


Horizon Scan 002: AI Governance

Artificial intelligence is already being ruled, in detail and at speed, by institutions that were never appointed to the job.



CONTENTS

—	Executive Summary	1	From Declarations to Consequences
2	Where AI Is Actually Governed	3	Rules Written, Rules Enforced
4	Implications for Institutional Leaders	—	Works Cited

RESEARCH AND ANALYSIS

Ashwin Telang
Gloria Chen
Ronan Amir

MANAGING EDITOR

Ashwin Telang

EXECUTIVE EDITOR

David Lovejoy

RESEARCH CONTRIBUTIONS

Hernando Liu

Published by Horizon Search Institute, the institutional author of this Scan.

Executive Summary

What the record since Bletchley shows about where AI is actually governed

AI is being governed through the arrangements that allow it to enter the economy. Firms need electricity and computing capacity, permission to trade, buyers willing to trust their products and a way to bear the cost when those products fail. Each need brings them into contact with institutions able to set conditions on the technology's use. As these conditions accumulate, decisions taken within familiar responsibilities acquire a wider significance. A grid connection helps determine where investment settles. The terms of insurance help establish which precautions a market expects.

These decisions also influence one another. A safeguard required by a major buyer can enter an insurer's assessment and, in time, inform a court's understanding of reasonable care. Firms organise their operations around those expectations, suppliers begin to accommodate them, and practices adopted for a particular purpose become established across a market. Through this process, institutions are helping to settle questions about AI that extend well beyond their immediate responsibilities.

This Scan follows that development from international commitments, through the institutions shaping deployment, to a comparison of legislation and enforcement across seven major regimes. We find considerable agreement on the principles that should guide AI, alongside pronounced differences in who can put them into effect. Understanding those differences requires attention to the relationships through which authority is exercised, the incentives that sustain it and the evidence available to judge its results.

Common principles need institutions that can carry them into practice.

The diplomatic record shows how much depends on what participants have agreed to do together. Within

AUKUS, cooperation on AI and autonomy gains practical purpose from shared operational needs and the requirement for systems to work together. BRICS supports coordination and joint projects through arrangements that leave members considerable discretion over implementation. Such differences matter because the work of implementation begins with decisions about responsibility. Someone must organise the activity, commit resources and establish whether the agreed result has been achieved. A declaration can make that work possible. Its value grows as participants build the means to act on it.

A standard gains influence as other institutions come to rely on it.

Once a buyer makes a contract conditional on meeting a recognised standard, the distinction between voluntary and mandatory becomes less useful to the firm seeking the work. The standards reviewed in this Scan acquire commercial and legal significance through procurement, insurance, regulatory guidance and the assessment of reasonable conduct. Their influence can survive changes in any one of those arrangements, because other institutions continue to find them useful. This helps explain why a standard may endure even as the legal protection attached to it changes. For firms, the lasting value lies in being able to show what precautions were taken, what they accomplished and where their limits remain.

Behind that reliance stands a chain of institutions that the Scan treats as a governance channel of its own. A legislature or board sets the obligation, a standards body writes the auditable requirements, a certifier attests that an organisation meets them, and an accreditor vouches for the certifier. Section 1 traces how voluntary standards acquire legal weight. Section 2 examines the chain directly under "The Fourth Point of Control", which covers ISO/IEC 42001, now certifiable under UK accreditation, and EN

18286:2026, the first European standard written in support of the AI Act.

Dependence gives governing power a geography.


The same process extends across borders. Countries that host computing infrastructure can set conditions on its connection and operation, while countries that rely on imported capacity must also contend with decisions made by foreign suppliers

and authorities. Trade controls, in turn, depend on customs agencies, licensing ministries and firms along the supply chain to make them effective. These relationships connect the Scan’s physical and geopolitical layers to its account of liability and assurance, where payment networks, insurers and courts shape the terms of use. Beneath them sits a further dependence on information. Institutions able to observe AI’s adoption and effects increasingly supply the evidence through which others understand it.

EXHIBIT 1

Suppliers lead in four of seven regimes and tie with the state in the United States. The state leads in China and courts in Canada.

Share of HSI’s assessed constraint score, by actor and regime (%)



Percentages reflect HSI’s assessment of binding instruments, not measured shares of costs or delays. The US state score primarily reflects federal export controls on advanced chips, largely affecting sales abroad. The US tie does not establish that domestic AI regulation and supply constraints impose equal burdens on US businesses.

Note: Each actor’s share is its assessed score (0–9) divided by the regime’s combined score. The label on the right names the leading instrument in that regime.

Source: Horizon Scan 002, Section 2

Decisions within one remit can distribute costs far beyond it.

A grid operator deciding how to accommodate new demand is also helping to determine which users bear the cost of expansion. An insurer specifying acceptable precautions can influence which firms can afford to enter a market. Each decision may be reasonable within the institution’s responsibilities, yet its wider effects may fall to people who had little part in making it. As these decisions become more

consequential, public policy must attend to their combined effects, including the distribution of costs and the opportunities to challenge a decision. The capacity to contain harmful AI behaviour also requires separate attention. Conditions on access can influence deployment while leaving unresolved what anyone can do once a system is acting.

Existing law carries most of the enforcement in this record.

The comparative register makes the role of established institutions especially clear. Of 38 selected enforcement actions from January 2023 through early August 2026, two relied on AI-specific law, both in China. Consumer protection and data protection accounted for 27 of the remaining cases. These are actions addressing recognisable problems, including misleading claims and misuse of personal information, through powers that already exist. The United States and Canada illustrate how such powers can support enforcement alongside limited dedicated AI legislation. In the European Union, extensive legislation is accompanied by staged implementation. China combines dedicated requirements with established enforcement institutions. The result is considerable variation in how a stated commitment becomes an obligation a firm must meet.

2 of 38

enforcement actions in the register relied on AI-specific law. Both cases were in China.

The passage of a law leaves consequential choices unsettled.

Implementation takes time, particularly where authorities need new expertise, procedures or resources. The Scan also records postponements, revisions and withdrawals that change what institutions are expected to enforce. These decisions deserve attention in their own right. Every extension of a transition period prolongs the role of the arrangements already shaping deployment, giving firms further reason to organise around them. What begins as an interim practice may become costly to replace once contracts, systems and expectations depend on it. Assessing a regime therefore requires an account of the capabilities being built and the choices being made after legislation passes. Its age alone tells us relatively little about where it is heading.

The evidence available shapes the policy that becomes possible.

To judge how AI is changing work, policymakers need to know whose experience the evidence represents and what else might explain the changes observed. Payroll processors, hiring platforms and model providers can offer timely information, each drawn from a particular population and collected for a particular purpose. Public statistics provide a broader basis for comparison, though several instruments examined in the Scan are weakening or lagging. Where that capacity is already thin, the experience of firms visible to private data providers can acquire disproportionate influence over the public account. Better measurement requires connecting these sources while preserving their differences, so that uncertainty about adoption, employment or productivity remains visible in the decisions built upon them.

. . .

Taken together, these findings describe a redistribution of governing power through the relationships on which AI depends. For institutional leaders, the practical task is to understand how those relationships affect the ability to operate, intervene and recover when conditions change. For policymakers, it includes deciding how responsibilities spread across several institutions can serve a coherent public purpose, especially where the costs of individual decisions are borne elsewhere.

The comparative scores offer a starting point for that inquiry. They describe legislation and enforcement within the Scan's scope, leaving open the larger question of how well either serves society. A narrower distance between commitments and action can result from stronger implementation or from reduced commitments. Judging the outcome requires evidence about the protections achieved, the costs imposed and the people affected. The governing arrangements taking shape around AI will ultimately have to answer to those tests.

1

SECTION 1

From Declarations to Consequences

When summit language converts into practice, and what it costs to make it do so

Across the post-Bletchley summit sequence, states have converged on the same vocabulary and diverged on nearly every mechanism that would give it consequence.

Since the first AI Safety Summit at Bletchley Park, successive international forums have sought to build common ground around the governance of advanced AI. States now routinely endorse the same terms, from transparency and accountability to human oversight. But the convergence is shallower than it appears. Most multilateral commitments fail to convert for a structural reason: they were never designed to carry legal or operational cost in the first place.

This section defines conversion as the point where governance language starts to carry a real cost: when an endorsed principle such as transparency or human oversight hardens into a procurement requirement, a statutory defence, or an enforceable penalty. A commitment can remain a declaration for years. It begins to shape conduct when a specific actor faces an immediate cost for failing to honour it.

Conversion is a cost-attachment event. Where misalignment carries an immediate price for a named actor, governance language hardens. Where it does not, the language accumulates and the enforcement structure underneath it stays where it was.

Three cases establish the pattern across three tiers of enforcement. Tier 1 is binding law with enforcement infrastructure still being assembled: the EU AI Act, where the legal framework exists while implementation remains uneven across Member States. Member States were required to designate market surveillance and notifying authorities by August 2025. A June 2026 implementation tracker counted nine of twenty-seven as having clearly designated both, twelve as partial, and six as having designated neither. Tier 2 is cost-driven operational convergence without formal law: AUKUS Pillar 2, where operational stakes force alignment even in the absence of a treaty. Tier 3 is vocabulary without a shared mechanism: BRICS, where no actor bears a direct cost for non-alignment and declarations remain declarations. A fourth finding runs outside the diplomatic track altogether: U.S. state legislatures and common law courts are attaching legal consequences to voluntary technical standards that were never negotiated as international agreements, and detaching them again.

EXHIBIT 2

Five major AI governance forums in three years.
None created a binding enforcement mechanism.

Selected international forums, November 2023 to July 2026; chronological order, schematic spacing



Note: New Delhi is included in the count. Its declaration advances voluntary, non-binding cooperation. The UN Dialogue is a governance forum, not a summit or a treaty-making body. The count covers the five forums shown.

Sources: UK Government; Seoul and Paris summit statements; Government of India, New Delhi Declaration (21 February 2026); United Nations. Supporting references and correction record: p. 46.

I • The Cost-Driven Case: AUKUS Pillar 2

The AUKUS partnership announced its first concrete AI-and-autonomy deliverable, a signature project on uncrewed undersea vehicle payloads, on 30 May 2026, nearly five years after AUKUS Pillar 2 launched in 2021. Capability delivery is not expected before 2027. UK Defence Secretary John Healey's assessment of the interval was direct: the partnership had "talked too much and delivered too little." (ASPI/The Strategist, June 2026)

For years before that announcement, independent analysts described AUKUS Pillar 2 as "a solution in search of a problem," with eight workstreams spanning AI, hypersonics, quantum, and cyber, widely seen as too diffuse to produce visible output. (USSC policy brief, January 2026) The timing of what finally concentrated attention coincides with the April 2026 exposure of a Russian cable-mapping operation off British shores and the recovery of a Chinese UUV in the Lombok Strait. The project that

emerged targets seabed protection and anti-submarine warfare, precisely the missions where misalignment between partners carries immediate operational consequence. The causal sequence is not definitively documented. The timing is consistent with a cost-logic reading, but the public record does not establish it.

Five years, under a trilateral defense pact with direct battlefield stakes, to produce one deliverable, and that is the fast case in this section. AUKUS Pillar 2 marks the upper bound of what cost-driven conversion currently looks like: narrow, delayed, and tied to a specific operational scenario.

The fastest conversion in this section took five years and one costed scenario. Everything measured against it is slower.

The same cost is visible closer to home, and has not yet converted. Cables carry roughly 99% of the UK's international data traffic and underpin about £1.4

trillion in daily transactions. Around 75% of UK-US cable capacity is in two cables landing at Bude, Cornwall. (JCNSS, September 2025, paras. 10, 51.) In April 2026, British armed forces exposed a covert Russian operation mapping this infrastructure, with the Defence Secretary describing a submarine deployed as a decoy while other vessels surveyed the cable network. (Geopolitical Monitor, April 2026) The government announced plans to consult on replacing 140-year-old subsea cable legislation on 29 May 2026, after the Joint Committee on the National Security Strategy called the existing posture "too timid." (GOV.UK, May 2026) The cost that moved AUKUS Pillar 2 toward a deliverable is the same cost UK domestic legislation has not yet fully absorbed. The operational exposure is immediate; the legal conversion is still lagging.

II • The Null Case: BRICS and the Multilateral Track

BRICS functions as a working example of non-binding cooperation: useful for vocabulary, coordination, and shared projects, but structurally weak for enforcement.

BRICS adopted its first standalone AI statement, the BRICS Leaders' Statement on the Global Governance of Artificial Intelligence, at the 17th Summit in Rio de Janeiro on 6 July 2025, with eleven countries listed by the Brazilian BRICS presidency: Brazil, Russia, India, China, South Africa, Saudi Arabia, Egypt, UAE, Ethiopia, Iran, and Indonesia. (Brazilian Presidential Office, January 2025; BRICS Brazil Presidency, July 2025) The statement calls for governance that must "mitigate potential risks and meet the needs of all countries, including those in the Global South," with the UN positioned "at its core." (BRICS Leaders' Statement on AI Governance, July 2025) More than a year on, BRICS has built no enforcement machinery beyond the language of the July 2025 package itself. The Data Economy Governance Understanding adopted alongside the statement carries no audit body, no shared compliance test, and no penalty for non-compliance. (IDEAS-BRICS)

The vocabulary convergence masks structural divergence that goes deeper than policy preference. A 2026 comparative study of BRICS+ members found five distinct regulatory logics operating under one

declaration. China regulates by binding sector-specific rules under strong state control. Brazil runs a rights-based participatory model built on democratic oversight. India anchors its approach in digital public infrastructure. Russia frames AI policy through sovereignty. Indonesia keeps its framework deliberately flexible, combining soft law with sandbox experimentation. (Soeparna & Sarli, SSRN, May 2026)

Eleven states signed the same declaration and went home to five incompatible regulatory logics. The signature was the cheapest part of the process.

Jaqueline Pigatto of Data Privacy Brasil attended the UN Global Dialogue in Geneva as a civil society delegate in July 2026. The Dialogue, she noted, was deliberately designed to produce non-binding outcomes. Its function was to establish minimum consensus on vocabulary and principles, rather than negotiate binding agreements. On the state of BRICS coordination:

"We don't think we are seeing coordinated efforts from the BRICS countries to make things actually happen." (Pigatto, interview, July 2026)

A joint statement from the Global Digital Justice Forum and the Global South Alliance, issued around the Geneva Dialogue, framed the same observation from inside the process: "The many summits and conversations about AI governance have failed to tackle these core issues." (GDJF/GSA Joint Statement, APC, July 2026) The assessment comes from actors who attend these processes, draft the language, and watch the outcomes.

Intan Soeparna, whose comparative study of BRICS+ regulatory profiles is among the most granular available, put it this way:

"BRICS+ is far more likely to succeed in building cooperative ecosystems, through standards, sandboxes, and shared infrastructure, than in creating a binding enforcement mechanism." (Soeparna, interview, July 2026)

Cooperation can occur around specific interests. Without institutional mechanisms, legal harmonization, external verification, and political will to accept collective enforcement over unilateral control, that cooperation stays at what she called "a coalition of convenience rather than a regulatory union." (Soeparna, interview, July 2026)

Asked whether countries might cooperate on individual projects while never converging on a shared governance model, Soeparna was direct: "They can cooperate, but it will be very difficult to converge." Technical cooperation follows national interest, and national interest is supported by philosophical framework. China's dual identity as rule-shaper and national-interest protector, Russia's sovereignty-first orientation, Brazil's rights-based democratic model, and Indonesia's adaptive positioning create what she described as a divergence of philosophical and legal frameworks that "will clash if there is a single idea that is not accommodating everything." She pointed to ASEAN's digital trade provisions as the working precedent: flexible by design, because they had to accommodate capacity disparities among members. That flexibility

made the framework workable and left it looser than enforcement requires. (Soeparna, interview, July 2026)

A coalition of convenience rather than a regulatory union. Cooperation is available on projects; convergence is not available on rules.

More than a year of inaction reflects the design working as intended. A process built to produce non-binding declarations has produced them. Participation is cheap precisely because obligation is optional. No member bears a direct cost when another fails to act on the declaration's language, and there is structurally little to convert.

EXHIBIT 3

Governance language turns into rules only where falling out of line costs someone something.

Three governance pathways compared on what gives them force, status as of August 2026

	Binding law	Cost of falling out of line	Shared enforcement	Where it stands
EU AI Act Tier 1 · Binding law, enforcement still being assembled	Yes In force since 2024; high-risk duties deferred to Dec 2027	Yes Fines and market exclusion for non-compliance	Decentralised 9 of 27 states had named both required authorities by June 2026	Binding law with uneven national enforcement
AUKUS Pillar 2 Tier 2 · Operational convergence without a treaty	No Defence pact, no AI treaty	Yes Battlefield consequences for misaligned partners	No No formal legal enforcement	One deliverable after five years, once a costed scenario emerged
BRICS Tier 3 · Shared vocabulary, no shared mechanism	No Leaders' statement, July 2025	No No member pays a price when another fails to act	No No audit body, compliance test or penalty	Eleven signatories, five incompatible regulatory logics

Note: Categories are qualitative classifications based on the cases analysed in Section 1, not numerical scores.
Source: European Commission; AUKUS official materials; BRICS Brazil Presidency; Horizon Scan 002, Section 1

The UK offers a domestic version of the same conversion problem. There is no dedicated UK AI statute; AI is governed contextually through

whichever sectoral law, whether financial services, healthcare, or consumer protection, already applies to the underlying activity. (House of Commons

Library, 2026) The February 2025 renaming of the AI Safety Institute to the AI Security Institute illustrates how that model has evolved: governance vocabulary has continued to shift across safety, security, and growth, while the institutional model remains centred on technical evaluation and stops short of broad regulatory enforcement. (DSIT and AI Security Institute, February 2025; Glacis, April 2026)

Eighty-one percent of UK finance and real estate firms cited unclear regulation as their top barrier to AI adoption, per government-commissioned research cited in the Interim Government Response to the AI Champions' AI Adoption Plans. The comparison with the EU and BRICS clarifies why. In the EU, binding law exists while enforcement infrastructure is still being assembled: the high-risk deadline moved to December 2027 under the AI Omnibus, and the delay reflects gaps in the standards, conformity assessment, and national enforcement infrastructure needed to make the law operational at scale. BRICS presents the opposite problem: common language without a shared mechanism for audit, compliance, or penalty. The UK follows a third path, relying on existing sectoral regulation while a developing assurance and accreditation architecture begins to give technical standards commercial consequence outside a comprehensive AI-specific statute. Section 2 examines that commercial route in detail.

A law on the books and a law that bites are two different instruments. The Omnibus is the gap between them made visible.

III • The Real Mechanism: Standards-to-Statute Conversion and Its Instability

The conversion from talk to binding rule that the diplomatic track has not produced is happening through a different channel. It runs through state legislatures, procurement expectations, enforcement guidance, and common law courts attaching legal consequence to voluntary technical standards that were never negotiated as international agreements.

NIST's AI Risk Management Framework is voluntary by design. The FTC, CFPB, FDA, SEC, and EEOC already cite it in enforcement guidance, and federal contractors face growing pressure to demonstrate alignment. (GAICC, April 2026) The mechanism that

converts this from informal expectation into legal consequence is traceable and specific: Texas's TRAIGA gives NIST-aligned internal review a role in statutory liability protection, creating a pathway through which a voluntary technical framework can affect legal exposure. A voluntary standard becomes legally consequential when a legislature attaches liability protection to compliance with it. (Maro, March 2026)

ISO/IEC 42001 is following a parallel path. It is increasingly expected in enterprise procurement, though it sits outside the EU harmonisation process and does not by itself establish conformity with the AI Act's high-risk requirements. It is precisely because ISO/IEC 42001 does not map directly onto the AI Act's quality-management requirements that a dedicated European standard was needed, and EN 18286 was written specifically for that purpose. (CEN-CENELEC, 2026)

A voluntary standard gains its force from the consequence an institution attaches to it, and loses that force the moment the institution takes the consequence away. No treaty is negotiated, and no summit is convened.

The mechanism is also more fragile than it appears. Colorado was the first state to codify this conversion into law under SB 24-205 in 2024. On 14 May 2026, Colorado repealed and replaced that statute with SB 26-189 before it ever took effect. The new law dropped the NIST/ISO safe harbor entirely. The repeal passed 34-1 in the Senate and 57-6 in the House, with almost no opposition on either side. (Glacis, June 2026)

Joe Braidwood, CEO of Glacis, whose firm sells AI compliance infrastructure and therefore has a commercial view on these questions, tracked companies through the transition. For serious companies, he argued, NIST was never the point of the safe harbor itself. "It was why their lawyers approved budget." Removing the safe harbor makes the financing case harder, but the new law still requires controls that rely on much of the same underlying compliance infrastructure. "NIST didn't get voted out," Braidwood said. "It went from being a legal defense to being the plumbing that makes those things possible." (Braidwood, interview, July 2026)

Braidwood read the near-unanimous margin as a verdict on regulatory design: legislators rejected the EU AI Act model, built on duty-of-care obligations and risk-classification machinery, while retaining the transparency and disclosure requirements. He expects that narrower model to prove more politically durable at the state level. The forecast comes from someone with a stake in a particular compliance architecture and warrants triangulation against independent analysis. (Braidwood, interview, July 2026)

A new federal dimension entered the picture in April 2026 when xAI sued the Colorado Attorney General in a case captioned xAI v. Weiser (filed 9 April 2026),

and the DOJ intervened on xAI's side on 24 April 2026. It was the first time the DOJ's newly established AI Litigation Task Force moved to invalidate a state AI law in federal court. The Colorado AG agreed not to enforce pending rulemaking.

“State legislatures will tell you what can pass. Federal courts are about to tell us what survives.” (Braidwood, interview, July 2026)

The conversion mechanism runs in both directions. A legislature can attach legal cost to a standard, and it can detach it before the statute ever operates.

EXHIBIT 4

Colorado gave firms that follow voluntary AI standards legal protection in 2024, then withdrew it in 2026 before the law took effect. Texas's version still stands.

Two US state laws that shield firms following recognised AI standards from some legal claims, from signing to August 2026



Note: A safe harbour protects a firm from certain legal claims if it can show it followed a recognised framework, such as the NIST AI Risk Management Framework. xAI challenged Colorado's law in federal court on 9 April 2026, and the US Department of Justice joined on its side on 24 April. The legislature replaced the law in May by votes of 34-1 in the Senate and 57-6 in the House.
Source: Colorado General Assembly; Texas Legislature; US Department of Justice; xAI v. Weiser; Horizon Scan 002, Section 1

NIST AI RMF is beginning to enter arguments about the standard of care in AI-related liability, creating a second potential conversion channel outside the legislature. Alignment with a voluntary standard may therefore begin to function as evidence of "reasonable conduct" through common law,

regardless of what any state has legislated. Braidwood grounded this in the 1932 TJ Hooper case: when new technology emerges, common law has long required organizations to pursue safety measures that make that technology reasonably safe in context. AI's inherent risks, from hallucination to

prompt injection, exist across deployments of transformer architecture, independent of jurisdiction or statute. The common law obligation to manage them does not wait for Congress. (Braidwood, interview, July 2026; Future of Privacy Forum, March 2026; TJ Hooper, 60 F.2d 737, 2d Cir. 1932) This claim, that courts are already treating NIST alignment as the standard of care, is presented here as an emerging trend; the caselaw is accumulating and not yet consolidated.

Three distinct, non-diplomatic conversion pipelines are now running simultaneously: Texas's statutory safe harbor, durable so far; Colorado's, reversed before it ever operated; and the judicial standard-of-care channel, accumulating without legislative action. None runs through a summit, and all three move faster and less predictably than the Bletchley-to-Geneva sequence. In Braidwood's formulation, all three are evidence problems. In every case, the operative question is whether anyone can independently verify what an organization actually did, regardless of which framework it claimed.

Standards-to-statute conversion is real and unstable. It gives voluntary standards legal consequence without giving operators a durable settlement.

Texas illustrates one route by which a voluntary technical standard can acquire legal consequence. The UK has developed substantial standards activity of its own, from DSIT's Blueprint and sector regulator principles to an emerging AI Growth Lab model, yet it has no comparable AI-specific statutory mechanism connecting those frameworks to liability protection. An advisory AI Growth Lab is now operating in legal services within existing regulatory frameworks, while the broader Growth Lab model involving targeted regulatory modifications remains tied to forthcoming legislation.

That absence of a statutory hook has not left the standards without consequence. BSI holds UKAS accreditation as the first UK certification body for ISO/IEC 42001, while DSIT's September 2025 roadmap treats third-party AI assurance as a market the state should help build. Procurement is increasingly asking suppliers to evidence AI risk management and human oversight, a control point techUK calls one of the most consequential in responsible AI adoption, and one through which standards such as ISO/IEC 42001 may acquire

commercial consequence. Investors are beginning to require assurance evidence in due diligence, and insurers to factor it into coverage. This is a mechanism developed further in Section 2. (techUK, 2026)

The sequencing is deliberate. DSIT's roadmap considered professionalisation, process certification, and firm accreditation as routes to improving quality in the assurance market, and prioritised professionalisation in the near term, judging that the market was not yet mature enough for the latter mechanisms to realise their full value.

Professionalisation does not itself create accountability, though some enforcement already exists at the level of the credential. Emma McGuigan, Chair of the AI Assurance Stakeholder Consortium, confirmed that BCS, the Institute for IT Professionals, can accept a complaint, investigate it, suspend someone from its register, and remove the chartered title. Unlike the position for lawyers and medics, this does not leave the person unlicensed to work, because that level of structure has never been created in the profession. (McGuigan, correspondence, September 2026)

Financial services shows the same split between sectoral regulation and AI-specific consequence. The FCA has explicitly chosen not to write AI-specific rules, relying instead on the Consumer Duty and the Senior Managers and Certification Regime, a framework that assumes a named human is accountable for each decision. The FCA's own Mills Review conceded in January 2026 that agentic AI may become "capable of independent decision-making" by 2030, but called it "premature" to change the rules, while a House of Commons Treasury Committee report warned the posture "risks serious harm to consumers and the wider financial system." (Freshfields, 2026; Baker McKenzie, 2026)

Over 75% of UK financial services firms now use AI, yet the same agentic systems documented at JPMorgan and Goldman are being deployed under a different accountability structure. The UK has no equivalent of California's Transparency in Frontier Artificial Intelligence Act (TFAIA) penalties or the EU AI Act's Article 14 human-oversight requirement, which for Annex III high-risk systems applies from December 2027, leaving existing sectoral regulation and the emerging assurance market to carry more of

the governance burden. (FCA, Mills Review, January 2026; California SB 53; Regulation (EU) 2024/1689, Art. 14)

Synthesis

Conversion from talk to rule is a cost-attachment event, not a maturation process.

AUKUS Pillar 2 demonstrates the ceiling: a trilateral defence pact with battlefield consequences for failure still took five years to produce one deliverable, with progress concentrating around a specific operational scenario in which the costs of non-alignment were immediate. BRICS demonstrates the floor: more than a year and a deliberate UN process in Geneva have produced a cooperative ecosystem without the rules to bind it, because the declarations themselves carry no legal obligation.

A BRICS-wide binding regime would require delegated authority, substantial legal harmonisation, and mechanisms for external verification. Rarer than any of these is the political will to accept collective enforcement over unilateral control. None of those conditions currently exists within BRICS+, while the UN Global Dialogue itself is designed as a non-negotiating forum for building shared understanding rather than producing binding agreements.

Even Tier 1 is provisional, and August was when it was tested. Regulation (EU) 2026/1744 moved the Annex III high-risk requirements to 2 December 2027, with requirements for high-risk AI systems embedded in regulated products running to 2 August 2028. Article 50 transparency obligations began applying on schedule on 2 August 2026, while GPAI obligations had already begun applying a year earlier. The high-risk deferral followed sustained industry pressure to pause parts of the timetable. The legal schedule is now differentiated, not simply delayed, and that differentiation sits on top of an uneven national enforcement architecture. Article 50 makes that fragmentation concrete: its transparency obligations apply across the Union, while

enforcement rests primarily with national market surveillance authorities.

Some of the clearest visible conversion mechanisms in AI governance today run through statehouses and courtrooms: state legislatures and common law courts attaching and detaching legal consequence to voluntary technical standards that were never negotiated as treaties. Two features of this moment distinguish it from the standard implementation-gap narrative. The federal government has now intervened on the side of a private company against a state AI law, rewriting the risk calculus for any operator building compliance programs on state-level safe harbors. And civil society from the Global South is documenting from inside these processes what this section argues from outside: non-binding design functions as a feature of multilateral participation, not an accident of it.

The newer assurance channel complicates the distinction between voluntary and binding governance. Colorado shows how quickly a legislature can detach a legal incentive from a technical framework without eliminating the framework itself. Braidwood's account helps explain why: the underlying evidence and control infrastructure can survive even when the legal incentive changes. The UK's developing accreditation and assurance architecture shows how that infrastructure may acquire a separate route to consequence through procurement and market access. That route raises the same question this section asks of statutes and courts: where does the cost attach? Commercial certification can matter through procurement or insurance, while professionalisation already carries some consequences at the credential level without yet reaching licensure-level control over the right to practise. Section 2 examines how far that route has developed.

For operators calibrating compliance to the wrong enforcement tier, the mispricing compounds with every agentic deployment that assumes a legal shield no longer in force.

2

SECTION 2

Where AI Is Actually Governed

Four layers where the binding constraints on AI now form, and who does the binding


The typical map of AI governance is a map of laws: the EU AI Act's deadlines, the American executive orders, the UK's regulator letters, China's filing regimes. Draw the map as a map of gaps instead and a different territory materializes.

The binding constraints on AI now sit in four layers, and in each of them the formal rulemakers arrive late, govern sideways, or cannot see the thing they are meant to govern.

EXHIBIT 5

Drawn as a map of laws, AI governance is a European and Chinese story. Drawn as a map of who actually binds AI today, it is mostly a story about suppliers.

Seven regimes: how much AI-specific law each has written, and which type of actor holds the strongest binding constraint on AI, August 2026



Note: Shading reflects how much AI-specific law each regime has written, from the Scan's comparative scores. Each marker shows the actor whose instrument binds most tightly in that regime today; suppliers means the utilities and firms that control grid connections, data-centre sites and chip supply.

Source: Horizon Scan 002, Sections 2 and 3; Natural Earth (map geometry)

The pattern across all four is that gaps do not stay empty. Where statute is absent or slow, governance accretes to whoever holds operational control: grid operators and utility commissions in the physical layer, customs agents, prosecutors, and ministries in the geopolitical one, card networks, insurers, certifiers, and courts in the accountability layer, and private data vendors in the evidence layer beneath them all.

The instruments of this improvised regime are subtle: a connection queue, a customs ledger, a kill

switch, a certificate, or an underwriting file. Each is doing work that a statute was expected to do, on a faster clock, with narrower accountability. The question for the next several years is less whether AI will be governed than by whom, with which instruments, and with what view of the system.

The published analysis that follows concentrates on seven regimes: the EU, the US federal system, the UK, China, South Korea, Singapore, and Canada. The wider sixteen-jurisdiction record remains in the evidence base and is drawn on selectively where a

single jurisdiction offers the clearest example of a supplier, a state, a court, or a certifier becoming the real point of control. Layers and points of control do different work throughout. The layers describe where governance forms, while the points of control describe who can make it binding.

Four points of control recur across the layers. The supplier binds through contract, the state through statute and ownership, and the court through judgment, while the certifier binds through standards that bodies such as ISO, IEC, CEN, and CENELEC write and that accreditors stand behind. Those standards bodies belong on the map as governance actors in their own right, and the section examines their machinery, including ISO/IEC 42001 and EN 18286:2026, under “The Fourth Point of Control,” which follows the evidence layer.

Layer 1. The Physical Layer

AI governance is quietly becoming utility governance: the constraints that bind are material, set in supply chains and substations before any statute speaks.

Governance of AI's physical substrate is forming in auction rules and interconnection queues, which already set the technology's price and pace. Laws fall years behind, leaving entire load bands with no owner of oversight.

The tightest chokepoint is high-bandwidth memory. One supplier, SK Hynix, has been effectively sold out

through 2026. That single order book throttles the supply of frontier accelerators more tightly than any export rule on the books. The dependency runs deeper than memory. The International Energy Agency notes that China refines 99 percent of the world's gallium, a metal essential to advanced chips and power electronics, and projects that data center demand alone will absorb more than a tenth of current global supply by 2030.


In China, a single national permitting decision, a single refinery outage, a single memory fab's order book each now moves the frontier more than most statutes do.

The energy numbers give the layer its scale. Global data center electricity consumption ran to roughly 415 terawatt-hours in 2024 and is projected by the IEA to more than double to about 945 TWh by 2030, slightly more than Japan's entire present consumption. In the United States, data centers account for nearly half of all electricity demand growth to 2030, and by decade's end the country is set to consume more electricity processing data than producing aluminum, steel, cement, chemicals, and every other energy-intensive good combined. The IEA estimates that around 20 percent of planned data center projects risk delay on grid constraints alone; transmission lines take four to eight years to build in advanced economies, and waiting times for transformers and cables have doubled in three years. Federal review tends to trigger around 100 megawatts and many state laws at 10, which leaves a wide band of facilities with no clear owner of oversight.

EXHIBIT 6

Data centres are on course to use more than twice as much electricity in 2030 as they did in 2024, more than Japan uses today.

Global data-centre electricity consumption, terawatt-hours per year



Nearly half

of all growth in US electricity demand to 2030 comes from data centres

1 in 5

planned data-centre projects risk delay because the grid cannot connect them in time

Note: The 2030 figure is the International Energy Agency's base-case projection. The Japan comparison is approximate: the IEA describes 945 TWh as slightly more than Japan's present total consumption.

Source: International Energy Agency, Energy and AI (2025); Horizon Scan 002, Section 2


Where no statute governs, the auction does. The clearest demonstration is PJM, the grid operator for 65 million people across thirteen states, where the capacity price (the amount paid to generators simply to exist and be available at peak) rose from \$28.92 per megawatt-day for 2024-25 delivery to \$269.92 for 2025-26, then to the regulatory cap of \$329.17 for 2026-27, a roughly tenfold repricing in two cycles.

PJM's independent market monitor attributes 63 percent of the 2025-26 increase, about \$9.3 billion, to data center demand. The July 2026 auction cleared at the cap again, locking elevated prices into 2028-29, and the December 2025 auction produced a record \$16.4 billion capacity bill while still missing its reserve target.

EXHIBIT 7

In America's largest power market, the price paid to keep generators available rose more than tenfold in two years. Data centres drove most of the increase.

PJM capacity auction clearing price, by delivery year, \$ per megawatt-day



Note: PJM serves 65 million people across thirteen states. The capacity price is what generators are paid to be available at peak. The 2027-28 and 2028-29 auctions (December 2025 and July 2026) cleared at the regulatory cap, shown at its 2026-27 level.
Source: PJM Interconnection; PJM Independent Market Monitor; Horizon Scan 002, Section 2

Residential customers in Washington DC, Maryland, and Ohio are already paying \$10 to \$21 more per month, with advocates projecting \$70 by 2028 absent reform, costs externalized to households one billing cycle at a time. Meanwhile the interconnection queue, the waiting line to plug new generation into the grid, has stretched from under two years in 2008 to more than eight, and developers cancelled 38 gigawatts of projects in 2025 alone. The auction rules and the queue design, both written by a grid operator and its regulator, are currently allocating the costs and the pace of American AI more directly than any AI law. That allocation is practically governance, whether or not anyone calls it that.

The same American instrument, a connection decision serving as a surrogate for a siting statute, now recurs across the advanced economies, and in each the regulator has quietly become the authority that decides where and whether AI compute can land.

Inside the EU, the clearest single instance sits in Ireland, where data centres consumed 22 percent of all metered electricity in 2024, up from 5 percent in 2015, more than the share taken by all urban households combined, according to the Central Statistics Office, and the share reached 23 percent in 2025. The body that governs that load is the Commission for Regulation of Utilities. On December 12, 2025, the CRU published its final Large Energy User connection policy, issued as a Section 34 direction to the system operators under the 1999 Electricity Regulation Act, which requires new data centres above 10 MVA to bring their own dispatchable generation and to procure renewable electricity equal to at least 80 percent of annual demand, and lets operators cut a facility's import capacity or move it to non-firm status if the associated generation underperforms. A utility regulator now sets the terms on which AI capacity connects to a national grid, and it acknowledges 5.8 GW of pending data-centre demand it must ration; the Oireachtas has yet to legislate on the question.

The world's new data-center siting authorities were never elected to zone anything. They are electricity regulators, and their instrument is the connection contract.

Singapore made the rationing explicit and turned it into a tender. After a de facto moratorium on new data centres from 2019, the Infocomm Media Development Authority reopened capacity through a Green Data Centre Roadmap in 2024 and a second Call for Applications on December 1, 2025, allocating at least 200 MW to operators that clear stringent efficiency and sustainability bars, including a power-usage-effectiveness target at or below 1.3 and at least 50 percent green power. Capacity is awarded by ministry selection, on criteria the ministry writes.

South Korea fuses the grid constraint to national industrial policy. The Special Act on the Promotion of Distributed Energy, effective June 14, 2024, imposes a power-system impact assessment on any data centre drawing 10 MW or more and requires facilities in metropolitan Seoul above 200,000 MWh a year to source a rising share of distributed energy. In practice the binding actor is KEPCO, the state utility, whose power-supply confirmation for a Seoul data centre has lengthened from two or three months to about twelve, alongside a policy restricting additional supply in the capital region and pushing new load to the provinces. The same national interest that rations Seoul's grid runs the other direction on memory: SK Hynix and Samsung hold the global HBM chokepoint that constrains Layer 1 everywhere else, so Korea both supplies the substrate and rations its domestic siting.

From the wider set, the Gulf states supply the clearest example of the state itself as the point of control, because there the state owns the compute outright. The 5GW UAE-US AI Campus in Abu Dhabi, anchored by G42's 1GW Stargate UAE cluster with OpenAI, Oracle, Nvidia, Cisco, and SoftBank, was announced in May 2025 under a bilateral US-UAE AI Acceleration Partnership, and the chips flow only through a "Regulated Technology Environment" that G42 built to satisfy US security conditions, including divestment from Chinese technology. Sovereign wealth sites the compute, and the siting terms are written jointly by a Gulf state and Washington; zoning law never enters the file.

Where the substrate concentrates, rulemaking follows: the countries that host compute are writing the terms for power, water, and siting. The countries without it will inherit those terms as importers.

The layer also has a geography, and the geography is lopsided. The World Bank's Digital Progress and Trends Report 2025 finds that high-income countries held 77 percent of global co-location data center capacity as of mid-2025, while low-income countries held less than 0.1 percent.

Half of the world's secure internet servers sit in the United States; on a per-capita basis the US has 200 times the servers of a typical middle-income country and 20,000 times those of a low-income one. The same report counts 2.6 billion people, a third of humanity, still offline in 2024, with per-capita data traffic of 1,400 gigabytes a year in rich countries against 5 in the poorest. For most of the world, the physical layer of AI is an import. The governance conversation about siting, cooling, water, and grid interconnection is happening only in the places that have the substrate, *while the places that lack it inherit whatever terms emerge*. That asymmetry surfaces again in Layer 2, because a country that cannot host compute must buy it as a service, and a service can be switched off.

In the Global South, the governing instrument changes from the kilowatt to the liter, and the venue changes from the auction to the court. Cooling is where AI meets the water table. The chips that train and run frontier models convert almost all the electricity they draw into heat, and the cheapest way to carry that heat away at scale is evaporation, which consumes fresh water rather than merely borrowing it. A campus of the size now under construction can evaporate millions of litres on a hot day. Where the grid is the binding constraint in the North, the aquifer is the binding constraint in much of the South, so the governing instrument shifts from the kilowatt to the liter and the venue shifts from the auction to the court.

The wider set's clearest example of a court as the point of control is Chilean, and the court is environmental: Google's Cerrillos data centre near Santiago, first approved in 2020 and designed to draw 7.6 million litres of potable water a day for cooling, was sent back for reassessment by the Second Environmental Court in 2024 over its reliance on a

drought-stressed aquifer, and the project was suspended and redesigned to air cooling. Community water governance forced the redesign before any AI statute could reach the question. Uruguay's Google project drew the same water objections.

Shaolei Ren, associate professor of electrical and computer engineering at the University of California, Riverside, argues that the gap is one of disclosure rather than authority. Asked what a jurisdiction should require first, he points to standardized reporting of both annual and peak draw, because "peak water and electricity demand often determines whether local infrastructure has enough capacity."

The division of labor he describes runs across two institutions: utilities and water authorities collecting the operational data, an environmental or public regulatory agency setting consistent reporting requirements. Neither was built to govern AI, and together they would decide where it can land. For jurisdictions without that reporting in place, Ren identifies a workable first pass, screening on-site water use, peak demand, water source, and the seasonal timing of use against local supply, then reading those against existing hydrological data, utility capacity, drought indicators, and watershed stress metrics.

In this way, the pattern holds: where electricity is the scarce input in the North, water is the scarce input in the South, and the institution that governs AI siting is whichever one controls the scarce input.

Where the North rations AI by the kilowatt, the South rations it by the liter, and the venue shifts from the auction to the courtroom.

Layer 2. The Geopolitical Layer: the contest over who controls it

Compute has become an instrument of statecraft, and the controls governing it are proving porous.


Restrictions on AI processors leak through cloud access and transshipment, arrangements that can be technically compliant and substantively evasive at the same time. The leakage now has a measured size.

Epoch AI's April 2026 study estimates that between 290,000 and 1.6 million H100-equivalent accelerators were smuggled into China through 2025, with a median of 660,000, roughly a third of China's total compute. The enforcement record supplies the case studies. In December 2025 the Justice Department's Operation Gatekeeper dismantled a network that moved at least \$160 million of restricted processors through straw purchasers, and in March 2026 federal agents arrested a Super Micro co-founder over a scheme alleged to have routed \$2.5 billion of AI servers through Taiwan and Southeast Asia. Set against roughly \$3 billion in diverted product across six prosecutions, the federal budget for policing export controls ran to about \$122 million for all of 2025.

EXHIBIT 8

Hundreds of thousands of restricted AI chips reached China. The budget to police the controls was \$122 million.

Estimated restricted accelerators smuggled into China through 2025, H100-equivalents



What was diverted, and what was spent to stop it, \$



Note: Smuggling estimates are Epoch AI's April 2026 range and median. The \$3 billion figure sums diverted product across six prosecutions, including Operation Gatekeeper (\$160 million, December 2025) and the Super Micro case (\$2.5 billion, March 2026). Bars are drawn to scale.
Source: Epoch AI (April 2026); US Department of Justice; Horizon Scan 002, Section 2

Compute controls are forming fastest at the border and in silicon, as Washington and Beijing converge on a shared regulatory grammar. They are failing at the ledger, where a \$122 million enforcement budget polices a multibillion-dollar leak.

Erich Grunewald of the Institute for AI Policy and Strategy, asked where a marginal enforcement dollar buys the most deterrence, points less to funding than to staffing and resolve. According to Grunewald, the highest-leverage moves are "hiring more overseas agents" and analysts to process leads, and the binding policies "are not bottlenecked by money but rather by political will."

Grunewald's point is that the tools to catch more diverted chips largely exist, and that the officials who could deploy them are held back by choices about priorities rather than by empty accounts. That reframes the enforcement gap as a decision the government is making, rather than a resource it

lacks, which is a harder problem to solve because no appropriation closes it.

Enforcement has consequently migrated to whoever sits closest to the cargo. Malaysia's trade ministry now requires a Strategic Trade Permit for any export, transshipment, or transit of US-origin AI chips, effective July 14, 2025, with 30-day advance notification even for unlisted items. Singapore's prosecutors have pressed a \$390 million fraud case over servers falsely declared for local end-use, and by July 2026 had charged corporate entities for the first time and seized a \$42 million bungalow from the alleged network. The busiest AI export-control officials of 2026 work in Kuala Lumpur and Singapore, applying trade statutes written for conventional dual-use goods, retrofitted now for frontier compute. That is the pattern this section keeps finding: the gap is filled by the nearest institution with operational reach, on the legal instruments it already has.

The second ring extends further, and each state joins the enforcement layer on the instrument closest to hand: an investment board in Bangkok turned transshipment cop, TSMC's compliance regime in Taipei, and the Dutch and Japanese licensing ministries that control the lithography and equipment exports Washington cannot restrict alone; the fuller second-ring record sits in the appendix. Even the allied core is an operational layer for a policy written elsewhere: the US diffusion framework names eighteen close partners, including Japan, the Netherlands, Korea, and Taiwan, as the trusted tier through which the regime actually runs.

A control regime designed in Washington and mirrored in Beijing is enforced in Kuala Lumpur, Bangkok, The Hague, and Taipei. The middle powers are the operational layer of both blocs.

Washington's own posture escalates unevenly, and the incoherence now has dates attached. On December 8, 2025 the administration announced it would permit H200 sales to China in exchange for a 25 percent export fee, implemented in a January 13, 2026 BIS licensing policy, in the same weeks its prosecutors were unveiling Operation Gatekeeper. By late March 2026 Congress had advanced the Chip Security Act, which would embed location tracking directly into advanced processors. The direction of travel matters more than the bill's fate, and it rests on a research base already several years deep: delay-based location verification, in which chips prove their position through cryptographic ping-timing against landmark servers, and offline licensing, in which chips require renewable cryptographic permits to operate at all. Enforcement is moving from paperwork into silicon, from licensing regimes administered at borders to telemetry welded into the product, which reads as a concession that the paper regime failed.

How much any of this can reach the compute already in China is a separate question, and Grunewald is skeptical there is a policy space at all: on the several hundred thousand H100-equivalents already landed, "I don't think there's much you can do," which leaves stopping the flow as the only lever. He adds a reason the installed base may matter less over time, since

production is ramping fast enough that "new compute quickly dominates old compute."

The practical upshot is that the border is the only place the controls really bite. Once a chip is running inside a Chinese data center, there is no switch to reach it and no permit to withhold, so the effort has to concentrate on the shipments that have not yet left. Grunewald's second observation softens the stakes of that limit: because so much new hardware is being built every year, the chips already inside China shrink as a share of the whole over time, which means the steady flow of new compute matters at least as much as the stock that has already slipped through.

The chips-for-commitments model shows the same devolution to bilateral arrangement. The Trump administration's late-2025 approvals let the UAE's and Saudi Arabia's national champions buy up to 35,000 Nvidia GB300-class chips each, conditioned on the validated-end-user assurances G42 gave the US government, including proof that chips could not be re-exported or remotely accessed by unapproved parties. Access to the substrate is now negotiated country by country, with a private national champion as the compliance surface.

The architecture is now mirrored from the other side. On October 9, 2025, China's commerce ministry issued six export-control notices asserting, for the first time, extraterritorial jurisdiction over foreign-made goods. Any product anywhere containing Chinese-origin rare earths above a 0.1 percent value threshold requires a Chinese export license, with semiconductor- and AI-related end uses explicitly singled out for case-by-case scrutiny and a 50 percent ownership rule that deliberately echoes the BIS Affiliates Rule adopted eleven days earlier. The two blocs are converging on the same design grammar, de minimis thresholds, foreign-direct-product rules, entity lists, while pointing the instruments at each other. The October measures sit suspended until November 10, 2026 under the Busan truce, which converts a regulatory regime into a countdown clock.

One firm now embodies the whole layer, and it is the wider set's clearest case of a supplier as the point of control. On September 30, 2025 the Dutch government seized control of Nexperia, a Netherlands-based, Chinese-owned maker of basic

automotive chips, under a 1952-era goods-availability statute, a statute that has lived many lives; on October 4 China banned exports from Nexperia's Dongguan plant, which packages more than 50 billion chips a year, and automakers on three continents began warning of production stops within weeks. Beijing eased the ban in November and The Hague suspended its order, yet the company's European and Chinese halves remain functionally separated as of July 2026, with the parent claiming \$8 billion from the Dutch state and the first Dutch trade-minister visit to Beijing since 2018 devoted largely to the file. Two ministries, invoking two emergency instruments, partitioned a single company and briefly governed the global auto supply chain between them. No AI statute anywhere played a role.

For every state except a handful of producers, sovereignty over AI has become contractual: capability arrives as a service, and a service can be revoked.

The cloud is the other half of the leak, and it cuts both ways. The World Bank finds that 87 percent of global cloud computing exports originate in the United States, with a share approaching zero delivered to low-income countries. Most nations therefore run their AI on rented American infrastructure governed by American policy, and the arrangement was stress-tested this June, when a Commerce directive took the most capable generally available model offline worldwide for nineteen days and gated its successors to approved customer lists. Allies discovered their access was as revocable as anyone's. Sovereignty over AI, for all except three or four states, is now a service-level agreement.

The absentees define the layer as much as the enforcers. India sits outside the allied control regime and has seen a GPU import surge, and its answer is procurement: the IndiaAI Compute Portal reports 38,231 deployed GPUs across three 2025 procurement rounds, far above the original 10,000-GPU target, with IndiaAI CEO Abhishek Singh confirming the figure in a September 2025 briefing and Sarvam AI receiving the largest single allotment of 4,096 Nvidia H100s. State purchasing puts a ministry in the allocator's chair, and Brazil, its comprehensive bill

PL 2338 still stalled, treats compute the same way; both records sit in the appendix.

Underneath the trade fight lies a deeper contest of philosophies. The EU treats AI as a certifiable product, the US treats it as a system to optimize, the UK leans on adaptive regulator-led oversight, and China governs it as state-managed socio-technical infrastructure. Each logic coheres on its own terms; set side by side, the four reconcile poorly. The contest also has a missing party. High-income countries, home to 17 percent of the world's population, hold 87 percent of notable AI models, 86 percent of AI startups, and 91 percent of cumulative venture funding; excluding China and India, the rest of the world holds under 1 percent of each. The four philosophies being reconciled are the philosophies of the producers. Most of humanity enters this layer as the governed, on terms negotiated elsewhere.

The UK's exposure in this layer is dependence purchased at scale. The state's own compute is real and operational: Isambard-AI at Bristol, built on 5,400 Nvidia GH200 superchips, anchors the AI Research Resource and now feeds a £500 million Sovereign AI Fund offering startups up to a million GPU-hours each. The private pledges dwarf it: Nvidia and partners committed to up to 120,000 GPUs in the UK by end-2026, Microsoft's Technology Prosperity Deal promised a 23,000-GPU supercomputer with Nscale, and Stargate UK was to host OpenAI workloads from 8,000 GPUs scaling toward 31,000. The ratio between the state's 5,400 chips and the market's 120,000 is the layer in miniature, and the risk it encodes stopped being hypothetical when OpenAI paused Stargate UK in 2026, a decision taken in San Francisco that repriced a British national capability overnight. The Sovereign AI Fund reads, in part, as the response. The country's parallel physical dependency, the subsea cables that land the transatlantic internet on British soil, is taken up in the London stress test that closes this section.

Layer 3. The Accountability Layer: who answers when an autonomous system acts

This is the fastest-widening gap and the least visible. Non-human identities are projected to exceed 45 billion by the end of 2026, roughly twelve times the human workforce, while only about a tenth of

organizations report any strategy to govern them. The blind spot has been measured: more than half of deployed agents operate with no security oversight or logging, only about a quarter of organizations report full visibility into how their agents interact, and 78 percent of executives concede they could not pass an independent AI governance audit. The supervisory frameworks that would ordinarily force the issue have stepped back at the decisive moment: the Federal Reserve's April 2026 rewrite of its model risk guidance, SR 26-2, excludes generative and agentic AI from scope entirely, leaving each institution to govern the newest class of systems with controls the letter declines to specify.

Accountability for autonomous action is forming in private hands, as card networks, underwriters, and courts allocate liability that agency and tort doctrine, built around a human principal, cannot yet assign.

Legacy doctrine struggles to attach. Agency law presumes a human principal, tort law presumes foreseeability, and harm from a multi-agent chain emerges across several autonomous steps with no clear hand on the wheel. Many of the new harms defy that framework.

Recourse is shifting toward strict liability: the EU shelved its dedicated AI Liability Directive and now leans on the revised Product Liability Directive, which treats software and AI as products and must be transposed into national law by December 9, 2026.

Into the doctrinal vacuum, two private rulebooks arrived, and together they show what accretion looks like in real time. The first came from the payments industry. Within eighteen months, Mastercard launched Agent Pay (April 2025), binding tokenized card credentials to a named agent, a merchant scope, and a consent policy; Visa followed with its Trusted Agent Protocol; and Google released AP2 (September 2025) with more than sixty partners, wrapping every agent purchase in cryptographically signed mandates that record what the user authorized, what the agent selected, and what was charged. By one count, six competing agent-payment protocols shipped within a single quarter. In April 2026 the FIDO Alliance, the biometrics standards body, stood up working groups to harmonize them, chaired by CVS Health, Google, and OpenAI on authentication

and by Mastercard and Visa on payments. The liability rules for autonomous commerce, who authorized what, who eats a disputed transaction, when an agent may act with no human present, are being drafted in card-network specifications and a standards consortium, years ahead of any legislature.

The protocols carry a second, quieter significance. They are the future evidence base of liability. Veronica Paternolli, a PhD researcher in computer science at the University of Verona who studies how agent protocols interact with tort doctrine, argues that the mandate records now being standardized are precisely what courts will need, because a workable regime must "spread accountability across the entire architecture of the AI agent." The allocation question, on her account, becomes tractable once the record shows who selected the protocol, who controlled the workflow, and who controlled the data. A signed mandate is, on this reading, a comparative-fault instrument waiting for its first case. The rulebooks forming in payments consortia may end up mattering less for what they permit than for what they make provable.

The second rulebook was taking shape more quietly, among underwriters, and it inverted the usual sequence of regulation. The Artificial Intelligence Underwriting Company, launched in July 2025 by alumni of frontier labs and safety organizations, publishes AIUC-1, a certification standard for AI agents spanning six risk families and roughly fifty controls, revised quarterly, audited by an accredited third party, and priced directly into an insurance policy. In February 2026, ElevenLabs became the first company insured under it, after its voice agents ran more than 5,800 adversarial tests; UiPath certified in March; Munich Re's HSB launched a standalone small-business AI liability product the same month. The loop is complete and entirely private: a standard, an audit, a premium. Where the state's model-risk guidance declined to specify controls, an insurer specified fifty-one and attached a price to each. This is the fourth point of control in miniature, a certificate, priced, standing between an agent and its market.

The private standards also have a doctrinal afterlife that their drafters may not intend. Ryan Calo, professor of law at the University of Washington,

whose work with computer scientists examines whether ordinary negligence can absorb agentic AI, observes that "custom has been a way to scaffold reasonableness inquiries," and certification regimes of this kind are how custom gets made. Once fifty-one controls circulate as the insured norm, a defendant who skipped them faces a doubled problem, since custom marks what the community treats as reasonable and its adoption by peers proves the controls were feasible all along. An impossibility defense collapses when the rest of the market is already complying. The premium, in other words, is writing the standard of care that the first wave of agent litigation will apply.

The underwriting instrument is going global, and the strictest agent rulebook in force is consequently being written by the same private hands in market after market. Tokio Marine established a group-wide "Basic Policy on AI Governance" in April 2025 built on transparency, mandatory human oversight, and bias controls, its two peers in Japan's 88-percent-concentrated property and casualty market are building comparable frameworks as they embed agentic systems in claims and underwriting, and Munich Re's aiSure product is expanding from the US into Europe and Asia. In each market the sequence is the same: the insurer specifies the controls and prices the risk before any statute defines the duty of care.

The accountability layer is also where certification will first acquire commercial force, because the instruments that move markets here are documents insurers and procurement desks already know how to read, and because the state is now deliberately building the supply side. The UK's Trusted Third-Party AI Assurance Roadmap commits public money and a skills framework to the third-party audit market, and accredited certification gives underwriters and buyers a signal they can price without opening the model. When an underwriting criterion or an eligibility clause first requires that signal, the fourth point of control will have hardened here, in contract and premium, before any statute names it.

The strictest agent regulation in force across the surveyed regimes is an insurance standard priced into a premium; the public frameworks meant to replace it remain unwritten.

Statute is late, and it is failing in three distinct ways at once. Each failure hands the accountability function to a different interim governor.

The first failure mode is death. Canada's Artificial Intelligence and Data Act, the AIDA in Bill C-27, died on the order paper when Parliament was prorogued in January 2025, and the industry minister confirmed in June 2025 it will not return in that form. What governs Canadian AI in its place is a patchwork of pre-existing instruments: the prudential regulator OSFI, whose revised model-risk Guideline E-23, published September 11, 2025 and effective May 1, 2027, explicitly extends to dynamic self-learning AI models; Quebec's Law 25 on automated-decision transparency; and the courts. The cleanest illustration is *Moffatt v. Air Canada* (2024), in which a tribunal held the airline liable for its chatbot's fabricated refund promise and rejected the argument that the bot was a separate entity, establishing through a small-claims decision the deployer-liability rule that the dead statute would have codified. Brazil's stalled PL 2338 runs the same track, with the courts and the data-protection authority governing in the interim.

The second failure mode is deliberate hollowness, and the wider set's exemplar is Japan, whose AI Promotion Act, enacted May 28, 2025 and in force from September 1, 2025, is a comprehensive national AI law by design and carries no penalties by design. Its only direct obligation on business is a "duty to cooperate," and governance runs through METI ministerial guidance and corporate self-regulation, which is a soft-law choice made in full knowledge of the harder alternatives next door.

The third failure mode is the enacted-without-machinery gap, and it mirrors the EU precisely. South Korea's AI Basic Act, effective January 22, 2026, is the first comprehensive national AI law in Asia, with transparency and high-impact obligations, extraterritorial reach, and a compute threshold set at 10^{26} FLOPS by its enforcement decree. Yet MSIT, the science ministry, is running a grace period of at least one year in 2026 that defers fines except in cases of "serious social harm, such as loss of life or human-rights violations," the maximum administrative fine is only KRW 30 million (about US\$20,163, for example for failing to label generative AI), and the enforcement decrees and institutional capacity are

still being assembled. A comprehensive statute exists; the machinery to make it bite does not yet, which is the same condition that pushed the EU's own high-risk deadline to December 2027.

Statute is failing on three tracks at once: dying in Ottawa, deliberately hollow in Tokyo, enacted without machinery in Seoul. In each, the interim governor is a court, a ministry, or an insurer.

The courts are the other interim governor, and the UK supplied the first major test, one that mapped the gap in detail and left it open. In *Getty Images v Stability AI* ([2025] EWHC 2863 (Ch), November 4, 2025), the first UK judgment on generative AI and intellectual property, Getty abandoned its primary copyright claims mid-trial because it could not show the training occurred in the UK, and the court then rejected the secondary infringement claim, holding that model weights do not contain or store the works they were trained on and so are not "infringing copies." The trademark finding that survived was, in the judge's words, "historic and extremely limited." Two consequences follow, and both echo the layers above. Territoriality means that training conducted abroad sits beyond the reach of UK copyright even when the outputs land in London, an accountability gap opened by geography before doctrine could reach the question. And the substantive question, whether training on protected works is lawful, remains undecided by any UK court, with academic criticism of the secondary-infringement reasoning already building toward appeal.

Where the statute is absent, the data-protection authority governs by improvisation, and inside the EU the instrument is the GDPR itself. Italy's Garante fined OpenAI €15 million in December 2024 for training ChatGPT without a lawful basis and without adequate age checks, the first GDPR penalty against a generative-AI company, and separately banned DeepSeek in Italy; the fine was later annulled by the Court of Rome in March 2026, which shows the improvised instrument is as reversible as the Colorado safe harbor in Section 1. India's courts supply a companion holding: in *ANI Media v. OpenAI* the Delhi High Court held in July 2026 that it had territorial jurisdiction over OpenAI even though its servers sit abroad, then declined an injunction on

the prima facie view that training can fall within fair dealing. The jurisdictional holding echoes *Getty v. Stability* directly: a court in an importing jurisdiction asserting authority over a US lab whose training happened elsewhere.

China is the case that proves the rule by inverting it. Where other states left the accountability layer to private rulebooks and improvising regulators, China built the machinery first. The Cyberspace Administration's algorithm-filing and generative-AI regime requires every public-facing model to register before launch, and by August 31, 2025 the CAC had processed 538 generative-AI service filings and 263 application registrations, 801 approved services in a publicly listed registry that no other country maintains. The courts moved in step: the Beijing Internet Court recognised copyright in an AI-generated image in November 2023, and the Guangzhou Internet Court held a generative-AI provider liable for infringing the Ultraman character in February 2024, with a Hangzhou judgment following in 2025. A registry plus a line of court judgments is functioning accountability machinery. The state built it early, accretion needs a vacuum, and China left none for the private rulebooks to fill.

China is the control case that makes the point: where the state built the registry and the courts first, private rulebooks did not have to.

The pattern to hold onto is who is doing the governing. The consequential accountability decisions of the past eighteen months were made by underwriters pricing exclusions and certifications, by card networks writing mandate schemas, by a national accreditation service vouching for the first AI certifiers, by plaintiffs' lawyers selecting defendants in disputes like *Mobley v. Workday*, and by a tribunal refusing to accept that a chatbot was a separate legal entity.

Legislatures are ratifying, at a distance, allocations the market and the courts have already made.

Layer 4. The Evidence Layer: whether anyone can see what is happening

Every layer of this section presumes a capacity the state is quietly losing, the capacity to see what AI is

doing. Grid planners need demand forecasts, export officials need diffusion estimates, and accountability regimes need a factual record of harms. That seeing is now the weakest layer of the four, and the clearest account of why comes from the economics of work, where the measurement effort is most advanced and the results are most contradictory.

The capacity to see what AI is doing is migrating from public statistics to private telemetry, and it is migrating fastest exactly where the policy questions are hardest.

The contradiction starts with the instruments. As the New York Times' Ben Casselman documented this summer, the standard measures of the economy predate the personal computer. The monthly jobs report has no category for the technology sector at all; the most recent detailed occupational breakdown dates to May 2025. The workhorse research tool is alarmingly sensitive to construction:

when economists at Northwestern and American University ran the same analysis under different AI-exposure measures, the choice of measure flipped the estimated effect of AI on jobs from negative to positive. The instrument decides the finding.

The signals point in opposite directions with equal confidence. Stanford's Digital Economy Lab, working from ADP payroll records, finds a 16 percent relative employment decline for workers aged 22 to 25 in the most AI-exposed occupations. Research from Ramp and Revelio Labs finds the firms adopting AI most intensively are adding jobs faster than the laggards. A National Bureau of Economic Research survey of more than 6,000 senior executives finds roughly 90 percent unable to attribute any measurable productivity change to AI after three years. The Yale Budget Lab finds the labor market's occupational mix shifting no faster than during the computer and internet eras.

EXHIBIT 9

AI's effect on jobs points in three different directions.

What four widely cited studies found about AI and employment, grouped by the direction of the effect each reports

SOURCE	Fewer jobs	No measurable effect	More jobs
Stanford Digital Economy Lab Payroll records from ADP	 16% relative employment decline for workers aged 22-25 in the most AI-exposed occupations		
Ramp and Revelio Labs Company spending and hiring data			 Firms adopting AI most intensively add jobs faster than those adopting least
National Bureau of Economic Research Survey of more than 6,000 senior executives		 About 90% cannot attribute any measurable productivity change to AI after three years	
Yale Budget Lab The mix of occupations across the labour market		 Occupations shifting no faster than in the computer and internet eras	

Note: The measure matters as much as the data. When economists at Northwestern and American University ran the same analysis with different definitions of AI exposure, the estimated effect on jobs flipped from negative to positive.

Source: Stanford Digital Economy Lab; Ramp and Revelio Labs; National Bureau of Economic Research; Yale Budget Lab; Horizon Scan 002, Section 2

The attribution problem compounds the measurement one: of the job cuts American companies announced through most of last year, about 4.5 percent were attributed to AI while ordinary market conditions accounted for nearly five times as many, *a gap economists have started calling AI-washing*. Erik Brynjolfsson's J-curve work on earlier general-purpose technologies, in which firms lose productivity while reorganizing around a new tool and gain it only afterward, would predict exactly this mixed picture.

Whether AI is destroying work is, on the present evidence, a question the instruments answer both ways. What makes this a governance problem in its own right is that the instruments were breaking before the question arrived, and the breakage is documented step by step.

In April 2025 the Bureau of Labor Statistics suspended consumer price collection entirely in three cities and trimmed roughly 15 percent of the sample elsewhere, citing a hiring freeze. In June it suspended researcher access to nearly all of its restricted-use datasets. On August 1, hours after a weak jobs report, the President fired the BLS commissioner, an act a former Trump-appointed commissioner called "totally groundless"; the annual benchmark revision the following month, 911,000 jobs, was the largest in two decades and triggered a Labor Department inspector-general review of the agency's capacity to collect its own flagship series. The cuts have landed on a load-bearing wall.

Then the 43-day autumn shutdown stopped collection outright: the October 2025 household survey was never conducted and cannot be reconstructed, a permanent hole in the employment record at the exact moment the AI-displacement question turned urgent. In June 2026, the Commerce Department banned noise infusion, a privacy technique statistical agencies rely on to publish granular data at all. The Economic Innovation Group's Nathan Goldschlag, whose July report is the sharpest inventory of the problem, states the stakes in one line: you cannot get the policy right if you do not know what is happening. It is premature, though, to write an obituary for the public instrument. Congress has begun to notice. The AI DATA Act,

introduced in June, would modernize the federal labor surveys and require an annual report on AI's workforce effects. It is a bill about statistical plumbing, which is precisely why it may be the most consequential AI legislation of the year: every other intervention inherits its blind spots.

The instrument decay is not confined to Washington. The public sensor is dimming across the advanced economies together, and the American case is simply the best documented.

The UK is running this experiment in an acute form, and the diagnosis is now official. Its flagship labour instrument, the Labour Force Survey, saw response rates fall so far that the ONS suspended detailed estimates outright from October 2023, and everything published since carries the downgraded badge of "official statistics in development." The Devereux Review, the independent inquiry into the ONS published in June 2025, concluded that most of the failures in core economic statistics were "the consequence of ONS's own performance," tracing them to funding choices that stripped the survey field force at precisely the wrong moment. The response rate ran to 21.5 percent in mid-2025, the two principal measures of employment currently disagree on direction, and the replacement Transformed Labour Force Survey faced its readiness assessment in July 2026 with transition aimed at November, which means the country's ability to answer the AI-and-work question was itself under evaluation this summer. The country deploying AI into three-quarters of its financial sector cannot presently say, with accredited statistics, what AI is doing to its workforce.

The Global South version is absence, and it is deepest exactly where AI adoption is leapfrogging. India, deploying AI at national scale and home to roughly 100 million weekly ChatGPT users by the Delhi High Court's own reckoning, has not conducted a census since 2011: the 2021 count was postponed and is now scheduled for 2027, which means every sample survey since 2020, including the labour-force survey PLFS and the consumer-expenditure survey, has run on a fifteen-year-old sampling frame. Brazil's IBGE and Indonesia's BPS face comparable capacity limits. The multilateral instruments meant to fill the gap, the ILO's AI-exposure work and the IMF's AI Preparedness Index, resolve poorly at the level

of a developing economy's actual labour market. The countries adopting AI fastest can measure its effects least.

The countries adopting AI fastest can see it least. Statistical invisibility is deepest exactly where the leapfrog is steepest.

The second structural fact is where the good data went. The most timely reads on AI's economic footprint now come from payroll processors, expense platforms, hiring-data firms, and the model providers themselves: ADP, Ramp, Revelio, Stripe, Anthropic. Each covers a nonrepresentative slice, none is obliged to publish, and together they out-resolve the state. The sharpest of these private sensors also delivers the layer's bleakest finding: Anthropic's own Economic Index, built from anonymized usage across 150-plus countries, reports that per-capita AI usage is strongly correlated with GDP and shows no sign of cross-country convergence, with no evidence that low-use countries are catching up.

The evidence layer, in other words, is privatizing along the same lines as the other three, and the privatized instrument can already see the divergence the public instruments cannot measure. For the countries on the far side of the World Bank's divide, running on 5 gigabytes of data per person a year, even that privatized visibility does not exist; policy there will be made on anecdote and imported narrative.

A state that cannot measure AI's effects cannot govern them; every intervention in the three layers above inherits the blind spots of this one.

The pattern, read across the layers

Set the four layers side by side and the section's finding sharpens into something close to a law of motion. Governance flows downhill to operational control, and it arrives carrying the instruments of whoever gets there first.

In the physical layer the instrument is the queue and the auction; in the geopolitical layer it is the customs ledger and the switch, thrown in both directions in the Nexperia affair; in the accountability layer it is the mandate schema, the certificate, and the underwriting file; in the evidence layer it is corporate telemetry. Each instrument is effective within its franchise, and each carries the priorities of its owner. An auction allocates cost, and by the market monitor's attribution it charged \$9.3 billion of data center demand to a single year's capacity bill before any legislature voted on the question. A card network allocates transaction liability, and it is already drafting the terms on which an agent may spend with no human present. An insurer's fifty-one controls are, functionally, the strictest agent regulation in force in any of the seven regimes surveyed. None of these governors was elected to the role, and none can see beyond its own layer. An accredited certificate allocates market access, and access is the quietest instrument of the four.

EXHIBIT 10

Suppliers already set firm limits on AI in all four areas we examined. Governments do so in only two.

How strongly each type of actor constrains AI today, scored from 0 (nothing in force) to 3 (the routine constraint)

0 No instrument in force 1 Present but deferred or partial 2 Binding in specific cases 3 Binding as the routine constraint

	Suppliers	The state	Courts	Certifiers and insurers
Physical Chips, power and data-centre sites	3 Supplier order books and allocation	3 Grid connection queues and auctions	2 Environmental review of data-centre sites	0 No certification instrument in force
Geopolitical Export controls and trade in chips	3 Contract terms passed down the supply chain	3 Customs records and export licences	1 Criminal prosecution for falsely declared servers	0 End-user assurances are contractual only
Accountability Who answers when AI causes harm	3 Payment networks' rules for AI agents	1 Supervisory guidance excludes generative AI	3 Ordinary liability law applied to AI deployers	3 Audited controls priced into insurance premiums
Evidence Measuring AI's economic footprint	3 Company data on payroll and spending	1 Degraded public statistics	0 No court route to the question	0 No assurance instrument in force

Note: Scores are a single-coder, indicative assessment. Each cell names the leading instrument observed in that area.
Source: Horizon Scan 002, Section 2

The Fourth Point of Control: the certificate between the statute and the system

On 15 January 2026, UKAS announced it had accredited BSI to certify AI management systems against ISO/IEC 42001:2023. BSI's own announcement is dated 17 November 2025. Neither notice specifies the grant date (see p. 46). The accreditation matters more than any single certificate it will produce, because it means the certifier has itself been assessed for competence and impartiality. That second layer of scrutiny allows a buyer who has never audited a vendor to treat the vendor's certificate as evidence. Policy was already in motion. The Department for Science, Innovation and Technology had published its Trusted Third-Party AI Assurance Roadmap on 3 September 2025, treating assurance as a market the state should deliberately build, with a

consortium and an innovation fund alongside work on the skills and professional standing of auditors. In July 2026, CEN and CENELEC, the European standards bodies, published EN 18286:2026, the first European standard developed in support of AI Act implementation, written to the quality-management obligations that Article 17 places on providers of high-risk systems. Its citation in the Official Journal of the European Union is still pending, and the pending step is the consequential one, because citation is what attaches a legal presumption of conformity, and presumption is what converts a voluntary text into the cheapest available route to compliance.

A sequence is now visible from stated duty to commercial consequence. A legislature or a board declares an obligation, a standards body translates it into auditable requirements, a certifier attests that an organization meets them, and an accreditor vouches for the certifier. Only the first step in that stack requires public authority. The more

procurement criteria and underwriting terms come to read the certificate, the less any regulator has to move for the standard to bind, and the seven regimes surveyed here sit at very different points on that curve. Certification of this kind insulates a firm from procurement risk well before it insulates anyone from harm. That distance is itself a gap.

This machinery changes the Scan's account of binding constraint. The decomposition that ran through the supplier, the state, and the court now carries a fourth point of control, the certifier, with the accreditor standing behind it. The four layers map where governance forms; the four points of control name who does the binding. A supplier binds vertically, through contract terms flowing down its own stack, while the certifier binds horizontally, since a body with no place in that stack decides whether an organization meets a standard, and the decision conditions entry to buyers, insurers, and procurement. Nobody sues, and nobody fines. The certificate is simply withheld, and the market does the enforcing. The channel also feeds the other three. Procurement contracts give suppliers one more document to demand. Citation in the Official Journal would let states bind through reference, a statutory thumb on the scale for whichever text earns the listing. Certificates will in time reach courtrooms as evidence of reasonable care, and judges will draw inferences from their absence. In this architecture, the keystone is accreditation, set by institutions most accounts of AI governance have yet to name. The fourth channel remains in formation, still short of full standing, both because no regime yet conditions market access on the certificate and because the assurance market beneath it is months old. Section 3 accordingly carries the certifier as an emerging fourth channel rather than a column in the 2026 index. A channel of this vintage cannot be weighed against three that already bind, and a score resting on so young a market would measure noise rather than practice. What would complete the channel is a requirement that makes the certificate a condition of doing business: a procurement clause or an underwriting criterion written so that no certificate means no contract. No regime in the set has yet written one, although the EU sits closest, since Article 17 already obliges providers of high-risk systems to run a quality-management system and since a cited EN 18286 would make certified conformity the easiest way to show it. The

accountability layer above already carries the loop's most complete private instance, an insurer's certification standard priced directly into premiums.

The fourth channel also imports a gap of its own kind into every layer above. A certificate under ISO/IEC 42001 attests that an organization maintains specified processes for governing its AI systems. The behavior of a particular deployed system remains a separate question, one the management-system instrument was never built to answer. Strong rules can therefore coexist with weak assurance, and a certified organization can field a system that misbehaves. An assurance signal can also travel faster and farther than the verification behind it, which is how a market that learns to read certificates could stop asking the questions certificates cannot answer, a perverse consequence for an infrastructure built to manufacture trust. That is reason to reread the four layers, once for the statutes and once for the certificates growing up beside them.

The wider jurisdictional record, preserved in the appendix, bears out the law of motion and adds five cross-jurisdiction findings. First, the utility regulator has become a de facto AI siting authority across the advanced economies: IMDA in Singapore and KEPCO in Korea allocate AI's physical footprint through tenders and connection confirmations, the CRU in Ireland, TenneT and the ACM in the Netherlands, and BAFA in Germany do the same through connection contracts and efficiency mandates across the appendix set, and the auction-and-queue instrument generalizes cleanly from PJM to all of them. Second, a second ring of middle powers, the Netherlands, Japan, Korea, Taiwan, Malaysia, Thailand, Singapore, and the UAE, now runs the operational layer of a control regime designed in Washington and mirrored in Beijing, each enforcing on the instrument closest to hand, from a Malaysian permit to a Thai investment-board rule to a Dutch 1952 statute. Third, statute is failing on three separable tracks, dying in Canada and stalling in Brazil, deliberately hollow in Japan, and enacted-without-machinery in Korea and the EU, while data-protection authorities, courts, and insurers govern in the interim, and China stands as the counter-case where the state built the registry and the courts first. Fourth, water joins electricity as a Layer 1 governance instrument, with Chilean and

Uruguayan courts rationing AI siting by the liter where Northern regulators ration it by the kilowatt. Fifth, a market-access mechanism is assembling in the assurance institutions: UKAS has accredited BSI, the first UK AI certifier; EN 18286 awaits its Official Journal citation in the EU, and an American insurer already prices certification into premiums, which makes the certificate a point of control in its own right across three regimes at once.

Two convergences give the pattern its edge. The first is architectural: Washington and Beijing are now building mirror-image control regimes, matching de minimis thresholds, foreign-direct-product rules, and ownership tests almost clause for clause, while the enforcement work devolves to Dutch ministries, Malaysian permit offices, Thai customs, and Singaporean prosecutors. The machinery of technology statecraft is standardizing even as its purposes polarize. The second is distributive, and it travels through every layer: the substrate, the models, the funding, the rulemaking venues, and now even the measurement all concentrate in the same small set of producer economies, and the private sensor that can actually see the resulting divergence reports that it is not closing. The countries that import their compute also import their governance and, increasingly, their self-knowledge, and the census-dark states are governing the fastest-moving technology in living memory on the oldest data.

The London stress test

The four layers of this section settle, with unusual concentration, on one city. London carries the

physical substrate of UK AI, lands the transatlantic internet, and runs the world's most explicit personal-accountability regime for financial firms, and each of those roles is governed by an institution built for something else. Three constraints frame the local stress test.

The grid. London is Europe's largest data centre market, with roughly two-thirds of UK facilities clustered within 20 miles of the city, which puts the physical substrate of UK AI directly on top of the country's most congested grid (House of Commons Library, May 2026). The average grid connection lead time for a new 50MW data centre in London now runs to seven years (JLL, via New Civil Engineer), and new housing in west London has stalled because data centres cannibalized the spare grid capacity (Electric Insights, Q1 2026). Ofgem reports roughly 140 projects seeking some 50 GW of connections, more than Britain's entire 45 GW peak demand (Ofgem, via Reinforce Technology). The oversight gap is structural: the UK designated data centres as Critical National Infrastructure in 2024 and passed the Planning and Infrastructure Act 2025 to bypass local planning, but no National Policy Statement for Data Centres has been finalised, leaving these CNI assets with no clear owner (Baker McKenzie, AI Growth Zones).

2 in 3

UK data centres sit within 20 miles of London, on one of the country's most congested grids.


EXHIBIT 11

London's three tightest constraints on AI are governed by rules between 10 and 141 years old, none written with AI in mind.

How close each constraint is to its limit, and how old the rule that governs it is, %

Power grid

Data-centre connection requests as a share of national peak electricity demand (about 50 GW sought against a 45 GW peak)



GOVERNING RULE


Electricity connection regime

Dates from 1989 · 37 years old

No national policy statement for data centres has been finalised

Undersea cables

Estimated share of UK-US cable capacity in two cables landing at Bude, Cornwall



GOVERNING RULE


Subsea cable legislation

Dates from about 1885 · 141 years old

Replacement under consultation; cable owners and maritime authorities oversee, with no AI mandate

Firms' grasp of their own AI

Share of 118 surveyed financial firms admitting only a partial understanding of the AI they use




GOVERNING RULE

Senior Managers and Certification Regime

Dates from 2016 · 10 years old

Overseen by the FCA and PRA; no AI-specific rules written

When each governing rule was written



Note: Grid requests (about 50 GW) exceed national peak demand (about 45 GW); cyan shows the excess. The cable figure is estimated capacity, not measured traffic or the share of UK internet activity lost if cables fail. Cable-law origin is approximate; ages are calculated against 2026.

Sources: JCNSS, Subsea telecommunications cables (September 2025), paras. 10 and 51; House of Commons Library; Ofgem; GOV.UK (May 2026); Bank of England/FCA survey of 118 firms. See p. 46.

The cables. The country's second physical dependency is coastal. The transatlantic internet lands on British soil through a concentrated set of subsea cables and landing points, and the compute the UK rents abroad, the cloud its firms run on, and the data its regulators watch all travel over them. The section's law of motion applies here too: the institutions with operational reach are cable owners, landing-station operators, and maritime authorities, none of which was constituted with AI dependence in mind. Section 1 carries the full treatment; the point for this box is that the dependency is physical before it is legal.

The attestation. The UK's accountability instrument for finance is the Senior Managers and Certification Regime, which makes named executives personally answerable for the systems their firms run, and the

regime is now absorbing AI on a thin evidence base. In the Bank of England and FCA's survey of 118 firms, 46 percent conceded only a partial understanding of the AI they use, which leaves senior managers attesting to systems their own institutions describe as partly opaque. The FCA's AI Live Testing cohort, running through end-2026, with evaluation due in early 2027, is a sandbox, and the rulebook comes later, if it comes. Section 1 examines the FCA's posture in full.

Read together, the three constraints supply the local lens: the city's AI capacity is rationed by a grid queue, its connectivity runs through a short stretch of coastline, and its accountability regime asks named individuals to vouch for systems their firms only partly understand.

3

SECTION 3

Rules Written, Rules Enforced

How far governance on paper runs ahead of governance that lands, across seven regimes

The preceding sections mapped the vocabulary and the gaps, jurisdiction by jurisdiction and layer by layer. This section sets the record side by side and


asks one comparative question: across the major regimes, how far does governance on paper run ahead of governance that actually lands, and what explains the distance?

EXHIBIT 12

Most jurisdictions have written more AI rules than they have enforced. The United States and Canada are the exceptions.

How much AI law each of seven jurisdictions has written, and how much it has enforced, each scored 0 to 100

● Rules written exceed enforcement ● Enforcement exceeds rules written ○ Circle size = number of enforcement actions in scope



Reading the chart

Each circle is one jurisdiction. The dashed line marks where enforcement matches the rules written. The further a circle sits from the line, the less its enforcement record could be predicted from its statute book.

35 points

United States: enforcement record ahead of rules written, the largest gap in either direction

30 points

European Union: rules written ahead of enforcement record

Note: Each circle is one jurisdiction; the further it sits from the dashed line, the less its enforcement record could be predicted from its statute book. The indices rank jurisdictions against one another and do not measure outcomes. Enforcement actions in scope: United States 18, China 6, European Union 6, South Korea 3, United Kingdom 3, Canada 2, Singapore 0.

Source: Horizon Scan 002, Section 3

Two regimes cross the identity line. The United States and Canada now enforce more than they have legislated, and nearly everything they enforce runs on law written for something else.

The geometry demonstrates that jurisdictions do not scatter randomly; they sort into a diagonal band and two departures from it. Along the band, restraint pairs with restraint and ambition with ambition, Singapore at one end and China at the other.

The departures tell the story: the EU standing furthest above the line, and the two common-law systems standing across it. Distance from the line, in either direction, measures the same thing, the degree to which a regime's enforcement record could not be predicted from its statute book.

Read with the underlying record, the scatter differs from the familiar map of statutes in five main ways.

The widest gap belongs to the most finished statute.

The European Union scores 91.7 on activity and 62.0 on bite. The AI Act is in force, its prohibitions have applied since February 2025, and its general-purpose

obligations since August 2025. The bite, though, is largely scheduled. Enforcement powers over general-purpose models activated on August 2, 2026, two days before this data closed; the AI Office marked the occasion by opening discussions with OpenAI and Anthropic, and it has yet to issue a fine. The high-risk obligations moved to December 2027 and August 2028 under the Digital Omnibus. The revised Product Liability Directive reaches national law through a byzantine transposition timetable: Hungary has completed it, roughly ten member states hold draft bills, France and Italy have yet to begin, and the directive touches only products placed on the market after December 9, 2026. The flagship fine the regime could point to, the Garante's €15 million against OpenAI, was annulled by the Court of Rome in March 2026. What bite the EU has today is borrowed from data-protection law: the Dutch authority's €30.5 million against Clearview stands, the €5 million against Replika stands, the DeepSeek ban stands.


Two regimes bite more than they legislate, and the pattern is this Scan's thesis in numbers. The United States scores 30.0 on activity and 65.0 on bite, a gap of minus 35, the largest in either direction. The paper side amounts to one narrow federal statute, the Take It Down Act of May 2025, plus executive orders that

revise one another. The enforcement side runs to eighteen in-scope actions: at least a dozen FTC cases on AI claims in 2025 alone, securities actions against advisers for AI-washing, the first AI hiring settlement at the EEOC, and judgments standing at \$48.6 million, \$18 million, and \$17 million. The register runs from a facial-recognition ban imposed on a pharmacy chain in December 2023, through a consent order over an AI weapons scanner that missed knives, to a criminal indictment of a founder whose shopping app marketed autonomous AI while employees completed the purchases by hand. None of it required an AI law. The devil is in the details, and the details are in the district courts, the consent orders, and the settlement schedules. Canada runs the same pattern at a fifth of the volume, with a gap of minus 9. Its statute died on prorogation in January 2025, and the deployer-liability rule that statute would have codified was set instead by a small-claims tribunal, in *Moffatt v. Air Canada*, for the price of a refund. A federal privacy investigation and a prudential guideline effective May 2027 round out the machinery. In both countries the general-purpose regulators are in the driver's seat, and the finding dovetails with Section 2's law of motion: governance accretes to whoever already holds an operational instrument.

EXHIBIT 13

Only the United States and Canada enforce more than they have written into law.

Rules written into law minus enforcement record, by jurisdiction, index points



Note: Gap = rules-written index minus enforcement-record index, each scored 0 to 100. Data to 3 August 2026. The scores rank jurisdictions; they do not measure outcomes.

Source: Horizon Scan 002, Section 3

One regime moves the two axes together. China scores 100 on activity and 78.7 on bite, the highest on both and the smallest gap among the regimes that legislated first. The registry came before the vocabulary: every public-facing generative service files before launch, 868 services stood registered with 530 applications on file by April 2026, and refusal is a big stick, because an unfiled model has no market to enter. The registry's reach extends past the generative tier: algorithm filings across all classes now exceed six thousand. The labelling measures in force since September 2025 produced named takedown cases within five months. The courts moved in step, from the Beijing Internet Court's 2023 image ruling through Guangzhou's provider-liability judgment to the Hangzhou line clarifying that ordinary fault liability reaches generative providers. China is the only regime in the set whose bite runs on AI-specific instruments. Everywhere else, the teeth were borrowed.

Deferral, more than defiance, manufactures the gap. South Korea enacted the most complete statute in

Asia, confirmed a compute threshold at 10^{26} FLOPS in its enforcement decree, capped the maximum administrative fine at KRW 30 million, about \$20,000, and then ran a grace period that insulates firms from even that through 2026. The EU has moved its high-risk deadline once and may be asked to move it again. Colorado repealed and replaced its AI law before the law ever took effect; New York's RAISE Act waits for January 2027; a December 2025 executive order stood up a Justice Department task force to litigate against the state laws that did pass. Occam's razor would chalk the gap up to youth, on the theory that statutes are new and machinery takes time. The register resists the comfort of that reading. The two regimes that cross the line never waited for a statute, and the regime with the most machinery built it before writing the vocabulary.

Restraint scores as balance, and the balance is thin. The United Kingdom sits at 43.3 and 40.7, a gap of 2.6, the closest to the identity line in the set. The symmetry is deliberate: little paper by design, and bite that keeps stalling short of a holding. The

Clearview fine, issued in 2022 and extinguished at the First-tier Tribunal, recovered its jurisdictional footing at the Upper Tribunal in October 2025 and now awaits a remitted hearing; the fine is four years old and still not final. *Getty v. Stability* reached the High Court and left the central training question undecided. Singapore anchors the bottom corner at 18.3 and 8.3, low on both axes by design: its frameworks carry no penalty anywhere in them, and its real leverage, the data-centre tender and the trade prosecutions, lives in Section 2's constraint mix.

2.6

index points separate what the UK has written into law from what it has enforced, the narrowest gap of the seven.


What the bite actually runs on

The register's 38 in-scope actions, read down the instrument column, reduce to one line: two ran on AI-specific law, and both are Chinese. Consumer-protection statutes carried fifteen, from the Rite Aid facial-recognition ban of December 2023 through the Air AI judgment of March 2026. Data-protection law carried twelve across four regimes. Copyright carried four, securities and fraud statutes three, employment law one, contract and tort the remainder. The common thread is age. The instruments doing the enforcing predate the technology by decades, and their operators needed no new mandate to act.

EXHIBIT 14

Almost every enforcement action against AI so far has relied on laws written before AI existed.

Consequential enforcement actions involving AI, by area of law, Jan 2023–Aug 2026, number of actions



Note: 38 consequential actions in scope, from the Rite Aid facial-recognition ban (December 2023) to the Air AI judgment (March 2026).
Source: Horizon Scan 002, Section 3, enforcement register

The composition holds within jurisdictions as well as across them. The European model wagers that an ounce of prevention is worth a pound of cure; the register shows the cure already dispensed daily, by agencies that never needed a new prescription, while the prevention waits on transposition. The record also cuts in the other temporal direction, because

what general law gives, general law can retract: the FTC vacated its own Rytr order in December 2025, and the Court of Rome erased Europe's flagship AI fine fifteen months after it was issued. Borrowed teeth can be recalled by the lender.


The middle band of the register has its own texture. Korean enforcement exists and stays small: a nominal fine against OpenAI in 2023, inspection letters to five model providers, a two-month suspension of DeepSeek downloads that ended on compliance. British enforcement exists and stays

unfinished: the Snap inquiry closed without a fine, and the Clearview matter has spent four years establishing who may decide it. Enforcement of this kind registers on the axis without moving markets, which is why the bite score weighs sanctions standing and liability live alongside the raw count.

EXHIBIT 15

The United States has brought the most AI enforcement actions, all under older consumer, securities and employment law. Only China has used a law written for AI.

Consequential enforcement actions involving AI, by jurisdiction and legal basis, Jan 2023–Aug 2026, number of actions



Note: 38 consequential actions in scope across seven jurisdictions, from the Rite Aid facial-recognition ban (December 2023) to the Air AI judgment (March 2026).

Source: Horizon Scan 002, Section 3, enforcement register

Of thirty-eight enforcement actions with real consequence, two ran on AI-specific law. Both are Chinese.

The scheduled future

It is premature, though, to write an obituary for statute. General-purpose enforcement in Brussels is live as of August 2, 2026. The Product Liability Directive begins to apply to products placed on the market after December 9, 2026, wherever transposition catches up with it. Korean fines resume

in 2027 when the grace period lapses, backed from September 2026 by data-protection penalties rising to 10 percent of turnover. Canada's prudential guideline takes effect on May 1, 2027; New York and Colorado arrive that January; the EU's high-risk obligations follow in December 2027 and August 2028.

If the schedule holds, the 2027 edition of this Scan will be measuring a different gap. Nearly every date on the exhibit has moved at least once in the past, and the single date that arrived on time is two days old. The ledger of retreat is the counterweight. The first stated American compute threshold left with the executive order that carried it, revoked in January

2025. Virginia's algorithmic-discrimination bill was vetoed. Colorado's first law never operated a day. The Korean grace period was announced before the statute it softens took effect. Deferral and revocation

are the same instrument at different settings, and both leave the activity score untouched while the bite score waits.

EXHIBIT 16

Most of the penalties written into AI law have yet to take effect. Nearly every date below has already moved at least once.

When AI rules and penalties take effect, August 2025 to August 2028

Already in force		Scheduled, not yet in force	
By 3 August 2026, when the evidence for this Scan closed		In date order	
Aug 2025	EU obligations on general-purpose AI models apply	Dec 2026	EU product-liability rules apply to new AI products
Jan 2026	South Korea's AI Basic Act takes effect FINES DEFERRED	Jan 2027	New York's RAISE Act and Colorado's replacement AI law take effect COLORADO LAW REPLACED ONCE
2 Aug 2026	EU gains power to enforce against general-purpose AI models	2027	South Korean fines begin when the grace period ends
		May 2027	Canada's model-risk guideline for banks and insurers takes effect
		Dec 2027	EU rules for high-risk uses such as hiring and credit scoring (Annex III) DEFERRED ONCE
		Aug 2028	EU rules for AI built into products already covered by EU safety law (Annex I)

Note: The EU product-liability directive applies wherever member states have transposed it into national law; only Hungary had done so by June 2026. The Korean grace period runs for at least one year from January 2026.
Source: Horizon Scan 002, Section 3; Jones Day transposition tracker, June 2026

Every stated compute threshold in the set is revoked, deferred, or capped below the cost of the conduct it polices. The threshold is the vocabulary; the deferral is the gap.

What the scores do not claim

Two checks discipline the finding, and one caution bounds it. The first check recomputes bite on AI-specific channels alone. The ranking collapses to a single bar, China's, and the American score falls toward zero, which is the point: the American bite is general-purpose law working at full occupancy. The second check strips volume and scores only

sanctions standing and liability live; the ordering of the seven holds. The caution concerns what a composite can honestly carry. An enforcement count rewards litigious systems, and the log scale damps that reward without erasing it. The implementation and machinery anchors are judgment calls, scored against written rubrics and colour-flagged in the Scan's research database so that any single call can be contested without the method collapsing. The scores rank; they do not measure welfare. A regime can close its gap from either side, by enforcing more or by promising less, and the chart is silent on which is wiser.

EXHIBIT 17

China and the European Union have written the most AI law. The United States leads on only one measure: the number of enforcement actions it has brought.

The six component scores behind each jurisdiction's two indices, 0 to 100

0-24 25-49 50-74 75-100

	Rules written into law			Enforcement record		
	Binding legal base	Implementation	Machinery in place	Enforcement actions	Sanctions standing	Liability live
China	100	100	100	66	100	70
European Union	100	75	100	66	70	50
South Korea	20	50	70	47	25	25
United Kingdom	60	30	40	47	50	25
Canada	60	30	40	37	50	70
United States	20	30	40	100	70	25
Singapore	0	15	40	0	0	25

Note: Each index is the average of its three components. Implementation and machinery are judgment calls scored against written rubrics; enforcement actions are counted on a log scale. The scores rank jurisdictions; they do not measure outcomes.

Source: Horizon Scan 002, Section 3

Set beside Section 2, the comparative layer sharpens the law of motion into a forecast. Governance flows downhill to operational control, and the scatter shows where the water has already pooled. The regimes above the identity line are waiting on

machinery. The regimes across it demonstrate what machinery does when nobody waits. Whether the paper catches up is the question for 2027. The answer for 2026 is that enforcement never waited.

4

SECTION 4

Implications for Institutional Leaders

Four priorities for acting before the interim regime hardens

Once we recognise how much of AI governance is being settled through procurement, infrastructure and the allocation of liability, several practical choices follow. Institutions can use these relationships to secure continuity and improve control over the systems they deploy. Public authorities can examine where the same relationships create risks that an individual firm has little reason, or insufficient influence, to address.

The four priorities below draw on familiar disciplines. Their application requires a fuller account of whose decisions the institution depends on, and what room it retains to act when those decisions change.

01

Map who can change the conditions under which you operate.

Begin with an activity the institution must be able to sustain, and follow the conditions that make it possible. A service may depend on a model provider whose capacity depends on a cloud operator, whose expansion depends on a connection decision. Following that sequence reveals how a decision several steps removed from the institution can change its costs or interrupt its work.

The Scan's infrastructure cases make the point tangible. Ireland's connection requirements and Singapore's allocation of data-centre capacity influence where investment can proceed and on what terms. An account of regulatory exposure should give such decisions a place alongside legislation. It should also reflect how established consumer and data-protection duties apply to the

service being offered, since those duties carry much of the enforcement documented here.

Responsibility for assembling this account needs to sit with someone able to bring the relevant functions together. Legal teams may see a licensing change, procurement a revised contract and operations a capacity constraint, with each holding part of the explanation for the same emerging risk. A named senior owner can ensure that these developments reach the people making investment and deployment decisions. Updates should follow consequential changes as they occur. The purpose is to give a decision-maker enough notice, and enough of the picture, to act.

02

Make the cost of losing access part of the investment decision.

As AI becomes embedded in services that customers and staff rely on, the cost of a supplier's withdrawal grows to include the work organised around it. Replacing a model may require changes to software, new testing, staff retraining and a period in which familiar tasks take longer. These costs belong in the decision to adopt the service, while alternatives remain relatively inexpensive to preserve.

Consider a planning exercise in which the primary model provider becomes unavailable for thirty days. The exercise should follow what happens to the work. Some activities may continue with a simpler system, others may return temporarily to staff, and some may need to stop. Establishing which is which gives the institution a basis for estimating the cost of interruption and deciding what continuity is worth.

30 days

without the main model provider: the planning exercise that shows what continuity is worth.

That estimate can inform negotiations over notice periods, access to records and assistance with migration, as well as investment in alternative capacity. The alternatives themselves require examination. Two suppliers may rely on the same underlying cloud infrastructure or be exposed to the same licensing decision, so apparent diversity can leave an important dependence intact.

Only once these connections are understood can a board judge whether the cost of preserving an alternative is proportionate to the disruption it would avoid. For public authorities, shared dependence across several important firms raises the further question of whether their individual continuity plans would remain workable during a common interruption.

03

Match a system's freedom to act with the ability to contain it.

When an AI system can initiate a payment or change a customer's account, oversight depends on what it is authorised to do and how quickly someone can intervene. The safeguards reviewed in this Scan give those questions practical form. Payment arrangements can specify whose permission an agent is acting on, while insurance requirements can make testing and documented controls a condition of cover.

These arrangements provide useful starting points for an institution deciding how much authority to delegate. Their value depends on the work they actually cover. A certificate concerned with management processes offers a different kind of assurance from a test of how a particular system behaves when instructions conflict or a transaction goes wrong. The institution needs to understand that scope before relying on the result.

Containment follows from the authority granted. A system permitted to spend needs enforceable

spending limits. A system able to make consequential changes needs a tested way to suspend further action, with a person who knows when to use it. Where intervention depends on a supplier, that dependence belongs in the contract and in the exercise used to test the response.

Keeping a record of these tests, permissions and interventions also makes later scrutiny more useful. An underwriter, supervisor or court can examine what the institution knew and did. The institution itself can use the same record to decide whether continued autonomy is justified as the system, its users and its operating conditions change.

04

Build evidence that can improve the next decision.

Claims about AI's effects on work often travel further than the evidence supporting them. A fall in hiring may reflect weaker demand, a change in the kinds of jobs being advertised or the introduction of a new technology. Establishing AI's contribution requires a basis for comparison and enough information about how work has changed to distinguish among those explanations.

Employers can begin by recording where AI is used, which tasks move between people and systems, and what happens to the time and quality of the work. Measures of redeployment and displacement should accompany productivity estimates, with the baseline and assumptions retained so that later results can be compared on consistent terms. Early findings may be mixed. Reorganisation carries costs, and a credible assessment needs to follow the adjustment long enough to establish whether improvements emerge.

An institution's own experience remains a partial view. Firms whose workers and customers appear in private datasets may differ substantially from those that remain unseen, particularly in economies with large informal sectors. Public statistical capacity is therefore part of the infrastructure needed to assess AI's consequences. Employers and financial institutions can contribute by supporting comparable measures and responsible research

access, while being explicit about whom their records include.

Such evidence gives workforce decisions a firmer basis and allows policymakers to examine claims across a wider range of experience. It also makes revision possible. An institution that retains the reasoning behind an early decision can return to it when the evidence changes, understand why expectations were disappointed and adjust with greater confidence.

New legislation will enter an economy already organised around many of these relationships. The work done now can give future rules something dependable to build on, while preserving the ability to change practices that experience proves inadequate. Institutions have an opportunity to make that capacity part of how they operate, through decisions whose effects they can explain and arrangements they remain able to revise.

REFERENCES

Works Cited

Sources consulted and interviews conducted, by section

Section 1

Primary and official documents

- BRICS Brazil Presidency. (2025, July 6). BRICS Leaders' Statement on the Global Governance of Artificial Intelligence. 17th BRICS Summit, Rio de Janeiro.
- BRICS Brazil Presidency. (2025, July 7). BRICS summit signs historic commitment in Rio for more inclusive and sustainable governance. brics.br.
- Brazilian Presidential Office. (2025, January 6). Brazil announces Indonesia as full member of BRICS. gov.br/planalto.
- California State Legislature. (2025). Senate Bill 53 (Wiener): Transparency in Frontier Artificial Intelligence Act.
- CEN-CENELEC. (2026, July 30). EN 18286:2026 — Artificial intelligence: Quality management system for EU AI Act regulatory purposes. European Committee for Standardization / European Committee for Electrotechnical Standardization.
- DOJ. (2026, April 24). Justice Department intervenes in xAI lawsuit challenging Colorado's algorithmic discrimination law. Department of Justice.
- DSIT and AI Security Institute. (2025, February). Tackling AI security risks to unleash growth and deliver the Plan for Change. Department for Science, Innovation and Technology.
- DSIT. (2025, September). Trusted third-party AI assurance roadmap. Department for Science, Innovation and Technology.
- FCA. (2026, January). The future of AI in UK financial services: The Mills Review. Financial Conduct Authority.
- GOV.UK. (2026, May). Plan to toughen protections for subsea internet cables amid heightened Russian activity. gov.uk.
- House of Commons Library. (2026). Artificial intelligence: Regulation (Research Briefing CBP-10003). parliament.uk.
- Regulation (EU) 2026/1744 of the European Parliament and of the Council amending Regulation (EU) 2024/1689 as regards the application dates of certain provisions (AI Omnibus). Published in the Official Journal of the European Union, July 24, 2026; entered into force July 27, 2026.
- UK Government. (2026). Interim government response to the AI Champions' AI Adoption Plans. gov.uk.

Court judgments and regulatory decisions

TJ Hooper, 60 F.2d 737 (2d Cir. 1932).

X.AI LLC v. Weiser, No. 1:26-cv-01515 (D. Colo. filed Apr. 9, 2026).

Trade press and analyses

ASPI/The Strategist. (2026, June 2). AUKUS Pillar 2 lunges for an operational capability underwater. Australian Strategic Policy Institute.

Baker McKenzie. (2026). AI governance update: FCA and agentic AI.

Data Privacy Brasil. (2026, July). Data Privacy Brasil participates in the first United Nations Global Dialogue on AI Governance.

Freshfields. (2026). FCA signals caution on AI-specific rules. Freshfields Bruckhaus Deringer.

Future of Privacy Forum. (2026, March). Incentives or obligations? The U.S. regulatory approach to voluntary AI governance standards. fpf.org.

GAICC. (2026, April). NIST AI Risk Management Framework: 2026 implementation playbook. gaicc.org.

GAICC. (2026). AI governance comparison: EU AI Act, NIST, ISO 42001. gaicc.org.

GDJF/GSA. (2026, July). What southern civil society wants from AI governance: Joint statement issued at the Global Dialogue on AI Governance. Association for Progressive Communications.

Geopolitical Monitor. (2026, April). Russian submarine operation targets UK undersea infrastructure.

Glacis. (2026, April). Guide to UK AI regulation. glacis.io.

Glacis. (2026, June). Colorado's AI Act repealed and replaced by SB 26-189. glacis.io.

IDEAS-BRICS. (2025). The Rio Declaration: Architecting a data economy for the Global South. ideas-brics.org.

Maro. (2026, March). NIST AI RMF vs ISO 42001: State safe harbor comparison. seekmaro.com.

Soeparna, I., & Sarli, A. C. (2026, May). BRICS AI governance: Comparative analysis of public-policy approaches of Brazil, Russia, India, China, and Indonesia. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6779294

techUK. (2026). DSIT's AI assurance roadmap takes its next step: techUK to lead the Industry Advisory Group of the AI Assurance Consortium. techuk.org.

techUK. (2026). Lessons from the AI procurement frontline: Beyond the contract and buying blind — Rethinking AI procurement as a governance function. techuk.org.

USSC. (2026, January). A menu of AI and autonomy options for AUKUS Pillar II. United States Studies Centre, University of Sydney.

Interview sourcing

Braidwood, J. (2026, July 1). CEO, Glacis. Interview by Gloria Chen.

McGuigan, E. (2026, September). Chair, AI Assurance Stakeholder Consortium. Correspondence with Gloria Chen.

Pigatto, J. (2026, July 2). Governance and Regulation Coordinator, Data Privacy Brasil. Interview by Gloria Chen.

Soeparna, I. (2026, July 7). Associate Professor of Law, Universitas Airlangga. Interview by Gloria Chen.

Section 2

Primary and official documents

Central Statistics Office (Ireland). "Data Centres Metered Electricity Consumption 2024" (Online ISSN 2811-5422). June 10, 2025. <https://www.cso.ie/en/releasesandpub>

lications/ep/p-dcmec/datacentresmeteredelectricity consumption2024/keyfindings

Commission for Regulation of Utilities (Ireland). "Large Energy Users Connection Policy," Decision Paper CRU/2025236. December 12, 2025.

TenneT TSO B.V. "No extra space on electricity grid in large part of Noord-Holland next decade." 2025. <https://www.tennet.eu/news/no-extra-space-electricity-grid-large-part-noord-holland-next-decade>

Hermans, S. (Minister for Climate Policy and Green Growth, Netherlands). Kamerbrief on the Grid Congestion Campaign (Aansluitoffensief netcongestie), letter to the House of Representatives. February 4, 2026 (via Stibbe).

Minister van Volkshuisvesting en Ruimtelijke Ordening (Netherlands). Voorbereidingsbesluit on hyperscale data centres (Art. 4.3, fourth paragraph, Wet ruimtelijke ordening). February 16, 2022.

Infocomm Media Development Authority / Economic Development Board (Singapore). "Green Data Centre Roadmap" (2024) and "Data Centre – Call for Application (DC-CFA2)." December 1, 2025.

Energieeffizienzgesetz (EnEfG), Federal Republic of Germany. In force November 2023; data-centre PUE and waste-heat (ERF) provisions effective July 1, 2026.

Special Act on the Promotion of Distributed Energy (대한민국 분산에너지 활성화 특별법), Republic of Korea. Effective June 14, 2024.

Ministry of Science and ICT (MSIT), Republic of Korea. Enforcement Decree to the Framework Act on AI Development (AI Basic Act). Effective January 22, 2026.

Cyberspace Administration of China (国家互联网信息办公室). Generative-AI service filings (538) and application registrations (263) as of August 31, 2025.

G42 / OpenAI. "Introducing Stargate UAE." May 22, 2025. <https://openai.com/index/introducing-stargate-uae/>

IndiaAI Mission (MeitY). IndiaAI Compute Portal, GPU deployment figures (38,231 across three 2025 rounds), CEO Abhishek Singh briefing. September 2025 (via AI CERTs News).

Norwegian Ministry of Finance (Finansdepartementet). Press release No. 52/2022, abolition of the reduced electricity-tax rate for data centres. October 6, 2022.

European Commission, Interoperable Europe. "Artificial intelligence standards and tools," status of EN 18286:2026 assessment for harmonised-standard citation. Accessed August 25, 2026.

European Commission. Implementing Decision C(2023) 3215 (M/593), standardisation request to CEN and CENELEC in support of Union policy on artificial intelligence. May 22, 2023.

BSI. "BSI becomes the first certification body accredited by UKAS and RvA to deliver certification for ISO/IEC 42001." November 17, 2025.

Court judgments and regulatory decisions

ANI Media Pvt. Ltd. v. OpenAI Inc., Delhi High Court, CS(COMM) 1028/2024, interim order (Bansal J.), July 24, 2026.

Moffatt v. Air Canada, 2024 BCCRT 149.

Beijing Internet Court, Li Yunkai v. Liu, (2023) Jing 0491 Min Chu No. 11279, November 27, 2023; Guangzhou Internet Court (Ultraman), February 8, 2024; Hangzhou Internet Court, February 10, 2025.

Autoridade Nacional de Proteção de Dados (Brazil). Preventive measure suspending Meta AI training, July 2, 2024; compliance plan accepted August 30, 2024.

Garante per la Protezione dei Dati Personali (Italy). Decision No. 755, OpenAI, €15M, December 20, 2024; Court of Rome annulment, March 18, 2026.

Segundo Tribunal Ambiental (Chile), Municipality of Cerrillos v. Coeva RM (Google Cerrillos Data Center), September 26, 2024.

Office of the Superintendent of Financial Institutions (Canada). Guideline E-23 (Model Risk Management), published September 11, 2025, effective May 1, 2027.

Trade press and analyses

DLA Piper; Philip Lee LLP; Mason Hayes Curran (CRU policy). Stibbe; NL Times (Dutch grid). Morgan Lewis; Reed Smith (Singapore). Columbia Climate Law Blog; Pinsent Masons (EnEfG). Cushman & Wakefield; King & Wood Mallesons; Lexology (Korea grid). Data Center Dynamics (UAE, Stargate). Cooley; White & Case; Future of Privacy Forum; Stimson Center; ITIF (Korea/Japan AI Acts). Chambers and Partners; The New Publishing Standard (ANI v. OpenAI). Cross-Border Data Forum (Garante). Global Network Initiative (STF Marco Civil). Trivium China; ChinaTalk; hellochinatech (CAC registry). Eco-Business; Asia Times; Nation Thailand

(Southeast Asian enforcement). Morningstar (Japanese P&C market share). Tokio Marine Holdings (insurer AI governance policy); Lumenova (EN 18286 approval date and Annex ZA scope); lawandtechnology.eu (Official Journal citation as the legally relevant event); Modulos (harmonized-standards pipeline and partial presumption); CMS Law (prEN 18286 enquiry-stage analysis).

Interview sourcing

Grunewald, Erich (Senior Researcher, Institute for AI Policy and Strategy). Email correspondence with Ashwin Telang, August 3, 2026.

Ren, Shaolei (Associate Professor of Electrical and Computer Engineering, University of California, Riverside). Written responses to questions from Ashwin Telang, August 2026.

Calo, Ryan, and Veronica Paternolli. Shadow Work and Protocols: A Contingent Case for Negligence in Agentic AI. Working paper, 2026.

Calo, Ryan (Professor of Law, University of Washington), and Veronica Paternolli (PhD researcher in Computer Science, University of Verona). Correspondence with Ashwin Telang on negligence, agent protocols, and the allocation of liability, August 2026.

Section 3

Primary and official documents

Stanford HAI, AI Index Report 2026: 47 jurisdictions with active AI-specific legislation and the enforcement-mechanism caveat (via Stanford Institute overview, July 2026).

European Commission / AI Office: GPAI enforcement powers effective August 2, 2026 (via Wilson Sonsini client alert and CNBC reporting, August 2–3, 2026); European Parliament adoption of the Digital Omnibus deferrals, June 16, 2026.

Cyberspace Administration of China: cumulative generative-AI registrations (868 services; 530 applications) as of April 30, 2026 (via Wedoany, May 13, 2026); labelling-measure enforcement notifications (via Bird & Bird TMT update, 2026).

Ministry of Science and ICT (Korea): AI Basic Act and Enforcement Decree in force January 22, 2026; grace-period terms and the March 2026 Institutional Improvement Task Force (via Cooley, KoreaTechDesk, Stimson Center).

Court judgments and regulatory decisions

Clearview AI Inc. v. Information Commissioner [2025] UKUT 319 (AAC), October 2025; ICO statement, October 8, 2025.

Court of Rome, annulment of Garante Decision No. 755 (OpenAI, €15M), March 2026; Garante decisions on Replika (€5M, May 2025) and DeepSeek (January 2025).

FTC v. Air AI (stipulated judgment, March 2026); FTC vacatur of the Rytr consent order (December 2025); FTC v. Growth Cave; Moffatt v. Air Canada, 2024 BCCRT 149.

Trade press and analyses

Jones Day, Product Liability Directive transposition tracker (Hungary complete; ~10 member states at draft stage), June 2026. Benesch and

National Law Review surveys of 2025–2026 FTC AI enforcement. Axis Intelligence and state trackers on EO 14365, Colorado SB 26-189, and the NY RAISE Act. Full row-level sourcing: Horizon Search Institute, Scan 002 original research database (methodology, enforcement register, activity and bite inputs, penalties and deferrals).

CORRECTIONS AND SOURCE CLARIFICATIONS

Revised 1 October 2026

Cable concentration (pp. 6–7 and Exhibit 11)

The figure is approximately 75% of estimated UK-US cable capacity, not measured traffic. The pre-read’s 79% has been corrected. The searchable text in Exhibit 11 has also been corrected. The parliamentary report notes that this does not imply the same proportion of UK internet activity would be lost if the cables failed.

Joint Committee on the National Security Strategy, 19 September 2025, paras. 10 and 51.

UKAS / BSI accreditation (pp. 11, 29 and 31)

BSI’s announcement is dated 17 November 2025. UKAS’s announcement is dated 15 January 2026. These are publication dates; neither notice specifies the date accreditation was granted. The text no longer treats January 2026 as the grant month.

BSI announcement, 17 November 2025.

UKAS announcement, 15 January 2026.

New Delhi added to Exhibit 2 (p. 6)

The chart and pre-read now cover five selected international AI governance forums. The New Delhi AI Impact Summit, 18–19 February 2026, produced voluntary, non-binding frameworks and principles. Its inclusion does not change the finding on binding enforcement.

Government of India, New Delhi Declaration, published 21 February 2026.

UN Dialogue: scope of the forum

The previous chart’s forecast of binding outcomes in 2027 has been removed. UN guidance describes the Dialogue as a discussion forum; it does not establish a date for binding enforcement.

United Nations, Countdown to the Global Dialogue on AI Governance: The Big Picture.