty in the short term, but this quickly becomes self-limiting when their utility diminishes as their user base shrinks. Open or strategically shared data sets maximise innovation and interoperability but reduce exclusivity.

AI developers are already shifting towards more diverse data sources, including proprietary or specialised scientific data sets, and extending beyond text to richer modalities such as video and sensor data, and synthetic data designed to fill data gaps. In parallel, countries are expanding efforts to digitalise public archives, build national language data sets and formalise public-private data partnerships that preserve trust while enabling innovation.²⁹


The question facing policymakers is not whether these data sets should be made open or closed, but how they can be governed in ways that ensure representation, preserve trust and unlock economic and scientific value, without isolating domestic ecosystems or ceding strategic agency. The appropriate approach will depend on the nature of the data set and a country's specific goals and objectives.

Models: Capability Versus Control Versus Alignment

AI models sit at the centre of national capability. Since the public release of ChatGPT-3.5 in 2022, the predominant AI paradigm has been one dominated by large language models (LLMs). Training a state-of-the-art LLM, however, now requires billions of dollars in compute, data and engineering talent, placing such efforts far beyond the reach of most countries. Recent examples, such as the new supercomputer “Colossus” built by xAI, illustrate the magnitude of the investment required.³⁰ The graphics processing unit (GPU) cost alone for Colossus 2 is likely to exceed the \$6 billion raised in its Series C funding in 2024.^{31,32} Over the next five years, trillions of dollars will be invested globally in frontier model development, with the US and China dominating both the capability and the capital.³³ An overview of the current distribution of the most capable models is shown in Figure 3.

FIGURE 3

Notable AI models are concentrated in the US, followed by China



Source: Stanford Institute for Human-Centered AI³⁴

Most countries will remain dependent on the US or China for their access to frontier AI models and foundational architectures. This dependence on foreign AI models carries clear risks: opacity around training, decision-

making and model evolution, misalignment with national norms and limited auditability. In sensitive domains, it creates potential security vulnerabilities, such as “sleeper agents” or models that are more likely to suggest insecure code to geopolitical adversaries.^{35,36} Accessing highly capable models without oversight or fallback options can create long-term dependencies.

In response, many countries are experimenting with open-weight and adaptable models, whose trained parameters are publicly available and can be run and fine-tuned locally, such as Mistral in France and Falcon in the UAE. These models allow countries to customise for local languages, regulatory environments and domain-specific needs. Countries unable to support their own open-weight or sovereign model initiatives can still assert meaningful agency by distilling, fine-tuning or layering capabilities on top of larger proprietary systems.

Ultimately, the key question here is not whether a country owns a model, but how it governs, adapts and applies the models it relies on.

Applications: Build Versus Buy Versus Hybrid

It is applications, rather than AI models themselves, that will create economic and public value. They determine whether AI improves productivity and accelerates economic growth or whether a country falls behind others that adopt and deploy the technology more quickly. As our report [*Governing in the Age of AI: A New Model to Transform the State*](#) emphasises, the transformative impact of AI on the state materialises only when it is embedded into real systems of delivery, decision-making and administration.

Governments therefore face a strategic trade-off between building applications domestically, buying them externally or pursuing hybrid approaches. Domestic development allows countries to embed national norms, tailor systems to local institutions and retain greater control over sensitive sectors such as health care, finance and public administration. Buying applications offers speed and lower cost, but risks reinforcing dependency, reducing future flexibility and embedding foreign standards.

Hybrid strategies, combining foreign-developed tools with local integration, domain-specific customisation and domestically governed data, often provide the most realistic path for countries seeking both capability and agency.

Sovereignty in applications ultimately rests on who captures value and who sets direction. Countries that already have strong sectoral institutions, including health-care systems, sophisticated financial markets, advanced manufacturing bases and research ecosystems, are uniquely positioned to create applications that competitors cannot easily replicate. Those that fail to adopt AI applications at scale risk losing competitiveness even if they possess strong models. Applications are where strategic intent becomes state capability, converting access to AI into national advantage.

Talent and Skills: Speed Versus Depth Versus Retention

The emerging AI economy will demand new skills from the workforce. A country's ability to design, deploy and govern AI systems will depend on its ability to cultivate, attract and retain talent. Building this foundation will require sustained investment in educational institutions and partnerships with the private sector, combined with robust research and innovation systems that link academic excellence with practical deployment.

As global demand for AI skills accelerates, and with two-thirds of employers planning to hire workers with AI expertise and 40 per cent expecting to automate some roles, countries must race not only to expand technical training, but to ensure their public institutions, firms and research ecosystems have the expertise needed to adopt AI at scale.^{37,38}

The trade-off in talent strategy lies between scaling quickly, building deep expertise and retaining skilled workers in a fiercely competitive global market. Kazakhstan's partnerships with foreign universities, Estonia's

integration of AI into the core of its education system and the UAE's nationwide use of generative tools such as ChatGPT each illustrate different models of translating AI skills ambition into capability.^{39,40,41}

Governments must rethink how education, research and labour systems interact. The challenge is not only to produce more technical specialists, but to ensure that institutions, from schools to universities to research centres, can adapt fast enough to match the pace of technological change. As TBI has previously set out, education systems must evolve to prepare citizens for a world in which AI is both a tool and a competitor, while research ecosystems need the autonomy, funding and continuity to anchor national expertise.⁴² Doing both can create a talent ecosystem capable of sustaining and governing AI over the long term.

Governance: Innovation Versus Assurance Versus Influence

Governance determines how AI is developed, deployed and trusted within a country, and how effectively a state can shape the external environment in which it competes. Domestic governance frameworks influence adoption, innovation, safety and accountability, while international governance shapes the norms, standards and market conditions that govern access to frontier systems. Together these dimensions determine how much agency a state retains as AI becomes embedded across its economy and public institutions.

The strategic trade-off in governance lies between flexibility and assurance, domestic control and global interoperability, and innovation and regulation. In previous research TBI has argued that decision-makers will have to craft governance strategies that support innovation and AI adoption, and provide guardrails and assurance to their domestic markets.⁴³ Overly restrictive regulation risks slowing innovation and deterring investment. Too little oversight undermines trust and leaves systems vulnerable to misuse. Governance also has geopolitical implications: countries whose regulatory

choices diverge sharply from major AI powers may find themselves constrained in accessing frontier capabilities or excluded from emerging global standards.

Getting this right will require policymakers to build credible institutions that can evolve with technology. This requires regulatory capacity, technical expertise, stable legal frameworks and active participation in international norm-setting processes. States that balance innovation with assurance and align governance with industrial, diplomatic and societal priorities will be best positioned to shape the terms of their engagement with global AI ecosystems, rather than simply absorbing rules and architectures built elsewhere.

04

A New Framework for Sovereign Posture: Control, Steer, Depend

As the trade-offs show, each AI-related decision a country makes strengthens some dimensions of sovereignty while constraining others. These tensions cannot be fully eliminated, only managed. The scale, speed and complexity of AI development mean that each state must navigate these trade-offs and make strategic choices about which capabilities to hold domestically, where to shape markets and norms, and where to rely deliberately on others.

To illustrate these variations in practice, this paper proposes a new framework for understanding how governments exercise agency in the AI era: **Control, Steer, Depend (CSD)**. This new taxonomy captures how governments exercise agency within an interdependent AI ecosystem, differentiating between postures of direct control, influence and managed dependence that, collectively, define the full spectrum of AI sovereignty. Importantly, countries rarely adopt a single posture across all layers of the AI stack. Instead, sovereignty emerges from their unique configuration of postures across AI systems.

The framework delineates three core postures:

Control: Direct Command Over Strategic Systems

Governments adopt a **Control** posture when they judge certain AI systems as too critical to outsource. This strategic orientation is often characterised by ensuring direct ownership, legal authority or exclusive operational command over AI infrastructure, data sets or governance mechanisms.

Control provides the strongest form of insulation against external shocks, but it is also capital-intensive and operationally demanding. More importantly, control is not synonymous with nationalisation or autarky. Often, it means ensuring that critical systems cannot be unilaterally withdrawn by foreign providers, or that governments retain domestic jurisdiction over data.

Yet overreaching for control can create inefficiency, duplication and isolation. Japan's investment in its FugakuNEXT supercomputer or France's trusted cloud initiative under European Union law illustrate control strategies for compute and data respectively.^{44,45}

Steer: Shaping Outcomes Without Full Ownership

A **Steer** posture allows states to influence AI outcomes without owning or directly controlling the underlying system. Governments steer through regulation, procurement, partnerships and the development of standards that help shape markets and norms. Many countries also amplify their steering capacity by collaborating through regional organisations, enabling them to pool leverage and coordinate positions.

Steering is often the most practical approach. However, its effectiveness will significantly depend on a state's institutional credibility, market power and geopolitical influence. Without the ability to enforce rules or shape demand, steering risks becoming symbolic rather than strategic.

Depend: Managing Reliance on External Actors

A **Depend** posture entails reliance on foreign providers for core elements of the AI stack. For many governments today, dependence remains the most practical way to access frontier capabilities at speed and scale, and at a fraction of the cost of developing them domestically. Dependence is not inherently negative, nor does it automatically make a country worse off. Instead, dependence should be shaped and managed strategically, ensuring it enhances sovereignty and agency rather than undermining them.

This posture can be further divided into **maker dependence** and **taker dependence**. The former is active and negotiated, with governments co-developing systems or securing technology transfer. Brazil's joint AI research labs with China exemplify this approach by leveraging external expertise while fostering domestic research ecosystems.⁴⁶ India's semiconductor strategy similarly blends dependence with negotiated co-production and

technology transfer, aiming to move from full reliance toward onshoring the semiconductor supply chain.⁴⁷ In contrast, the latter is passive, with systems adopted wholesale, with little to no domestic capacity to modify, audit or govern them. This creates a risk of vendor lock-in especially where long-term contracts govern mission-critical data and services.

Operationalising CSD Across the AI Stack

Applied across the seven core layers of the AI stack, the CSD framework enables a granular, actionable analysis of where states hold influence, where they depend on others and how they can improve resilience. Every layer requires customised responses, not only in what is built or adopted but also in the leverage and resilience it incorporates. Figure 4 illustrates this application and offers a diagnostic lens for countries to evaluate their postures across the stack and identify where interventions can most effectively expand national agency.

FIGURE 4

Operational definitions of CSD across the AI stack

Layer	Control	Steer	Depend - maker	Depend - taker
Compute infrastructure	The state owns or has jurisdiction over domestic compute infrastructure, with full fallback capacity (e.g., public clouds, national data centres, local high-performance computing (HPC).	The state uses strategic procurement, investment incentives or public-private partnerships to shape infrastructure, with partial fallback capacity.	The state co-develops or adapts foreign infrastructure with limited influence; fallback capacity is weak or externally reliant.	The state relies on foreign-owned infrastructure with no influence or fallback capacity; loss of access would disrupt AI operations.
Energy	The state has domestic energy generation and regulation policies aligned to AI infrastructure needs.	The state shapes energy strategies through incentives or partnerships, with some reliance but strategic control over cost and direction.	The state co-develops or accesses energy infrastructure via regional/foreign agreements without fallback or pricing autonomy.	The state fully relies on imported energy or infrastructure with no control over source, pricing or reliability.
Data	The state governs, stores and uses domestic data with high-quality data sets, strong institutions and legal enforcement. Data flows are nationally directed.	The state shapes data access via localisation, interoperability or strategic data-sharing agreements. Partial control exists but not full sovereignty.	The state adapts or negotiates access to external platforms or shared systems, but lacks domestic infrastructure or governance to be self-reliant.	The state depends on foreign platforms for data storage, access or governance with no influence or fallback.
Models	The state (or domestic actors) can fully develop, fine-tune and deploy AI models with local compute and human capacity. Local models exist as fallback.	The state has influence over models (e.g., co-development, local fine-tuning, regulatory input) and partial fallback options via strategic partnerships.	The state customises or co-develops with foreign actors but lacks domestic capacity to operate models independently.	The state uses foreign models "as-is" via APIs or licenses, with no ability to govern, audit or substitute.
Applications	The state (or local actors) builds, deploys and controls AI tools for public use, ensuring adaptability, security and relevance to national needs.	The state defines standards and sets terms for application use through audits, procurement rules or modular development partnerships.	The state co-develops applications with external actors, providing input but lacking the capacity to build or maintain them independently.	The state adopts off-the-shelf foreign applications with limited adaptation or control.
Talent and skills	The state has strong local education, research and professional ecosystems for AI. It retains and grows a domestic talent pipeline.	The state actively shapes talent strategies (e.g., incentives, diaspora, joint R&D). Some gaps exist, but national priorities are evident.	The state engages in talent partnerships (e.g., joint universities, scholarships) but lacks robust domestic ecosystems or retention capacity.	The state depends on external pipelines (talent trained in foreign educational and skills systems, initiatives led by multinational tech firms or donors, and imported human capital) without a strategy for local capability development.
Governance	The state has robust domestic regulatory institutions, AI strategies and enforcement mechanisms. It sets and implements national norms and safeguards.	The state participates actively in international norm-setting or creates adaptive national policy frameworks with partial enforcement and influence.	The state is involved in multilateral or donor-led rulemaking with limited agenda-setting power or local enforcement.	The state adopts foreign rules or standards without tailoring or influence, lacking institutional capacity to govern independently.

Source: TBI analysis

Control, Steer and Depend are not rungs on a ladder, nor a linear progression from “less sovereign” to “more sovereign”. They are strategic postures that vary across the AI stack and evolve over time, enabling states to expand agency and build resilience. A state’s ability to adopt any of these postures, however, is shaped by its institutional capacity, credibility, alliances and structural constraints. Sovereignty therefore emerges from how effectively a country configures and progressively expands the set of postures available to it, using them to strengthen bargaining power, manage dependence and negotiate interdependence on its own terms.

05

From Trade-Offs to Strategy: Navigating Choices for Sovereignty in the Age of AI

AI sovereignty is determined by the strategic choices governments make under real constraints including fiscal realities, technological capacity and geopolitical trade-offs. The following matrix translates these principles into action. It helps policymakers identify the key questions to consider at each layer of the AI stack, the spectrum of available options and the conditions under which each becomes viable. Rather than prescribing a single path, it maps the decision space within which countries can exercise agency.

FIGURE 5

Key questions and options: Compute infrastructure

COMPUTE INFRASTRUCTURE

Trade-offs	Critical questions	Options	Feasibility and risks
<ul style="list-style-type: none"> Build versus buy compute capacity Centralised versus distributed compute infrastructure Chip resilience versus chip sovereignty 	<ul style="list-style-type: none"> How much compute must be domestic for continuity? What mix of high-performance computing (HPC), cloud and sovereign inference (the deployment of trained AI models to produce real-time outputs) do we need? Should we pursue chip resilience, and through which levers? Do we have steady demand to keep national capacity at or greater than 60–70 per cent utilised? Is regional compute pooling credible given governance, latency and export-control risks? Do we have institutional capacity to govern contracts and enforce safeguards? 	<ul style="list-style-type: none"> External cloud with safeguards: Rely on commercial cloud providers while preserving exit, portability and control through audit rights, local key management (encryption keys are held and managed domestically), multiple cloud design and flexible access to GPU markets; schedule heavy jobs in clean/low-cost windows Build national capacity: Develop government-controlled/dedicated national compute (for example via shared data centres or national HPC) with sovereign key management, priority access for research and small and medium-sized enterprises (SMEs), and siting aligned to energy availability Regional hub: Pool demand through regional-economic-community-level GPU clusters anchored to firm clean baseload power, with shared governance, clear Service Level Agreements (SLAs) and latency-aware workload placement to minimise lag 	<ul style="list-style-type: none"> External cloud: High feasibility, fast deployment <i>Risk:</i> Creeping lock-in if safeguards are weak or untested National capacity: Viable with clear demand and fiscal space <i>Risk:</i> Under-utilisation, talent/operational gaps, higher total cost of ownership Regional hub: Best for small/mid-sized states with shared energy baseload <i>Risk:</i> Dependency on cross-border trust, exposure to export controls, latency issues

Source: TBI analysis

FIGURE 6

Key questions and options: Energy

ENERGY			
Trade-offs	Critical questions	Options	Feasibility and risks
<ul style="list-style-type: none"> Renewable versus firm low-carbon power Domestic generation versus imported power Grid capacity for AI versus domestic affordability for households/industry 	<ul style="list-style-type: none"> Can renewable and low-carbon sources meet 24/7 AI energy needs, and at what cost given local geography and grid strength? How can AI-driven power demand be prevented from distorting domestic supply for households and industry? Should compute be co-sited with renewables and/or new firm-capacity sources (for example nuclear, gas, hydro, geothermal) and under what conditions? To what degree should countries rely on regional energy markets versus building national capacity for AI-related demand? Is there sufficient public and political acceptance for large-scale energy projects? What geopolitical dependencies or leverage does each energy technology create? 	<ul style="list-style-type: none"> Optimise existing supply: Manage AI demand through improvements to data-centre efficiency, flexible scheduling, demand response and off-peak or renewable-rich operating windows Co-site and upgrade: Align compute with renewable or low-carbon generation through targeted grid upgrades, power-purchase agreements and proximity to hydro, geothermal or solar resources Add firm clean capacity: Invest in dedicated firm power such as hydro, geothermal, nuclear or renewables with storage to support sustained AI/HPC demand; re-use waste heat Pool regionally: Meet AI-related energy demand through cross-border power pools, green corridors and long-term off-take guarantees linked to shared compute infrastructure 	<ul style="list-style-type: none"> Optimise existing supply: Quick/low-cost <i>Risk:</i> Limited headroom for future AI scale Co-site: Best for resource-rich areas <i>Risk:</i> Financing gaps, permitting, local opposition New capacity: Countries with abundant scalable energy resources and access to affordable capital <i>Risk:</i> Financing gaps, stranded assets if demand does not materialise, limited social buy-in from the public Regional: Best for small/mid-sized states <i>Risk:</i> Dependency on cross-border trust

Source: TBI analysis

FIGURE 7

Key questions and options: Data

DATA			
Trade-offs	Critical questions	Options	Feasibility and risks
<ul style="list-style-type: none"> • Fragmentation versus whole-of-government interoperability • Openness versus sovereignty and trust • Privacy versus analytical value • National control versus cross-border scale 	<ul style="list-style-type: none"> • Which data sets must remain sovereign (ID, health, security)? • Where is localisation essential and where does openness unlock more value? • Are ministries coordinating or operating in silos, weakening data value? • Can we align cross-border data rules to enable regional scale? • Does consent/portability exist by design? 	<ul style="list-style-type: none"> • Open and interoperable: Maximise value through API-first government systems, common schemas, shared data catalogues and open non-sensitive data sets • Controlled and localised: Retain national control over sensitive data via consent by design, trusted research environments, national data exchanges/markets and role-based access • Cross-border sharing: Enable regional scale through data commons, federated analytics (with data remaining in place), model-to-data approaches, and shared certification and assurance frameworks 	<ul style="list-style-type: none"> • Open and interoperable: Quick ecosystem lift <i>Risk:</i> Weak privacy or uneven enforcement • Controlled/localised: Needed for sensitive domains <i>Risk:</i> Siloing, slower innovation, duplication • Cross-border: Best for small/mid-sized states and representativeness <i>Risk:</i> Misaligned laws across regions, political trust gaps, need for shared assurance rules

Source: TBI analysis

FIGURE 8

Key questions and options: Models

MODELS			
Trade-offs	Critical questions	Options	Feasibility and risks
<ul style="list-style-type: none"> • Cultural fit versus cost/complexity • Open-weight models versus proprietary commercial models • Adaptation speed versus sovereign control • Single-vendor simplicity versus gateway diversity 	<ul style="list-style-type: none"> • Are adapted foreign models sufficient for our language and domain needs? • Which data sets require sovereign or semi-sovereign models? • Can we afford to build, maintain and continuously update domestic models? • Do we have the ability to audit and evaluate external models? • Would regional pooling of data/model development reduce cost and increase representativeness? 	<ul style="list-style-type: none"> • Adapt foreign models: Fine-tune external models for local language and domains, supported by transparency requirements, evaluation access and human-in-the-loop oversight for high-risk use cases • Build domestically: Develop national or sector-specific models where sovereignty or sensitivity demands it; potentially release open-weight models and support them with national safety labs and red-teaming/evaluation capacity • Build regionally: Pool data, compute and expertise to develop shared regional models with common safety, evaluation and governance arrangements 	<ul style="list-style-type: none"> • Adapt: Affordable, fast <i>Risk:</i> Dependence on external roadmaps, vendor licensing changes, limited fallback rights • Domestic: Feasible only with capital, compute and talent <i>Risk:</i> High operational expenditure, performance-parity challenges, long-term support burden • Regional: Realistic for mid-sized states <i>Risk:</i> IP/data-sharing disputes that can be solved with clear governance

Source: TBI analysis

FIGURE 9

Key questions and options: Applications

APPLICATIONS			
Trade-offs	Critical questions	Options	Feasibility and risks
<ul style="list-style-type: none"> Speed of rollout versus long-term flexibility Domestic innovation versus turnkey imports Vendor lock-in risk versus modular interoperability 	<ul style="list-style-type: none"> Do we need quick service wins or sustainable modular systems? Are procurement rules preventing lock-in by enforcing open standards and portability? Can we sustain or exit critical applications if vendor terms change? Could shared modules reduce costs and improve scale? 	<ul style="list-style-type: none"> Turnkey adoption: Deploy off-the-shelf applications rapidly through commercial vendors, supported by strict SLAs and clear exit and portability provisions Modular build: Develop interoperable applications using open standards in procurement, reference architectures and reusable components such as forms, enabling local ecosystems and long-term flexibility Shared modules: Co-develop cross-border digital building blocks (for example ID, payment rails or registries) to reduce costs and achieve scale 	<ul style="list-style-type: none"> Turnkey adoption: Best for urgent service delivery <i>Risk:</i> Vendor lock-in and high exit/portability costs later Modular build: Best where local ecosystem exists <i>Risk:</i> Slower rollout, higher upfront integration costs Shared modules: Best for small states seeking economies of scale <i>Risk:</i> Alignment delays, cross-border policy friction

Source: TBI analysis

FIGURE 10

Key questions and options: Talent and skills

TALENT AND SKILLS			
Trade-offs	Critical Questions	Options	Feasibility and risks
<ul style="list-style-type: none"> • Domestic pipelines versus overseas training • Public-private partnerships versus over-reliance on vendors • Long-term academic reform versus rapid upskilling • Retention versus global mobility (brain drain) 	<ul style="list-style-type: none"> • How do we reduce dependence on foreign expertise? • Do our institutions have absorptive capacity for trained talent? • Can we retain talent against global salaries? • Should we create mobility schemes to stabilise supply? 	<ul style="list-style-type: none"> • Baseline AI literacy: Establish public-sector and economy-wide AI literacy programmes as foundational capability • Embedded specialists: Rapidly build state capability through secondments and fellowships as well as targeted pay flexibility for critical roles • Skills and talent pipelines: Invest in long-term talent through research hubs, academic reforms and collaboration with leading global institutions • Knowledge transfer between industry and academia: Embed skills through structured collaboration between industry, academia and government • Regional mobility: Stabilise talent supply via regional centres of excellence, skilled talent visa schemes and mutual recognition of credentials 	<ul style="list-style-type: none"> • Baseline literacy: Universally feasible <i>Risk:</i> Low engagement without incentives, inconsistent standards • Embedded specialists: Fast capability gains; knowledge transfer <i>Risk:</i> Low retention without career paths • Pipelines: Needed for long-term sovereignty <i>Risk:</i> Slow returns on investment, poor domestic absorption leading to brain drain • Knowledge transfer: Boosts skills relevance <i>Risk:</i> Misalignment with public interest if governance is weak • Regional mobility: Stabilises skills supply <i>Risk:</i> Requires coordinated funding and political trust

Source: TBI analysis

FIGURE 11

Key questions and options: Governance

GOVERNANCE			
Trade-offs	Critical questions	Options	Feasibility and risks
<ul style="list-style-type: none"> • Domestic priorities versus international alignment • Rule-taker or rule-shaper • Regulatory ambition versus enforcement capacity • Reactive versus anticipatory governance 	<ul style="list-style-type: none"> • How do we balance global norms with domestic flexibility? • Where must we align with blocs or countries like the EU or US for market access, where must we diverge? • Do we have the capacity to enforce the rules we set? • Can a regional position amplify our positions? • Should we regulate proactively or only after incidents? 	<ul style="list-style-type: none"> • Align and localise: Adopt core global AI norms while tailoring implementation to national priorities and institutional capacity • Targeted sovereignty: Apply stronger rules and controls in high-risk or strategic domains while maintaining broad interoperability for trade and innovation • Shape selectively: Lead norm-setting where national strengths exist, while adopting global standards for efficiency and market access in other areas • Regional alignment: Amplify influence through mutual recognition, shared audits and regional governance mechanisms 	<ul style="list-style-type: none"> • Align and localise: Quickly achievable with some adaption <i>Risk:</i> Weak enforcement, shallow adaptation • Targeted sovereignty: Protects key domains <i>Risk:</i> Market access frictions, regulatory burden • Shape selectively: Focuses limited capacity <i>Risk:</i> Country's governance efforts spread too thin: passive rule-taker • Regional alignment: Amplifies voice through regional blocs or global initiatives <i>Risk:</i> Countries only agree on lowest-common-denominator standards

Source: TBI analysis

06

Sovereignty in Practice: National Pathways Through the AI Stack

The following case studies (Japan, the UAE, Kenya, France, India and Brazil) illustrate that there is no single path to AI sovereignty. Instead, states exercise sovereignty across the AI stack through distinct configurations of control, steer and managed dependence that reflect their political economies, resource endowments, institutional strengths and strategic objectives. Each country covered in Figure 12 has followed its own unique pathway to maximise the opportunity offered by AI, while blending domestic capacity building with carefully structured partnerships.

FIGURE 12

A comparative overview of CSD across the AI stack

	Compute infrastructure	Energy	Data	Models	Applications	Talent and skills	Governance
JAPAN	Steer	Depend - taker	Steer	Steer (niche Control)	Steer	Depend - maker	Steer
UAE	Steer	Control	Control	Control	Steer	Depend - maker	Steer
KENYA	Depend - taker	Depend - maker	Steer	Steer	Depend - taker	Depend - taker (in transition to Depend - maker)	Steer
FRANCE	Steer	Control	Control	Steer	Steer	Steer	Steer
INDIA	Steer	Depend - maker	Control	Steer	Steer	Control	Steer (in transition to Control)
BRAZIL	Steer	Control	Control	Steer	Depend - maker	Steer	Control

Source: TBI analysis

Sovereignty in AI is not secured through autarky or full-stack dominance, but through a state's ability to diversify options, manage its interdependencies proactively and align technological choices with its national purpose.

Distinctive sovereignty pathways appear across these case studies showing how states translate AI ambition into architecture:

- First, advanced economies such as Japan and France assert sovereignty through fallback capacity and governance credibility. Both invest in domestic compute and modelling ecosystems while embedding themselves within allied supply chains. Their strategies prioritise building resilience via fallback capacities, ensuring national continuity if foreign access is disrupted, and exercising regulatory influence through global bodies such as the EU and G7.
- Second, capital-rich countries such as the UAE pursue sovereignty through state-led investment and ownership, building direct control over data, model development and energy, while striking high-value partnerships with frontier actors on favourable terms. This model uses capital and centralised coordination to secure proximity to the frontier, ensuring that global capability is anchored within national jurisdiction and aligned with domestic authority.
- Third, large developing economies such as India and Brazil exercise sovereignty through steering power, using their sizeable digital markets, growing technical talent bases and assertive state direction to shape AI standards, incentivise domestic model ecosystems and attract strategic investment. India's leadership in AI talent and digital public infrastructure, and Brazil's combination of renewable-energy strength and regulatory enforcement, illustrate how states can steer global value chains without total self-sufficiency.
- Fourth, developing digital economies such as Kenya advance sovereignty through negotiated dependence, leveraging alliances with major technology actors to build and accelerate foundational capability while embedding local governance, data protection and national priorities. This approach shows how states with limited domestic compute or talent can still expand agency by shaping the terms on which external capacity is deployed, rather than passively absorbing it.

What sets states apart in the age of AI is not the lack of interdependence but their ability to manage it. Countries that are best able to build selectively, steer where influence is possible and structure dependence so that foreign capacity supports rather than hinders their national goals will bolster their sovereignty.

How Six Countries Configure Control, Steer, Depend to Strengthen Strategic Agency

The case studies that follow provide a detailed overview of how these sovereignty pathways are constructed in practice. Using the CSD framework introduced in earlier chapters (with full findings [set out in the Annex](#)), they trace how governments combine and recombine these levers across the AI stack to build sovereign capability. The case studies highlight distinct pathways: selective fallback capability in Japan, frontier acceleration in the UAE, negotiated interdependence in Kenya, regulatory power in France, digital-public-infrastructure-led scaling in India and state-guided co-creation in Brazil. Together, they demonstrate that sovereignty without full independence is both feasible and strategically advantageous, reflecting a field of viable strategies shaped by national constraints and comparative advantages.

JAPAN: RESILIENCE-FIRST APPROACH THROUGH SELECTIVE INVESTMENT AND PARALLEL CAPABILITY STRENGTHENING

Japan's approach to AI sovereignty centres on targeted investments in domestic capacity building while remaining deeply integrated into global supply chains. Rather than pursuing full-stack autonomy, Japan concentrates on areas of the AI stack where it can meaningfully differentiate itself, such as by developing fallback compute capability (Fugaku), building culturally aligned open-weight models (Fugaku-LLM, Rakuten, Takane) and strengthening trusted data infrastructure.⁴⁸ This selective approach gives it both frontier access and credible continuity in the event of external disruption, reflecting a posture that blends control in priority areas with steering and negotiated dependence elsewhere.

A central pillar of this strategy is resilient technical foundations. Domestic supercomputing and emerging national GPU clusters, plus a growing suite of open-weight Japanese models, form a technical backbone that can support continued operations even if external access is impeded. Initiatives such as the Generative AI Accelerator Challenge (GENIAC) and improved national language training data also strengthen the foundations for culturally aligned model development.^{49,50} These investments are complemented by a governance approach that emphasises responsible, interoperable and transparent AI systems, underpinned by the Act on Protection of Personal Information (APPI), the Act on Promotion of Research and Development, and Utilization of Artificial Intelligence-related Technology (AI Promotion Act) and Japan's leadership of the Hiroshima AI Process, a G7-led initiative to develop shared principles, risk-mitigation measures and governance approaches for advanced AI.⁵¹ This combination strengthens Japan's international credibility, positions the country as a credible norm-setter and enhances its attractiveness as a base for AI innovation.

Japan faces demographic challenges and shortages in digital talent, as well as structural dependencies in import-reliant energy systems. It confronts them through active capacity building, such as through grid upgrades, energy-transition strategies and structured talent pipelines. This dual capability-strengthening model expands domestic capacity while simultaneously leveraging global ecosystems to accelerate innovation and assure frontier access.

Japan's strategy contrasts sharply with faster, investment-led ecosystems such as the UAE and differs from Kenya's sequenced capacity building. Unlike France, which leans on regulatory influence, Japan's posture emphasises resilience and continuity. Its approach shows how an advanced economy with structural vulnerabilities can enhance sovereignty by selectively reinforcing critical layers while leveraging international ecosystems for scale.

UAE: FRONTIER ADJACENCY THROUGH ACCELERATED STATE INVESTMENT

The UAE has pursued one of the world's fastest-moving AI-sovereignty strategies, driven by strong state capacity, centralised coordination and deep partnerships with global technology firms. Its posture combines direct control in areas relevant to national security with aggressive co-development and frontier collaboration, accelerating its geopolitical positioning as a regional AI powerhouse. The resulting hybrid posture means that the UAE retains direct control where AI intersects with national security, data governance and public-service delivery, while steering and co-developing in areas that tap into foreign expertise or the ability to scale innovation.

The country's sovereign capacity is reinforced by highly centralised coordination and state-led investments, most visibly through the establishment of a dedicated AI ministry and flagship institutions such as technology group G42 and the Mohamed bin Zayed University of Artificial Intelligence (MBZUAI).⁵² What distinguishes the UAE is its willingness to act early and invest heavily, turning its sovereign energy base, domestic cloud ecosystem and political agility into strategic assets. Its open-weight Arabic LLMs, such as Falcon, Jais and K2 Think, demonstrate how a country can shape linguistic and cultural representation in global AI.⁵³ These exemplify a strategy built on speed and scale to attain frontier capability while maintaining local jurisdiction and narrative control.

However, despite its ambition and strong capacity, the UAE remains dependent on foreign chips and frontier models. These dependencies are mitigated through long-term negotiated access, co-production agreements and US-aligned security partnerships, which collectively secure reliable compute and model access. By pairing strong domestic investment with structured interdependence, the UAE mitigates risks while accelerating frontier capabilities, turning external reliance into a managed asset rather a vulnerability.

Compared to Japan's resilience-first posture, Kenya's sequenced dependence or France's governance-led model, the UAE prioritises acceleration and global positioning that seeks frontier adjacency through scale and ambition. Its trajectory illustrates how a small state with high capacity can expand sovereignty through decisive investment and structured interdependence.

KENYA: PROGRESSIVE CAPACITY BUILDING THROUGH SEQUENCED, NEGOTIATED INTERDEPENDENCE

Kenya's AI-sovereignty strategy acknowledges its structural constraints (limited compute, uneven energy availability and nascent model ecosystems) but treats them as the basis for sequenced capability building. Instead of pursuing premature autonomy, Kenya has adopted a phased approach: securing negotiated access to frontier systems through a partnership with Cassava Technologies, strengthening regulatory authority and expanding domestic capability incrementally in partnership with global firms.^{54,55} The core strategy is to build sovereignty through sequencing – a deliberate progression from access to capability to influence.

The most visible expression of this model is the landmark \$1 billion Microsoft-G42 initiative to develop a geothermal-powered data-centre campus and a new Azure cloud region for East Africa.⁵⁶ By anchoring compute to its geothermal-energy advantage, Kenya converts a domestic strength in energy into a key enabler for sovereign AI capacity. Strengthened data protection, as well as policy signals reflecting the country's prioritisation of domestic training capacity and open-data curation, and co-development models across health, agriculture and climate illustrate a posture of directed, strategic dependence.^{57,58}

Kenya manages its constraints through directed, strategic dependence, by securing frontier access via long-term partnerships, retaining oversight of data governance, compliance and procurement standards, and using regulatory strengths to shape how foreign infrastructure operates domestically.

Kenya's pathway contrasts with Japan's fallback model, the UAE's frontier proximity and France's regulatory strength. It shows how an emerging economy can expand agency not through immediate capability, but through negotiation, sequencing and leveraging unique national assets. Kenya's example demonstrates that sovereignty is attainable even for states with limited resources.

FRANCE: SOVEREIGNTY THROUGH PUBLIC-LED INFRASTRUCTURE AND REGULATORY POWER

France's approach to AI sovereignty reflects a long tradition of state-led industrial strategy combined with regulatory authority, creating a model that prioritises domestic compute, enforceable data governance and nuclear-backed energy sovereignty. This combination enables the state to steer AI development from a position of infrastructural and institutional strength while remaining interconnected with global ecosystems.

What distinguishes France is the depth and coherence of its state-anchored foundations. High-performance computing systems, from the upgraded Jean Zay supercomputer to the forthcoming joint MGX, Bpifrance, Mistral and NVIDIA AI campus near Paris (with a capacity of about 1.4 GW), anchor frontier compute within French jurisdiction while maintaining interoperability with global supply chains.⁵⁹ Sovereign-cloud initiatives such as Bleu and the SecNumCloud certification standard reinforce this control, ensuring that sensitive data remain under national oversight and stewardship. France's open-data infrastructures (data.gouv.fr, Recherche Data Gouv) further support research and innovation while enabling the state to steer data flows with confidence grounded in strong enforcement of the General Data Protection Regulation (GDPR) by its national data regulator, the National Commission on Informatics and Liberty (CNIL). These measures together give France meaningful leverage over how AI systems are built, governed and deployed domestically.

Although France benefits from strong institutions and access to nuclear energy, it operates within a highly globalised AI ecosystem and relies on foreign semiconductor supply chains and multi-national cloud providers. France manages these dependencies by maintaining open channels for

innovation while imposing conditions such as data protection and cloud security. By combining domestic compute expansion with regulatory tools, France turns interdependence into a source of leverage.

France's strategy contrasts with acceleration-led models such as the UAE's or resilience-first approaches such as Japan's. Unlike Kenya's or Brazil's partnership-driven capacity building, France's posture derives from long-standing institutional strength, abundant low-carbon nuclear energy and regulatory credibility. This combination enables France to shape global norms where scale requires openness, while exercising control in domains where domestic institutions confer durable leverage. France illustrates how a state with robust governance capacity can convert regulatory power and infrastructure investment into meaningful technological influence without pursuing full-stack technological independence.

INDIA: SOVEREIGNTY THROUGH SCALE, DIGITAL PUBLIC INFRASTRUCTURE AND LAYERED CONTROL

India's AI-sovereignty model leverages its population's scale, its domestic talent reservoir and state-backed digital public infrastructure (DPI) to forge a distinctive model of AI capability. Through the IndiaAI Mission and the broader India Stack, the government asserts strategic control over data, digital identity and foundational digital services, while partnering with domestic and international firms to bridge gaps in compute, semiconductors and energy reliability. This produces a layered sovereignty posture: control where the state has strong institutional capacity (data governance, standards, India Stack and regulatory norms), steering where it can direct investments and innovation (compute, model development and public-sector applications) and managed dependence where structural constraints remain (chips, energy and frontier-scale compute).

What distinguishes India is the depth of its DPI, which gives the state unparalleled leverage over identity, payments, data sharing and digital public-service delivery. Targeted industrial strategy complements these foundations. The India Semiconductor Mission, with more than \$10 billion committed to catalyse domestic fabrication, assembly and design, exemplifies this effort to transition from taker-dependence to maker-

dependence in the compute supply chain.⁶⁰ India's growing ecosystem of sovereign-relevant models (BharatGPT, Sarvam-1 and Sarvam M), its prioritisation of vernacular-language models, public-service use cases and large-scale multimodal systems in policy, and its role in global governance forums (the Global Partnership on Artificial Intelligence, G20 and the AI Action Summit) further demonstrate how it converts scale, data and institutional coherence into strategic leverage. Co-investment arrangements with AWS, Microsoft, AMD and others expand national training capacity while reinforcing India's ability to adapt foreign technology rapidly to local needs.

India manages its constraints in semiconductor fabrication, energy reliability and frontier compute through structured partnerships such as co-investment in cloud and compute infrastructure, joint model development and localisation requirements. Regulatory tools covering data protection, accountability and standards provide the state with additional leverage to shape how foreign technology providers operate in India.

India's pathway differs from regulation-centred approaches such as in France and from Kenya's sequenced model of capacity-building through negotiated dependence. Instead, India benefits from a unifying national digital architecture that few economies can replicate, providing a common architecture across a large and diverse federal system. Its approach shows how a large federal democracy can expand agency by governing key levers of data, identity and digital infrastructure while leveraging partnerships to fill capability gaps.

Rather than seeking full control across the AI stack, India uses scale, market depth and state-coordinated DPI to generate strategic leverage. This layered posture constitutes a sovereignty model rooted in population scale, institutional coordination and deliberate prioritisation rather than technological self-sufficiency.

BRAZIL: SOVEREIGNTY THROUGH REGULATORY STRENGTH AND STATE-GUIDED CO-CREATION

Brazil's AI-sovereignty strategy is shaped by its position as Latin America's largest digital, energy and industrial market. Its approach is one of strategic hybridisation, combining a state-guided model of ecosystem co-creation with regulatory authority and a clean and abundant energy matrix to build domestic capability while remaining open to global partnerships. Brazil is anchoring key elements of the AI stack, from data governance to energy security and emerging compute infrastructure, while using public policy to steer both domestic actors and foreign investors towards national priorities.

Brazil's capability-building model is rooted in co-creation rather than isolation. Public institutions, universities and domestic firms are developing Portuguese-language and regionally aligned models such as SoberanIA, Amazônia 360 and GovBERT-BR. These initiatives are supported by significant public R&D commitments (approximately \$4 billion under the AI Plan 2024–2028), and increasing international cooperation, including the China–Brazil joint AI lab.⁶¹ Brazil also leverages its regulatory power and energy sovereignty to attract hyperscale investments.

While Brazil depends on foreign providers for cloud services, advanced chips and foundation models, it manages these dependencies through regulatory strength and investment conditionalities. Major commitments including Microsoft's \$2.7 billion commitment, AWS's \$1.8 billion expansion and Scala's 4.75 GW AI City all follow reinvestment, localisation and compliance requirements that ensure global compute capability is physically and economically anchored on Brazilian soil.^{62,63,64} This combination of enforcement capability and infrastructural advantage allows Brazil to attract and condition global capability while building local industry.

Brazil's posture differs from India's DPI-anchored model of layered control, the UAE's frontier acceleration strategy driven by capital and global partnerships, and Japan's resilience-first approach centred on fallback capacity. Its distinctiveness lies in its combination of regulatory leverage, supported by a favourable energy base and large internal market, which together allow it to shape how foreign capability enters and operates

domestically. This configuration enables Brazil to embed foreign investment within national safeguards and convert dependence into domestic innovation and capability building, demonstrating how large emerging economies can expand agency without pursuing full-stack autonomy.

07

Policy Levers for Expanding Agency in the Age of AI

The following recommendations set out practical actions that decision-makers can take to exercise sovereignty in the age of AI. They show how accelerating AI diffusion across sectors strengthens domestic competitiveness and reduces single-point dependence, how aggregated demand and pooled infrastructure can secure more predictable access to frontier capacity, and how targeted investments in data governance, talent, energy and assurance institutions can anchor long-term capability.

Each lever represents a core trade-off between openness and control, speed and stability, and access and autonomy. Taken together, they shape how countries position themselves within the global AI ecosystem, ensuring external capabilities reinforce national priorities rather than constrain them. This provides a pathway for countries to remain connected to frontier innovation while developing the institutional, technical and market foundations required for resilience and sustained competitiveness.

Secure Access to Frontier AI Models and Compute

Frontier AI models and large-scale compute are becoming essential infrastructure for national competitiveness in the new AI-enabled economy. The strategic challenge for governments is not to own these systems outright, but to secure reliable, predictable and governable access to these capabilities. Each country must make deliberate choices about how much compute it requires, which function can rely on remote or foreign infrastructure, and which require domestic hosting to ensure resilience, continuity and security in the event of disruption to global compute access or frontier model capabilities. The same is true for models. Policymakers should identify which use cases can rely on globally available frontier systems and those that demand higher levels of control, auditability or security where options such as customised open-weight models or secure domestic deployments offered by industry may be more appropriate.

- **Develop a national strategy that defines frontier model and compute requirements, access pathways and long-term needs.** Building domestic AI infrastructure requires coordination among various government functions, from planning and grid management to implementation. Governments should develop a clear, cross-government strategy, led from the prime minister's or president's office, that sets out sector-specific compute and model requirements, identifies which workloads must be sovereign and which can be external, and defines expectations around security, availability and resilience. Central coordination is essential to align planning, energy, procurement and international negotiations.
- **Build sufficient domestic compute for deployment and inference.** Governments should concentrate their resources and efforts on securing domestic inference (the ability to run trained AI models and generate real-time outputs) capacity, rather than frontier training. Ensuring local availability of national GPU clusters, sovereign cloud regions or high-availability inference infrastructure operated under domestic governance enables AI systems to be deployed reliably across public services. This domestic capacity underpins continuity in critical services such as health care, justice, border management and national security during global outages, export-control disputes or vendor policy shifts.
- **Negotiate multi-year sovereign access agreements that secure predictable, high-priority access to frontier capability.** Governments should establish structured, multi-year access agreements with leading international frontier AI firms to guarantee predictable access to frontier AI capabilities. These agreements should include minimum compute allocations for priority sectors, stable long-term pricing, non-interruption clauses, emergency access provisions and audit rights, reducing exposure to external supply shocks while strengthening bargaining power.
- **Form regional bargaining blocs to negotiate frontier access collectively.** Pooling demand across neighbouring or allied states allows governments to negotiate lower prices, guaranteed capacity slices and shared access pools, including coordinated training runs for regional

open-weight models. These kinds of consortia further reduce dependence on any single provider or jurisdiction, and potentially increase resilience to geopolitical or climate shocks.

Accelerate AI Adoption and Diffusion Across Sectors

Developing or possessing large-scale frontier AI models does not, by itself, create economic or strategic value. Sovereignty is strengthened when AI becomes a widely diffused capability, used by firms, public institutions and citizens to increase productivity, improve services and enable new forms of innovation. For most countries, this means that their focus should be oriented towards accelerating the creation and deployment of AI applications across the economy and ensuring that labour, capital and institutions can absorb and adapt to these gains. This requires governments to focus simultaneously on use and absorption: supporting the development of applications that can scale across sectors and creating the economic conditions that allow AI-driven productivity to translate into new firms, industries and higher living standards rather than displacement or concentration alone.

- **Make AI-driven transformation an explicit pillar of economic strategy.**

Governments should embed AI adoption into national growth, industrial and sector strategies, with clear goals for how AI will modernise manufacturing, health care, finance, agriculture and public services. This also requires aligning regulatory, skills and investment policies to enable firms to adopt AI at scale and to support the creation of new industries that emerge from AI-driven innovation. By taking this approach, governments can accelerate economic renewal, capture higher-value segments of global supply chains and secure long-term national prosperity.

- **Treat national data sets as strategic assets to accelerate adoption.**

Governments should develop trusted national data institutions or shared data spaces that curate high-quality, representative and machine-readable data sets in priority domains such as health, agriculture, climate and finance. These assets enable more accurate fine-tuning and

application development, reduce dependence on foreign data sets that may not reflect local contexts, including languages, and create reusable building blocks for research and innovation.

- **Design sector-specific, use-case-oriented regulatory frameworks that enable innovation.** Horizontal, one-size-fits-all AI regulations often slow deployment in highly regulated fields such as health, energy, manufacturing and finance. Instead, governments should adopt sectoral frameworks that set clear guardrails around high-risk uses while enabling experimentation and faster deployment in low-risk contexts.

Aggregate and Signal National Demand to Shape the AI Market

Private investment, talent and infrastructure flow to places where demand is credible, coordinated and sustained over time. When AI demand is fragmented across ministries, agencies or regions, ecosystems struggle to emerge: vendors face inconsistent requirements, solutions fail to scale and startups lack predictable pathways to growth. Aggregating demand, within government and where possible across regions, transforms small or uneven markets into structured opportunities that attract investment, justify national or regional AI infrastructure and create meaningful leverage in negotiating frontier access. Clear demand signals enable governments to influence vendor behaviour, set interoperability expectations and crowd in private capital towards national priorities. In a global environment where every country is competing for frontier access and investment, demonstrating coordinated domestic demand is often the decisive factor in attracting hyperscalers, investors and partners.

- **Publish multi-year AI roadmaps and national compute demand forecasts.** Governments should produce and update annual forecasts of national compute demand over one, three and five-year horizons, covering both public sector and anticipated private-sector use. These roadmaps should indicate which proportions of demand require

sovereign capacity and which can be met through external providers, reducing uncertainty for investors and enabling coherent planning across ministries.

- **Strengthen scale-up finance and investment pathways for AI-enabled firms.** Domestic AI ecosystems require access to growth capital to scale. Governments should review capital-market rules, institutional-investor mandates and development-finance instruments to ensure pension funds, insurers, sovereign-wealth funds and public banks can invest in growth equity and venture funds supporting AI-driven companies.
- **Use targeted incentives, co-investment guarantees and pooled procurement to de-risk early deployment.** In sectors where public value is high but commercial returns are uncertain, governments can deploy subsidies, co-investment guarantees or minimum-use agreements to anchor early demand. Pooling procurement across public-sector organisations further reduces risk for suppliers, delivers better value for money and ensures products are developed in line with national needs rather than fragmented local requirements.

Treat Interoperability as a Core Component of Sovereignty

As governments invest in new AI infrastructure, systems and services, interoperability must be treated as a sovereign capability and not a technical afterthought. In a rapidly evolving technological landscape, no state can predict which models, platforms or vendors will dominate in the future. Without interoperability, early architecture and procurement choices can hard-wire long-term dependencies that become difficult and costly to unwind. By contrast, interoperability underpins strategic autonomy by ensuring that public systems can switch providers, combine models and integrate new capabilities without wholesale rebuilds. It allows governments to benefit from global innovation while retaining the freedom to reconfigure their AI stack over time. This is especially important in an environment where frontier capabilities are concentrated, and vendor incentives may diverge

from public priorities. Governments should therefore focus on modular architectures, open standards and clear interoperability requirements across all major AI systems and services.

- **Define and enforce a national AI systems architecture based on modularity and open standards.** Governments should establish clear architecture principles for all new public-sector AI systems, including common data formats, open APIs and modular components. A dedicated architecture function can coordinate standards across ministries, prevent siloed development and ensure systems remain composable as technology evolves.
- **Use public procurement to make interoperability the default.** Public-procurement rules should explicitly favour AI-enabled systems that demonstrate interoperability across providers, clouds and orchestration frameworks. Contracts for major AI systems should require open interfaces, multi-cloud compatibility, and full portability of data and models in widely used formats, ensuring workloads can migrate without prohibitive switching costs or loss of functionality.
- **Embed interoperability obligations into regulation and oversight.** Regulators should make interoperability a baseline expectation in sectoral and cross-sector AI frameworks rather than an optional feature. This includes mandating switching rights, disclosure of interface specifications and ongoing compatibility with alternative providers. These obligations should be paired with structured knowledge transfer such as embedded engineering support, secondments and training so that public institutions can maintain and adapt systems independently over time.

Build and Scale Smaller, Efficient and Contextually Relevant Models

Frontier-scale models are costly, compute-intensive, linguistically narrow and often misaligned with local contexts. For most countries, training such models is neither necessary nor feasible. Strategic value instead lies in smaller, efficient and adaptable models that can be tailored to national languages, regulatory environments and priority sectors such as health,

finance, energy and science. As open-weight models rapidly narrow the performance gap with closed frontier systems, these approaches offer governments more cost-effective, governable and contextually aligned options for deploying AI at scale. By focusing on model adaptation rather than model ownership, countries can strengthen agency without entering the frontier model race. Fine-tuned open-weight models and small language models can deliver reliable public-sector performance, keep sensitive data within national governance frameworks and support sector-specific innovation. Used alongside selective access to frontier systems, they enable governments to deploy AI where it matters most while retaining control over data, risk and operational continuity.

- **Invest in fine-tuning open-weight foundation models for national languages, domains and regulatory requirements.** Rather than building models from scratch, governments should support the fine-tuning of existing open-weight models on nationally relevant data sets – for example, for linguistic alignment or domain-specific use in health care, agriculture or public administration. As demonstrated by Japan’s Fugaku LLM, among other similar emerging efforts, these models can become sovereign assets that are culturally aligned, contextually relevant and capable of efficient deployment.
- **Develop regional model consortia for shared linguistic and cultural contexts.** Countries with overlapping languages or use cases can co-develop open-weight models, similar to the Southeast Asian Languages in One Network (SEA-LION) initiative to share costs, expand data set diversity and expand applicability across borders.⁶⁵ Regional model alliances would enable smaller states to access systems they cannot develop on their own while strengthening cross-border innovation ecosystems.
- **Adopt hybrid model architectures that combine local and frontier capability.** Hybrid designs allow governments to combine the low-cost, locally governed capabilities of small models for sensitive or routine tasks with the specialised strengths of frontier systems for complex reasoning or multimodal workloads. This hybrid approach can reduce operational

costs and exposure to external policy shifts or outages while preserving access to state-of-the-art performance where required.⁶⁶ Governments can incentivise hybrid pilots for high-impact public services.

- **Establish domain-specific data trusts and national data libraries under clear governance frameworks.** Instead of negotiating training-data access on an ad-hoc basis, governments can create data trusts that securely pool anonymised data sets for model fine-tuning and evaluation. These can be sector-specific, as in education or health care, or as TBI has previously set out, they can be national endeavours that support private-public collaboration through Data Biome partnerships that enable model development aligned with national priorities.^{67,68,69}

Invest in Talent and State Capacity

Sovereignty is rooted not only in technical capability, but in human and institutional capacity. A country's ability to harness AI's potential will ultimately depend on the capacity of its people and institutions to understand, deploy and govern it. Access to models or compute is insufficient without a workforce that can use AI productively and public institutions that can evaluate risks, procure systems and enforce rules with confidence. Governments must act on two fronts. First, they need to reskill and upskill the existing workforce, so AI adoption translates into productivity gains rather than displacement or concentration. Second, they must attract, develop and retain world-leading technical talent while embedding that expertise within the state itself. Without this internal capacity, governments remain dependent on external actors for interpretation, implementation and oversight, regardless of how much infrastructure they access. By strengthening talent at every level, governments can create the conditions necessary for innovation, resilience and broad-based economic gains.

- **Accelerate national reskilling and upskilling for an AI-enabled economy.** Governments should launch large-scale national programmes to build AI literacy and practical skills across sectors, from manufacturing and agriculture to health care and public administration. This includes

modular training pathways, sector-specific curricula developed with industry, and incentives that encourage firms and workers to adopt AI tools in everyday workflows.

- **Reform universities and research institutions to compete for AI talent.** Universities must be equipped to support a modern innovation economy rather than reproduce rigid academic hierarchies. Governments should modernise governance and incentives, enable competitive remuneration, and create flexible career paths that allow movement between academia, industry and government, helping to retain talent and accelerate knowledge transfer.
- **Build deep public-sector technical capacity to govern AI credibly.** States should raise baseline AI literacy across the civil service, including senior leaders, policymakers, technical teams and frontline officials, through mandatory training and continuous learning programmes. At the same time, governments should establish dedicated technical career tracks with competitive pay, clear progression and rotation opportunities, complemented by fellowships and secondments that embed expertise within core public institutions.

Align AI Infrastructure with Sustainable Energy Planning

The availability, reliability and cost of electricity are emerging as critical bottlenecks on a country's ability to deploy and scale AI. As AI adoption accelerates, AI workloads will place growing pressure on national grids and long-term energy systems. Because energy infrastructure is shaped by structural constraints and long time horizons, most countries cannot expand supply at the speed or scale required for compute-intensive AI. Sovereignty therefore depends less on building more energy than on aligning AI infrastructure with the energy systems countries can reliably expand and sustain. Countries that fail to integrate AI needs into their energy planning risk higher operational costs, grid instability, increased reliance on imported power and, ultimately, an inability to host sufficient compute domestically. By contrast, countries that plan proactively, co-locating compute with reliable, low-carbon energy and integrating AI demand into grid strategy, can attract

investment, negotiate from a position of strength with hyperscalers and ensure that critical digital services remain resilient. In all cases, future-oriented energy planning is fundamental to sustaining national competitiveness in an AI-driven world.

- **Steer AI infrastructure towards energy-secure, low-carbon regions.**

Governments can designate AI energy zones near hydro, geothermal, nuclear or major solar corridors, and use targeted incentives such as concessional finance, long-term power-purchase agreements (PPAs) and expedited permitting to attract compute to grid-stable regions. This approach reflects practices seen in France's nuclear-enabled AI campus, Kenya's geothermal-backed cloud deployments and Brazil's hydro-powered data-centre ecosystem.

- **Integrate AI demand forecasting into national grid and infrastructure planning.**

Ministries of ICT and energy should jointly model AI workloads and incorporate them into transmission, storage, cooling and substation planning. Aligning grid upgrades with projected compute growth reduces the risk of localised stress, prevents crowding out of households or SMEs, mitigates the risk of local blackouts, and ensures AI expansion matches real energy availability and future supply.

- **Adopt green-compute standards and transparent energy reporting.**

Governments should introduce standards tied to efficiency, utilisation and emissions, such as waste-heat recovery, low-carbon PPA requirements and reporting on energy use, emissions intensity and environmental-mitigation measures, for all significant AI infrastructure projects. These measures strengthen accountability, push vendors towards energy-efficient model and infrastructure designs, reduce environmental impact, and ensure that the expansion of AI infrastructure aligns with national climate targets and long-term energy resilience.

08

Conclusion: The Long Game of Strategic Resilience

Building resilience in the age of AI is less about controlling a single technology than developing the capacity to adapt as technologies evolve. AI will continue to advance faster than policy cycles, and this pace of change means sovereignty cannot be treated as an end state. It is an ongoing practice bringing investment, governance, diplomacy and public purpose into coherence so that states can shape innovation rather than merely absorb it.

There is no single ideal sovereign posture. What is feasible for large industrial economies may be inefficient or unattainable for smaller or energy-constrained states. Sovereignty instead emerges from how governments sequence and align decisions across the AI stack: asserting control where necessary, steering markets and standards where they have leverage, and managing interdependence where capability gaps are structural. Dependencies are inevitable, but they can be designed to be diversified, reciprocal and capable of being renegotiated as national priorities shift.

AI will continue to reconfigure rather than resolve geopolitical competition. The same systems that promise inclusion, efficiency and economic growth can also concentrate power in compute, energy and talent supply chains. Smaller, efficient open-weight models may broaden participation, but the upper layers of the AI stack will remain highly centralised. Navigating these asymmetries will require governments to balance openness with security, leverage with prudence and ambition with realism.

The measure of success for sovereignty in the age of AI will not be total self-sufficiency, but strategic resilience. This is defined by the capacity of a country to anticipate interdependence, learn rapidly from disruption and steer innovation towards public value. Countries that treat AI as a central pillar of state capability, deploy it widely across the economy and public

services, and build institutions capable of evaluating and governing it will be best placed to preserve agency in a world of accelerating technological change.

AI sovereignty, therefore, cannot be pursued in isolation; it must be built through deliberate interdependence. States must ground their approach in their comparative strengths and invest in the layers of the AI stack that matter most for national priorities. The most competitive countries will be those that adopt AI at scale, align their investments with clear strategic intent, and negotiate their place in the global AI ecosystem from a position of purpose and confidence. Maintaining sovereignty in the AI era will depend on pragmatic decision-making, strategic prioritisation and continuous negotiation across a rapidly evolving technological landscape.

Annex

[Download the Annex as a PDF.](#)

Acknowledgements

The team would like to extend its thanks to the following experts who offered their advice and guidance in the development of this report:

Yuko Harayama, Tokyo Centre of the GPAI Expert Community

Caroline Meinhardt, Stanford HAI

Mariam Razak and Glen Robinson, Microsoft

Brittany Smith, OpenAI

Lawrence Wee, Infocomm Media Development Authority of Singapore

Suzy Wild, Anthropic

Endnotes

- 1 <https://trt.global/world/article/d68fb4b84c55>
- 2 <https://avpn.asia/about/press/avpn-and-google-org-announces-usd-15-million-ai-opportunity-fund-asia-pacific-to-empower-workers-in-apac-for-the-ai-driven-future/>
- 3 <https://www.cigionline.org/articles/chinas-ai-governance-initiative-and-its-geopolitical-ambitions/>
- 4 <https://www.csis.org/analysis/united-arab-emirates-ai-ambitions>
- 5 <https://www.bennettschool.cam.ac.uk/blog/aseans-geopolitical-strategy/>
- 6 <https://www.brookings.edu/wp-content/uploads/2025/08/GS%5F08252025%5FAISA%5Freport.pdf>
- 7 <https://policyreview.info/concepts/digital-sovereignty>
- 8 <https://www.cer.eu/insights/can-eu-afford-drive-out-american-cloud-services>
- 9 <https://www.economist.com/business/2025/07/13/can-nvidia-persuade-governments-to-pay-for-sovereign-ai>
- 10 <https://www.dhs.gov/publication/fact-sheet-and-report-dhs-advances-efforts-reduce-risks-intersection-artificial>
- 11 <https://www.economist.com/business/2025/07/13/can-nvidia-persuade-governments-to-pay-for-sovereign-ai>
- 12 <https://news.microsoft.com/source/emea/2025/11/microsoft-expands-digital-sovereignty-capabilities/>
- 13 <https://www.atlanticcouncil.org/blogs/geotech-cues/the-sovereignty-trap/>
- 14 <https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses>
- 15 <https://www.ft.com/content/05206915-fd73-4a3a-92a5-6760ce965bd9>
- 16 <https://www.csis.org/analysis/understanding-us-allies-current-legal-authority-implement-ai-and-semiconductor-export>

- 17 <https://merics.org/en/report/chinas-drive-toward-self-reliance-artificial-intelligence-chips-large-language-models>
- 18 <https://epoch.ai/blog/trends-in-ai-supercomputers>
- 19 <https://www.nytimes.com/interactive/2025/06/23/technology/ai-computing-global-divide.html>
- 20 <https://epoch.ai/data-insights/ai-supercomputers-performance-share-by-country>
- 21 <https://www.iea.org/reports/energy-and-ai/executive-summary#abstract>
- 22 <https://www.goldmansachs.com/what-we-do/investment-banking/insights/articles/powering-the-ai-era/report.pdf>
- 23 <https://news.microsoft.com/source/emea/features/the-port-town-in-norway-emerging-as-an-ai-hub/>
- 24 <https://www.bbc.co.uk/news/articles/cvg8z2mx508o>
- 25 <https://www.politico.eu/article/emmanuel-macron-answer-donald-trump-fossil-fuel-drive-artificial-intelligence-ai-action-summit/>
- 26 <https://epoch.ai/blog/will-we-run-out-of-data-limits-of-lm-scaling-based-on-human-generated-data>
- 27 <https://commoncrawl.github.io/cc-crawl-statistics/plots/languages>
- 28 <https://www.nature.com/articles/d41586-024-03990-2>
- 29 <https://blogs.microsoft.com/on-the-issues/2025/07/20/eudigitalunlock/>
- 30 <https://www.rdworlondonline.com/how-xai-turned-a-factory-shell-into-an-ai-colossus-to-power-grok-3-and-beyond/>
- 31 <https://siliconangle.com/2025/12/30/elon-musk-reveals-plan-expand-xais-colossus-data-center-2-gigawatts/>
- 32 <https://x.ai/news/series-c>
- 33 <https://www.ft.com/content/efele350-62c6-4aa0-a833-f6da01265473>
- 34 https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf
- 35 <https://arxiv.org/abs/2401.05566>
- 36 <https://www.crowdstrike.com/en-us/blog/crowdstrike-researchers-identify-hidden->

vulnerabilities-ai-coded-software/

37 <https://www.sciencedirect.com/science/article/pii/S004873323001828>

38 <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>

39 <https://www.timeshighereducation.com/news/kazakhstan-has-enough-branch-campuses-says-he-minister>

40 <https://e-estonia.com/ai-leap-2025-estonia-sets-ai-standard-in-education/>

41 <https://www.businesstoday.in/technology/news/story/uae-makes-chatgpt-plus-subscription-free-for-all-residents-as-part-of-deal-with-openai-477948-2025-05-27>

42 <https://institute.global/insights/public-services/generation-ready-building-the-foundations-for-ai-proficient-education-in-englands-schools>

43 <https://institute.global/insights/tech-and-digitalisation/how-leaders-in-the-global-south-can-devise-ai-regulation-that-enables-innovation>

44 <https://blogs.nvidia.com/blog/fugakunext/>

45 <https://www.idc-a.org/news/France-is-building-its-own-Cloud-de-Confiance-for-government-agencies-and-critical-infrastructure-pl>

46 <https://www.brasildedefato.com.br/2025/05/28/china-brazil-ai-agreement-reinforces-joint-researches-and-infrastructure-development/>

47 <https://carnegieendowment.org/research/2025/08/indias-semiconductor-mission-the-story-so-far?lang=en>

48 <https://www.datacenterdynamics.com/en/news/japans-riken-partners-with-nvidia-and-fujitsu-for-zettascale-fugakunext-supercomputer/>

49 <https://www.meti.go.jp/english/policy/mono%5Finfo%5Fservice/geniac/index.html>

50 <https://arxiv.org/abs/2404.17733>

51 <https://www.gov-online.go.jp/en/assets/hj%5FNovember%5F2025%5Fp24-25.pdf>

52 <https://www.reuters.com/world/middle-east/microsoft-g42-announce-200-mw-data-centre-capacity-expansion-uae-2025-11-05/>

53 <https://www.g42.ai/resources/news/mbzuai-and-g42-launch-k2-think-leading-open-source-system-advanced-ai-reasoning?>

- 54 <https://www.businessdailyafrica.com/bd/corporate/technology/uk-based-firm-to-deploy-kenya-s-first-rentable-ai-servers-5263850>
- 55 <https://techafricanews.com/2025/11/12/google-launches-gemini-pro-in-kenya-to-revolutionize-education-with-ai/>
- 56 <https://www.reuters.com/technology/microsoft-g42-invest-1-billion-kenya-build-data-center-2024-05-22/>
- 57 <https://www.odpc.go.ke/wp-content/uploads/2024/02/TheDataProtectionAct%5F%5FNo24of2019.pdf?>
- 58 <https://ict.go.ke/sites/default/files/2025-03/Kenya%20AI%20Strategy%202025%20-%202030.pdf>
- 59 <http://mgx.ae/en/news/mgx-bpifrance-mistral-ai-and-nvidia-launch-joint-venture-build-europes-largest-ai-campus>
- 60 <https://economictimes.indiatimes.com/industry/cons-products/electronics/india-eyes-chip-parity-with-us-china-by-2032-as-10-bn-semiconductor-push-gains-momentum/articleshow/125454590.cms?from=mdr>
- 61 <https://www.gov.br/g20/en/news/brasil-launches-a-usd-4-billion-plan-for-ai-and-prepares-global-action>
- 62 <https://www.reuters.com/technology/microsoft-make-27-billion-cloud-ai-investments-brazil-2024-09-26/>
- 63 <https://www.reuters.com/technology/amazons-aws-unit-invest-18-bln-brazil-through-2034-2024-09-11/>
- 64 <https://datacentremagazine.com/data-centres/scala-data-centers-ai-city-secures-5-gw-approval-in-brazil>
- 65 <https://sea-lion.ai/our-story/>
- 66 <https://arxiv.org/pdf/2502.15964>
- 67 <https://ai.gov.uk/knowledge-hub/tools/content-education-store/>
- 68 <https://trustplatform.sg/>
- 69 <https://institute.global/insights/tech-and-digitalisation/governing-in-the-age-of-ai-building-britains-national-data-library>

Follow us

[facebook.com/instituteglobal](https://www.facebook.com/instituteglobal)

x.com/instituteGC

[instagram.com/institutegc](https://www.instagram.com/institutegc)

General enquiries

info@institute.global

Copyright © January 2026 by the Tony Blair Institute for Global Change

All rights reserved. Citation, reproduction and or translation of this publication, in whole or in part, for educational or other non-commercial purposes is authorised provided the source is fully acknowledged Tony Blair Institute, trading as Tony Blair Institute for Global Change, is a company limited by guarantee registered in England and Wales (registered company number: 10505963) whose registered office is One Bartholomew Close, London, EC1A 7BL.