Standard Intelligence

White paper

The Tripartite Convergence

Fragmentation, crisis, and the staggered emergence of three global AI regulatory blocs

PUBLISHED

June 2026

WEB

standardintelligence.com

AUDIENCE

For enterprise, regulatory and policy communities

CONTENTS

Summary 3

1. Fragmentation is transitional, and it resolves on three timelines 5

2. The forecasts are structured judgements, rated for confidence 6
- 2.1 Decisions cannot wait for outcomes 6
- 2.2 Three methods, each with a stated limitation 6
- 2.3 Three tiers of evidence 7
- 2.4 Three levels of confidence 7

3. Between 2022 and 2026, the major jurisdictions took divergent paths 8
- 3.1 The EU moved first, and its reach is extraterritorial 8
- 3.2 The US left AI to its states, and the states diverged 9
- 3.3 China regulated AI as industrial policy and exported its norms through infrastructure 10
- 3.4 The UK diverged in rhetoric more than in practice 11
- 3.5 Elsewhere, jurisdictions kept their options open 11

4. Fragmentation's costs fall unevenly, and they build pressure for coordination 12

5. Technical standards are narrowing the gaps before legislation does 13

6. Security crises will add a layer of emergency law between 2026 and 2028 13
- 6.1 AI-accelerated vulnerability discovery has outpaced remediation 14
- 6.2 The post-quantum transition compounds the risk 14
- 6.3 The result will be a surge of poorly coordinated emergency law 15

7. US federal legislation is the hinge of the forecast 16
- 7.1 The obstacles are well documented, and three dynamics will erode them 16
- 7.2 Federal law will build on existing agencies, not a new regulator 17
- 7.3 Federal law would give the EU a counterpart for mutual recognition 17

8. Three blocs will form, each held together by a different mechanism 18
- 8.1 A bloc is a shared regulatory grammar, not a treaty 18
- 8.2 The EU-anchored bloc: rights-based, recognisable by 2026 to 2027 18
- 8.3 The China-anchored bloc: sovereignty-first, recognisable by 2027 to 2029 19
- 8.4 The Americas bloc: liability-led, recognisable by 2030 to 2032 20

9. Some major states will stay outside all three blocs, and weaker states risk being aligned by default 21

10. Between 2027 and 2031, two blocs will operate without the third 22

11. The forecast rests on eight assumptions, and the Americas bloc is its weak point 23
- 11.1 The eight assumptions 23
- 11.2 How each assumption could fail 24
- 11.3 The composite forecast 25

12. The likeliest alternative is that US federal legislation fails 26

13. Operating across three blocs creates a compliance trilemma 27

14. Blocs will become instruments of power, and the largest safety gaps sit between them 28

Where to start, and what it costs 29

Basis and Standing 30

Sources 31

Summary

Global AI regulation fragmented sharply between 2022 and 2026, as the EU, the United States and China went separate ways and a long tail of jurisdictions took positions somewhere between them.

We expect this fragmentation to be transitional rather than stable. Its costs accumulate, it leaves governance gaps, and it builds political pressure for resolution. Our forecast is that it resolves into three durable regulatory blocs (sets of jurisdictions that share a regulatory grammar and accept audit and conformity regimes across the bloc, without needing a treaty), and that they form on staggered timelines rather than on a single convergence date.

The forecasts are structured judgements with stated confidence, not statements of fact. They combine regulatory trajectory analysis, geopolitical framing and scenario modelling, grade the evidence in three tiers, and rate each claim High, Medium or Medium-Low (composite forecasts may fall between two levels). The probabilities given for alternative scenarios are judgements, to be read as rankings rather than point estimates.

Between 2022 and 2026, the major jurisdictions took divergent paths. The EU built an extraterritorial, risk-tiered Act, now amended by the Digital Omnibus. The US federal government revoked its main AI executive order and moved to challenge state AI laws, while the states legislated on incompatible definitions; China regulated application by application and exported its norms through infrastructure.

Fragmentation's costs fall hardest on the smallest actors, and they build pressure for coordination. AI fragmentation differs from GDPR-era fragmentation because the regimes disagree about what is being regulated, not only about implementation. We anticipate at least one politically visible cross-border AI harm event between 2026 and 2028.

Technical standards are narrowing the gaps before legislation does. ISO/IEC 42001, the NIST AI Risk Management Framework and the European harmonised standards give enterprises a shared vocabulary, and procurement rather than mandate is driving much of ISO/IEC 42001's adoption.

Security crises will add a layer of emergency law between 2026 and 2028. Anthropic's April 2026 disclosure that one of its models had found thousands of high-severity vulnerabilities, and the slow migration to post-quantum cryptography, are likely to provoke fragmented emergency legislation. We assess the likelihood of this surge at 60%; it complicates the path to the three blocs without preventing it.

US federal legislation is the hinge of the forecast. We expect it in 2029 to 2030, driven by the cost of the state patchwork, a competitiveness argument and a deadlock-breaking incident. It would create the counterpart the EU needs for mutual recognition.

Three blocs will form, each held together by a different mechanism:

- an EU-anchored, rights-based bloc, held together by market access and recognisable by 2026 to 2027 (High confidence);

- a China-anchored, sovereignty-first bloc, held together by infrastructure dependency and recognisable by 2027 to 2029 (between Medium and High);
- an Americas, liability-led bloc, held together by bilateral relationships with the United States and recognisable by 2030 to 2032 only if US federal legislation passes (Medium-Low).

The complete structure follows by 2032 to 2034 (Medium). The underlying direction, consolidation into a small number of blocs, is rated High.

Some major states will stay outside all three blocs, and weaker states risk being aligned by default. India is the most consequential non-aligned state. States without the capacity to build their own frameworks are likely to be drawn into whichever bloc their infrastructure or export markets belong to.

Between 2027 and 2031, two blocs will operate without the third. In this bipolar interregnum, US enterprises face the AI Act's full extraterritorial reach with no domestic framework to negotiate from. That pressure makes US federal legislation more likely, not less.

The forecast rests on eight assumptions, and the Americas bloc is its weak point. Each assumption carries a confidence rating, an evidence tier and a stated consequence if it fails.

The likeliest alternative is that US federal legislation fails. We assess that scenario at 25%; the interregnum would then harden into a durable bipolar world with a large non-aligned zone that includes the United States itself.

Operating across three blocs creates a compliance trilemma. The blocs' philosophies conflict at the level of design, their conformity evidence does not transfer, and security-layer obligations sit outside all three general frameworks.

Blocs will become instruments of power, and the largest safety gaps sit between them. No body has the mandate to coordinate a binding response to a transnational AI harm event.

For enterprises, the first step is to adopt the AI Act as the global baseline and map it to the NIST AI RMF through ISO/IEC 42001. Design for compatibility with a future US federal framework now, and treat security-layer obligations as a separate compliance domain. For policymakers, the interregnum is the window to negotiate EU-Americas mutual recognition, agree a G7 or OECD minimum floor and build regulatory capacity in non-aligned states.

This paper states the position as at 11 October 2026. Its forward-looking statements are our forecasts, not statements of fact. Basis and Standing, at the end, sets out which propositions are binding law, which are policy or guidance, and which are our own analysis.

1. Fragmentation is transitional, and it resolves on three timelines

Artificial intelligence has produced one of the most compressed and uneven regulatory responses in modern legislative history. In less than four years, several major jurisdictions moved from ad hoc, sector-specific interventions to economy-wide frameworks, for a technology whose properties change faster than legislative cycles. Each framework reflected its own legal tradition rather than a shared template.

The usual framing of this period, innovation against regulation, obscures the structural question. AI is likely to be regulated almost everywhere in the end; what matters is whether the world's regulatory architecture converges enough to be governable at a global level. In our view unbounded fragmentation is not stable, because it accumulates compliance costs, leaves governance gaps and generates political pressure for resolution.

Our forecast is that the resolution takes the form of three durable regulatory blocs, which form on staggered timelines. A bloc, as we use the term, is a set of jurisdictions that share a regulatory grammar (compatible definitions, comparable risk criteria and overlapping obligations), accept audit and conformity regimes across the bloc, and exert a pull on third countries; it needs no treaty (section 8.1). A bloc is "recognisable" when it has a stable core membership, a functioning conformity architecture and a pull on states outside it.

The three blocs reach that point at different times:

- the EU-anchored bloc is forming now and will be recognisable by 2026 to 2027;
- the China-anchored bloc is also well advanced and will be recognisable by 2027 to 2029;
- the Americas bloc depends on US federal legislation and the mutual recognition that would follow it, and will be recognisable by 2030 to 2032 only if that legislation passes.

The complete structure follows by 2032 to 2034. A single convergence date would assume that the three blocs form on comparable timescales, which the evidence does not support; it would overstate certainty about the Americas and understate how far the other two have already advanced.

Between about 2027 and 2031, the world will be in a bipolar interregnum. Two blocs will be operating while the Americas bloc has not yet formed. We treat this as a configuration with its own dynamics and risks (section 10), not as a gap between two states of the world.

2. The forecasts are structured judgements, rated for confidence

2.1 Decisions cannot wait for outcomes

Regulatory scholarship has been cautious about forward-looking claims, because elections, incidents and technological discontinuities shape regulatory trajectories and resist systematic prediction. A forecast of legislative timelines can be wrong in ways that retrospective analysis cannot.

Compliance functions, regulators and policy practitioners still have to make consequential decisions before outcomes can be observed. Investment in compliance infrastructure, standards engagement and regulatory advocacy has to be calibrated to expected futures. A forecast that states its assumptions, rates its confidence, names the conditions under which it would be revised and separates confidence in direction from confidence in timing is more useful than either unsupported assertion or silence.

2.2 Three methods, each with a stated limitation

The central claims rest on three methods, and on their agreement rather than on any one alone.

Regulatory trajectory analysis is the primary method. It identifies the direction already built into legislative programmes (the momentum of policy commitments, the institutional interests that sustain them and the compliance pressures they create) and extends it forward. It draws on path dependency: once a framework has enough stakeholder investment, reversing it usually costs more than continuing.

The AI Act illustrates the point. Behind its text sits a six-year investment by the Commission, the Parliament, twenty-seven governments, a High-Level Expert Group, an AI Office and a generation of compliance professionals, which makes reversal within the horizon very unlikely. The method's limitation is that it handles discontinuous events less well than incremental change, which the scenario analysis in section 12 partly addresses.

Geopolitical framing is the second method. It places regulatory trajectories within great-power competition, trade management and the strategic use of regulatory power, drawing on Strange's account of structural power (Strange, 1994), Drezner's work on international regulatory regimes (Drezner, 2007) and Bradford's analysis of the Brussels Effect (Bradford, 2020; defined in section 3.1). It explains features that domestic institutional analysis cannot: why the AI Act has broad extraterritorial scope, why US federal legislation is more likely under competitiveness pressure than under domestic harm alone, and why China exports its model through infrastructure dependency rather than legislative harmonisation.

Its limitation is that great-power dynamics are among the least tractable elements of any forecast.

Scenario modelling is the third method, used in sections 6 and 12. It builds internally coherent alternative futures that depart from the base case along identified uncertainties, assigns each a probability based on the assessed likelihood of its trigger

conditions, and traces the consequences. The probabilities are not derived from formal modelling; they are structured judgements calibrated against historical base rates for analogous events, and they should be read as rankings rather than point estimates.

2.3 Three tiers of evidence

The evidence varies in strength across the claims it supports, so we grade it in three tiers.

Tier 1 is enacted law, published standards and disclosed data. It includes the AI Act's text, the NIST AI RMF and post-quantum standards, ISO/IEC 42001's requirements, and Anthropic's published disclosure of what Project Glasswing found (section 6.1). We treat claims grounded here as established, with the qualification that a company's disclosure is a record of what it reported rather than independent verification.

Tier 2 is peer-reviewed literature, official records and the history of comparable episodes. It includes official impact assessments and legislative records. Claims about the Brussels Effect and the conditions for US federal legislation draw mainly on this tier.

Tier 3 is practitioner assessment, industry submissions, think-tank analysis and informed commentary. We use it for claims about compliance burden, lobbying dynamics and how frameworks not yet in force are likely to work. It is more open to motivated reasoning, and we identify the claims that rest on it.

Every forecast draws on all three tiers by extrapolation, and its confidence rating reflects both the quality of the evidence and the distance of the forecast from it.

2.4 Three levels of confidence

We use three levels, High, Medium and Medium-Low, and define each by four conditions:

- **High:** the claim rests on Tier 1 or Tier 2 evidence, analogous precedents exist, it describes a direction rather than a specific timeline (or a near-term window for a process already under way), and no identified counterfactual is likely enough to reverse its direction.
- **Medium:** the claim rests on Tier 2 evidence with significant Tier 3 support, precedents are imperfect, it describes both direction and approximate timing, and at least one counterfactual has a non-trivial chance of altering the outcome.
- **Medium-Low:** the claim rests mainly on Tier 3 evidence and extrapolation, precedents are limited, it needs a specific conjunction of conditions, and several counterfactuals together carry a significant probability. Its base-case direction nonetheless remains the single most probable outcome.

There is deliberately no "Low" level. Assumptions we hold with low confidence do not carry weight in the base case; they appear instead as the trigger conditions for the alternative scenarios in section 12. Composite forecasts in section 11.3, which combine assumptions rated at different levels, may be rated between two adjacent levels where neither fits, and are described as such.

3. Between 2022 and 2026, the major jurisdictions took divergent paths

3.1 The EU moved first, and its reach is extraterritorial

The AI Act is a deliberate attempt at durable governance. Regulation (EU) 2024/1689 entered into force on 1 August 2024, completing a project that began with the Commission's 2018 AI strategy and the High-Level Expert Group's 2019 ethics guidelines, and took legislative form in the Commission's April 2021 proposal. Its architecture is tiered by risk: it prohibits certain practices, requires conformity assessment for high-risk systems, sets transparency duties for a wider range of systems and imposes a separate set of obligations on providers of general-purpose AI models.

Its extraterritorial reach is the source of its global significance. Under Article 2(1), the Act applies to providers placing AI systems on the EU market or putting them into service there, wherever they are established; to deployers established or located in the EU; and to providers and deployers in third countries where the system's output is used in the EU. Combined with the size of the EU market, that scope creates the conditions for the Brussels Effect, the process by which EU rules become global standards through the commercial choices of firms rather than through treaties (Bradford, 2020).

The mechanism is commercial. An enterprise that wants to sell in the EU must comply, and if it also sells elsewhere, maintaining separate configurations often costs more than complying with EU rules everywhere, especially where EU revenue is substantial. Aggregated across thousands of enterprises, those decisions produce a de facto global standard.

The GDPR is the precedent. Before the GDPR applied in May 2018, commentators predicted that its extraterritorial scope would provoke resistance and a fragmented global picture. In practice several major US platforms extended GDPR-derived controls globally, and many jurisdictions enacted data protection laws that drew on the GDPR, creating the de facto adequacy club whose structure we expect the EU-anchored AI bloc to repeat (section 8.1).

Assumption 1 (section 11.1), which we rate Medium, is that the AI Act's Brussels Effect will follow a similar path. The main uncertainty is its speed and completeness.

Enforcement will set the practical reach of that effect. The AI Office, established within the Commission in January 2024 and given its tasks by the Act, supervises general-purpose AI models. Under Article 75 as amended by the Digital Omnibus (Regulation (EU) 2026/1744), it is also exclusively competent for AI systems based on a general-purpose model developed by the same provider or undertaking (subject to exceptions) and for AI systems that are, or are integrated into, very large online platforms and search engines; national market surveillance authorities supervise most other high-risk systems.

The distributed architecture has yet to prove itself. GDPR enforcement was less consistent in its early years than its design intended, with national authorities varying in resources, interpretation and appetite. If AI Act enforcement starts as unevenly, its extraterritorial reach may take longer to materialise than the base case assumes.

3.2 The US left AI to its states, and the states diverged

Federal inaction in the US has been a policy choice, not only a failure. It reflected a coherent set of political and economic interests, even as it produced outcomes that are increasingly hard to sustain.

The first Trump administration (2017 to 2021) emphasised US leadership, a deregulatory posture and resistance to binding international frameworks. Executive Order 13859 of February 2019 launched an American AI Initiative focused on research, workforce and standards leadership, and asked the Office of Management and Budget to guide agencies' regulatory and non-regulatory approaches rather than proposing new regulation. The Biden administration took a more interventionist line in Executive Order 14110 of October 2023, which directed NIST guidelines and standards, reporting by developers training frontier models above a compute threshold, and sector guidance across the FDA, FTC, CFPB and EEOC.

The second Trump administration revoked Executive Order 14110 on 20 January 2025 (Executive Order 14148). In December 2025 it issued Executive Order 14365, which set up an AI Litigation Task Force to challenge state AI laws and called for a legislative recommendation for a uniform federal framework, which the White House published in March 2026 as a non-binding recommendation favouring targeted federal preemption of state laws. As at 11 October 2026, the US had no binding horizontal AI legislation.

State legislation filled the vacuum. California legislated most actively. SB 1047, the Safe and Secure Innovation for Frontier Artificial Intelligence Models Act, would have imposed pre-deployment safety testing on frontier developers, but Governor Newsom vetoed it in September 2024 on the ground that it regulated by model size and cost rather than by the risk of the deployment.

The veto did not end California's ambition. SB 53 (2025) imposed transparency, safety-framework and critical safety incident reporting duties on frontier developers, and further bills addressed definitions, content disclosure and automated decisions.

Colorado's AI Act (SB 24-205), signed in May 2024, was the most extensive horizontal state framework, requiring impact assessments, transparency disclosures and consumer appeal mechanisms. In May 2026 Colorado repealed and replaced it with SB 26-189, which drops impact assessments in favour of documentation, notice, disclosure after adverse decisions and human review, from January 2027.

Texas considered a similar structure in HB 1709 but enacted a much narrower law, HB 149 (2025), focused on intentional discrimination and government use. New York City's Local Law 144 of 2021, requiring bias audits of automated employment decision tools, was an early and influential sector-specific measure.

The states multiplied obligations and fragmented the vocabulary beneath them. State frameworks used different definitions of "AI system", "high-risk", "automated decision" and "deployer", different thresholds, different conformity methods and different enforcement arrangements. For an enterprise operating nationally, the burden approached that of several foreign jurisdictions at once, without a single authoritative counterpart.

Federal sector frameworks did useful work within their domains. The FDA's action plan for AI- and machine-learning-based software as a medical device addressed device

risks with technical sophistication, and the FTC used its Section 5 powers against consumer harm. The NIST AI Risk Management Framework, published in January 2023, is voluntary but influential: its four functions (Govern, Map, Measure and Manage) supply a vocabulary used in federal agency guidance and some state legislation, and we expect its risk-based, outcome-oriented approach to shape any federal law.

3.3 China regulated AI as industrial policy and exported its norms through infrastructure

Chinese AI regulation serves several purposes at once. It is consumer protection, industrial policy, national security governance and political-control infrastructure, operating in a state-directed economy in which regulator and regulated enterprise work on a national project together rather than as adversaries. Reading it as an attempt to reach Western objectives by other means misreads its architecture.

China regulated application by application. Rather than build a horizontal framework first, it issued rules at the frontier of deployment: the Provisions on the Management of Algorithmic Recommendations (in force March 2022), the Provisions on Deep Synthesis Internet Information Services (January 2023), addressing synthetic media, and the Interim Measures for Generative Artificial Intelligence Services (August 2023), which set a filing regime for public-facing generative AI in China. The framework is still being extended, with Measures for Labelling AI-Generated Synthetic Content in force from September 2025 and an AI provision added to the amended Cybersecurity Law from January 2026.

Each rule addressed a real problem while building in requirements that serve political control: transparency that gives the state visibility into recommendation systems, labelling of synthetic content, and security assessments that examine training data for politically sensitive content. Providers of services with public-opinion attributes or social-mobilisation capacity must file their algorithms with the Cyberspace Administration of China and pass a security assessment before launch. All providers must ensure their systems do not generate content that endangers national security.

The difference from the EU is constitutional, not technical. In our reading, Chinese AI ethics frameworks use the vocabulary of fairness, transparency and human oversight, but within a conception of the relationship between individual and state in which individual rights do not constrain the state's authority to manage AI risk as EU fundamental rights law constrains the AI Act. For bloc formation, this means the China-anchored bloc will not converge with the others on fundamental rights.

China exports its model through infrastructure dependency. The Digital Silk Road, announced in 2015 as part of the Belt and Road Initiative, is a sustained programme of digital infrastructure investment across Central Asia, the Middle East, Sub-Saharan Africa and Southeast Asia. Huawei telecommunications infrastructure, ZTE network equipment and Chinese-built cloud and surveillance systems embed Chinese choices about standards, protocols and security architecture, and migrating away from them is costly.

As a result, recipient states find Chinese requirements (data localisation, security assessment and algorithmic transparency to the regulator) easier to accommodate than Western alternatives that require different technical architectures.

3.4 The UK diverged in rhetoric more than in practice

The UK matters both as a major AI economy and because its choices shape the EU-anchored bloc. Its March 2023 white paper, A Pro-Innovation Approach to AI Regulation, set a deliberately divergent course: no horizontal equivalent of the AI Act, but existing sector regulators applying AI principles within their mandates. The AI Safety Institute, later renamed the AI Security Institute, was to address frontier risks through research and voluntary engagement.

Commercial reality keeps UK and EU governance closer than the rhetoric suggests. UK enterprises whose AI systems are placed on the EU market, or whose outputs are used in the EU, are subject to the Act extraterritorially, and, in our assessment, UK financial institutions work under model risk expectations that increasingly reference EU standards. UK life sciences companies whose AI is a medical device must meet the EU Medical Devices Regulation to sell in the EU. We expect, with Medium confidence, that sector-by-sector application of existing UK frameworks will over time produce obligations broadly compatible with the AI Act without formal alignment.

The UK-EU data protection adequacy arrangement survived Brexit and was renewed in December 2025, and it is the template for alignment without harmonisation. Our forecast that UK-EU AI alignment deepens between 2027 and 2031 does not require the UK to adopt the Act. It requires only that the cost of divergence in high-value sectors, together with a normalising relationship, produces convergence at the operational level.

3.5 Elsewhere, jurisdictions kept their options open

Between the three poles, a large group of jurisdictions combined alignment, divergence and deliberate ambiguity. Canada's Artificial Intelligence and Data Act, part of Bill C-27, was the most complete attempt to replicate the AI Act's risk-tiered framework in a common law jurisdiction. It died on the order paper when Parliament was prorogued in January 2025, and as at 11 October 2026 Canada had no federal AI statute.

Japan has favoured the human-centred framing of its Society 5.0 framework and agile governance, combining voluntary guidelines with a 2025 AI Promotion Act that sets a framework for government measures without penalties. Its membership of the G7, which produced the Hiroshima AI Process and its Code of Conduct, and of the OECD links it to Western frameworks while preserving domestic flexibility. Singapore's Model AI Governance Framework typifies the ASEAN preference for voluntary, principle-based frameworks compatible with several international standards.

India's Digital Personal Data Protection Act 2023 passed only after years of delay and three drafts, which shows how hard consensus on AI-adjacent data governance is in a large federal democracy. Its AI governance took the form of non-statutory AI Governance Guidelines, released in November 2025, and its place outside the blocs is taken up in section 9.

4. Fragmentation's costs fall unevenly, and they build pressure for coordination

Fragmentation distorts markets, and it shifts market power towards incumbents. The overhead of complying with several incompatible frameworks falls unevenly. Large enterprises with dedicated regulatory functions can absorb it; smaller enterprises and start-ups usually cannot.

Three kinds of actor bear the heaviest share. High-growth start-ups entering international markets spend a disproportionate share of limited capital on compliance, which entrenches incumbents, and public bodies face the complexity without the commercial incentive that drives private investment in expertise. Civil society organisations that use AI for advocacy and service delivery may find the overhead excludes them altogether, concentrating the benefits of AI in well-resourced commercial actors.

AI fragmentation is costlier in kind than the data protection fragmentation the GDPR replaced. When it proposed the GDPR in January 2012, the Commission estimated that replacing divergent national rules with a single regulation would save businesses around EUR 2.3 billion a year, mainly by removing duplicated administrative requirements such as notifications. Those national laws were different implementations of a shared framework. The AI regimes rest on different philosophies, with different definitions of the regulated object, different risk criteria, different conformity methods and different institutional counterparts.

Coordination across legal cultures rarely happens without a precipitating event. The Basel Committee was set up after the failure of Bankhaus Herstatt in 1974; the Montreal Protocol followed the discovery of the Antarctic ozone hole; and the NIS2 Directive was shaped in part by the WannaCry and NotPetya attacks of 2017, which struck after the original NIS Directive had been adopted. The Cambridge Analytica revelations of March 2018, two months before the GDPR came into application, sharpened attention to data protection enforcement.

We anticipate at least one significant cross-border AI harm event between 2026 and 2028 that exposes the governance gap in a politically visible way. The forecast is probabilistic rather than specific. In our assessment the base rate of significant AI harm events is rising as systems are deployed more deeply into consequential domains, and divergence creates real opportunities for harm.

The likeliest candidates exploit divergence structurally, where a design or deployment choice is rational precisely because no single jurisdiction can close the gap. Financial AI operating across borders is one: a system that complies in each jurisdiction individually can still create instability through aggregate effects that no single regulator has the mandate or information to assess. AI-generated disinformation timed to a major election, with a traceable cross-border origin, would create the conditions for emergency coordination.

5. Technical standards are narrowing the gaps before legislation does

Convergence is under way at the level of technical standards, and it will prepare the ground for legislative alignment. ISO/IEC 42001:2023, the first international standard for AI management systems, provides a vocabulary and process framework that crosses jurisdictional boundaries. It specifies requirements for establishing, maintaining and improving an AI management system, including risk assessment, impact assessment, data management and human oversight, in a form that maps onto the requirements of all three blocs without matching any of them.

Procurement is driving ISO/IEC 42001 adoption more than mandate. In our experience, large enterprises running vendor due diligence, public procurement frameworks and supply-chain risk programmes are beginning to treat certification as a baseline qualification. The result is a population of certified enterprises whose compliance posture is partly harmonised across jurisdictions before any legislative alignment.

The NIST AI RMF spreads by a different route. As a NIST product it carries credibility with US enterprises and federal procurement, and its use in state bills gives it a de facto regulatory status at home. Its international influence comes from NIST's role in ISO/IEC standards development and from the investment non-US enterprises have made in NIST-aligned programmes to reach the US market.

European harmonised standards add a third layer. Under Article 40 of the AI Act, a high-risk system that complies with a harmonised standard whose reference has been published in the Official Journal is presumed to conform with the requirements that standard covers. CEN and CENELEC published the first standard developed for this purpose, on quality management (EN 18286), in July 2026, but as at 11 October 2026 we had found no AI Act harmonised standard cited in the Official Journal.

Once references are cited, the presumption will give enterprises inside and outside the EU a strong reason to adopt them. Alignment between the European standards, ISO/IEC 42001 and the NIST AI RMF is incomplete but developing, narrowing the gaps between frameworks even while legislation diverges. We think technical standards receive less weight in governance analysis than they deserve, and the record from 2022 to 2026 supports that view.

6. Security crises will add a layer of emergency law between 2026 and 2028

AI systems that can discover software vulnerabilities autonomously and at scale are a distinct driver of regulation. They differ from the cost and incident pressures in section 4

in three ways: the capability has been demonstrated rather than theorised, it is spreading on a timescale of months (section 6.3), and the cybersecurity governance it meets was not designed to respond at that pace.

6.1 AI-accelerated vulnerability discovery has outpaced remediation

In April 2026, Anthropic disclosed that its Claude Mythos Preview model had found thousands of high-severity vulnerabilities, including some in every major operating system and web browser. The examples included a now-patched 27-year-old vulnerability in OpenBSD and a chain of Linux kernel vulnerabilities that escalated ordinary user access to full root control. The disclosure launched Project Glasswing, a defensive initiative with AWS, Apple, Broadcom, Cisco, CrowdStrike, Google, JPMorganChase, the Linux Foundation, Microsoft, NVIDIA and Palo Alto Networks, intended to apply the capability defensively before adversaries obtained equivalent tools.

Anthropic committed up to USD 100 million in usage credits and said it did not plan to make Mythos Preview generally available. As at 11 October 2026, it offered its later Mythos models with advanced cyber capabilities to qualifying security professionals through its Cyber Verification Program.

At disclosure, AI-assisted discovery had outpaced remediation, and it can, for now, outpace the capacity of enterprises, open-source maintainers and public bodies to patch. At announcement, Anthropic reported that more than 99% of the vulnerabilities it had found were not yet patched. Cisco, a launch partner, was quoted in the announcement as judging that AI capabilities had crossed a threshold that fundamentally changes the urgency of protecting critical infrastructure.

Glasswing also exposes an asymmetry with direct regulatory consequences. About 50 organisations took part at launch, and Anthropic added about 150 more in June 2026, bringing the total to some 200. That is still a small fraction of the enterprises, public bodies and infrastructure operators exposed to the same vulnerabilities. Rules that set minimum security standards for every participant cannot work as intended if the AI capability needed to meet those standards is available only to the largest and best-resourced organisations.

6.2 The post-quantum transition compounds the risk

NIST published its first post-quantum cryptography standards in August 2024: FIPS 203, a module-lattice key-encapsulation mechanism based on CRYSTALS-Kyber; FIPS 204, a module-lattice digital signature standard based on CRYSTALS-Dilithium; and FIPS 205, a stateless hash-based signature standard based on SPHINCS+. They provide a sound basis for migrating away from RSA and elliptic-curve cryptography, both of which a sufficiently powerful quantum computer running Shor's algorithm could break.

The governance problem is the pace of migration. NIST notes that deploying today's public-key infrastructure took almost two decades, and it plans to deprecate and remove quantum-vulnerable algorithms from its standards by 2035. The lag opens a "harvest now, decrypt later" window, in which state adversaries collect encrypted data today (financial records, health data, communications and AI model weights) to decrypt once quantum capability matures.

The risk meets AI governance in two places. AI systems process, generate and store data whose long-term confidentiality the quantum threat may compromise, so a framework that ignores the cryptographic protection of AI-processed data is incomplete.

AI is also increasingly deployed in critical infrastructure (grid management, transaction monitoring and healthcare diagnostics), the highest-priority targets for harvest-now, decrypt-later operations, and no single framework yet addresses that compound risk coherently.

6.3 The result will be a surge of poorly coordinated emergency law

We treat this surge of emergency law as part of the base case rather than as an alternative to it, and label it Scenario E for comparison with the alternatives in section 12. **We assess the likelihood that it unfolds alongside the base case at 60%.** It complicates the path to the three blocs without preventing it.

The proliferation window is likely to close before any jurisdiction legislates for AI-scale vulnerability discovery. Anthropic warned at launch that it would not be long before such capabilities proliferate, possibly beyond actors committed to deploying them safely, and that the fallout for economies, public safety and national security could be severe; in June 2026 it said it expected many other AI companies to have comparable models within six to twelve months. NIS2's requirements were calibrated to a pre-AI environment, CIRCIA's reporting regime does not yet address AI-discovered vulnerabilities in third-party software, and China's framework has no specific provision for disclosure of AI capability.

Legislative machinery in all three jurisdictions runs on timescales of eighteen months to three years. On Anthropic's own assessment, proliferation will be faster.

The post-quantum transition is a second emergency trigger. It could be set off by a public government assessment that a credible quantum decryption capability has been demonstrated, which would be enough to prompt mandatory migration timelines without any incident, or by disclosure that harvested financial, health or classified data had been decrypted. NIST has written the standards; the legislative questions are when governments make migration timelines mandatory, and whether they do so together or separately.

We expect intense, overlapping and poorly coordinated activity across all three jurisdictions between about 2026 and 2028:

- in the EU, an extension of the Cyber Resilience Act to address AI-specific requirements, with member states adding national emergency measures that go beyond the harmonised floor and fragment the internal market;
- in the US, executive orders setting post-quantum timelines for federal contractors and critical infrastructure, conflicting state AI security statutes, and NIST guidance binding only on federal contractors;
- in China, extension of CAC security assessment to the disclosure of AI-discovered vulnerabilities, and tighter domestic cryptographic requirements that in effect close regulated sectors to foreign providers.

Emergency legislation tends to outlast the emergency. Post-2001 surveillance statutes and data retention regimes show the pattern. Security-specific rules will harden into permanent law at different rates in different jurisdictions, creating a separate layer that has to be managed alongside the general AI frameworks rather than absorbed into them.

The capability asymmetry that Glasswing exposed is likely to become embedded rather than corrected. A rule requiring AI-assisted scanning of critical infrastructure mandates a capability that, in our assessment, most operators do not have and cannot acquire without significant investment. The likely result is compliance by well-resourced enterprises and structural non-compliance by the rest.

Post-quantum migration will create differences that resist mutual recognition. Cryptographic incompatibility is built into hardware security modules and protocol implementations, which governance-level agreements cannot abstract away. A system that must support both the NIST standards and China's SM-series algorithms needs different architectures, not merely different documents.

A revived EU-US Trade and Technology Council, or a successor dialogue on cybersecurity, could by the early 2030s fold the security layer into the wider alignment process. That timeline runs to the end of our main horizon, so enterprises should treat security-layer compliance as a distinct domain from 2026 to 2035.

7. US federal legislation is the hinge of the forecast

7.1 The obstacles are well documented, and three dynamics will erode them

The obstacles to federal AI legislation are structural. The US has failed to enact federal data privacy legislation for more than twenty years despite bipartisan agreement in principle, and AI is structurally harder. The division of authority between federal and state government exposes any preemptive legislation to litigation.

Acting through existing agency powers is slow, uncertain and legally fragile. Rulemaking demands substantial procedure, and rules face review under the major questions doctrine reinvigorated in West Virginia v. EPA (2022) and, since Loper Bright Enterprises v. Raimondo (2024) ended Chevron deference, the courts' own reading of the statute.

Three dynamics will gather force between 2026 and 2029. The first is the exhaustion of the state-level experiment. By 2028, enterprises operating under the laws of Colorado, California, New York, Texas and a growing roster of other states will have documented enough cost to change industry's lobbying position.

Privacy offers a precedent. Major technology companies moved to supporting comprehensive federal privacy legislation with preemption once state laws multiplied, because the patchwork had become more costly and less predictable than a federal floor.

The second is a competitiveness argument. The deregulatory instinct that has blocked legislation draws on a narrative in which regulation constrains American leadership. By 2028 that narrative will face a credible counter: that US enterprises facing incompatible state obligations, and the AI Act's extraterritorial reach without a federal counterpart, are disadvantaged in standards negotiations and trade discussions and lack a coherent interlocutor for regulatory cooperation.

The third is an incident. We do not predict a particular event, but we expect a politically attributable harm event to complete the conditions for federal legislation rather than

to create them. The analogy is not with disasters that create law from nothing but with incidents that break deadlocks in processes already substantively advanced, as the 2008 financial crisis released years of accumulated reform work into the Dodd-Frank Act.

7.2 Federal law will build on existing agencies, not a new regulator

US federal legislation will carry the imprint of American administrative law, constitutional constraint and a decade of accumulated regulatory activity. It will not be modelled on the AI Act. We expect it to combine three elements.

A risk-based horizontal classification. Shared definitions, above all a single definition of "AI system" that supersedes the state-level variants, and a framework of high-risk categories derived largely from the NIST AI RMF's risk mapping. That common vocabulary would let sector regulators implement requirements within their existing mandates without definitional conflicts.

Coordinated sector-specific implementation. Rather than a horizontal regulator like the AI Office, which would need new enabling legislation and would face constitutional challenge, the framework would designate lead agencies within existing domains: the FDA for medical AI, the CFPB for consumer financial AI, the FTC for consumer-facing systems and the SEC for investment AI. A coordinating body, most likely NIST's Center for AI Standards and Innovation (formerly the AI Safety Institute) or a successor, would provide horizontal guidance without primary enforcement authority.

Federal minimum obligations for high-risk deployments. These would include mandatory incident reporting to the sector regulator and the coordinating body, documentation requirements compatible with (though not identical to) the AI Act's technical documentation under Article 11, and a list of prohibited applications narrower than the AI Act's. We expect that list to cover facial recognition in public spaces without judicial authorisation, social scoring and real-time biometric surveillance in defined contexts, areas where we think bipartisan consensus is achievable.

Preemption will be the most contested element. The most likely resolution is conditional preemption. States could keep requirements in areas the federal framework does not cover, but could not impose stricter requirements in defined high-risk categories where federal authority is clear.

7.3 Federal law would give the EU a counterpart for mutual recognition

The EU-US relationship on AI is asymmetric. The EU has an economy-wide binding framework; the US has executive orders, sector guidance and a voluntary risk management framework. The Trade and Technology Council, which has not met at ministerial level since April 2024, produced shared commitments on principles (transparency, accountability, human oversight and non-discrimination) but could not progress to operational mutual recognition without comparable US domestic obligations.

Federal legislation would change this. It would create a US counterpart with authority to negotiate mutual recognition, and a basis for comparing the two frameworks obligation by obligation rather than principle by principle. The negotiation of the EU-US Data Privacy Framework is the closest precedent, showing both that transatlantic alignment is achievable without full harmonisation and that it is fragile under judicial challenge.

The two frameworks would share much and differ on fundamentals. They would converge on compatible risk classification, overlapping transparency and documentation duties, a shared commitment to human oversight, comparable incident reporting and broadly similar prohibitions. The tensions are real: the EU's fundamental rights framing creates individual rights enforceable in court that the US liability model does not replicate, notified-body conformity assessment (where the AI Act requires it) would not be mirrored, and GDPR-based data governance would be handled differently.

We therefore expect a mutual recognition framework rather than a harmonisation agreement, with significant carve-outs on both sides for defence, intelligence and law enforcement.

8. Three blocs will form, each held together by a different mechanism

8.1 A bloc is a shared regulatory grammar, not a treaty

The closest precedent for an AI bloc is the GDPR adequacy club. As at 11 October 2026, the EU had recognised adequate protection in Andorra, Argentina, Brazil (January 2026, reciprocated by Brazil), Canada (commercial organisations), the Faroe Islands, Guernsey, Israel, the Isle of Man, Japan, Jersey, New Zealand, South Korea, Switzerland, the United Kingdom (renewed in December 2025), Uruguay and the United States (limited to organisations in the Data Privacy Framework), and, as an international organisation, the European Patent Organisation (July 2025). The EEA countries apply the GDPR itself through the EEA Agreement.

No multilateral treaty created this club. It emerged from bilateral decisions, domestic reform and the incentive of market access.

We use "bloc" in the same functional sense. A bloc is a set of jurisdictions that share three things:

- a common regulatory grammar: compatible definitions, comparable risk criteria and overlapping obligations, whether written into law or carried by norms built into shared infrastructure;
- audit and conformity regimes accepted across the bloc, through mutual recognition or through members accommodating the anchor jurisdiction's regime;
- a pull on third countries, through market access or infrastructure dependency.

The value of the definition is that it picks out the forces that make some alignments self-reinforcing and some non-alignments costly to sustain.

8.2 The EU-anchored bloc: rights-based, recognisable by 2026 to 2027

The bloc's core is the EU27 and the EEA states, and its orbit is already wide. By 2027 we expect its orbital membership (jurisdictions whose frameworks are compatible enough for mutual recognition or de facto equivalence) to include:

- the United Kingdom, driven by market access and digital trade;

- Canada, through its existing GDPR adequacy for commercial organisations, at least until an Americas bloc forms (section 8.4);
- Japan, through bilateral digital partnership agreements and the Hiroshima AI Process;
- South Korea, through its AI Basic Act and its adequacy decision;
- Australia, through privacy reform and trade relationships.

Four features define the bloc:

- a rights-based framing that makes fundamental rights (dignity, privacy, non-discrimination and due process) enforceable constraints on AI design and deployment;
- a precautionary approach, in which the provider must show conformity before placing a high-risk system on the market and deployers carry separate duties of use and oversight;
- conformity assessment that is largely internal for Annex III systems but involves notified bodies for certain biometric systems and for products covered by sectoral legislation (Article 43);
- data governance obligations grounded in the GDPR and referenced throughout the AI Act.

The Act's high-risk obligations apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems (Article 113, as amended by Regulation (EU) 2026/1744). Under Article 111(2) as amended, high-risk systems already on the market before those dates are caught only if their design changes significantly afterwards, although systems intended for use by public authorities must comply by 2 August 2030.

Distributed enforcement will produce variation, and with it room for arbitrage. National authorities enforce most of the Act, with the AI Office coordinating and directly supervising general-purpose AI models and the systems within its Article 75 competence (section 3.1). That arrangement creates predictable variation in rigour across member states, as GDPR enforcement did, producing both compliance uncertainty and opportunities to locate in lighter-touch jurisdictions.

The bloc's main political vulnerability is the competitiveness argument. Critics argue that conformity assessment imposes costs that disadvantage EU developers against US and Chinese competitors. The concern has some empirical basis, but in our view it is overstated: it leaves out the market-access benefits of compliance, and consistently enforced rules can push firms to build compliance capability that becomes an advantage in risk-sensitive sectors.

8.3 The China-anchored bloc: sovereignty-first, recognisable by 2027 to 2029

This bloc holds together differently from the other two. The EU-anchored and Americas blocs cohere around a shared regulatory grammar. The China-anchored bloc coheres mainly around shared infrastructure and the norms built into it, with formal alignment secondary, and its members accommodate China's security assessment rather than negotiate mutual recognition (section 3.3). It meets the definition in section 8.1 by that second route.

China's domestic framework is coherent on its own terms and incompatible with the others. Its layers (application-specific regulations, sector security requirements, and the overarching Cybersecurity Law and Data Security Law) form a technically

sophisticated whole. For regulated categories of systems and data, it requires domestically approved cryptographic standards (SM2, SM3 and SM4), storage of regulated data on domestic servers and security assessment, creating a governance zone structurally incompatible with EU and US requirements in ways mutual recognition cannot resolve.

Its orbit spans three regions:

- in Southeast Asia, Cambodia, Laos and Myanmar have deepened their dependence on Chinese providers with little counterbalancing Western investment;
- in the Gulf, Saudi Arabia and the UAE pursue diversified strategies that combine significant Chinese investment with US cloud investment, and are best described as hedgers that lean towards the China-anchored bloc in specific domains without full alignment;
- in Sub-Saharan Africa, Chinese telecommunications and data centre investment is most extensive in countries including Ethiopia, Kenya, Nigeria and South Africa.

Dependency does not determine alignment, but it sets the path of least resistance. Some states will invest in migrating away as pressure to align grows. For states with limited technical capacity and constrained finances, however, migration is costly, and accommodating the norms already built into deployed infrastructure is the easier course.

8.4 The Americas bloc: liability-led, recognisable by 2030 to 2032

The Americas bloc depends on US federal legislation and would form around it. Its members would be drawn in through bilateral relationships with Canada (through the USMCA digital trade chapter and any successor to the lapsed Artificial Intelligence and Data Act), with Mexico (through the USMCA and supply-chain dependency), and with a group of OECD-aligned Latin American states including Brazil, Chile, Colombia and Argentina. Canada thus appears in both blocs: it sits in the EU-anchored orbit while no Americas bloc exists, and we expect it to move towards the Americas bloc once one forms.

The bloc's character reflects a different legal culture. Its main enforcement mechanism would be accountability after the event, through tort law, FTC enforcement and sector agency action, rather than conformity assessment before it. Its built-in headroom for innovation (lighter obligations for non-critical applications, reliance on voluntary standards and safe harbours for systems that comply with technical standards) reflects a preference for permitting deployment and managing harm over withholding deployment until safety is shown.

Compliance would cost less on average but vary more. Without mandatory conformity assessment before deployment, absolute compliance costs are lower. An enterprise that deploys a high-risk system without adequate testing faces large liability exposure, while one that invests in a thorough internal programme can comply more cheaply than under the EU's model.

EU-Americas relations would be managed coexistence rather than harmonisation. A revived Trade and Technology Council, or a successor EU-US dialogue, would provide the machinery for managing differences, and mutual recognition would reduce but not remove the compliance differential (section 7.3).

Brazil is the swing state. It is the largest potential member outside North America, and since January 2026 it has held mutual adequacy with the EU; its Lei Geral de Proteção de Dados already aligns closely with the GDPR, which pulls it towards the EU-anchored bloc. We expect its ties across the Americas, including with the United States, to draw it into the Americas bloc in the end, but the call is close, and Brazil's choice will influence AI governance across Latin America.

9. Some major states will stay outside all three blocs, and weaker states risk being aligned by default

India is the most consequential non-aligned state. With 1.4 billion people, a fast-growing AI sector and a tradition of strategic autonomy rooted in the Non-Aligned Movement, it has both the capacity and the will to refuse membership on terms set by others. Its governance has been designed to preserve policy space: the Digital Personal Data Protection Act 2023 is conceptually aligned with the GDPR but structurally different in ways that create friction for cross-border data operations, and the November 2025 AI Governance Guidelines prefer a principles-based approach to the EU's prescriptive conformity model.

India prefers to shape norms through multilateral forums rather than join a bloc. Its 2023 G20 presidency produced the AI section of the G20 New Delhi Leaders' Declaration, and it hosted the AI Impact Summit in February 2026.

We expect India's non-alignment to hold within the 2032 to 2034 horizon, with caveats. Its non-alignment is not neutrality; it engages with all three blocs on different terms in different sectors. Indian pharmaceutical companies in EU markets will need to comply with the AI Act, Indian financial services firms serving US clients align with US expectations, and some government infrastructure deployments may align with Chinese technical standards.

That multi-directional, partial alignment is functionally different from membership. It is sustainable as long as India keeps enough bargaining power in each bilateral relationship.

ASEAN is internally divided. Singapore and Vietnam are the most likely to align closely with the EU-anchored or Americas blocs, given their positions as a financial centre and, in our assessment, a US-aligned manufacturing hub. Indonesia, with some 285 million people and a developing AI sector, may have the scale for a domestically anchored framework.

The governance gap is concentrated in states that lack the capacity to build their own frameworks. Such states default to whichever external framework costs least in the short term. Where infrastructure is mainly Chinese-built, that default may be the China-anchored bloc's norms; where enterprises mainly serve EU markets, the AI Act's extraterritorial reach may pull the state towards the EU framework without any domestic action.

This passive alignment, drawn into a bloc through dependency rather than choice, is the condition most likely to leave lasting gaps in regulatory capacity.

10. Between 2027 and 2031, two blocs will operate without the third

The staggered timelines create an intermediate configuration with its own logic and risks. This is the bipolar interregnum introduced in section 1: the period, assessed at about 2027 to 2031, in which the EU-anchored and China-anchored blocs are both recognisable while the Americas bloc has not yet formed. It begins when both blocs are recognisable (2027 is the earliest plausible date for both) and ends when the Americas bloc is (2030 to 2032 in the base case, and only if US federal legislation passes).

Our base case is an interregnum of three to five years. Delayed US legislation could extend it into the mid-2030s.

The interregnum makes US federal action more likely, not less. During it, US enterprises in EU markets face the AI Act's full extraterritorial obligations with no domestic framework to provide a baseline or to anchor a mutual recognition negotiation. Large US technology enterprises in effect comply with the Act globally, the Brussels Effect at full strength, while having no domestic equivalent to offer in return for relief from its most burdensome provisions.

That is the condition in which industry lobbying shifts from resisting federal regulation to asking for it. The calculation changes from "no regulation is better than federal regulation" to "federal regulation we can shape is better than EU regulation imposed from outside".

The China-anchored bloc confronts US enterprises with a binary choice. They can invest in maintaining access under Chinese technical standards or accept effective exclusion. The choice is sharpest during the interregnum, because there is no Americas mutual recognition framework to point to.

Third countries face a sharper choice under two blocs than they will under three. In a three-bloc world, a state trading with both Western blocs can keep compatible alignment with each. In the interregnum, a state that cannot afford real non-alignment must choose between the EU's rights-based architecture and China's sovereignty-first model. That pressure will sort some hedging states into firmer alignment, hardening the configuration the eventual Americas bloc has to accommodate.

The interregnum carries a specific safety risk. In a three-bloc world, EU-US mutual recognition partly bridges the two sides' safety assessment methods, and shared Five Eyes and OECD relationships provide channels for cross-bloc incident response. In the interregnum there is no such bridge on the Americas side: US enterprises deploy under state laws and voluntary standards without a federal safety framework, and the AI Office has no coherent US interlocutor in a significant cross-border incident.

The interregnum calls for a two-part enterprise strategy. The operational priority is AI Act compliance as the de facto global standard. The strategic task is to design systems now for compatibility with a future US federal framework, so that adapting to US requirements is efficient rather than a retrofit; the closing section sets out both steps.

11. The forecast rests on eight assumptions, and the Americas bloc is its weak point

11.1 The eight assumptions

Each assumption below carries a confidence rating on the scale in section 2.4, the evidence for it and its main vulnerability.

Assumption 1: the Brussels Effect holds under geopolitical strain (Medium). We assume the Brussels Effect works for AI as it has for data protection, with EU rules gaining extraterritorial influence through market access and leading multinationals to adopt EU-compliant architectures globally. It could weaken if US-EU trade tensions escalate after 2026, if early AI Act enforcement proves inconsistent, or if a large alternative market bloc changes the cost-benefit of EU compliance. We rate it Medium rather than High because the Act's extraterritorial mechanics, graduated by risk tier, are more complex than the GDPR's single horizontal standard, and because deliberate strategic divergence is a real possibility in the late 2020s.

Assumption 2: US federal AI legislation passes in 2029 to 2030 (Medium-Low). This rests on three conditions: that state-level proliferation raises compliance costs enough to move industry from resisting federal preemption to preferring it; that a competitiveness framing opens bipartisan support; and that no major geopolitical event disrupts the calendar. The rating reflects the base-rate problem set out in section 7.1, the failure over more than twenty years to pass federal privacy legislation. Its failure is Scenario B, the main alternative (section 12).

Assumption 3: China's regulatory orbit remains coherent (Medium). We assume that China's export model, which combines algorithmic regulation, mandatory domestic standards and infrastructure dependency through Belt and Road partnerships, produces a coherent bloc with identifiable and fairly stable membership. It is vulnerable to US export controls that make China's ecosystem less attractive, to domestic economic pressure that reduces outward infrastructure investment, and to partner states facing binary alignment pressure. Scenario C, at 20%, explores the orbit shrinking.

Assumption 4: post-quantum risk materialises as a near-term catalyst (Medium-Low). We treat post-quantum risk as a near-term pressure point, particularly through harvest-now, decrypt-later attacks on AI-processed data. This rests on two nested claims: that state quantum capability advances faster than enterprise migration to the NIST standards, and that a demonstrable incident or a credible intelligence assessment creates political urgency for intervention rather than continued voluntary migration. Expert estimates of when a cryptographically relevant quantum computer will exist vary by more than a decade, so the timing is uncertain even if the direction is sound.

Assumption 5: Glasswing-type disclosures produce binding regulatory responses (High). We assume that AI-accelerated vulnerability discovery, of the kind Anthropic disclosed for Claude Mythos Preview (section 6.1), will produce binding regulatory responses rather than being absorbed by voluntary frameworks. The precedents are extensive: GDPR breach notification, NIS2's obligations for critical infrastructure, CIRCIA and existing disclosure frameworks all show that governments do not leave significant security risks to voluntary management. The assumption is stated at the level of category (some binding response will follow), not of specifics, because its nature, timing and coordination remain uncertain.

Assumption 6: standards bodies pre-harmonise but do not substitute for legislation (High). We treat ISO/IEC 42001, the NIST AI RMF, the CEN and CENELEC harmonised standards and ITU work as precursors to legislative alignment rather than self-sufficient governance. Two decades of data protection experience support this: ISO/IEC 27001 was published more than a decade before the GDPR applied, and it did not prevent binding obligations. Standards lower the cost of compliance once legislation arrives and provide a shared technical vocabulary, but they do not replace the political act of making binding law.

Assumption 7: India remains non-aligned within the horizon (Medium). We expect India to join none of the three blocs by 2032 to 2034, given its domestic AI capacity, its tradition of strategic autonomy and its deep dependence on both Western and Chinese ecosystems. The real alternative is that a much deeper US-India TRUST initiative (the successor to the Initiative on Critical and Emerging Technology, iCET) pulls India towards the Americas bloc before the structure completes. We expect the bilateral relationships to deepen without producing membership, but the margin is not large.

Assumption 8: a bloc does not need a treaty (High). Defining a bloc as a shared regulatory grammar rather than a treaty structure is deliberate (section 8.1). The closest analogue is the GDPR adequacy framework, through which some sixteen jurisdictions formed a de facto club of data protection equivalence without a multilateral treaty. Readers who require a treaty-based definition will find the forecast overstated; our answer is that treaty-based international AI governance is not achievable within the horizon, so a useful framework has to work with the softer definition.

11.2 How each assumption could fail

Table 1 compares the eight assumptions on five dimensions. "Directional risk" is the likely direction of error: "optimistic" means that if the assumption is wrong, the outcome is probably less favourable to the forecast, and "pessimistic" means it is probably more favourable.

Table 1: Core assumptions

No.	Assumption	Confidence	Evidence tier	Directional risk	Sensitivity of forecast	Depends on
A1	Brussels Effect holds	Medium	1 to 2	Optimistic	High	A8
A2	US federal AI legislation passes 2029 to 2030	Medium-Low	2 to 3	Optimistic	High	A1, A3

No.	Assumption	Confidence	Evidence tier	Directional risk	Sensitivity of forecast	Depends on
A3	China's orbit remains coherent	Medium	2 to 3	Optimistic	Medium	A7
A4	Post-quantum risk as near-term catalyst	Medium-Low	2 to 3	Pessimistic	Low to Medium	A5
A5	Glasswing-type disclosures produce binding responses	High	1 to 2	Pessimistic	Low to Medium	A4
A6	Standards bodies as pre-harmonisers	High	1 to 2	Neutral	Medium	A1, A2
A7	India non-aligned within horizon	Medium	2 to 3	Optimistic	Low	A3
A8	A bloc needs no treaty	High	1 to 2	Neutral	High	A1

Table 2 sets out how the forecast would change if each assumption failed in the direction shown.

Table 2: The forecast if an assumption fails

No.	If wrong	Revised outcome	Severity
A1	Weaker	Smaller orbital membership; timeline slips	Moderate
A2	Fails or comes later	No Americas bloc; bipolar world; completion moves to 2035 or later	Severe
A3	Orbit shrinks	China-anchored bloc becomes near-domestic; bipolar world with a large non-aligned zone	Moderate
A4	Slower	Scenario E less intense; convergence slightly faster	Low
A5	Weaker	Security fragmentation less severe; Scenario E about 30%	Low to moderate
A6	Standards substitute for law	Timeline extends; blocs form around standards bodies	Moderate
A7	India aligns	Fourth-bloc scenario eliminated; Americas bloc larger	Low
A8	Treaty required	Blocs reclassified as proto-blocs; substantive forecasts unchanged	Low

11.3 The composite forecast

The three-bloc forecast is a composite of sub-forecasts, each with its own confidence and timing. The staggered structure follows from that. The direction beneath it, that unbounded fragmentation gives way to a small number of blocs, we rate High, because each alternative in section 12 also ends in a small number of blocs, even where a large non-aligned zone remains.

The EU-anchored bloc, recognisable by 2026 to 2027, is rated High. Its formation is under way: the Act is in force, its extraterritorial mechanics are operating and the adequacy precedent is established, so the main uncertainty is which states orbit it at the edges, not whether the bloc exists. The China-anchored bloc, recognisable by 2027 to 2029, is rated between Medium and High, because the domestic framework is extensive and still growing and the infrastructure model is established, while the stability of its orbit under geopolitical pressure is uncertain.

The Americas bloc, recognisable by 2030 to 2032, is rated Medium-Low. It depends on a legislative event that has not happened and faces documented obstacles, which makes it the forecast's weak point.

The bipolar interregnum of about 2027 to 2031 is the expected near-term configuration rather than a risk; we rate it between Medium and High, the same as the China-anchored bloc on which it depends. It does not depend on US legislation, since it continues if that legislation fails.

Table 3: Composite forecast

Sub-forecast	Date	Assumptions	Confidence
EU-anchored bloc recognisable	2026 to 2027	A1, A6, A8	High
China-anchored bloc recognisable	2027 to 2029	A3, A8	Between Medium and High
Bipolar interregnum operative	2027 to 2031	A1, A3, A8	Between Medium and High
Americas bloc recognisable	2030 to 2032	A2, A6, A8	Medium-Low
Near-term fragmentation surge (Scenario E)	2026 to 2028	A4, A5	Between Medium and High
Complete tripartite structure	2032 to 2034	A1 to A8	Medium

Three points follow from the table. First, the complete structure is rated higher than the Americas bloc because its later window absorbs some slippage in US legislation. Second, the interregnum is a configuration in its own right, and it puts structural pressure on US legislative action (section 10).

Third, the dependency between Assumptions 1 and 8 is the main structural vulnerability. Both concern how the EU-anchored bloc coheres, and both would fail together if the EU's market power weakened so far that extraterritorial pull could no longer substitute for treaty alignment; we assess the probability of that joint failure at about 15%.

12. The likeliest alternative is that US federal legislation fails

Four scenarios depart from the base case; a fifth, the near-term surge (Scenario E), unfolds within it and is described in section 6.3. Each probability is our judgement of how likely that departure is, assessed separately and read as a ranking (section 2.2).

Scenario B: deeper fragmentation (25%). US federal legislation fails, or passes in a form so hollowed out by carve-outs, preemption limits and enforcement constraints that it produces no meaningful harmonisation. The state patchwork persists into the early

2030s, the absence of a federal counterpart prevents EU-US mutual recognition from advancing, and the Americas bloc does not form; the result is a bipolar structure with a large, internally differentiated non-aligned zone that includes the United States itself.

This is the main alternative. On horizontal technology regulation, the US legislative record has consistently been more conservative than its international counterparts and more prone to coalitions breaking up in the final stages.

Scenario C: the China-anchored bloc shrinks (20%). Accelerating US export controls on advanced AI hardware, combined with domestic economic pressure, slow China's outward infrastructure investment and weaken the dependency that holds its orbit together, and partner states facing binary pressure choose Western alignment on cost-benefit grounds. China's domestic framework stays intact and continues to govern the world's largest single-country AI market. Its pull on other states weakens substantially, however, moving the structure towards two blocs or one and a half.

Scenario A: faster convergence (15%). A catastrophic AI-attributable event before 2029, clearly traceable to a cross-jurisdictional governance failure, accelerates coordination in the way the discovery of the Antarctic ozone hole accelerated the Montreal Protocol. For this to happen, the event must be attributable to systems operating across borders, the attribution must be technically demonstrable and politically legible, and the international climate must permit US-EU-China coordination rather than mutual blame. The historical base rate for rapid treaty-level agreement on technology governance is low, and the conditions for comparable speed are not present.

Scenario D: a fourth bloc emerges (15%). India, drawing on its capability, its population of 1.4 billion and its diplomatic networks across the Global South, cultivates a fourth alignment built on digital sovereignty, data localisation and bilateral rather than multilateral arrangements. This would need substantial investment in regulatory capacity in partner states and in domestic institutions able to act as credible international interlocutors. The scenario is plausible given India's G20 ambitions and its TRUST partnership with the US, but current trajectories do not yet show the sustained investment it requires.

13. Operating across three blocs creates a compliance trilemma

Three stable blocs would be an improvement on unbounded fragmentation, but only a partial one. Three coherent counterparts, each developing mutual recognition, are tractable in a way that twenty or more independent frameworks are not. Operating across all three at once is nevertheless a new kind of challenge.

The trilemma arises from three conditions:

- **The philosophies conflict at the level of design, not only documentation.** A system built to respect the AI Act's restriction on real-time remote biometric identification in public spaces for law enforcement may not serve a use that is

lawful under Chinese rules; a training-data approach that satisfies the GDPR's data minimisation may not pass China's security assessment; and a transparency architecture that satisfies US FTC disclosure expectations may not meet the depth of the AI Act's documentation.

- **Conformity evidence does not transfer.** The EU conformity file is not what a US sector regulator expects, and China's assessment format is incompatible with both.
- **Security-layer obligations sit outside all three general frameworks.** Enterprises in critical infrastructure must manage obligations derived from NIS2 in the EU, CIRCIA reporting in the US once its final rule takes effect, and Chinese security assessments as a separate domain (section 6.3).

Enterprises will respond with regulatory passporting. By that we mean internal architectures designed to show conformity with one bloc in a form that translates efficiently to another, supported by automated monitoring, AI-assisted documentation and specialist advisers in each jurisdiction. The capability will give early movers a competitive advantage, and it will concentrate markets by raising barriers to entry for smaller enterprises.

14. Blocs will become instruments of power, and the largest safety gaps sit between them

The blocs will shape geopolitical competition, not merely reflect it. Regulatory power, the ability to set the terms on which others take part in a market, is a form of structural power that works independently of military and economic coercion (Strange, 1994). Each bloc will use it as a geopolitical tool, deepening dynamics already visible between 2022 and 2026.

The EU's use of extraterritoriality is the most mature example. By conditioning market access on compliance, the EU extends its norms to enterprises in jurisdictions that have not adopted them, and the influence is more durable than diplomatic pressure because commercial incentives enforce it. Beneath the legislation, the contest over who defines "high-risk", whose criteria become the ISO default and whose conformity method is recognised is both commercial and geopolitical, with China's investment in ITU work, the EU's use of CEN and CENELEC standards and the US's NIST-led framework each used as an instrument.

Blocs could reinforce technological decoupling. If enterprises must choose between architectures built for the EU-anchored and China-anchored blocs, and maintaining both is prohibitive for all but the largest, the rational response is to focus on one market and leave the other. That kind of decoupling would harm AI safety (divergent testing methods and no shared incident reporting), economic welfare (lost efficiency) and stability (less interdependence, and so fewer structural incentives to cooperate).

In our view the tripartite structure is the best outcome achievable within the horizon, but it leaves serious safety gaps. The most important concern systemic risks from the aggregate behaviour of systems operating across all three blocs, which no

single framework is designed to address. Divergent risk thresholds allow regulatory-void deployment: an application prohibited in the EU-anchored bloc can be deployed from a non-aligned or China-anchored jurisdiction while its effects cross borders.

A disinformation system that is prohibited in the EU but lawful outside it can still target EU citizens. A system barred from EU financial applications but lawful in less regulated markets can still affect global stability through interconnected markets.

No body can coordinate a binding response to a transnational AI harm event. None currently combines the mandate, authority, technical capacity and legitimacy. The Financial Stability Board, the Basel Committee and IOSCO offer partial precedents, but the diversity of applications, the speed of capability development and the tensions between blocs make effective cross-bloc machinery very hard to build.

The most realistic route to a global minimum floor runs through the G7 and the OECD. By a minimum floor we mean baseline prohibitions and incident reporting obligations that apply across all three blocs. The processes that produced the OECD AI Principles, the Hiroshima AI Process and its Code of Conduct are voluntary and remain at the level of principle, but they provide institutions and relationships that could support a more substantive binding instrument in the 2030s.

Where to start, and what it costs

On our forecast, the fragmentation of 2022 to 2026 is the disorder that precedes consolidation, but the three blocs will not consolidate on the same timeline. The two that already exist in outline, the EU-anchored and China-anchored blocs, are current trajectories more than forecasts. The Americas bloc is the forecast's genuinely future element, and the period before it forms is the one enterprises and policymakers have to plan for.

First: adopt the AI Act as the baseline, and map it once. Treat AI Act compliance as the de facto global baseline for rights-based markets, and document its equivalence to NIST AI RMF outcomes using ISO/IEC 42001 as the bridge. Most of the cost is work an enterprise selling into the EU must do anyway; the addition is the mapping, which is done once per system and maintained.

Next: design for a US federal framework before it exists. Make compatibility with a federal framework of the kind described in section 7.2 an architectural requirement now, so that the eventual standard does not force a costly retrofit. Manage US state law as a fragmented residual in the meantime, and add China-anchored compliance only where Chinese infrastructure dominates a market.

Ongoing: run security-layer obligations as a separate domain. NIS2- and CIRCIA-derived AI security duties, Chinese security assessments and post-quantum migration timelines will not map cleanly onto general AI governance frameworks (section 6.3), and they will need their own owners and their own monitoring from 2026 to 2035.

For policymakers, the interregnum is the window. Negotiate EU-Americas mutual recognition while it is open, so that an Americas bloc brings immediate interoperability rather than a further multi-year negotiation. Pursue a G7 or OECD minimum-floor instrument with baseline prohibitions and incident reporting, and invest in regulatory capacity for non-aligned states before pressure forces them into a bloc they are not yet equipped to evaluate.

On timing. The AI Act's high-risk obligations apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems, subject to Article 111(2) for systems already on the market, and as at 11 October 2026 no fine under the AI Act had been confirmed by any regulator or official source. The direction we are most confident of, that unbounded fragmentation gives way to a small number of blocs, is rated High; that the result is three blocs, complete by 2032 to 2034, is rated Medium, with the Americas bloc as the binding constraint.

A poorly managed interregnum is a real possibility. If US federal legislation fails, mutual recognition is never negotiated, security-layer obligations harden into permanent technical barriers and non-aligned states become governance voids, the interregnum becomes a durable bipolar world; that is Scenario B, at 25%, and the default outcome of continued US inaction. The technical foundations are in place or under construction; what is missing is political will, in Washington above all.

An enterprise that has adopted the AI Act as its baseline, mapped it to the NIST framework and designed for a federal standard is prepared for either outcome: a third bloc, or a two-bloc world that hardens.

This paper is provided for educational and informational purposes and is not legal, investment or policy advice. It describes the position as at 11 October 2026 and contains forecasts of regulatory developments that may not occur as described; confidence ratings express analytical judgement, not assurance. Take qualified professional advice before acting on it.

Basis and Standing

Binding law. All AI Act references are to Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744 (Digital Omnibus on AI, in force 27 July 2026): Article 2(1) on territorial scope; Article 11 on technical documentation; Article 40 on the presumption of conformity from harmonised standards; Article 43 on conformity assessment; Article 75 on the AI Office's supervisory competence; Article 113 on application dates (2 December 2027 for Annex III high-risk systems, 2 August 2028 for Annex I); and Article 111(2) on high-risk systems already on the market. The NIS2 Directive, the Cyber Resilience Act and the Medical Devices Regulation. The GDPR (Regulation (EU) 2016/679) and the Commission's adequacy decisions. US: Colorado SB 24-205 and SB 26-189, California SB 53, Texas HB 149, New York City Local Law 144 of 2021, CIRCIA; West Virginia v. EPA (2022) and Loper Bright Enterprises v. Raimondo (2024). China: the Cyberspace Administration of China's provisions on algorithmic recommendation, deep synthesis and generative AI, the

labelling measures, and the Cybersecurity Law as amended with effect from 1 January 2026. Other jurisdictions' laws (India's Digital Personal Data Protection Act 2023, Japan's AI Promotion Act, South Korea's AI Basic Act) are cited for their existence and general character only.

Policy, guidance and voluntary frameworks. US Executive Orders 13859, 14110 (revoked by Executive Order 14148) and 14365, and the White House's March 2026 legislative recommendations, which are not law. The NIST AI Risk Management Framework, which is voluntary; NIST's post-quantum standards (FIPS 203, 204 and 205), which are standards for US federal information systems. ISO/IEC 42001:2023 and EN 18286, which are voluntary standards. The UK's 2023 AI white paper, India's AI Governance Guidelines, Singapore's Model AI Governance Framework, and the G7 Hiroshima AI Process and OECD AI Principles, none of which is binding.

Disclosure. The facts given about Project Glasswing and Claude Mythos Preview are as Anthropic reported them in April and June 2026. They are a company's account of its own work, not independent verification.

Our own analysis. The definition of a bloc and of recognition; the three blocs, their membership and their recognition windows; the bipolar interregnum; the fragmentation surge; the eight assumptions and their ratings; the evidence tiers and confidence scale; the scenario probabilities; the forecast architecture of US federal legislation; the compliance trilemma; and regulatory passporting are our own and are not drawn from any official source. Statements that procurement drives ISO/IEC 42001 adoption rest on our experience.

Research, and what it supports. Bradford (2020) supports the account of the Brussels Effect as unilateral regulatory globalisation through market access; Strange (1994) supports the concept of structural power; Drezner (2007) supports the role of great-power interests in international regulatory coordination. None of them forecasts the bloc structure described here.

Enforcement position. As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source, and we had found no AI Act harmonised standard cited in the Official Journal.

Sources

1. Regulation (EU) 2024/1689 (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024, consolidated version of 27 July 2026: Articles 2, 11, 40, 43, 75, 111 and 113. https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng
2. Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), OJ L, 2026/1744, 24.7.2026; in force 27 July 2026. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
3. Regulation (EU) 2016/679 (General Data Protection Regulation), OJ L 119, 4.5.2016, p. 1. https://eur-lex.europa.eu/eli/reg/2016/679/oj

4. European Commission, Adequacy decisions (list of countries and organisations recognised, including Brazil, 26 January 2026; United Kingdom renewal, 19 December 2025; European Patent Organisation, 15 July 2025). https://commission.europa.eu/law/law-topic/data-protection/international-dimension-data-protection/adequacy-decisions_en
5. European Commission, "Commission proposes a comprehensive reform of data protection rules to increase users' control of their data and to cut costs for businesses", press release IP/12/46, 25 January 2012. https://ec.europa.eu/commission/presscorner/detail/en/IP_12_46
6. European Commission, AI Act standardisation. https://digital-strategy.ec.europa.eu/en/policies/ai-act-standardisation; CEN-CENELEC, EN 18286, AI quality management system for regulatory purposes (July 2026).
7. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, January 2023. https://doi.org/10.6028/NIST.AI.100-1
8. National Institute of Standards and Technology, Post-Quantum Cryptography: FIPS 203, FIPS 204 and FIPS 205 (13 August 2024), and the migration timeline in NIST IR 8547. https://csrc.nist.gov/projects/post-quantum-cryptography
9. ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system. https://www.iso.org/standard/42001
10. Executive Order 13859, Maintaining American Leadership in Artificial Intelligence (11 February 2019). https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence
11. Executive Order 14110, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (30 October 2023). https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence
12. Executive Order 14148, Initial Rescissions of Harmful Executive Orders and Actions (20 January 2025), revoking Executive Order 14110. https://www.federalregister.gov/documents/2025/01/28/2025-01901/initial-rescissions-of-harmful-executive-orders-and-actions
13. Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence (11 December 2025), sections 3 and 8. https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/
14. The White House, National Policy Framework for Artificial Intelligence: legislative recommendations (20 March 2026).
15. California SB 1047 (vetoed 29 September 2024) and SB 53 (signed September 2025); Colorado SB 24-205 (signed May 2024) https://leg.colorado.gov/bills/sb24-205 and SB 26-189 (signed 14 May 2026) https://leg.colorado.gov/bills/sb26-189; Texas HB 149 (signed June 2025); New York City Local Law 144 of 2021.
16. West Virginia v. EPA, 597 U.S. 697 (2022). https://www.supremecourt.gov/opinions/21pdf/20-1530_n758.pdf; Loper Bright Enterprises v. Raimondo, 603 U.S. 369 (2024). https://www.supremecourt.gov/opinions/23pdf/22-451_7m58.pdf
17. Cyberspace Administration of China: Provisions on the Management of Algorithmic Recommendations in Internet Information Services (2022); Provisions on the Administration of Deep Synthesis Internet Information Services (2023); Interim Measures for the Management of Generative Artificial Intelligence Services (2023); Measures for Labelling AI-Generated Synthetic Content (2025).

Cybersecurity Law of the People's Republic of China, as amended on 28 October 2025, in force 1 January 2026.
18. Department for Science, Innovation and Technology, A Pro-Innovation Approach to AI Regulation (white paper, March 2023). https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach
19. Bill C-27, Digital Charter Implementation Act, 2022, Part 3 (Artificial Intelligence and Data Act), Parliament of Canada (died on prorogation, January 2025).
20. Japan, AI Promotion Act (2025); India, Digital Personal Data Protection Act 2023 and India AI Governance Guidelines (November 2025); Republic of Korea, AI Basic Act (2025, in force January 2026).
21. Anthropic, Project Glasswing (7 April 2026, with later updates). https://www.anthropic.com/glasswing
22. Anthropic, Claude Mythos Preview research post (7 April 2026), source of the 27-year-old OpenBSD vulnerability and the more than 99% unpatched figure. https://www.anthropic.com/research/mythos-preview
23. Anthropic, Expanding Project Glasswing (2 June 2026). https://www.anthropic.com/news/expanding-project-glasswing
24. Anu Bradford, The Brussels Effect: How the European Union Rules the World (Oxford University Press, 2020).
25. Daniel W. Drezner, All Politics Is Global: Explaining International Regulatory Regimes (Princeton University Press, 2007).
26. Susan Strange, States and Markets, 2nd edn (Pinter, 1994).
27. Directive (EU) 2022/2555 (NIS2 Directive), OJ L 333, 27.12.2022, p. 80; Regulation (EU) 2024/2847 (Cyber Resilience Act), OJ L, 2024/2847, 20.11.2024; Regulation (EU) 2017/745 (Medical Devices Regulation), OJ L 117, 5.5.2017, p. 1.
28. Commission Implementing Decision (EU) 2023/1795 of 10 July 2023 on the adequate level of protection of personal data under the EU-US Data Privacy Framework.
29. Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA), 6 U.S.C. 681 to 681g; Texas HB 1709 (2025, not enacted).
30. Data Security Law of the People's Republic of China (2021); Brazil, Lei Geral de Proteção de Dados (Law No. 13.709 of 2018).
31. ISO/IEC 27001, Information security management systems (first published 2005).
32. G7, Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems (30 October 2023); G20 New Delhi Leaders' Declaration (September 2023); OECD, Recommendation of the Council on Artificial Intelligence (OECD AI Principles, 2019, revised 2024).
33. Personal Data Protection Commission Singapore, Model Artificial Intelligence Governance Framework, 2nd edn (January 2020).
34. US Food and Drug Administration, Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan (January 2021).

Provisions are cited by article, section or paragraph in the text above. Forward-looking statements are our forecasts, not statements of fact.

Standard Intelligence

GET IN TOUCH

Plan for three blocs, not twenty regulators.

EMAIL

hello@standardintelligence.com

WEB

standardintelligence.com

LOCATION

Manchester, United Kingdom

AI Use: AI tools were used in research, editorial review and drafting support. All sources were retrieved and checked against the published record by the author, who holds editorial responsibility for the contents and for the conclusions drawn. This disclosure is made voluntarily. To the extent the Article 50(4) obligation in Regulation (EU) 2024/1689 applies to this paper as text published to inform the public on matters of public interest, the exception for content that has undergone human review and for which a person holds editorial responsibility applies.

STANDARD INTELLIGENCE LIMITED · REGISTERED IN ENGLAND AND WALES · COMPANY NO. 16737890