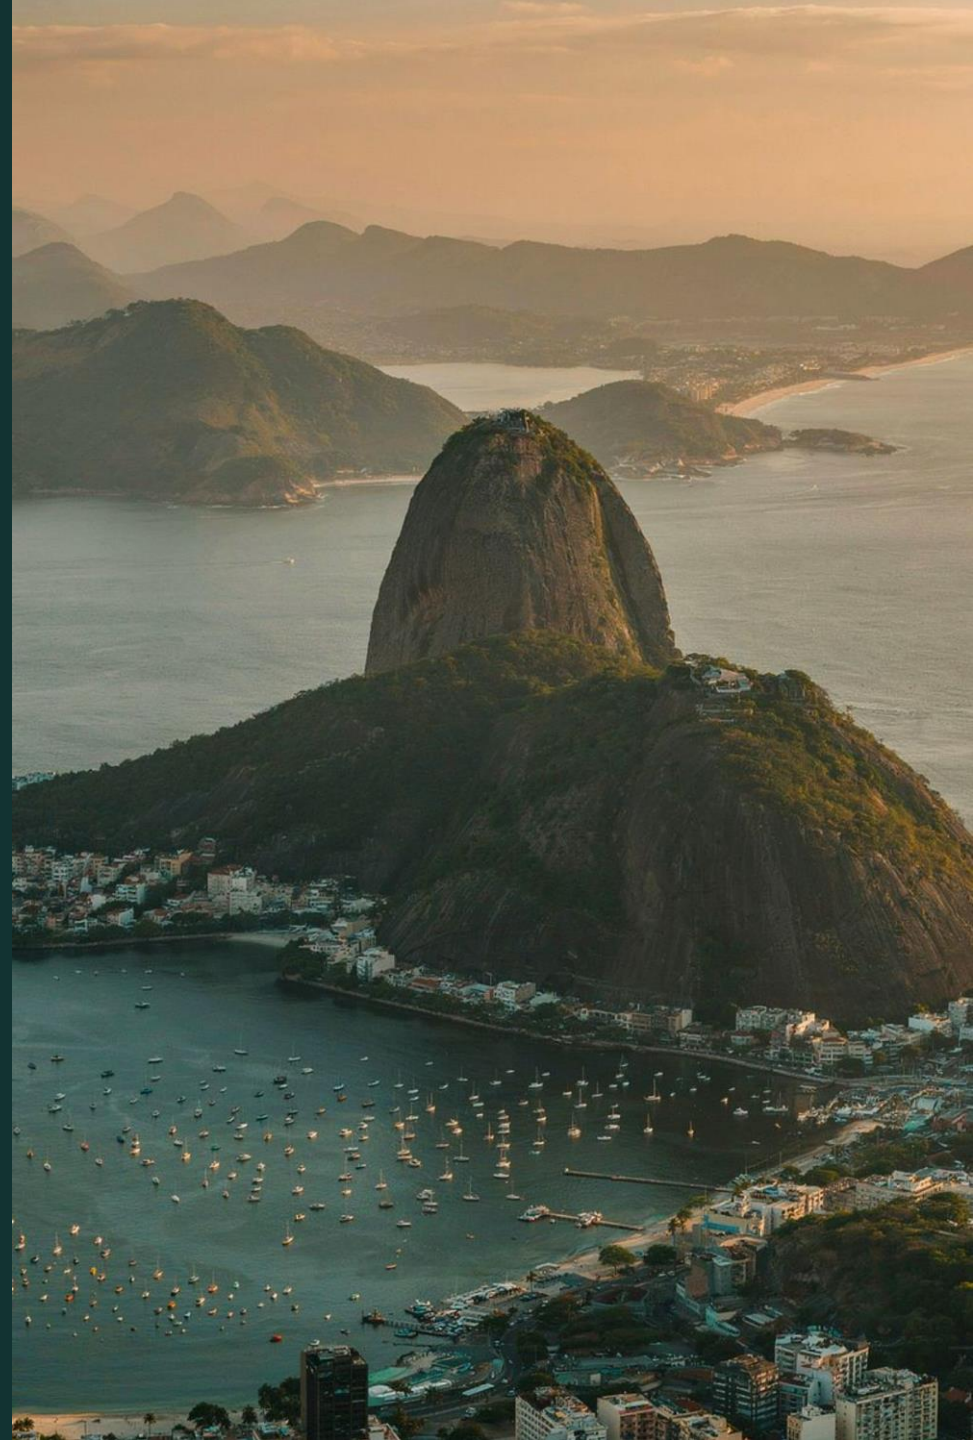
AI Sovereignty and the Challenge for Middle Powers

Nate Persily
James B. McClatchy Professor of Law
Co-Director, Stanford Law AI Initiative
Stanford AI Governance Lab
Stanford Law School





The Big Picture

Governance

How can public policy and private rules mitigate the risks and maximize the benefits of AI?

Sovereignty & Diffusion

How can countries retain autonomy as AI pushes toward dependence?

- Most countries in the world face an unprecedented challenge: Rapidly and responsibly integrate a revolutionary and unpredictable technology developed by a select few foreign corporations.



In the AI era, sovereignty is not the ability to build everything yourself.

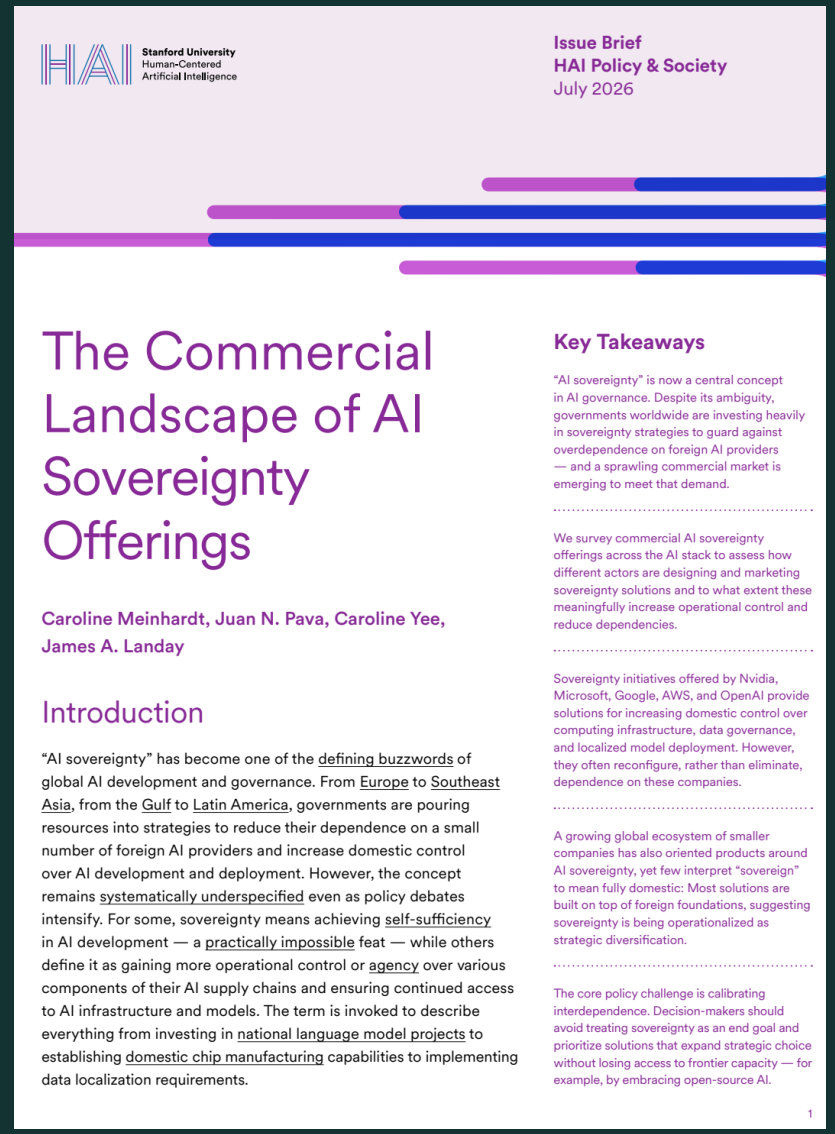
It is the ability to choose what you depend upon —

and to change that choice when your national interests require it.



AI Sovereignty

The concept and the concern



Stanford University
Human-Centered
Artificial Intelligence

Issue Brief
HAI Policy & Society
July 2026

The Commercial Landscape of AI Sovereignty Offerings

Caroline Meinhardt, Juan N. Pava, Caroline Yee, James A. Landay

Introduction

"AI sovereignty" has become one of the defining buzzwords of global AI development and governance. From Europe to Southeast Asia, from the Gulf to Latin America, governments are pouring resources into strategies to reduce their dependence on a small number of foreign AI providers and increase domestic control over AI development and deployment. However, the concept remains systematically underspecified even as policy debates intensify. For some, sovereignty means achieving self-sufficiency in AI development — a practically impossible feat — while others define it as gaining more operational control or agency over various components of their AI supply chains and ensuring continued access to AI infrastructure and models. The term is invoked to describe everything from investing in national language model projects to establishing domestic chip manufacturing capabilities to implementing data localization requirements.

Key Takeaways

"AI sovereignty" is now a central concept in AI governance. Despite its ambiguity, governments worldwide are investing heavily in sovereignty strategies to guard against overdependence on foreign AI providers — and a sprawling commercial market is emerging to meet that demand.

We survey commercial AI sovereignty offerings across the AI stack to assess how different actors are designing and marketing sovereignty solutions and to what extent these meaningfully increase operational control and reduce dependencies.

Sovereignty initiatives offered by Nvidia, Microsoft, Google, AWS, and OpenAI provide solutions for increasing domestic control over computing infrastructure, data governance, and localized model deployment. However, they often reconfigure, rather than eliminate, dependence on these companies.

A growing global ecosystem of smaller companies has also oriented products around AI sovereignty, yet few interpret "sovereign" to mean fully domestic: Most solutions are built on top of foreign foundations, suggesting sovereignty is being operationalized as strategic diversification.

The core policy challenge is calibrating interdependence. Decision-makers should avoid treating sovereignty as an end goal and prioritize solutions that expand strategic choice without losing access to frontier capacity — for example, by embracing open-source AI.

Six Competing Conceptions

Self-Sufficiency

Own and produce the full stack domestically, from chips to applications.

Strategic Autonomy

Keep real, viable alternatives so no single supplier holds veto power.

Regulatory Authority

Apply domestic law and institutions to AI systems operating within a country's territory.

Economic Capacity

Capture the economic value created by AI adoption, not merely enable foreign firms to extract it.

Cultural & Linguistic

Ensure native language, law, and local institutions are embedded in AI systems.

Democratic Agency

Ensure that societies — not just governments or firms — retain meaningful choice over AI's role.

The Central Paradox

1

Chips

Designed & fabricated abroad

2

Models

Trained in the U.S. and China

3

Cloud

A handful of multinationals

4

Deployment

Data, users, institutions, harms

How can a country be sovereign over a technology that no country controls?

FIVE QUESTIONS FOR ANY TECHNOLOGICAL RELATIONSHIP

The Exit Test

Access — Can we obtain it?

Is supply subject to foreign government discretion, commercial terms that can be revised, or export control classification?

Understand — Can we evaluate it?

Do we have the technical capacity to audit what the system does, how it fails, and whether it meets our standards?

Adapt — Can we modify it?

Can we fine-tune, align, or distill the system to serve native legal, linguistic, and institutional requirements?

Govern — Can our law reach it?

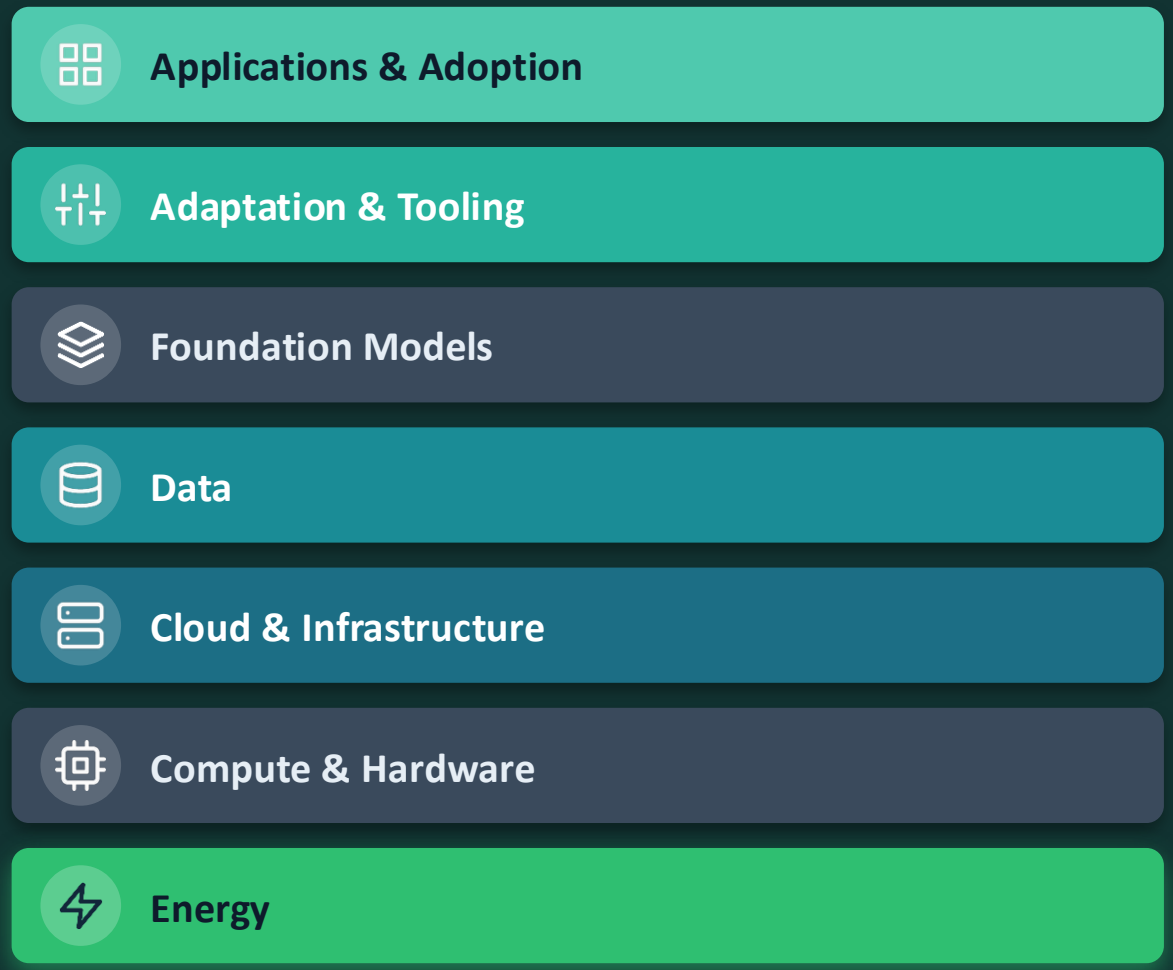
Is the deployer, developer, or supplier subject to Brazilian regulatory jurisdiction and court orders?

Exit — Can we replace it?

Is there a credible alternative, and what is the cost of switching? Dependence becomes a sovereignty problem when exit cost becomes prohibitive.

Sovereignty is less like ownership and more like optionality.

Sovereignty across the stack



BUILDABLE
Run foreign models across the economy

BUILDABLE
Integration, fine-tuning, local talent

GATED
Frontier labs — capital, data, compute

BUILDABLE
Local, language & domain datasets

BUILDABLE
Host data centers on home soil

GATED
Fabrication concentrated; export-controlled

BUILDABLE
Often import-dependent — powers every layer

SEA-LION = Southeast Asian Languages in One Network

A family of open-source large language models developed by AI Singapore under the nation's National Multi-Modal LLM Program

SEA-LION.AI

Models Leaderboard Blog Case Studies Community Resources About Collaborate

Catalysing AI Innovation for Southeast Asia

SEA-LION is a family of efficient, open-source, multilingual, multimodal language models designed to understand Southeast Asia's diverse languages, cultures, and contexts.

SEA-LION Telegram ↗ SEA-LION Playground ↗



Why we built SEA-LION

Existing LLMs often display strong biases in cultural values, political beliefs, and social attitudes. This is due to their training data, especially those scraped from the Internet, having heavily influenced content from Western, industrialized, rich, educated, and democratic (WIRED) societies. People from non-WIRED societies are less likely to be literate, use the internet, or have their output easily accessible, leading to a significant imbalance in the data.

SEA-LION addresses this by being trained on a greater volume of content produced in Southeast Asian languages like Thai, Vietnamese, and Bahasa Indonesia. This ensures better data representation and alignment compared to Western or Chinese models. SEA-LION models understand the nuances in SEA languages and demonstrate a greater awareness of the region's cultural context.

This lowers the barrier to adoption for governments, enterprises, academia, and end-users while effectively expanding Southeast Asian languages and cultural representation in mainstream LLMs which are currently dominated by models with a low representation of data from, about and by Southeast Asia.

Sovereignty over development vs. deployment

THE SAME AI STACK CREATES DIFFERENT KINDS OF NATIONAL LEVERAGE

DEVELOPMENT: constrained

Frontier models, advanced chips and hyperscale compute remain concentrated abroad.

Domestic capacity is expensive; access depends on suppliers and export controls.

DEPLOYMENT: more governable

Governments can regulate domestic applications, procurement, data, oversight and consequential uses.

Local adaptation can increase practical autonomy.

Middle powers have more control over how AI is deployed than over how frontier AI is developed.

AI governance strategies for middle powers

Who regulates what—and with which legal tools?

REGULATORY GOAL	APPLICATION
Legal architecture	Promotion statutes; binding horizontal AI laws; voluntary assurance; sectoral regulation; public-sector controls
Regulatory target	Frontier model developers versus downstream deployers and government agencies
Strategic leverage	Market access, public procurement, technical testing, industrial specialization and coalitions
Core limitation	Jurisdiction over domestic uses does not automatically yield access to foreign frontier models

Who is regulated? A second axis of comparison

MANDATORY VS. VOLUNTARY IS ONLY ONE DIMENSION

Frontier developers

Model safety and documentation

Example: South Korea's frontier-scale trigger; EU general-purpose AI rules

High-impact deployers

Risk controls for consequential uses

Examples: South Korea, Vietnam, Peru; EU high-risk applications (from Dec. 2027)

Government agencies

Public procurement, administrative use and impact assessments

Examples: Peru public sector; public-sector governance elsewhere

A country can combine all three targets—and mix binding rules with voluntary guidance.

Middle powers' approaches to AI governance

#	GOVERNANCE APPROACH	DEFINING FEATURE	ILLUSTRATIVE COUNTRIES
1	**Adopt the EU AI Act in form or content**	Horizontal, risk-based AI law; mandatory duties for high-risk applications	Vietnam*; Brazil & Chile (proposed)
2	**Mandatory AI-specific regulation, different model**	Binding duties for high-impact systems or frontier-model safety	South Korea · Peru
3	**Sector-specific regulation through existing law**	Existing privacy, consumer, employment and sectoral regulators apply law to AI	UK · Canada · South Africa* · (All countries)
4	**Voluntary governance and self-regulation**	Model frameworks, standards, evaluations and assurance without general mandates	Singapore · Japan · Australia · India · Malaysia · Indonesia
5	**Government-centered AI governance**	Public-sector procurement, automated decisions and administrative oversight	Taiwan (AI Basic Act, 2025) · Canada
6	**State-directed AI development and governance**	Industrial strategy, investment, infrastructure and selected binding controls	UAE · Saudi Arabia

Categories describe dominant AI-specific approaches, not exclusive legal regimes. Countries may appear in more than one category; existing binding laws can apply even where AI-specific guidance is voluntary.

VIETNAM



A regulatory approach with EU and Chinese elements

Risk classification (Art. 9)	Distinguishes high-, medium- and low-risk AI systems based on foreseeable uses and harms.
Provider classification (Art. 10)	Providers self-classify systems before use; classification dossiers required for medium- and high-risk systems.
Pre-service notice	Medium- and high-risk classifications are notified to the Ministry of Science and Technology via the national AI portal.
High-risk conformity assessment (Art. 13)	**Required before deployment and after significant changes. Systems on the Prime Minister's high-risk list need assessment by a registered or recognized independent body; others may use provider self-assessment.**
What assessment verifies (Art. 14)	Risk management; data governance and quality; technical dossiers and logs; human oversight; transparency; incident handling; and information for regulatory accountability.
Deployers and modifications	Coordinate reassessment when modifications, integration or changes in use increase risk.

Source: Vietnam Law on Artificial Intelligence No. 134/2025/QH15 (passed Dec. 10, 2025; effective Mar. 1, 2026), Arts. 9–14. Conformity assessment and technical implementation are distinct from the EU AI Act.

PERU



RISK-BASED AI LAW AND REGULATIONS: DIFFERENT PUBLIC / PRIVATE DUTIES

Prohibited uses	Certain manipulative, unlawful-surveillance and other harmful uses are prohibited.
High-risk applications	Specified consequential uses: employment, education, health, justice, credit and public benefits.
Transparency	Notify users of AI interaction and explain consequential decisions where required.
Private-sector safeguards	Maintain system records, governance policies, and meaningful human oversight.
Public-sector safeguards	Pre-deployment impact assessments, security audits and responsible-AI governance.
Private impact assessments	Voluntary; not a general EU-style conformity assessment.
Implementation / enforcement	Phased by sector: health, education, justice and finance from Sept. 2026; other sectors through 2029. Existing privacy and sectoral laws also apply.

HOW PERU DIFFERS FROM THE EU

No general provider conformity assessment; private-sector rights impact assessments are voluntary; public-sector duties are more explicit.

Sources: Peru Law 31814; Supreme Decree 115-2025-PCM, Arts. 22–36; EU AI Act, Arts. 9–15, 27, 43.

SOUTH KOREA: TWO DISTINCT REGULATORY TRIGGERS



HIGH-IMPACT USES VERSUS FRONTIER-SCALE SAFETY

CATEGORY A

High-impact AI

Trigger: the system's use and potential consequences.

Examples: healthcare, employment, credit, transportation and other sensitive contexts.

Principal provisions: Arts. 2, 33–35

CATEGORY B

Frontier-scale AI

Trigger: training compute, technological sophistication and potential serious harm.

All three required: over 10^{26} training FLOPs; state-of-the-art AI technology; risk of broad, significant harm to fundamental rights.

Principal provisions: Art. 32

The categories can overlap, but neither automatically implies the other.

SOUTH KOREA: HIGH-IMPACT AI



Similar to EU regulation of high risk systems

Definition (Art. 2)	Designated consequential uses, including healthcare, employment, credit, transportation and critical infrastructure.
Classification (Art. 33)	Operators review whether the system is high-impact; they may seek an official determination.
Risk management (Art. 34)	**Establish and implement measures to identify and manage risks from high-impact AI.**
Explanation and oversight (Art. 34)	**Provide explanation measures, user protection and appropriate human supervision.**
Documentation (Art. 34)	Maintain documentation of implementation of required safeguards.
Impact assessment (Art. 35)	**"Best efforts"** to assess effects on fundamental rights; not an unconditional mandatory assessment.
Conformity assessment	**No general EU-style mandatory premarket conformity-assessment** or independent certification regime.

SOUTH KOREA: FRONTIER-SCALE AI SAFETY



COMPUTE + TECHNOLOGY + RISK: ALL THREE CONDITIONS MATTER

Training threshold (Art. 32)	More than **10^{26} cumulative training FLOPs** under the implementing decree.
Additional conditions	**State-of-the-art AI technology AND risk of broad, significant harm to fundamental rights**.
Lifecycle risk management	**Identify, assess and mitigate safety risks across the AI lifecycle.**
Incident readiness	Establish **monitoring and response systems for AI safety incidents**.
Government reporting	Submit implementation results for safety measures to the Ministry of Science and ICT.
Conformity assessment	No general requirement for independent pre-deployment certification.
Transparency (Art. 31)	Separate AI-use notice and generative-content disclosure rules apply regardless of frontier status.

Voluntary AI governance

Singapore · Japan · Australia · India · Malaysia · Indonesia

COMMON FEATURE	HOW THE VOLUNTARY MODEL WORKS
Government-issued guidance	Governments publish principles, codes, frameworks, or guardrails rather than a general AI-specific licensing regime.
Risk-sensitive safeguards	Organizations are encouraged to identify foreseeable harms and calibrate testing, human oversight, and controls to risk.
Organizational accountability	Guidance emphasizes responsible personnel, documentation, data governance, transparency, monitoring, and incident response.
Voluntary adoption, legal backstop	The AI frameworks are generally nonbinding; privacy, consumer, sectoral, and other existing laws remain enforceable.

KEY DISTINCTION	Singapore and Japan offer especially detailed guidance; Japan's AI Promotion Act (2025) is binding but carries no penalties; Malaysia has proposed a binding AI Governance Bill (2026).

Sources: Singapore PDPC Model AI Governance Framework; Japan AI Guidelines for Business & AI Promotion Act; Australia Guidance for AI Adoption (2025); India AI Governance Guidelines; Malaysia MOSTI National Guidelines on AI Governance & Ethics (2024); Indonesia Circular 9/2023.

GULF STATES: STATE-LED AI DEVELOPMENT & GOVERNANCE



COMMON MODEL: PUBLIC INVESTMENT + GOVERNMENT ADOPTION + SELECTIVE REGULATION

State as builder	Sovereign-backed investment in AI firms, compute, infrastructure and domestic capabilities.
State as user	National strategies prioritize AI adoption in government services and strategic industries.
Governance approach	AI ethics, risk-management and adoption frameworks; not a general EU-style conformity regime.
Binding obligations	Privacy, cybersecurity and sectoral rules remain enforceable; scope differs by country.

COUNTRY VARIATIONS

SAUDI ARABIA	SDAIA-led coordination; AI ethics and risk frameworks; national data governance.
UAE	State-backed AI investment; federal and emirate-level institutions; AI Charter.
QATAR	National AI strategy; QCB requirements for regulated financial institutions.
BAHRAIN · KUWAIT	Government-led strategies and guidance; less extensive AI-specific regimes.

Analytical distinction: a shared development strategy, not a single uniform legal regime.

Sources: UAE AI Strategy 2031 & AI Charter; Saudi SDAIA ethics / risk frameworks (2026); Qatar Central Bank AI Guideline.

The Middle Power Dilemma

THREE GOALS THAT CAN REINFORCE ONE ANOTHER—OR COLLIDE

ACCESS

Use the most capable foreign models and infrastructure to capture AI's benefits.

AUTONOMY

Retain authority to impose domestic legal, cultural and democratic priorities.

RESILIENCE

Reduce dependence on any one foreign provider, government or supply chain.

Practical sovereignty means preserving the ability to choose, govern and switch—not building every layer at home.

A Menu for AI Governance

1. **Predeployment testing**
   Red-teaming, evals, and benchmarking before release
2. **Transparency & documentation**
   Model cards, data sheets, disclosure to counterparties
3. **Incident reporting & monitoring**
   Logging, notification duties, post-deployment watch
4. **Accountability & governance**
   A named owner and board-level oversight of AI risk
5. **Risk-management process**
   A repeatable, framework-mapped assessment cycle
6. **AI system inventory**
   A catalog of deployed systems and their risk tiers
7. **Data governance & provenance**
   Training-data lineage, IP rights, privacy compliance
8. **Human oversight & control**
   Escalation paths, rollback, and kill-switch capability
9. **Third-party / supply-chain risk**
   Governing vendor and foundation-model dependencies
10. **Security controls**
    Access management, robustness, misuse prevention
11. **Auditability & evidence**
    Audit trails and attestations a third party can inspect

EXISTING LAWS THAT GOVERN AI

Adapt existing legal regimes—or enact AI-specific rules?

ELECTIONS & SPEECH
Political deepfakes; campaign disclosure; election interference

CHILDREN
Online safety; age-appropriate design; exploitation

PRIVACY & DATA
Consent; profiling; biometric and sensitive data

CIVIL RIGHTS
Employment; housing; credit; disparate impact

CONSUMER PROTECTION
Deception; unfair practices; AI-generated fraud

INTELLECTUAL PROPERTY
Copyright; trademarks; training and outputs

TORT & PRODUCT LIABILITY
Negligence; defective products; personal injury

HEALTHCARE
Medical devices; clinical decisions; patient safety

FINANCIAL SERVICES
Lending; insurance; securities; model risk

CRIMINAL LAW & CYBERSECURITY
Impersonation; hacking; evidence; malicious use

ADMINISTRATIVE LAW
Government decisions; due process; procurement

ANTITRUST
Market power; collusion; vertical integration

CROSS-CUTTING: Labor • Education • Energy, environment & data centers

Key choice: Update existing law, or create new AI-specific obligations for developers and deployers?



Conclusion

- AI governance constitutes an increasingly important attribute of sovereignty.
- All countries must navigate integration of domestic interests with dependence on foreign technology.

STANFORD LAW SCHOOL | AI INITIATIVE

https://law.stanford.edu/ai-initiative/

Nate Persily, Faculty Co-Director
npersily@stanford.edu

RegLab · LiftLab · CodeX · Rock Center · Rhode Center · LDL · HEAL-AI · AI Learning Hub · LST