SenecAI CONSULTING

GLOBAL AI POLICY REPORT

ISSUE NO. 01

2026

What Awaits Organizations Deploying AI Globally

A Comparative Index of AI Regulatory Intensity Across 22 Jurisdictions

Matei Ștefan

Founder, SenecAI Consulting

First Edition · Issue No. 01

August 2026

SenecAI
CONSULTING

Global AI Policy Report

What Awaits Organizations Deploying AI Globally

First Edition · August 2026

A Comparative Index of AI Regulatory Intensity Across 22 Jurisdictions

Matei Ștefan, Founder, SenecAI Consulting

ISSUE NO. 01

SenecAI CONSULTING

Global AI Policy Report 2026

What Awaits Organizations Deploying AI Globally

PUBLISHER	SenecAI Consulting
AUTHOR	Matei Ștefan, Founder, SenecAI Consulting
EDITION	First Edition, Issue No. 01
PUBLICATION DATE	August 2026

CONTACT	matei.stefan@senecai.eu
WEBSITE	senecai.eu

CONTENTS

Table of Contents

01 Executive Summary 6

02 The Question 11

03 Methodology 15

3.1 The scoring framework 16

3.2 Rules developed during scoring 18

3.3 The EU member-state delta model 19

3.4 The US federal-plus-state model 20

3.5 Limitations 21

04 The Global Picture 22

4.1 The full ranking 23

4.2 Reading the table 26

05 Region by Region 28

5.1 Overview 29

5.2 The European Union and its member states 29

5.3 The United States 32

5.4 China 35

5.5 Singapore, Japan, and South Korea 36

5.6 The Gulf states 38

5.7 Rest of the world: Canada, Brazil, and South Africa 40

CONTENTS

Table of Contents (continued)

06 Who Enforces AI Policy, and What Deployers Must Actually Do 42

07 Does the Law Match the Enforcement? 51

08 Potential Criticisms and Responses 57

09 What to Look Out For Next 61

10 What This Means in Practice for Deployers 64

11 Conclusion 67

A Why This Study Differs from Others in the Market 70

B List of Abbreviations 72

C Source Appendix 75

D Country Scorecards 78

— How to Cite This Report 85

01

CHAPTER 01

Executive Summary

How intense is AI regulation, really — and what does that mean for an organization deploying across 22 jurisdictions?

EXECUTIVE SUMMARY

Executive Summary

This study analyzes the most relevant 22 jurisdictions globally and answers the question of how intense and dense the regulatory requirements are in each — in other words, how burdensome it is for a deployer to place an AI product into that market.

Governments are regulating AI in very different ways. Some have passed one wide-reaching AI statute. Some rely on existing laws that were not written with AI in mind. Some are relying on strong vertical and use-case laws issued directly from by the executive branch. Some govern mostly through the contracts a company signs to do business with the state. Reading the statutes side by side does not tell a deployer what will actually happen on the ground, because two regimes that look equally strict on paper can produce very different outcomes in practice, depending on whether a regulator has the staff, budget, and legal authority to act on what the law allows.

This study exists to answer a more practical question: how much regulatory weight will an organization actually feel if it puts an AI system into a given market? To measure that, each of the 22 jurisdictions is scored across five weighted categories, each broken into three sub-criteria and scored on a 0–100 scale against a fixed rubric. The five category scores are combined into a single composite score per jurisdiction. Jurisdictions are then placed into one of three tiers based on that composite: Tier 1 (composite of 70 or above) is high intensity, Tier 2 (50 to 70) is moderate intensity, and Tier 3 (below 50) is emerging. The full scoring framework, the weighting of each category, and the rules developed while scoring are set out in the Methodology section later in this report.

The Index compares jurisdictions across five dimensions:

- **Governance approach** — comprehensive legislation, sectoral absorption, or the sovereign/centralized model
- **Legal density versus enforcement capacity** — the gap between what a law says and what a regulator can actually do
- **Instrument type and count** — binding statute, soft law and guidance, sectoral rules, or jurisprudence-driven exposure
- **Penalty architecture** — administrative fines, criminal liability, injunctive power
- **Extraterritorial reach** — whether the regime binds foreign providers and imposes data localization

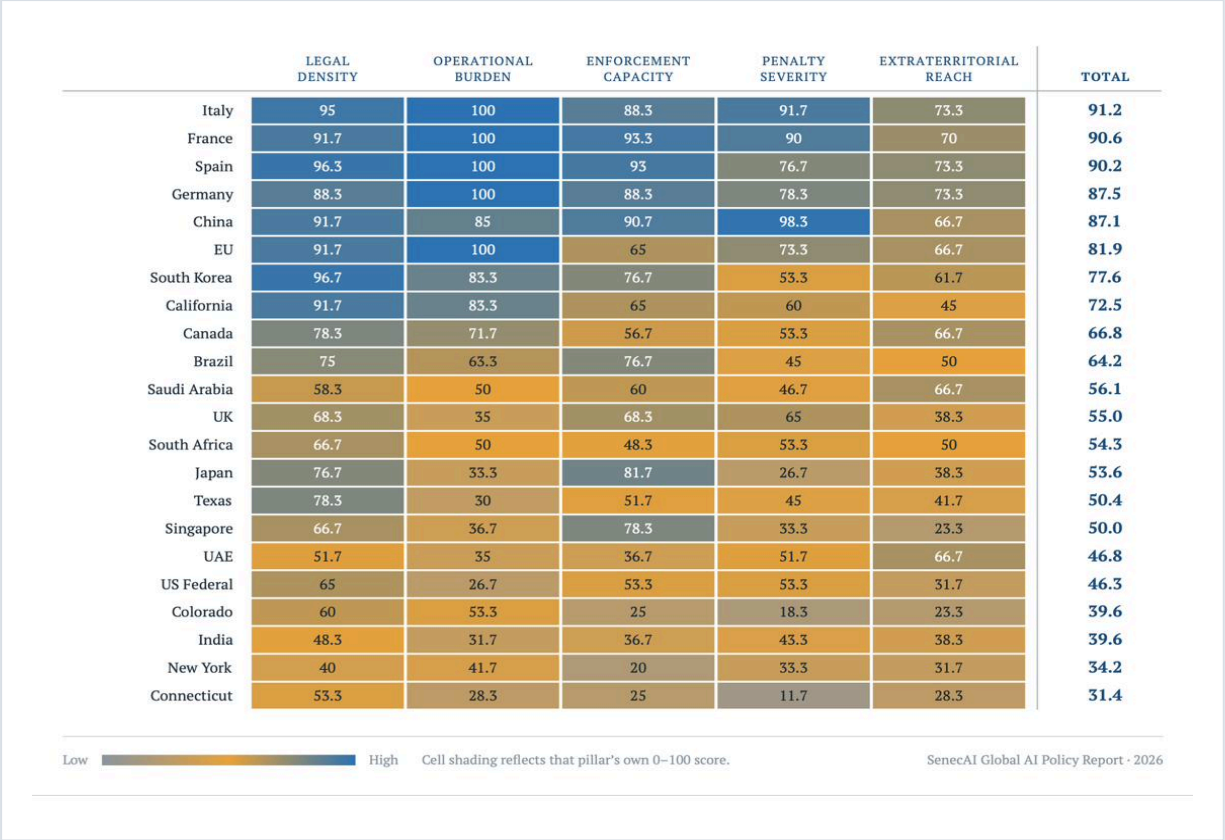

	LEGAL DENSITY	OPERATIONAL BURDEN	ENFORCEMENT CAPACITY	PENALTY SEVERITY	EXTRATERRITORIAL REACH	TOTAL
Italy	95	100	88.3	91.7	73.3	91.2
France	91.7	100	93.3	90	70	90.6
Spain	96.3	100	93	76.7	73.3	90.2
Germany	88.3	100	88.3	78.3	73.3	87.5
China	91.7	85	90.7	98.3	66.7	87.1
EU	91.7	100	65	73.3	66.7	81.9
South Korea	96.7	83.3	76.7	53.3	61.7	77.6
California	91.7	83.3	65	60	45	72.5
Canada	78.3	71.7	56.7	53.3	66.7	66.8
Brazil	75	63.3	76.7	45	50	64.2
Saudi Arabia	58.3	50	60	46.7	66.7	56.1
UK	68.3	35	68.3	65	38.3	55.0
South Africa	66.7	50	48.3	53.3	50	54.3
Japan	76.7	33.3	81.7	26.7	38.3	53.6
Texas	78.3	30	51.7	45	41.7	50.4
Singapore	66.7	36.7	78.3	33.3	23.3	50.0
UAE	51.7	35	36.7	51.7	66.7	46.8
US Federal	65	26.7	53.3	53.3	31.7	46.3
Colorado	60	53.3	25	18.3	23.3	39.6
India	48.3	31.7	36.7	43.3	38.3	39.6
New York	40	41.7	20	33.3	31.7	34.2
Connecticut	53.3	28.3	25	11.7	28.3	31.4

Figure 1. Index heatmap across all five pillars, all 22 jurisdictions.

KEY TAKEAWAY Composite scores split into three clear bands, with Italy leading and Connecticut trailing across all 22 jurisdictions.

AI governance has split into three durable approaches — comprehensive legislation, sectoral absorption, and the sovereign/centralized model — and within each, it is enforcement capacity, not statutory ambition, that determines what a deployer actually experiences on the ground.

THE CENTRAL FINDING

Every jurisdiction in the Index has taken one of three approaches to governing AI, and in every case, what a regulator can actually enforce matters more than what the statute says on paper.

- **Comprehensive legislation** means a government has passed one wide-reaching AI-specific law, or, in the EU's case, a central law plus country-level additions on top of it. This is the group made up of the EU baseline, Italy, France, Spain, Germany, South Korea, and Brazil. It produces the highest scores in the Index, but only where a real regulator with staff and powers stands behind the law. A comprehensive law with no working authority behind it does not score as high as one that is fully staffed and active.
- **Sectoral absorption** means a government has not passed one broad AI law, and instead lets existing rules, such as data protection law, consumer protection law, and rules specific to individual sectors, stretch to cover AI. This is the largest group in the Index by number of jurisdictions, covering California, Texas, Colorado, New York, Connecticut, the UK, Canada, Japan, Singapore, South Africa, and India. It is also the most varied group, because how strict this approach feels in practice depends entirely on how willing existing regulators are to apply their existing powers to AI. That willingness is a choice regulators make, not something guaranteed by the law itself.
- **The sovereign/centralized model** splits into two distinct mechanisms, but shares one common trait: a powerful state apparatus that sets AI policy from the top down, rather than through an independent regulator answering to a legislature. In Saudi Arabia and the UAE, that top-down authority expresses itself lightly in statute and concentrates instead in the state's own commercial and procurement relationships — a deployer feels real pressure only where it needs the state as a customer, licensor, or gatekeeper, and comparatively little where that relationship can be avoided. China does not fit the procurement pattern: its administrative rules bind private-sector deployers directly, with penalties that apply regardless of any relationship with the state, and its enforcement-capacity score is among the highest in the Index. What all three still share is that the policy itself, in each case, is set centrally rather than negotiated through the pluralistic rulemaking seen elsewhere in this Index.

Supporting findings

The legal-density-to-enforcement-capacity gap is arguably the sharpest diagnostic this data produces, because it separates jurisdictions where the law and the enforcement match from jurisdictions where they do not.

- Japan shows a near-zero gap: its enforcement activity tracks its statutory ambition closely, even though it imposes no monetary penalties at all.
- Colorado shows one of the widest gaps in the Index: a comprehensive statute that is currently unenforceable in practice pending litigation.
- Texas shows the same gap in the opposite direction: a comparatively narrow AI-specific statute paired with a general-law enforcement record that outweighs what the text alone would suggest.

The EU's real intensity sits below the baseline, inside its member states, not at the EU level itself.

- Italy, France, Spain, and Germany each score above the EU baseline.
- Together, these four occupy four of the Index's top six positions.
- This shows the AI Act working as a floor, while national criminal law, data-protection enforcement, and co-determination rules set the ceiling a deployer actually experiences in each country.

The United States is better understood as a set of separate markets than as one jurisdiction.

- The US federal row sits in Tier 3, below both Saudi Arabia and South Africa.
- California alone sits in Tier 1, in the same band as South Korea.
- The gap between California and the lowest-scoring US jurisdiction is wide enough that a single national compliance program is unlikely to serve the US market well.

The sovereign/centralized model works through at least two different mechanisms, not one.

- China reaches Tier 1 through a large, actively enforced body of administrative rules.
- Saudi Arabia and the UAE reach materially lower composite scores while relying on comparatively little formal statute, with procurement and licensing relationships carrying the weight that a Western deployer would normally expect to find in legislation.

02

CHAPTER 02

The Question

Why a binary yes/no catalogue of AI laws can't answer the question deployers are actually asking.

AT A GLANCE

This chapter sets out the question the Index answers — not whether a jurisdiction has an AI law, but how much regulatory weight a deployer actually feels there — and introduces the three governance approaches that organize every jurisdiction in this report.

Every AI-policy tracker built to date answers a narrower question than the one deployers are actually asking. A binary "has an AI law: yes/no" catalogue would treat a country with a single, toothless statute the same as a country with a comprehensive law backed by an active regulator. It would also treat a narrow statute reaching only a handful of frontier developers the same as a broad statute reaching every deployer in the market. Neither comparison survives contact with an actual compliance budget.

This report exists to move past that binary cataloguing. In its place, it offers a weighted measure of what the answer looks like in practice. It asks a more specific question. Given a jurisdiction's statutes, its regulators, and its enforcement record to date, how intense and how dense is the regulatory burden an organization can expect to carry there? And how does that burden compare across the 22 markets where AI deployment is most concentrated and most consequential?

That question could not have been asked credibly at an earlier stage of this technology's regulation. The EU AI Act was the pioneer here. Over the last two years, it has helped push AI governance from a voluntary, nice-to-have set of principles toward a binding legal obligation. In the period right after the AI Act passed, most of the jurisdictions in this Index were still at the proposal stage: green papers and draft bills, not enacted statutes or appointed regulators. Governments are no longer in that position. A critical mass of the jurisdictions surveyed here now have an operational record. Statutes are in force. Authorities have been designated. In a meaningful subset, enforcement actions have already been completed.

In the last two years, AI governance has shifted from a voluntary, nice-to-have practice to a binding legal obligation. Given that shift, governments have had to choose an approach to regulate.

Faced with that choice, governments have not landed on a spectrum of individually distinct national styles. Instead, their responses cluster into three durable approaches.

The first is comprehensive legislation. A jurisdiction passes one wide-reaching, AI-specific statute. In the EU's case, that statute sits underneath a stack of member-state additions. This approach accepts more compliance friction in the near term. In exchange, it offers legal certainty, and a first-mover's influence over how the rules get shaped.

The second is sectoral absorption. A jurisdiction does not pass a broad AI statute at all. Instead, it lets existing rules, on data protection, on consumer protection, on specific sectors, stretch to cover AI. This preserves flexibility for the regulator. It comes at a cost for the deployer, though, who has a harder time mapping their exposure in advance.

The third is the sovereign/centralized model. It concentrates AI rule-making, enforcement, and commercial gatekeeping within the executive branch, unconstrained by the separation of powers that shapes the other two approaches. It reaches deployers through one of two mechanisms. China represents an administrative-registry variant: a dense, actively enforced body of executive rules binds private-sector deployers directly, and no government-contract relationship is required at all. Saudi Arabia and the UAE represent a procurement-leverage variant: obligations attach mainly to licenses, approvals, and the state's own commercial relationships, rather than to one law applying evenly to every deployer, and a deployer feels the least pressure where it can avoid needing the state as a customer or licensor. The common thread across both variants is institutional concentration, not the specific channel through which enforcement runs.

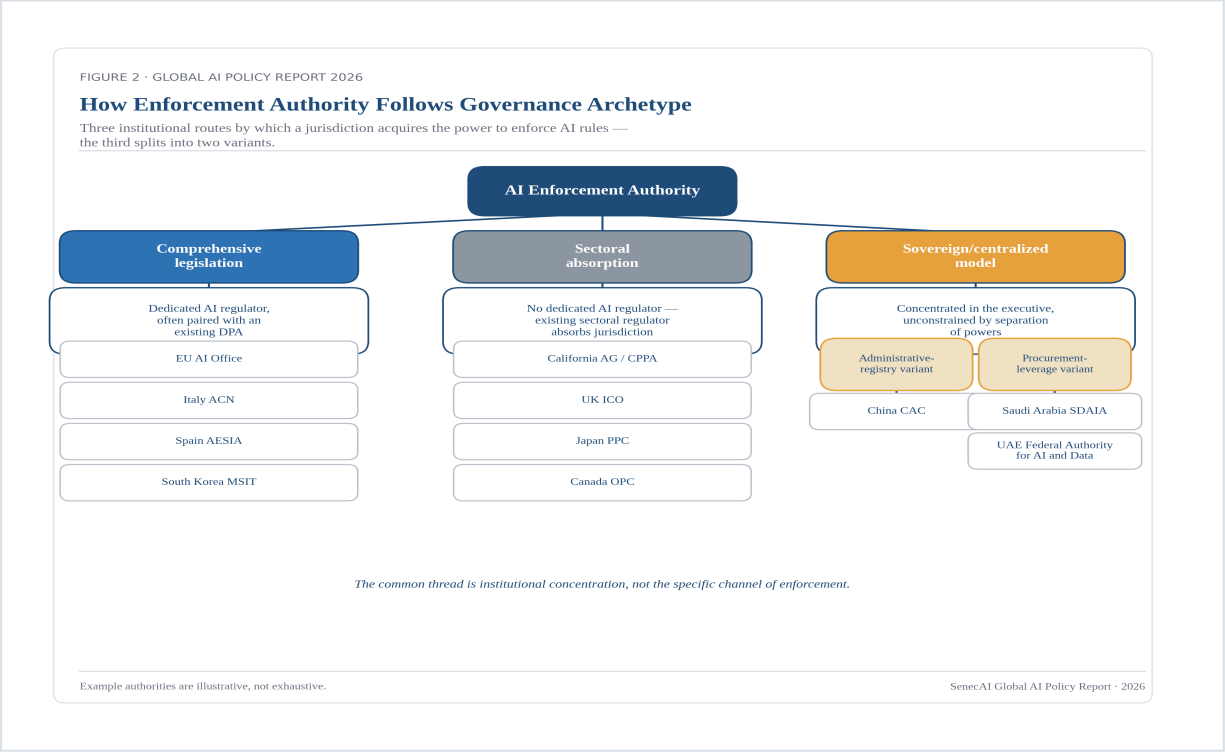

Figure 2. Enforcement pattern by governance approach — the sovereign/centralized model splits into two variants.

KEY TAKEAWAY Three routes to enforcement power, but the third splits into two genuinely different mechanisms.

None of the three approaches is presented here as superior to the others in the abstract. Each is a defensible response to the same underlying uncertainty about how to govern a fast-moving technology. And each produces a different distribution of risk for the organization that has to operate under it. What differs, and what the remainder of this report sets out to measure, is how much of each approach's apparent ambition is actually backed by the capacity to enforce it. That is a question the statutory text alone cannot answer.

This report covers 22 jurisdictions: the EU baseline together with its four largest AI-relevant member states (France, Germany, Italy, Spain); the United States as a federal row plus five states that have enacted AI-specific statutes (California, Colorado, Texas, New York, Connecticut); the United Kingdom; Canada and Brazil; China, South Korea, Japan, and Singapore; India; Saudi Arabia and the UAE; and South Africa. It is written for the audience that has to act on it, not just track it: in-house counsel, compliance officers, and policy professionals at organizations deploying AI systems across borders. And it measures a single variable on their behalf. Not how sophisticated, fair, or innovation-friendly a jurisdiction's AI policy is. What awaits an organization that puts an AI system into that market.

03

CHAPTER 03

Methodology

Fifteen sub-criteria, five weighted pillars, and the rules developed while scoring 22 very different legal systems.

AT A GLANCE

This chapter documents the scoring framework itself: the five weighted pillars, the fifteen underlying sub-criteria, the ten rules adopted during scoring, and the limitations the Index does not attempt to overcome.

3.1 The scoring framework

Each jurisdiction is scored against fifteen sub-criteria, grouped into five weighted pillars. Every sub-criterion gets a score from 0 to 100 against a fixed rubric. The rubric anchors at 0, 25, 50, 75, and 100 are reference points, not the only allowed values. Scores can land on any integer between anchors. Each pillar score is the average of its three sub-criteria. The composite is the weighted sum of the five pillars.

PILLAR	WEIGHT	WHAT IT MEASURES
Legal density	25%	Whether binding AI law exists, how broad it is, and how mature it is. Comprehensive statutes score above sectoral rules that happen to cover AI.
Operational burden	25%	What deployers actually have to do: pre-deployment assessments, documentation, human oversight.
Enforcement capacity	20%	Whether a competent authority exists, is staffed, has published guidance, and has a track record of enforcement.
Penalty severity	15%	Maximum fines, criminal or personal liability, and injunctive or market-removal powers.
Extraterritorial reach	15%	Whether the regime binds foreign providers, requires data localisation, and cooperates on cross-border enforcement.

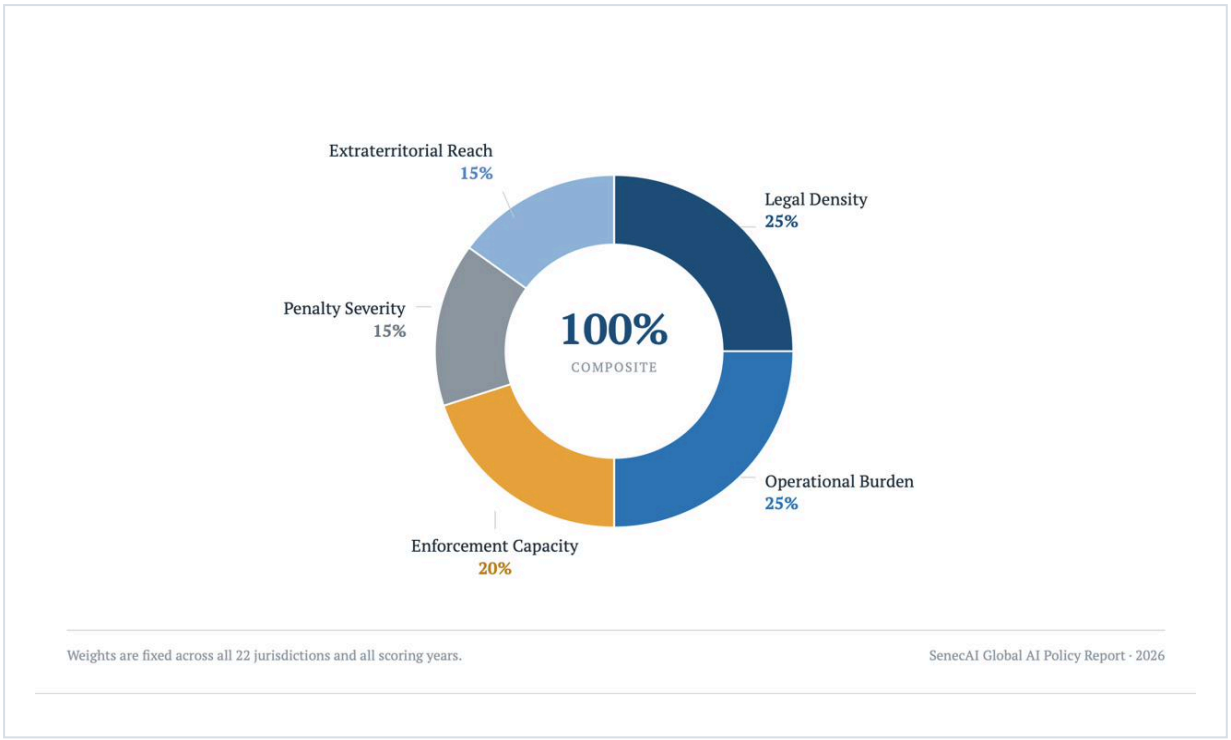

Figure 3. Five-pillar weighting.

KEY TAKEAWAY Legal density and operational burden together account for half the composite score.

The fifteen sub-criteria break down as follows. Legal density: statutory scope, binding force, maturity (enacted, partially in force, or still at bill stage). Operational burden: pre-deployment obligations, documentation and transparency, human oversight and individual rights. Enforcement capacity: whether a competent authority exists and is staffed, whether it publishes guidance, and its enforcement track record. Penalty severity: maximum fines, criminal or personal liability, injunctive or market-removal powers. Extraterritorial reach: whether foreign providers are bound, whether data localisation is required, and whether the jurisdiction cooperates cross-border.

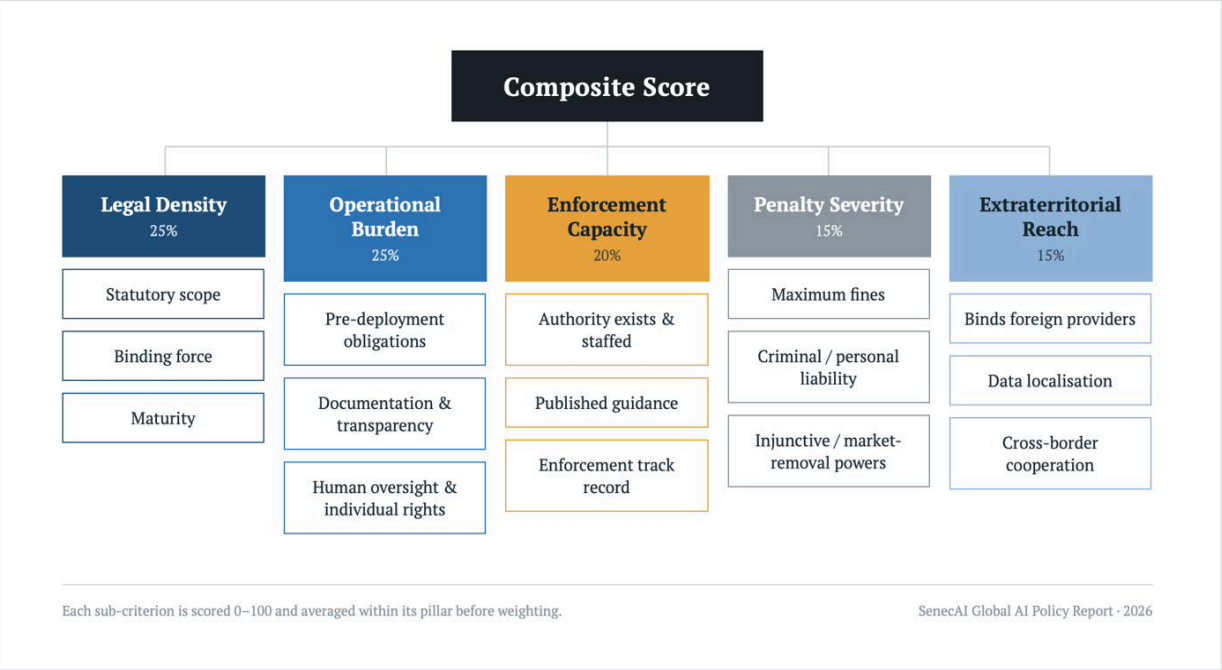

Figure 4. The fifteen sub-criteria.

KEY TAKEAWAY Every pillar breaks into exactly three measurable sub-criteria, keeping the rubric consistent across jurisdictions.

Composites sort into three tiers: Tier 1 at 70 or above, Tier 2 at 50 or above, Tier 3 below 50.

DEFINITION A **composite score** is the weighted sum of a jurisdiction's five pillar scores. A **tier** is simply the band that composite falls into — it carries no meaning beyond the number itself.

Every score traces to a source in the jurisdiction-specific research brief behind that row. One scorer worked against the locked rubric, to keep scoring consistent across 22 very different legal systems. Each completed row was then checked cell by cell against both the rubric and the underlying brief before being confirmed final.

3.2 Rules developed during scoring

No rubric can anticipate every case a 22-jurisdiction comparison will throw up. A set of rules was adopted as specific jurisdictions forced a decision, then applied back to earlier rows where consistency required it.

Enforcement scope (Rule 1). Track-record credit counts enforcement by any national authority acting on AI-related conduct, not just actions under an AI-specific statute. AI-statute-specific enforcement counts for more than general-law enforcement that happens to touch AI. A deployer fined under data-protection law for an AI system still experiences that as AI regulatory exposure, regardless of which statute the order cites.

Anchors are reference points, not fixed values (Rule 2). The 0/25/50/75/100 anchors describe reference states. Scores interpolate to the integer that best reflects where a jurisdiction actually sits, rather than snapping to the nearest anchor.

AI-specific weighting, with a stacking exception (Rule 3). Instruments that specifically target AI weight above general statutes that happen to reach AI. Exception: several stacked AI-specific instruments can score as high as one comprehensive statute, if their combined obligation on deployers is comparable. China's five-instrument regime established this. California's stack of disclosure statutes confirmed it also applies outside a centralized-authority system.

Pre-enactment credit goes to maturity only (Rule 4). An advanced bill that hasn't passed gets credit only through the maturity sub-criterion, never spread across obligation or penalty cells that assume an operative statute. Brazil is scored on the LGPD regime actually in force, plus a maturity bump for the pending AI bill. This is flagged for revision once the bill passes.

Authority calibration (Rule 5). A single consolidated super-regulator scores 80–100 on authority. A strong two-body arrangement scores 60–65. DPA-led or fragmented supervision scores 40–55. This is why Saudi Arabia's SDAIA scores well above the UAE's distributed model, and part of why China's CAC scores above purely sectoral jurisdictions.

Soft law counts under guidance, not obligation (Rule 6). Ethics charters and voluntary frameworks that become binding through government procurement are scored under guidance, not operational burden. The obligation is real, but it depends on who the deployer chooses to sell to, not on the general population of deployers.

Four more rules apply specifically to the five US state and federal rows. See 3.4.

> **KEY INSIGHT**
>
> The stacking exception in Rule 3 is what allows California, a jurisdiction with no single comprehensive AI statute, to reach Tier 1. Several AI-specific instruments stacked together can equal one comprehensive law, if the combined burden on deployers is comparable.

3.3 The EU member-state delta model

France, Germany, Italy, and Spain are each scored as the EU baseline plus a national delta, not independently from scratch. Nine cells reflect AI Act provisions that apply identically across the EU, and are capped at the EU baseline for every member state. Four cells, criminal liability, data localisation, guidance, and authority, can score above the baseline where a country's own instruments add something real. An anti-double-counting rule stops a national regulator's enforcement record, the CNIL's or Italy's Garante, from being credited in both the EU row and the member-state row. This is why member states can and do outscore the EU baseline: the floor is the same for all four, and what they add on top is earned nationally.

3.4 The US federal-plus-state model

The US isn't scored on the EU's baseline-plus-delta structure, because there's no federal AI statute for states to build on. Each US row, Federal, California, Colorado, Texas, New York, and Connecticut, is scored independently against the same rubric. That's why several states outscore the federal government: they've enacted a statute the federal row lacks entirely, not because they're adding a delta to a shared floor. Because there's no federal baseline to protect, the double-counting rule runs the other way: general federal law that a state deployer also faces is counted once, in the US Federal row. State rows only get credit for state-specific instruments and state AG enforcement of state law.

Four rules specific to this set:

State AG enforcement counts (Rule 7). Enforcement by a state Attorney General under general consumer-protection or data-privacy law counts toward that state's track record, even without an AI-specific statute behind it. This is why Texas's enforcement record counts despite TRAIGA's narrow scope.

Repeal-and-replacement and litigation stays (Rule 8). If a statute is repealed and replaced mid-project, the successor is scored as the current law. An active federal stay is reflected by lowering maturity and injunctive-power scores, not by dropping the row. Colorado's replacement of CAIA, stayed in *X.AI LLC v. Weiser,* with successor SB 26-189 is the precedent.

Frontier-only scope (Rule 9). A statute that only covers frontier-model developers, not deployers generally, is capped below comprehensive deployer-facing regimes on every scope-sensitive cell, no matter how strict it is for the developers it does cover. New York's RAISE Act set this precedent; it also applies to Connecticut's narrower frontier provisions under SB 5.

Stacked disclosure can equal a single statute (Rule 10). California confirmed that Rule 3's stacking exception works outside a centralized authority. A jurisdiction with no single horizontal statute can still reach Tier 1 through stacked disclosure rules, multiple enforcement bodies, and dense documentation requirements, if the total burden on deployers is comparable to a comprehensive regime.

3.5 Limitations

Three things this Index can't do.

LIMITATION 1

It measures how much regulatory burden a deployer experiences, not how well a regime is implemented or how sound the policy is. A jurisdiction can score high while its enforcement apparatus is still inexperienced; that's a different question from what this Index measures.

LIMITATION 2

Every score is a snapshot as of mid-2026. A single court ruling, effective date, or bill passage can move a row a full tier within months. The Colorado stay, South Korea's fine activation in January 2027, and Brazil's pending bill are three live examples, flagged elsewhere in this report.

LIMITATION 3

One person did all the scoring. That trades some inter-rater reliability for consistency across jurisdictions. The safeguards here are a rubric locked before scoring started, rules applied the same way once adopted, a documented source for every cell, and a second pass on every row against both the rubric and the research brief before it was finalized.

04

CHAPTER 04

The Global Picture

The full ranking of all 22 jurisdictions, and the three patterns that organize it.

AT A GLANCE

This chapter presents the complete 22-jurisdiction ranking and reads the table for the patterns that hold regardless of any single jurisdiction's story: a European top, a sectoral-absorption middle, and a bottom shaped by deferral rather than neglect.

Section 2 set out the three approaches governments have taken to AI governance and the question this report asks of each jurisdiction: given its statutes, its regulators, and its enforcement record, how intense a regulatory burden should an organization expect to carry there. This section presents the answer as a single ranked list.

4.1 The full ranking

Composite scores are the weighted sum of five pillar averages, each pillar itself an average of three sub-criteria scored 0 to 100 against a fixed rubric:

- **Legal density (25 percent)** — binding instrument, scope breadth, maturity
- **Operational burden (25 percent)** — pre-deployment obligations, documentation, human oversight
- **Enforcement capacity (20 percent)** — authority, guidance, track record
- **Penalty severity (15 percent)** — maximum fine, criminal liability, injunctive powers
- **Extraterritorial reach (15 percent)** — foreign providers, data localization, cross-border enforcement

The resulting composite sorts each jurisdiction into one of three tiers: Tier 1 (composite of 70 or above, high intensity), Tier 2 (50 to 70, moderate intensity), and Tier 3 (below 50, emerging).

Figure 5. Average pillar profile by tier.

KEY TAKEAWAY Tier 1 jurisdictions separate from Tier 2 and 3 mainly on legal density and operational burden, not on penalty severity.

RANK	JURISDICTION	APPROACH	COMPOSITE	TIER
1	Italy	Comprehensive legislation	91.2	Tier 1
2	France	Comprehensive legislation	90.6	Tier 1
3	Spain	Comprehensive legislation	90.2	Tier 1
4	Germany	Comprehensive legislation	87.5	Tier 1
5	China	Sovereign/centralized model	87.1	Tier 1
6	EU (baseline)	Comprehensive legislation	81.9	Tier 1
7	South Korea	Comprehensive legislation	77.6	Tier 1
8	California	Sectoral absorption	72.5	Tier 1
9	Canada	Sectoral absorption	66.8	Tier 2
10	Brazil	Comprehensive legislation	64.2	Tier 2
11	Saudi Arabia	Sovereign/centralized model	56.1	Tier 2
12	United Kingdom	Sectoral absorption	55.0	Tier 2
13	South Africa	Sectoral absorption	54.3	Tier 2
14	Japan	Sectoral absorption	53.6	Tier 2
15	Texas	Sectoral absorption	50.4	Tier 2
16	Singapore	Sectoral absorption	50.0	Tier 2
17	UAE	Sovereign/centralized model	46.8	Tier 3
18	US Federal	Sectoral absorption	46.3	Tier 3
19	Colorado	Sectoral absorption	39.6	Tier 3
20	India	Sectoral absorption	39.6	Tier 3
21	New York	Sectoral absorption	34.2	Tier 3
22	Connecticut	Sectoral absorption	31.4	Tier 3

Eight jurisdictions sit in Tier 1, eight in Tier 2, and six in Tier 3.

KEY INSIGHT

Tier 1 is not a European or Western club: China and California sit alongside Italy, France, Spain, and Germany at the top of the ranking, reached through two entirely different governance approaches.

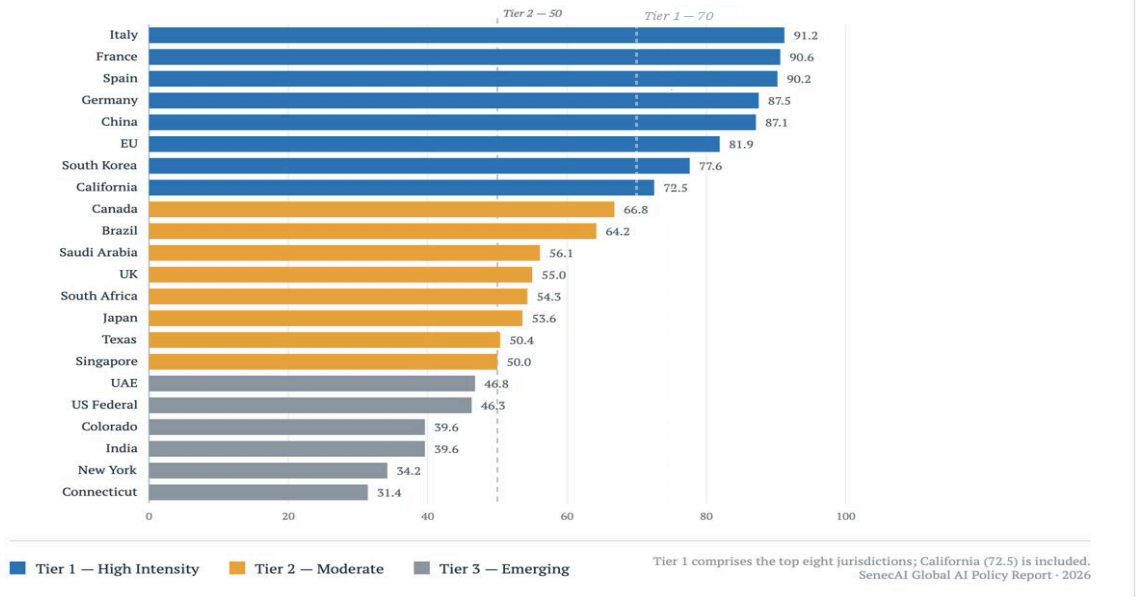

Figure 6. Full ranking by composite score.

KEY TAKEAWAY The top of the ranking is dominated by comprehensive-legislation jurisdictions, but China breaks that pattern at rank five.

4.2 Reading the table

Three patterns organize the list, independent of any single jurisdiction's story.

The top of the table is European, but nationally rather than supranationally so. Italy, France, Spain, and Germany each score above the EU baseline itself, and together occupy four of the top six positions. The EU AI Act functions as a floor common to all five rows; what separates the member states from the baseline, and from each other, is the national law layered on top of it, not the Regulation.

The middle of the table is dominated by jurisdictions with no horizontal AI statute at all. Canada, the United Kingdom, South Africa, Japan, Texas, and Singapore each reach Tier 2 through some combination of sectoral rules, general consumer or data protection law, and active regulators applying that law to AI conduct, rather than through an AI-specific statute reaching every deployer. Their position on this list depends less on what has been legislated than on how forcefully existing law is being applied.

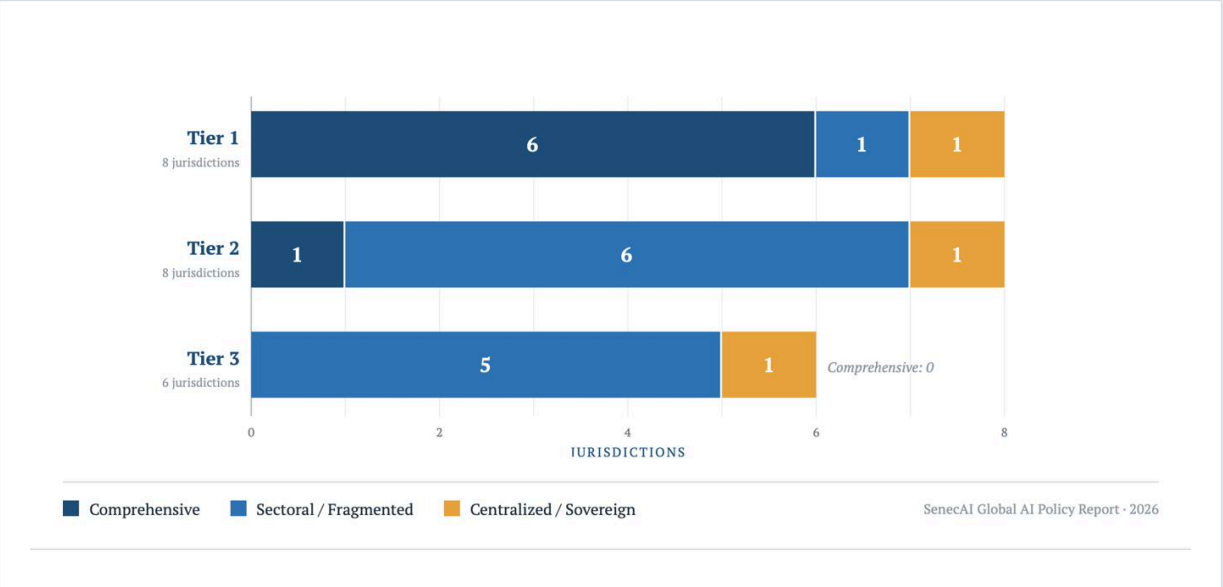

Figure 7. No archetype owns a tier.

KEY TAKEAWAY All three governance approaches appear in every tier — the approach a jurisdiction takes does not by itself predict its intensity.

The bottom of the table is not, on close reading, a list of laggards. It is a list of jurisdictions where intensity has been deliberately narrowed or deferred: by a statute that names only frontier developers rather than deployers generally, by a litigation stay, by a scope limited to a specific harm, or by a legislative process still in progress. What produced each jurisdiction's position at the bottom differs case by case, a distinction the Region by Region section takes up directly.

One further point is worth stating plainly at this stage, since the table invites the comparison on its own: the United States does not appear as a single row. It appears as six, spanning from California in Tier 1 to Connecticut in Tier 3, a spread wider than the distance separating most pairs of countries elsewhere in this Index. A national compliance program built around a single assessment of "US risk" will not reflect the reality this table shows.

05

CHAPTER 05

Region by Region

Working through each grouping in turn: what actually awaits an organization that deploys an AI system there.

AT A GLANCE

This chapter works through each regional grouping in turn — the EU and its member states, the United States, China, East Asia, the Gulf, and the rest of the world — answering the same question each time: what actually awaits a deployer there.

5.1 Overview

The 22 jurisdictions in this Index vary more within regions than most readers expect. The United States spans a wider range on its own than most pairs of countries do. China ranks in the global top five without a single horizontal AI statute. Two Gulf states, built on nearly identical instruments, land nine points apart because of one thing: how consolidated their institutions are. What follows works through each grouping in turn, answering the same question each time: what actually awaits an organization that deploys an AI system there.

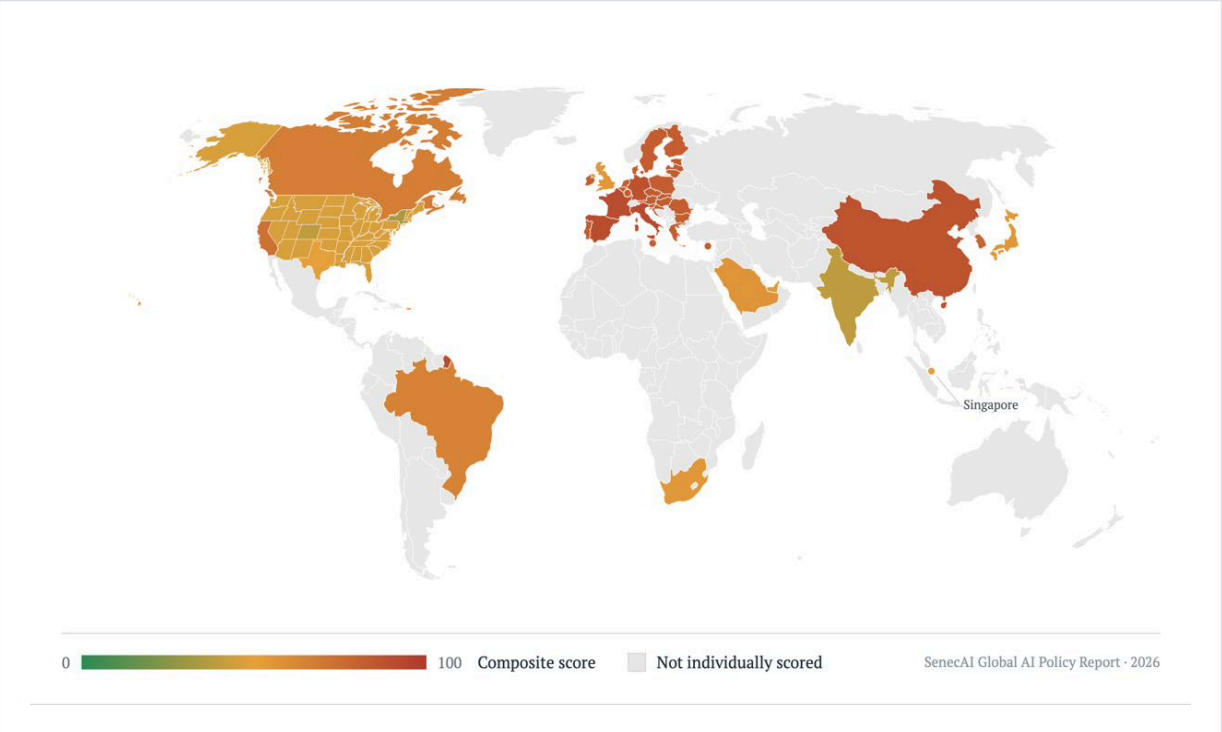

Figure 8. World map by tier.

KEY TAKEAWAY High-intensity jurisdictions cluster in Europe and East Asia, but California and China break the geographic pattern.

5.2 The European Union and its member states

Membership in the EU sets a common floor for AI regulation. It does not set a common ceiling. Every EU member state included in this Index lands above the EU baseline itself, each for a different reason.

JURISDICTION	COMPOSITE	RANK	WHAT DRIVES THE DIFFERENCE
Italy	91.2	1	Law 132/2025, the first national AI statute in the EU; a deepfake criminal offense; the bloc's heaviest workplace-AI regime; dual AgID/ACN authority
France	90.6	2	Loi SREN deepfake offenses; CNIL as a de facto AI regulator with a dedicated AI unit and Europe's most active DPA enforcement record
Spain	90.2	3	AESIA operational since September 2023, an eighteen-month head start on every other national AI authority; mandatory workplace transparency
Germany	87.5	4	Works-council co-determination on workplace AI; a fragmented but deep supervisory architecture spanning BNetzA, BfDI, and the state DPAs
EU (baseline)	81.9	6	The AI Act, GDPR, the DSA, the CRA, and the recast Product Liability Directive: the shared floor every member state inherits

Nine cells tied to the AI Act itself are fixed at the same level for every member state. Four cells, criminal liability, data localization, guidance, and authority, allow a national overlay on top of that floor, and this is where the ranking above comes from. Italy added a national statute with a deepfake offense and the bloc's toughest workplace-AI rules. France's edge is institutional: the CNIL runs a dedicated AI unit and enforces more actively than any other DPA in Europe. Spain got there first, with AESIA running since 2023, well before any other national AI authority in this Index. Germany's edge is unique to it: works-council co-determination gives employees a binding say over workplace AI, something no other jurisdiction here has.

> **THE PRACTICAL TAKEAWAY**
>
> "The EU" is not one compliance target. The AI Act sets the floor, but the real ceiling in any given member state comes from national law that operates independently of Brussels, and that is usually where enforcement shows up first.

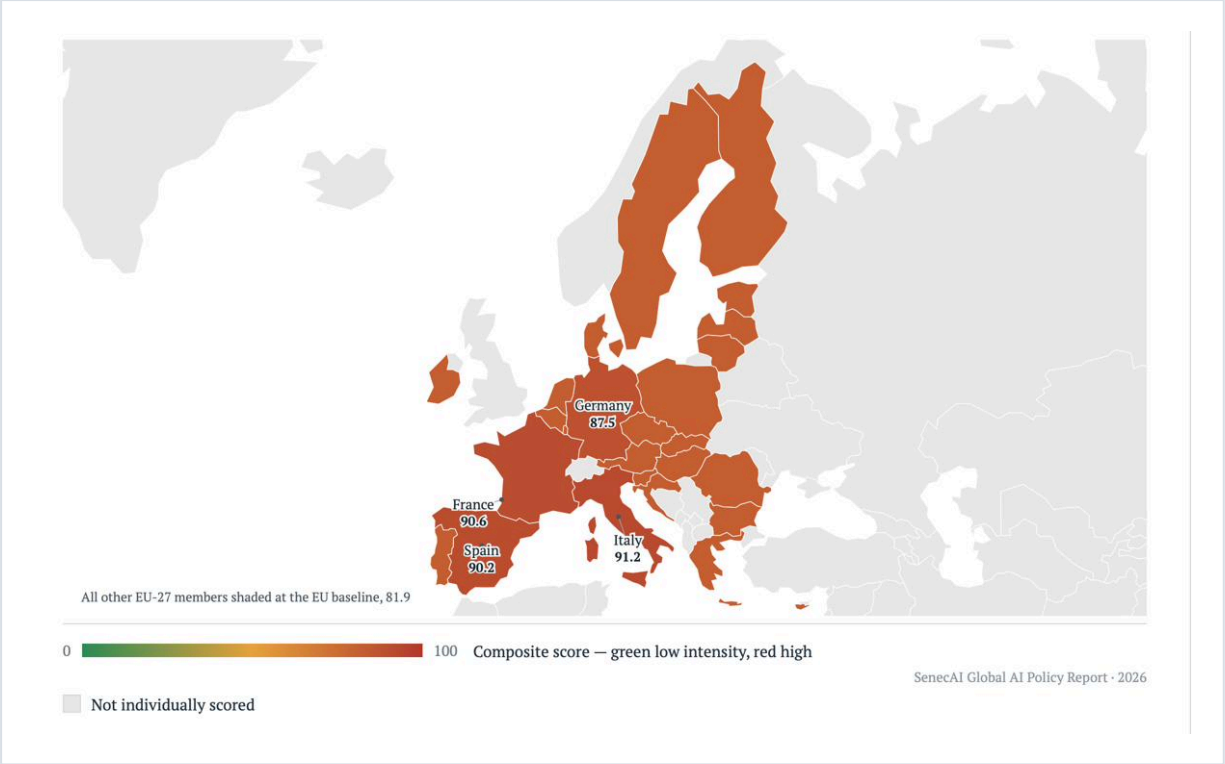

Figure 9. EU and member states.

KEY TAKEAWAY Every scored member state outranks the EU baseline itself.

Figure 10. EU baseline-plus-delta model.

KEY TAKEAWAY Nine cells are fixed at the EU baseline for every member state; only four cells allow a national overlay.

5.3 The United States

The United States enters this Index as six separate rows, spanning 41 points from top to bottom, wider than the gap between the EU baseline and India. US AI policy is genuinely fragmented: there's no federal law for the states to build on, so each row is scored on its own merits.

JURISDICTION	COMPOSITE	TIER	CHARACTER OF THE REGIME
California	72.5	T1	SB 53 (TFAIA) frontier transparency, AB 2013 training-data disclosure, SB 942, and CCPA automated-decision rules. Disclosure-based, not pre-approval
Texas	50.4	T2	TRAIGA in force since January 2026: a narrow prohibited-practices catalogue, a 60-day cure period, AG-exclusive enforcement. Real exposure comes from other state statutes
US Federal	46.3	T3	No horizontal statute; FTC Section 5, sectoral regulators, and the TAKE IT DOWN Act
Colorado	39.6	T3	SB 26-189 replaced the earlier Colorado AI Act with a considerably lighter law: no disparate-impact duty, narrower scope, in force January 2027
New York	34.2	T3	The RAISE Act reaches only frontier developers above specific revenue and compute thresholds, effective January 2027, leaving most deployers untouched
Connecticut	31.4	T3	SB 5 covers the widest range of domains of any 2026 state law but is enforced through a general consumer-protection statute

On Colorado, specifically: the original Colorado AI Act was the only US law to import disparate-impact liability into AI. It faced a federal constitutional challenge, *xAI v. Weiser*, with the Justice Department intervening on the state's side of the argument against it. Colorado's Attorney General agreed to hold off on enforcement while the case proceeded. Before the court ruled, Colorado repealed the law and replaced it with SB 26-189, so the constitutional question was never actually decided. The replacement drops the disparate-impact duty and narrows the law considerably, a real retreat from the original.

Three things stand out about the US picture. First, nearly all binding AI-specific law sits at the state level; the federal government relies on general consumer-protection authority, sectoral regulators, and one narrow statute on non-consensual intimate imagery. Second, the states are pulling in different directions on purpose: California keeps adding disclosure obligations without adopting the stricter, vetoed SB 1047 model; Texas pairs a light statute with some of the toughest state enforcement in the country under other laws; New York limited its law to frontier developers meeting specific thresholds; Connecticut covers the most ground but enforces it gently. Third, the federal government is actively working to lower state-level intensity, not just watching from the sidelines, and Colorado is the clearest example of that.

For a deployer, there's no single US compliance program to plan around. California's obligations are as heavy as South Korea's. Texas is light on paper and heavy on enforcement risk. New York and Connecticut, passed the same year, cover almost entirely different companies.

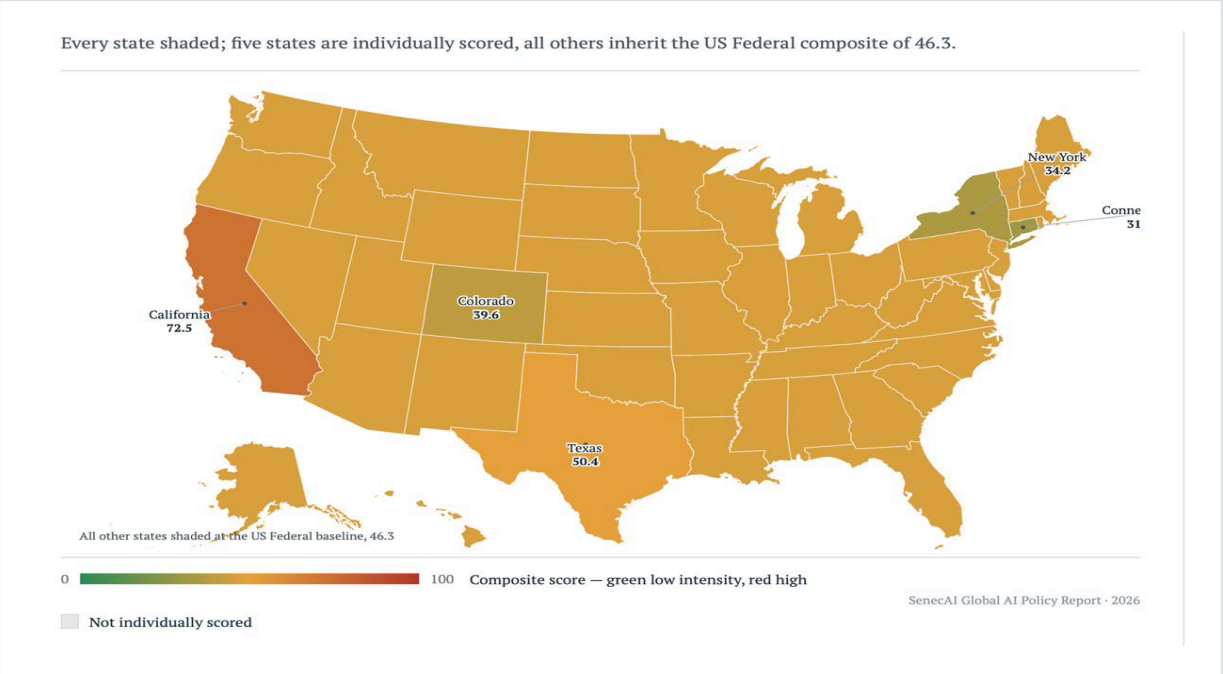

Figure 11. United States.

KEY TAKEAWAY US intensity varies more by state than by any single federal policy.

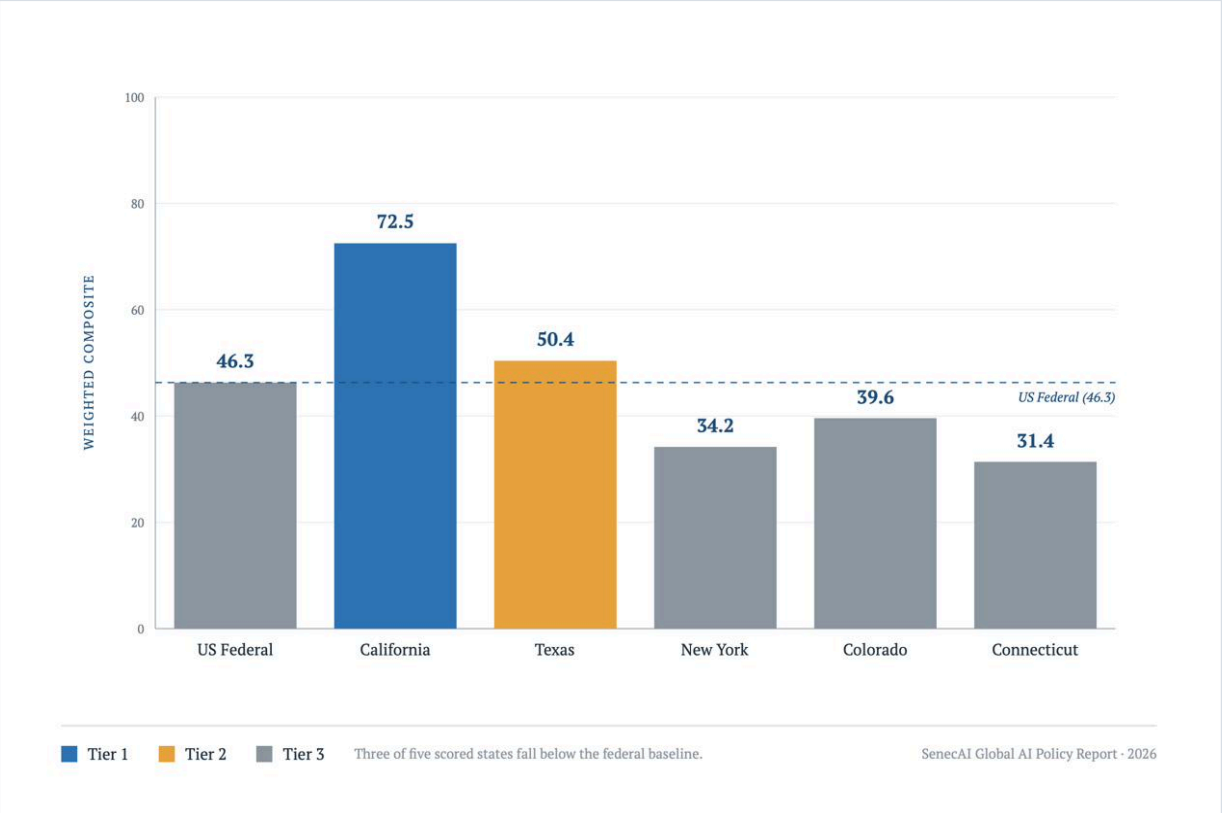

Figure 12. US state fragmentation.

KEY TAKEAWAY California and Connecticut sit at opposite ends of a 41-point spread within one country.

5.4 China

China is the cleanest example in this Index of the sovereign/centralized model: a regime that governs AI through binding, use-case-specific instruments, fast, narrow, and stacked, rather than through a single horizontal statute. As of mid-2026 there is no comprehensive AI Act in mainland China; a unified AI Law sits on the National People's Congress's 2026 agenda as a "legislative research" item, putting realistic adoption at 2027 or later. That does not make China a soft-law jurisdiction. The five-layer stack already in force, the Cybersecurity Law as amended, the Data Security Law and the Personal Information Protection Law, the Algorithm Recommendation, Deep Synthesis, and Generative AI trilogy, the AIGC Labelling Measures and their supporting national standard, and the algorithm registry itself, is materially more demanding on most operational metrics than several jurisdictions with headline-grabbing comprehensive laws.

The Personal Information Protection Law is the fine vector to plan around: up to RMB 50 million or 5 percent of turnover, the same order of magnitude as the EU AI Act's 7 percent ceiling once stacked with the Cybersecurity Law's 2025 tenfold increase to fines on critical information infrastructure operators. A single data-protection fine issued in China remains the largest ever imposed globally. The algorithm registry, mandatory pre-launch filing within ten working days, a publicly searchable filing list, and a requirement to display the model name and filing number, has no equivalent anywhere else in this Index. The AIGC Labelling Measures, in force since September 2025, add one of the world's most comprehensive whole-pipeline content-labelling regimes.

The doctrinal frame matters as much as the fine schedule. China's instruments use "service provider" as the universal duty-bearer, without the EU's clean separation between provider and deployer, so a multinational integrating a Chinese-filed model still carries its own filing, labelling, content-security, and training-data duties. For counsel mapping mainland China, the Cyberspace Administration of China is the single most consequential regulator to track, the Personal Information Protection Law is where the cash exposure lives, and the algorithm registry is the gateway condition for market access. The stack kept growing through mid-2026: dedicated rules for anthropomorphic companion services and for AI agents both took effect July 15, 2026, and the Ministry of Commerce began enforcing export-license requirements on model weights and training infrastructure the same month, evidence that the regime is still densifying rather than settling. Section 6 sets out the mechanics of each.

5.5 Singapore, Japan, and South Korea

Korea, Japan, and Singapore share geography and overlapping diplomatic commitments, and almost nothing else. Korea has enacted Asia's first comprehensive horizontal AI law. Japan built the reference point for an innovation-first statute that binds government more than it binds private companies. Singapore built the benchmark for voluntary, assurance-led governance. The 27-point spread between Korea and Singapore is wider than the distance between either of them and any other jurisdiction in this Index.

JURISDICTION	COMPOSITE	TIER	CHARACTER OF THE REGIME
South Korea	77.6	T1	AI Framework Act in force since January 2026; an “indirect impact” trigger broader than the EU’s own; AI-specific fine is trivial, real exposure runs through data-protection law
Japan	53.6	T2	AI Promotion Act binds government, not deployers; no fines of its own; Digital Agency procurement rules function as de facto compliance requirements
Singapore	50.0	T2	No binding AI statute; the Model AI Governance Framework is backed by a well-resourced regulator and real data-protection penalties; first jurisdiction with a dedicated agentic-AI framework (January 2026)

South Korea is built on the AI Framework Act, in force since January 2026, and pairs a genuinely binding statute with an “indirect impact” jurisdictional trigger broader than the EU’s own output-based test. The headline reads like the EU’s; the fine print doesn’t. The AI-specific administrative fine is trivial by comparison, and the real financial exposure runs through Korea’s data-protection law rather than the AI statute itself.

Japan operates the AI Promotion Act, a strategic statute that binds government agencies rather than private deployers, and carries no administrative fines of its own. Real compliance risk runs through the country’s data-protection law and through what amounts to procurement-as-de-facto-law: Digital Agency procurement rules and the AI Guidelines for Business bind any vendor pursuing public-sector or large enterprise work, regardless of what the AI statute itself requires.

Singapore has no binding AI statute at all, relying instead on the Model AI Governance Framework, but backs that voluntary framework with a fully resourced regulator and real data-protection penalties, which is why it still clears Tier 2. In January 2026 Singapore became the first jurisdiction to publish a dedicated governance framework for agentic AI, ahead of the EU, the UK, and Korea.

Across all three, the same lesson recurs: the AI-specific statute is rarely where the money is. In Korea, Japan, and Singapore alike, the real financial exposure runs through general data-protection law applied to AI conduct, not through the AI Act, the AI Promotion Act, or the Model Framework by name.

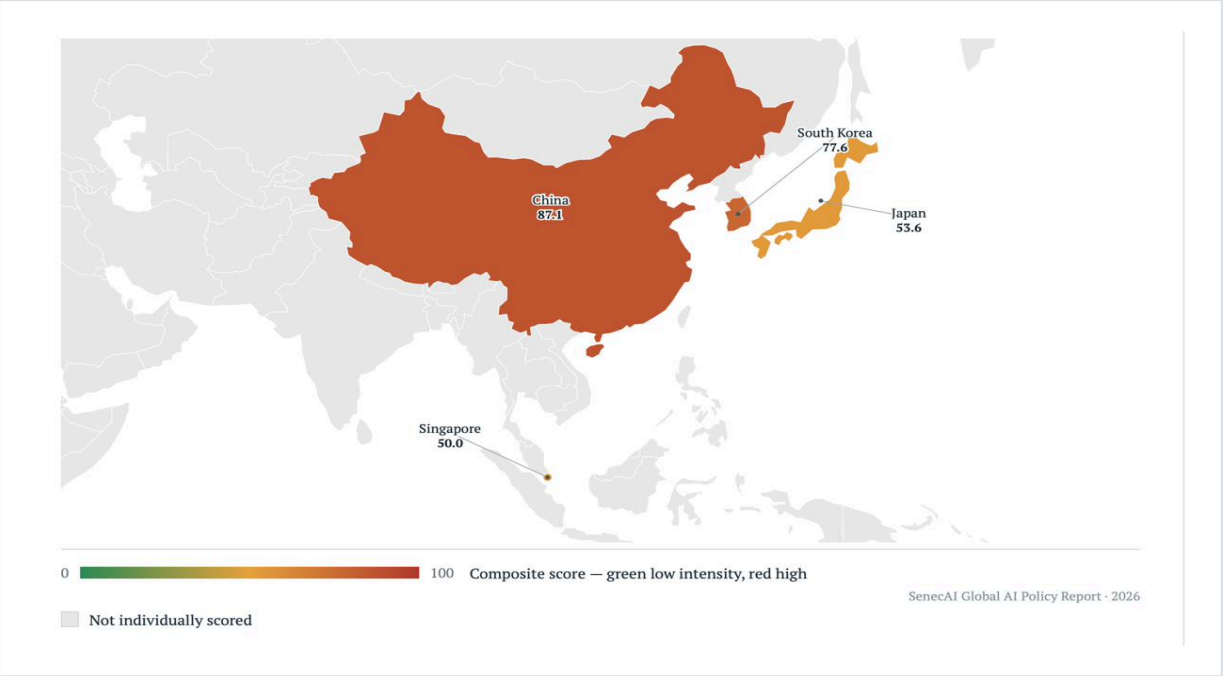

Figure 13. Singapore, Japan, South Korea.

KEY TAKEAWAY Geographic neighbors, but a 27-point spread separates Korea's binding statute from Singapore's voluntary framework.

5.6 The Gulf states

Saudi Arabia and the UAE form this Index's cleanest matched pair: nearly identical instruments producing a nine-point gap that comes down to one thing, how consolidated their institutions are.

JURISDICTION	COMPOSITE	TIER	STRUCTURE
Saudi Arabia	56.1	T2	SDAIA, the most consolidated national AI authority in this Index, chaired at the highest level of government; a data-protection law plus an anti-cybercrime law; ethics principles made binding through procurement
UAE	46.8	T3	Functions split across an AI office, several ministries, and free-zone authorities; a data-protection law plus a cybercrimes law; an AI charter made binding through procurement; heavy industrial-policy weight

Both ground AI obligation in a data-protection law with a fine ceiling well below the EU, Korea, or Singapore. Both run a non-binding AI ethics charter that becomes contractually binding the moment a deployer signs with the government. Both pair this light regulatory touch with major industrial-policy weight, sovereign investment, special economic zones, state-anchored compute, that makes the government the largest AI customer in the market.

The nine-point gap is institutional, not textual. Saudi Arabia's SDAIA does everything, data regulator, ethics authority, planner, procurement gatekeeper, in one body reporting to the top of government. The UAE splits the same functions across an AI office, several ministries, and free zones, which lowers its enforcement score even with the same underlying ambition. In either country, the contract, not the statute, is the real AI law, and the sanction that matters is exclusion from the government relationship that anchors the market. Both jurisdictions added a formal instrument in mid-2026 without changing this basic picture: Saudi Arabia's new Copyright Law, in force from August 1, 2026, carves out an AI-training exception from author permission; the UAE's DIFC free zone has run its own AI-specific Regulation 10 since January 2026, a narrower binding layer that sits alongside, not instead of, the procurement relationship described above.

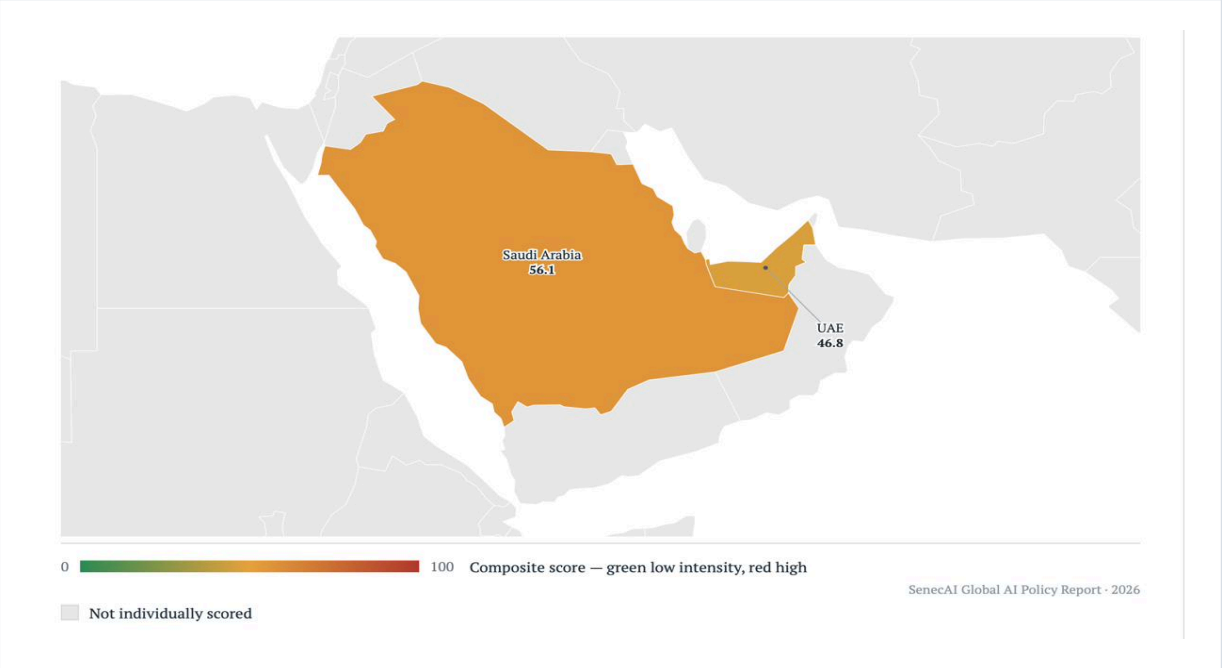

Figure 14. The Gulf states.

KEY TAKEAWAY Nearly identical instruments still produce a nine-point gap, driven entirely by institutional consolidation.

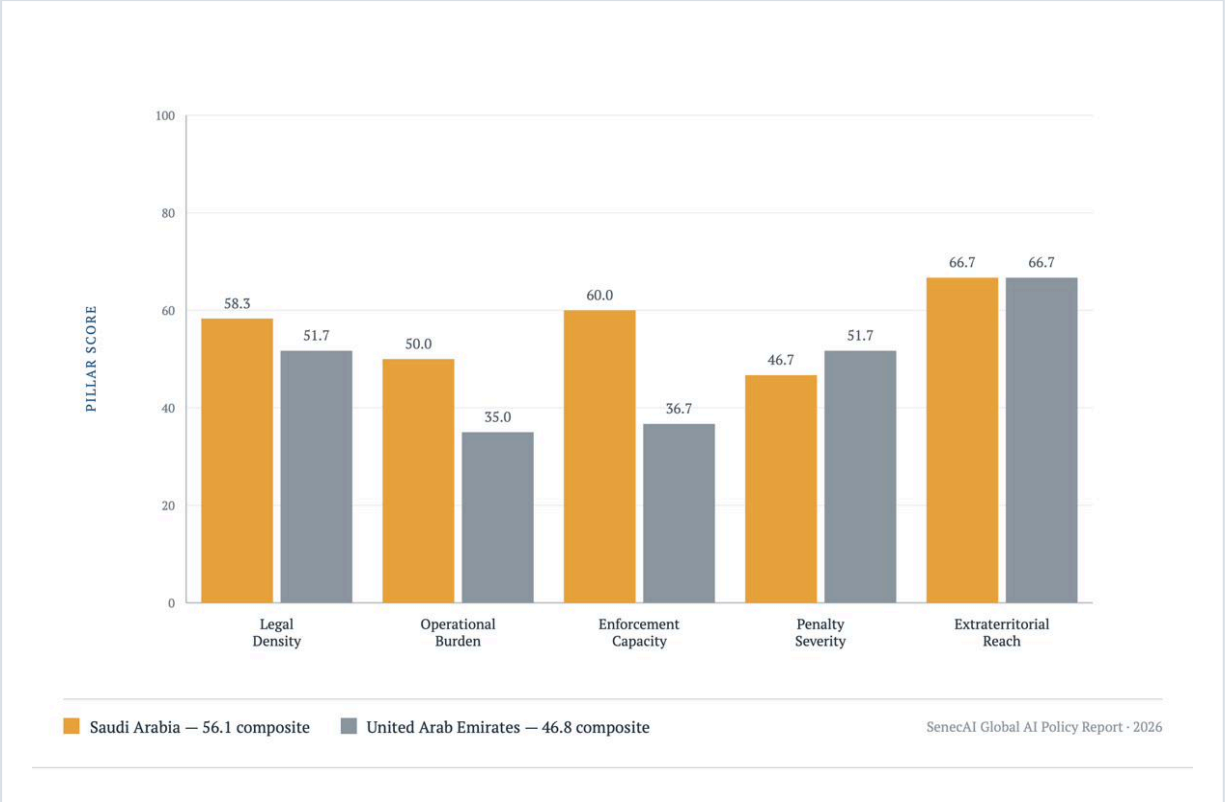

Figure 15. Gulf comparison.

KEY TAKEAWAY Saudi Arabia's single consolidated authority outscores the UAE's split structure on enforcement capacity alone.

5.7 Rest of the world: Canada, Brazil, and South Africa

Canada, Brazil, and South Africa share a profile more than a region: each has meaningful AI-adjacent enforcement capacity today, but none has a binding horizontal AI statute in force yet. Brazil is the closest to changing that: its comprehensive bill, PL 2338, is well advanced and could pass within the life of this report.

Canada (66.8, Tier 2) is the only G7 country whose comprehensive federal AI bill died on the floor: the Artificial Intelligence and Data Act lapsed in January 2025 and hasn't been re-tabled. Regulatory weight has shifted to the provinces and to sectoral regulators instead. Quebec's Law 25 is the de facto AI rulebook, with the broadest reach in the country and penalties up to CAD 25 million or 4 percent of worldwide turnover. Ontario adds workplace AI-hiring disclosure from January 2026. The banking supervisor OSFI has extended its risk guideline to cover AI directly, giving Canada real sectoral teeth even without a national statute.

Brazil (64.2, Tier 2) pairs an advanced comprehensive bill (PL 2338) with a data-protection law, the LGPD, that already produces real enforcement, and an electoral AI regime with no peer in this Index: Brazil's electoral court runs a deepfake and synthetic-content enforcement operation that goes fully live again for the country's general elections, backed by cooperation agreements with every major platform in the country. Once PL 2338 passes, Brazil's score is likely to move well into Tier 1.

South Africa (54.3, Tier 2) is the only African jurisdiction in this Index, chosen over Nigeria, Kenya, and Egypt because it currently combines the continent's most active data-protection enforcement with its most advanced formal AI-policy instrument. Nigeria and Kenya both have AI-specific bills before their legislatures, but neither had cleared the process as of mid-2026; Egypt has a national AI strategy but no enacted AI-specific law. South Africa's Draft National AI Policy, gazetted and Cabinet-approved in April 2026, is further along than either, and its data-protection regulator has actually issued fines rather than guidance alone. That said, South Africa shows the widest gap in this cluster between paper ambition and enforcement reality: its regulator has issued only a handful of fines to date, none yet collected, and the country has no AI-specific statute and no AI safety institute as of mid-2026. For a deployer, the immediate exposure runs through the data-protection law and sector regulators; the AI-specific statute is a 2027-or-later prospect.

06

CHAPTER 06

Who Enforces AI Policy

And what deployers must actually do, jurisdiction by jurisdiction.

AT A GLANCE

This chapter names the actual enforcing authority for every one of the 22 jurisdictions and sets out the concrete obligations a deployer must meet in each, organized by governance approach.

Enforcement authority in this Index follows three patterns. Comprehensive-legislation jurisdictions designate a specific AI regulator, typically alongside an existing data-protection authority. Sectoral-absorption jurisdictions have no dedicated AI regulator; enforcement falls to whichever existing regulator has jurisdiction over the underlying conduct. Sovereign/centralized jurisdictions concentrate AI rule-making, enforcement, and commercial gatekeeping within the executive branch, unconstrained by the separation of powers that shapes the other two patterns. Enforcement reaches deployers through one of two mechanisms: direct administrative enforcement backed by a dense body of executive rules, or state-controlled commercial leverage exercised through procurement and licensing. The common thread across both variants is institutional concentration, not the specific channel through which enforcement runs.

Across all three patterns, enforcement in this Index is mostly administrative rather than judicial. Courts function as a review mechanism in most cases, not the primary point of contact. China is a partial exception: it combines administrative fines with an active civil court docket on copyright and training-data disputes.

Comprehensive legislation

The EU

REGULATOR The AI Office, part of the European Commission, has direct enforcement authority over providers of general-purpose AI models. National market-surveillance authorities supervise high-risk systems in each member state.

KEY OBLIGATIONS Providers of high-risk systems must complete a conformity assessment before placing the system on the market. Deployers must operate high-risk systems under human oversight, follow the provider's instructions, and monitor for risk.

DISTINCTIVE FEATURE For certain infringements, fines can reach 3% of global annual turnover or €15 million once the relevant enforcement provisions apply.

Italy

REGULATOR ACN supervises the AI Act; AgID handles conformity-body notification. The Garante, Italy's data-protection authority, has issued the country's AI-related fines to date.

KEY OBLIGATIONS Deployers face the EU baseline obligations, plus one Italian addition: in healthcare, public administration, the judiciary, and the regulated professions, a human professional must make the final decision.

France

REGULATOR France has no dedicated AI authority. The CNIL, the data-protection authority, handles enforcement and has already fined an AI vendor the GDPR maximum.

KEY OBLIGATIONS Deployers must follow CNIL guidance on training-data use. AI-generated content delivered to French consumers must be available in French, under the Loi Toubon.

Spain

REGULATOR AESIA, the first AI Act authority designated in the EU, supervises alongside the AEPD, the data-protection authority.

KEY OBLIGATIONS Deployers meet the AI Act baseline through AESIA. Employers using AI in hiring, evaluation, or dismissal decisions must inform worker representatives in advance.

Germany

REGULATOR BNetzA, the telecoms regulator, was designated to supervise the AI Act through a new AI Service Centre. The BfDI and sixteen state data-protection authorities handle AI matters touching fundamental rights.

KEY OBLIGATIONS Deployers face the EU baseline, plus Germany's state-level data-protection enforcement.

South Korea

REGULATOR MSIT enforces the Framework Act. The PIPC enforces the larger privacy exposure under it.

KEY OBLIGATIONS Operators of high-impact AI must maintain a risk-management plan, keep a human in the loop, and give users prior notice. Generative AI providers must label AI-generated output. Large foreign providers must appoint a domestic representative.

Brazil

REGULATOR The ANPD, the data-protection authority, functions as the AI regulator and would coordinate the SIA framework if PL 2338 passes.

KEY OBLIGATIONS Deployers processing personal data must honor the LGPD right to human review of automated decisions and complete an impact assessment for high-risk processing. Generative AI providers must meet the TSE's election-period deepfake labelling and disclosure rules.

> **KEY INSIGHT**
>
> Even within the same governance approach, who actually enforces varies widely — from a single dedicated AI authority (Spain's AESIA) to a data-protection regulator standing in for one entirely (France's CNIL). Knowing which pattern applies matters more than knowing the statute's name.

Sectoral absorption

California

REGULATOR No single regulator. The Attorney General and the CPPA share enforcement.

KEY OBLIGATIONS Businesses using automated decision-making technology for a significant decision must complete a risk assessment and provide consumers a pre-use notice and opt-out. Frontier model developers must publish a safety framework and report critical incidents.

Texas

REGULATOR The Texas Attorney General is the sole enforcer.

KEY OBLIGATIONS TRAIGA prohibits AI deployed with intent to discriminate, manipulate, or produce unlawful content.

DISTINCTIVE FEATURE The Attorney General's enforcement record under the state's biometric-privacy and consumer-data statutes is larger, including a $1.4 billion settlement with Meta.

Colorado

REGULATOR The Colorado Attorney General enforces.

KEY OBLIGATIONS The current statute, SB 26-189, is stayed pending litigation. It would impose a duty of reasonable care on deployers of high-risk systems to prevent algorithmic discrimination, with consumer transparency notices.

New York

REGULATOR The Attorney General enforces. A new office, DIGIT, housed at DFS, is not yet operational.

KEY OBLIGATIONS The RAISE Act applies only to large frontier developers, who must publish a safety framework before deployment and report critical incidents within 72 hours.

Connecticut

REGULATOR The Attorney General enforces most provisions as unfair trade practices under CUTPA.

KEY OBLIGATIONS Employers using automated employment-decision technology must notify affected applicants and employees. Frontier developers face separate whistleblower-protection duties.

The UK

REGULATOR No dedicated AI regulator exists. The ICO functions as the default regulator, alongside the FCA, Ofcom, and other sectoral bodies.

KEY OBLIGATIONS Deployers processing personal data through AI must meet the UK GDPR standard on automated decision-making. AI-generated harmful content falls under the Online Safety Act.

Canada

REGULATOR No federal AI statute is in force. The OPC enforces federal privacy law. Quebec's CAI enforces the binding ceiling in this Index.

KEY OBLIGATIONS Deployers serving Quebec must complete a privacy impact assessment before deploying an AI or automated decision-making system and provide a right to human review.

Japan

REGULATOR METI and MIC maintain the non-binding AI Guidelines for Business. The PPC enforces the binding exposure, under the APPI.

KEY OBLIGATIONS The guidelines expect deployers to classify their role, run bias and safety testing, and maintain transparency notices, though none of this is mandatory. Personal data used in AI is subject to APPI's training-data and cross-border-transfer rules.

Singapore

REGULATOR IMDA maintains the voluntary Model AI Governance Frameworks. The PDPC enforces the binding exposure, under the PDPA.

KEY OBLIGATIONS Deployers processing personal data through AI must meet PDPA training-data and consent requirements.

DISTINCTIVE FEATURE Selling to Singapore's government or major enterprises generally requires alignment with the voluntary frameworks as a procurement condition.

South Africa

REGULATOR The Information Regulator enforces POPIA, the only binding instrument reaching AI directly.

KEY OBLIGATIONS Deployers must register an Information Officer and complete an impact assessment for automated decision-making. Breaches must be reported through the mandatory eServices portal.

India

REGULATOR No single regulator exists. MeitY coordinates policy. The Data Protection Board, not yet staffed, will enforce the DPDP Act once operational.

KEY OBLIGATIONS Deployers training AI on personal data must obtain informed consent once the DPDP Act's substantive provisions take effect, in 2027. 2026 rules already require synthetic content labelling and rapid takedown.

US Federal

REGULATOR No horizontal AI statute exists. The FTC enforces under its general unfair-and-deceptive-practices authority. Sectoral regulators cover employment, health, finance, and securities.

KEY OBLIGATIONS Deployers may not make unsubstantiated AI capability claims or conceal AI use from consumers. Under the TAKE IT DOWN Act, covered platforms must maintain a notice-and-takedown system for AI-generated non-consensual intimate imagery, with removal required within 48 hours of a valid request.

DEFINITION **Sectoral absorption** means a jurisdiction has not passed a broad AI-specific law; existing data-protection, consumer-protection, or sector-specific regulators simply extend their existing authority to cover AI conduct.

The sovereign/centralized model

China

REGULATOR The CAC leads AI enforcement, with sectoral ministries, through an administrative registry rather than statute passed by a separate branch of government. Enforcement rests on a suite of existing instruments rather than a single AI law: the Cybersecurity Law, the Data Security Law, the Personal Information Protection Law, the Algorithm Recommendation Provisions, the Deep Synthesis Provisions, and the Generative AI Measures together set the obligations a deployer must meet.

KEY OBLIGATIONS Deployers must register generative AI services with the CAC before launch and label AI-generated content across all formats and channels. Enforcement runs mainly through administrative fines, rectification orders, and warnings. Civil courts handle copyright and training-data disputes separately.

DISTINCTIVE FEATURE The apparatus continues to expand. Interim Measures for the Administration of AI Anthropomorphic Interactive Services, issued jointly by the CAC and four other departments, took effect July 15, 2026: they require algorithm filing, a security assessment, and clear AI-disclosure for companion and chatbot services, bar virtual intimate-relationship services for minors, and require parental consent for other anthropomorphic services offered to under-14s. Implementation Opinions on intelligent agents, jointly issued by the CAC, NDRC, and MIIT and effective the same day, create China's first dedicated regulatory category for AI agents, treated separately from generative AI. The Ministry of Commerce has also begun enforcing export-license requirements on model weights, training pipelines, and inference infrastructure developed by Chinese companies before transfer to foreign entities, issuing administrative warnings to three companies in the first week of enforcement.

Saudi Arabia

REGULATOR SDAIA, chaired by the Crown Prince, sets AI policy, runs data protection through the NDMO, and administers the procurement-linked ethics standards, within the same executive structure that directs state AI investment vehicles including HUMAIN.

KEY OBLIGATIONS Deployers seeking government or giga-project business are expected to show ISO 42001-grade governance and alignment with SDAIA's AI Adoption Framework. Formal statutory exposure is modest: the PDPL penalty ceiling is roughly $1.3 million.

DISTINCTIVE FEATURE A new Copyright Law, in force from August 1, 2026, adds a text-and-data-mining exception permitting reproduction of copyrighted works for AI training or algorithm development without the author's permission or compensation, subject to a proportionality test limiting use to what serves that purpose; the scope of the exception awaits forthcoming Implementing Regulations.

The UAE

REGULATOR As of June 2026, the UAE consolidated its AI Office, data-protection function, and the digital-government arm of its telecoms regulator into a single Federal Authority for Artificial Intelligence and Data, reporting to Cabinet and chaired by the Minister of State for AI. The onshore PDPL Executive Regulations remain pending; industry accounts describe the underlying data-authority function as never fully operational under the former Data Office, with the new Authority now the likely catalyst to finally move them forward. A National AI System sits as an advisory Cabinet member.

KEY OBLIGATIONS Deployers processing personal data must meet the PDPL's impact-assessment and breach-notification requirements. Deployers selling to federal or federally adjacent customers are expected to align with the UAE Charter's twelve principles, now administered by the new Authority.

DISTINCTIVE FEATURE Within the DIFC free zone, a separate AI-specific instrument, DIFC Regulation 10, has been in force since January 2026: it requires AI impact assessments, transparency obligations for AI-driven decisions, and documentation of high-risk use cases, with fines of USD 25,000 to 50,000 per violation — a distinct free-zone layer relevant to multinationals structuring UAE entities through DIFC or ADGM.

07

CHAPTER 07

Does the Law Match the Enforcement?

Comparing legal density against enforcement capacity across all 22 jurisdictions.

AT A GLANCE

This chapter compares each jurisdiction's legal density against its enforcement capacity, showing where the gap is widest, where it is negative, and what that gap means for compliance planning.

Legal density and enforcement capacity do not always move together. A jurisdiction can score high on one and low on the other. Comparing the two pillars shows where a statute is backed by institutional capacity and where it is not.

Legal density alone is an incomplete measure of deployer exposure. A broad statute behind an under-resourced regulator produces less practical exposure than a narrow statute behind an active one.

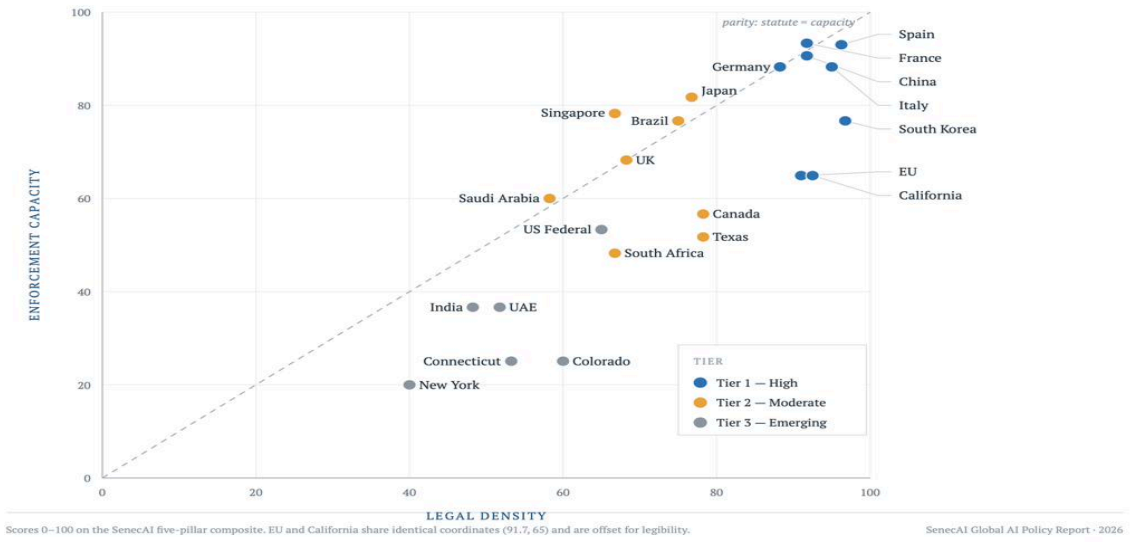

Figure 16. Does the law match the enforcement? Legal density vs. enforcement capacity.

KEY TAKEAWAY Points above the line have more statute than enforcement muscle; points below have more muscle than statute.

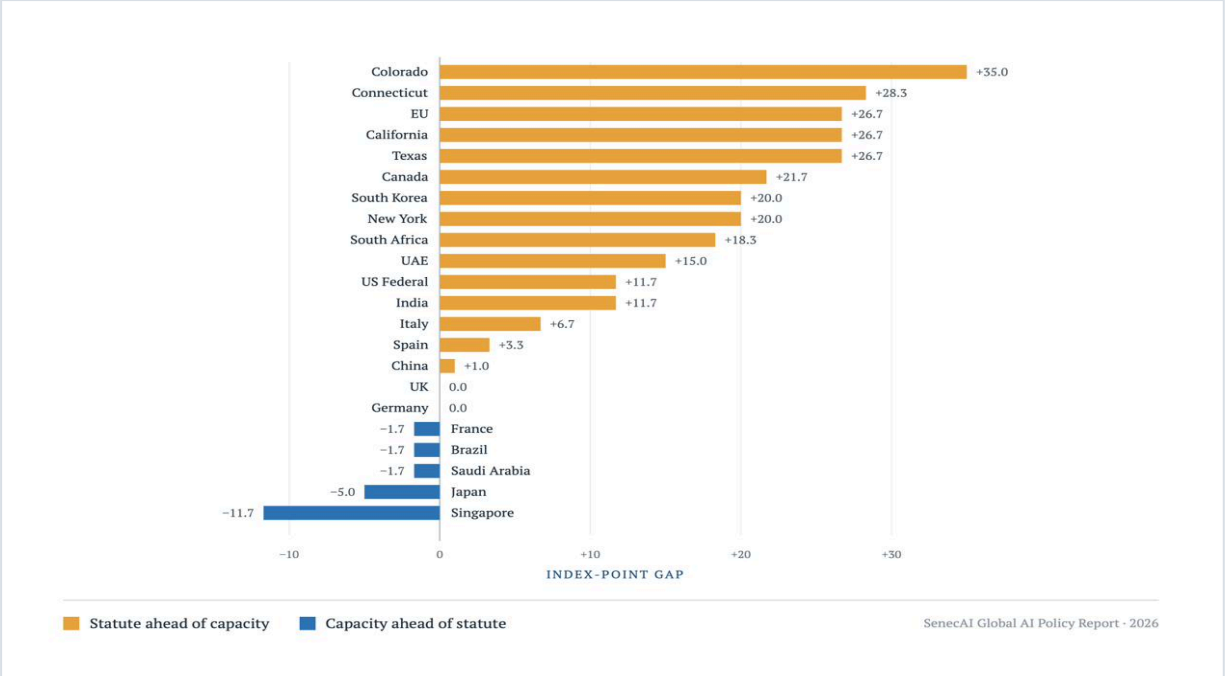

Figure 17. Diverging gap, all 22 jurisdictions.

KEY TAKEAWAY Colorado shows the widest gap in the Index; Singapore shows the widest gap in the opposite direction.

The pattern across all 22 jurisdictions

JURISDICTION	LEGAL DENSITY	ENFORCEMENT CAPACITY	GAP
Colorado	60.0	25.0	+35.0
Connecticut	53.3	25.0	+28.3
EU	91.7	65.0	+26.7
California	91.7	65.0	+26.7
Texas	78.3	51.7	+26.7
Canada	78.3	56.7	+21.7
South Korea	96.7	76.7	+20.0
New York	40.0	20.0	+20.0
South Africa	66.7	48.3	+18.3
UAE	51.7	36.7	+15.0
US Federal	65.0	53.3	+11.7
India	48.3	36.7	+11.7
Italy	95.0	88.3	+6.7
Spain	96.3	93.0	+3.3
China	91.7	90.7	+1.0
UK	68.3	68.3	0.0
Germany	88.3	88.3	0.0
France	91.7	93.3	−1.7
Brazil	75.0	76.7	−1.7
Saudi Arabia	58.3	60.0	−1.7
Japan	76.7	81.7	−5.0
Singapore	66.7	78.3	−11.7

What the pattern shows

The gap does not align with the three governance approaches. Each approach can produce a wide gap, a narrow one, or a negative one. Governance approach determines what gets written into law. Enforcement capacity is a separate, later determination.

Colorado, Connecticut, the EU, California, and Texas show the widest gaps in the Index, for different reasons. Colorado's SB 26-189, the successor to the original Colorado AI Act, which was stayed in *X.AI LLC v. Weiser* and repealed before taking effect, does not itself take effect until January 2027. Connecticut's AI-specific provisions are recent, and CUTPA enforcement against them has no track record yet. The EU's AI Act is the densest single instrument in the Index; the AI Office is well staffed, but not all national market-surveillance authorities have been designated, and enforcement actions have not begun. California spreads disclosure duties across several agencies, with no single regulator carrying the enforcement weight seen in China or Italy. Texas is the exception in this group: TRAIGA is narrow, but the Attorney General's enforcement record under the state's consumer-protection and data-privacy statutes, including a $1.4 billion settlement with Meta, already exceeds what the statute's text implies.

Canada, South Korea, New York, South Africa, UAE, US Federal, and India show a moderate positive gap. In most cases, enforcement capacity is still being built rather than absent. South Korea's Framework Act and Canada's sectoral rules are backed by regulators still developing toward their full mandate. New York and India show the same pattern at a lower overall level, on narrower statutes. South Africa and the UAE combine substantial statutory text with authorities that remain under-resourced relative to their mandate.

Italy, Spain, China, the UK, and Germany show a gap near zero. In each case, the regulator behind the statute is largely built out. Italy and Spain pair national additions to the EU baseline with active enforcement bodies. China backs its suite of AI-related instruments, the Cybersecurity Law, the Data Security Law, the Personal Information Protection Law, and the algorithm- and generative-AI-specific rules built on top of them, with matching administrative capacity. The UK and Germany reach the same alignment through existing sectoral regulators rather than a dedicated AI authority.

France, Brazil, Saudi Arabia, Japan, and Singapore show a negative gap: enforcement capacity exceeds legal density. France and Brazil pair moderate statutory reach with regulators, the CNIL and the ANPD, that are already active beyond what the AI-specific text requires. Saudi Arabia's procurement and licensing relationships carry more weight in practice than its formal statute. Japan's AI Promotion Act carries no monetary penalties, but the ministries and data protection commission behind it are fully staffed and active. Singapore's frameworks remain voluntary by design, while its enforcement bodies, the IMDA and the PDPC, already function above what the voluntary text implies.

Legal density describes what a jurisdiction has written down. Enforcement capacity describes what a regulator is prepared to act on. A compliance program based on the statute's text alone will misjudge exposure wherever the two diverge.

08

CHAPTER 08

Potential Criticisms and Responses

The objections a careful reader is likely to raise, answered directly.

AT A GLANCE

This chapter sets out the eight most legitimate objections a careful reader is likely to raise about this Index's scope, methodology, and authorship, and answers each directly.

A comparative index of this kind can, of course, be subjected to a number of legitimate criticisms. This section sets out the objections a careful reader is likely to raise, and answers them directly, rather than leaving them for the reader to discover on their own.

8.1 Why these 22 jurisdictions, and not others.

It may seem obvious why the European Union, the United States, China, and South Korea belong in a study of this kind. What is less obvious is why the selection stops where it does within each of those categories, and a careful reader is likely to ask a more specific version of the question: why these four EU member states and not a fifth; why these five US states and not a sixth; why Brazil and not Argentina; why South Africa and not Nigeria or another African jurisdiction. The honest answer is that any line drawn at 22 jurisdictions is, at the margin, somewhat arbitrary. Inclusion in this Index reflects a combination of market size and AI deployment concentration, and the existence of a binding instrument or enforcement record substantial enough to score meaningfully. Italy, France, Spain, and Germany were selected among EU member states, and California, Texas, Colorado, New York, and Connecticut among US states, on that same basis. Brazil was included over Argentina, and South Africa over Nigeria or other African jurisdictions, for the same reason: at the time of scoring, each had the more developed binding instrument or enforcement record to measure. Argentina, Nigeria, Australia, Ireland, and Illinois were considered and set aside as close calls, not overlooked.

8.2 Why China scores higher than the EU.

China's composite score, the highest in the Index among sovereign/centralized jurisdictions and above the EU baseline itself, is not one that most deployers would describe as more rights-protective, more procedurally predictable, or more comfortable to operate under than the European Union. That is not a contradiction, once the variable being measured is kept in view: this Index scores regulatory intensity as experienced by a deployer, not rule-of-law quality, procedural fairness, or institutional legitimacy. China's score reflects administrative rules that are dense, bind private-sector deployers directly, and are consistently enforced, which is precisely what this Index measures. The European Union's score reflects a genuinely binding, rights-protective framework whose enforcement is distributed across member states and moves more slowly from statute to sanction. A high score in this Index is not an endorsement of the regime that produced it, and a comparatively lower score is not a criticism of the regime that produced it either. The Index measures burden and enforcement reality, not desirability.

8.3 Why five pillars, and not others.

Two additional pillars were considered while this framework was being built, and set aside. The first was a standalone pillar measuring a jurisdiction's alignment with international frameworks such as the OECD AI Principles or the Council of Europe Framework Convention on AI. This was set aside because that alignment is already captured indirectly, through the Legal Density and Guidance sub-criteria, and adding a separate pillar for it would have double-counted the same underlying fact without changing what a deployer actually experiences on the ground. The second was a pillar measuring public trust and civil-society engagement in a jurisdiction's AI governance process. This was set aside because it measures something closer to institutional legitimacy or regulatory readiness than to deployer-facing burden, and this report treats those as distinct variables that this Index does not attempt to score together. The five pillars that remain, Legal Density, Operational Burden, Enforcement Capacity, Penalty Severity, and Extraterritorial Reach, were chosen because each translates directly into a cost or a risk a deploying organization actually bears.

8.4 This Index is the work of one author.

This First Edition reflects the analysis and scoring judgment of a single author, working alone rather than through a research team. That was a deliberate choice for this edition, made in the interest of maintaining a single, consistent point of view across all 22 jurisdictions, and it carries the limitation that comes with any single-author work: no second, independent perspective checked the result before publication.

8.5 Comparing legal systems that are not, in a strict sense, comparable.

This Index places common-law US states, civil-law EU member states, China's administrative-registry model, and the Gulf states' contract-and-licensing regimes on a single numeric scale, and a reader may object that these are fundamentally different kinds of legal systems that cannot be reduced to one number without losing something important. This objection has real force and is not fully answerable. What this Index compares is not legal systems in the abstract, and it does not claim that a civil-law statute and an administrative circular are the same kind of instrument. It compares a single, narrower variable: the intensity of AI regulatory exposure an organization can expect to experience by deploying a product into a given market. That variable is measurable across different legal traditions because compliance obligations, enforcement actions, and financial penalties all translate into comparable cost and risk for the deployer facing them, whatever legal tradition produced them. It is a comparison of outcomes a deployer will feel, not of legal doctrine, and the report does not claim to offer more than that.

8.6 A snapshot that will already be somewhat out of date.

Every score in this Index reflects the regulatory landscape as of August 2, 2026, in a field where a single court ruling, a statute's effective date, or a bill's enactment can move a jurisdiction a full tier within months. Colorado's statute remains under litigation stay, South Korea's penalty provisions activate in January 2027, and Brazil's comprehensive bill remains pending, three examples discussed later in this report. There is no better time to publish a study of a field moving this quickly than yesterday; the second-best time is today. This Index does not claim to be permanent. It claims to be accurate as of its publication date, and it will be updated on a regular basis to keep pace with the jurisdictions it tracks.

8.7 Enforcement Capacity often reflects general-law enforcement, not AI-specific enforcement.

In several jurisdictions, the enforcement record behind a high Enforcement Capacity score comes from data-protection regulators, sectoral regulators, or criminal prosecutors applying general law to AI-related conduct, rather than from enforcement under a dedicated AI statute. A reader could argue that this measures the strength of a jurisdiction's general regulatory and law-enforcement apparatus more than its AI-specific regulatory maturity, and gives a boost to jurisdictions with assertive privacy or criminal regulators regardless of how developed their AI-specific framework actually is. This is a deliberate choice rather than an oversight: an organization sanctioned under a general data-protection statute for AI-enabled conduct experiences that sanction as AI regulatory risk, regardless of which statute technically authorized it, and a score that counted only AI-specific enforcement would understate the environment most deployers actually face. The two are not fully conflated: enforcement under an AI-specific statute is weighted more heavily than enforcement under general law applied to AI. But the underlying judgment, that general-law enforcement should count at all toward an "AI regulatory intensity" score, is a real one, and a reader who disagrees with it would score several jurisdictions in this Index differently.

8.8 This Index has not been peer-reviewed.

Unlike comparable work published by academic institutions or organizations such as the IAPP or the OECD, this First Edition has not been through an external peer-review process, which compounds the single-author limitation described above. This is an accurate description of where the study stands today. This First Edition is published as a working index, intended to establish the framework and invite scrutiny of its methodology and results. Future editions are intended to incorporate external peer review before publication.

09

CHAPTER 09

What to Look Out For Next

Seven developments most likely to move rows, or whole tiers, before the next edition.

AT A GLANCE

This chapter flags the seven specific developments — from the Colorado stay to China's next regulatory update — most likely to move a jurisdiction a full tier before the next edition of this Index.

Seven developments are the most likely to move rows, or whole tiers, before the next edition of this Index.

The Colorado stay. X.AI LLC v. Weiser is the first federal constitutional challenge to reach this stage against a comprehensive, deployer-facing state AI statute. As SB 26-189 approaches its January 2027 effective date, Colorado's maturity and injunctive-power scores will rise, closing the largest legal-density-to-enforcement gap in the Index. If the challenge succeeds on the merits instead, the precedent constrains comparable provisions in Texas, California, and Connecticut.

South Korea's fine activation. The Framework Act's administrative fines take effect in January 2027 — the first time a comprehensive AI statute outside the EU moves from in-force-but-unpenalised to fully operative. Korea's first enforcement actions after that date will show whether comprehensive legislation converts into real enforcement without the EU's institutional scaffolding, or whether it depends on features — the AI Office, cross-border cooperation — that Korea doesn't have.

The EU's high-risk milestone. As the AI Act's high-risk obligations reach full application, member-state maturity scores are already capped at the EU baseline. Whether enforcement begins promptly through the national authorities — AESIA, ACN, BNetzA, the CNIL — will determine whether those enforcement-capacity scores hold up or turn out to have been early.

Brazil's bill. PL 2338 is the single largest potential mover in the Index. Enactment would move credit currently confined to the maturity sub-criterion into the full set of obligation and penalty cells, plausibly pushing Brazil from mid-Tier 2 toward the Tier 1 boundary in one step.

Japan's data-protection amendment. A pending amendment would give Japan its first administrative-fine mechanism for AI-adjacent conduct. It's the development most likely to close the gap between Japan's already well-built institutional capacity and its current penalty severity score, which sits near zero.

US federal preemption. The federal government is currently enforcing AI-adjacent conduct through general law while litigating against state statutes — an unstable position. A court ruling establishing broad preemption would pull the US spread down across every state row; a federal AI statute would pull it up. Either would reshape roughly a quarter of the Index's 22 rows at once.

China's Generative AI Measures update. The CAC is drafting an update to the 2023 Generative AI Measures, expected to open for public comment in August 2026, likely adding model capability disclosure, red-teaming documentation, and stricter content-moderation thresholds above a user-volume ceiling. Coming alongside the Interim Measures on anthropomorphic AI services and the new intelligent-agents framework, both effective July 15, 2026, it would extend a regulatory apparatus that is still densifying rather than settling, and is the item most likely to move China's already near-ceiling Legal Density score higher still.

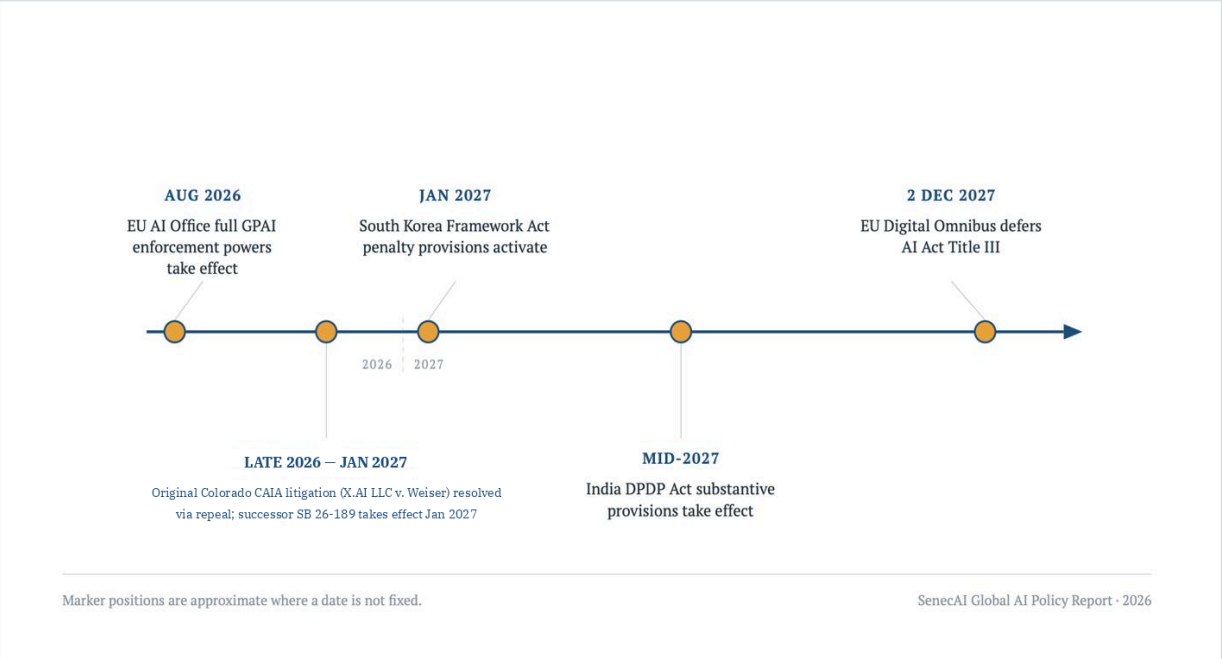

Figure 18. Forward-looking timeline.

KEY TAKEAWAY Seven separate developments could each move a jurisdiction a full tier before the next edition.

10

CHAPTER 10

What This Means in Practice for Deployers

Six practical rules for deciding where, and how, to deploy AI systems across these 22 markets.

AT A GLANCE

This chapter translates the findings of Chapters 6 and 7 into six practical rules for building a compliance program across these 22 markets.

The findings in Sections 6 and 7 point to a few practical rules for deciding where, and how, to deploy AI systems across these 22 markets.

1. Build to the strictest jurisdiction in each governance approach, then check for local add-ons. A program built to satisfy the EU AI Act's high-risk obligations, China's filing and labelling regime, and California's disclosure rules will clear most other requirements in this Index by default. The efficient approach is one baseline program plus a per-market check for add-ons — Italian criminal exposure, German works-council rules, Korean extraterritorial triggers, Texas biometric rules — since these are exactly what a global baseline won't catch.

2. Know who actually enforces before worrying about what the statute says. In comprehensive-legislation jurisdictions, that's usually the AI-specific regulator. In sectoral-absorption jurisdictions — the UK, Canada, Japan, Singapore, South Africa, India, the Gulf — it's whichever existing authority already has jurisdiction, most often the data-protection regulator, so compliance should sit with the privacy function rather than a new AI team. In the sovereign/centralized model, it's often the body that also controls market access. A program built around the statute alone will miss who shows up.

3. Weigh enforcement capacity separately from legal density when setting budget. The gap analysis in Section 7 is the most useful operational tool in this report: Texas acts faster than Colorado despite a lighter statute; France acts faster than Germany despite a thinner national layer. A risk assessment based only on whether a statute exists will misallocate resources away from the markets most likely to actually enforce.

4. In Saudi Arabia and the UAE, the contract functions as the practical regulatory gateway. Binding obligation attaches at the point of government contracting or platform licensing, not general legislation. What matters legally is the counterparty contract review, not a compliance read of a statute. What matters as a sanction is losing that relationship, not a fine. A deployer that never bids for government work or seeks a domestic license carries much less exposure than the law on paper suggests.

5. Treat the United States as six separate markets, not one. Run California as its own Tier 1 program, comparable in weight to South Korea. Track Texas Attorney General enforcement the way you'd track EU regulator guidance. Put Colorado, New York, and Connecticut on a watchlist tied to specific dates — the Weiser stay's resolution, New York's January 2027 effective date, Connecticut's phased rollout — rather than treating them as active programs today, unless the organization builds frontier models, in which case New York moves to the front of the line now.

6. Assume extraterritorial reach applies by default. Most of the 22 jurisdictions bind foreign providers whose systems are used locally or marketed in from abroad, and South Korea's "indirect impact" standard goes further than the EU AI Act's own trigger. Operating from outside a market isn't a compliance strategy in Tier 1, and it doesn't work in most of Tier 2 either.

11

CHAPTER 11

Conclusion

AI regulation is entering a phase of permanent divergence rather than convergence.

AT A GLANCE

This chapter closes the Index by restating its central finding — that enforcement capacity, not statutory ambition, is what a deployer actually experiences — and sets out the Index's plans for future editions.

AI regulation is entering a phase of permanent divergence rather than convergence. While jurisdictions increasingly share common objectives, transparency, accountability, and risk management, they are pursuing them through fundamentally different legal architectures. For organizations deploying AI globally, compliance will therefore become less about understanding a single AI law and more about navigating multiple regulatory models simultaneously.

This Index set out to answer a narrower question than most AI-policy trackers: not how strict a jurisdiction's law reads, but how much regulatory weight an organization actually feels once it deploys there. The answer splits into three patterns rather than one. Comprehensive legislation, sectoral absorption, and the sovereign/centralized model reflect genuinely different relationships between states and the companies operating in their markets, and each has produced a Tier 1 result on its own terms — Italy, California, and China respectively. Statutory ambition alone doesn't predict this. Enforcement capacity, scored separately in Section 7, is what actually separates a jurisdiction's paper position from what a deployer experiences on the ground.

The fragmentation this report documents isn't a phase before everyone converges on one global standard. The practical task for an organization deploying AI across these 22 jurisdictions is to build compliance around the pattern each jurisdiction actually follows, not the statute it happens to have published.

This First Edition reflects the regulatory landscape as of August 2, 2026, in a field that Section 9 shows can shift by a full tier within months. It isn't a permanent ranking. It's accurate as of its publication date, and later editions of this Index will track the same 22 jurisdictions, plus any that join them, on a regular basis.

LOOKING AHEAD

This Index is intended to be updated regularly rather than published once. The objective is to release a new edition annually, as Issue No. 02 and beyond, so that the ranking keeps pace with a regulatory landscape that continues to move quickly. Future editions are expected to expand both the number of jurisdictions covered and the analytical depth applied to each one, while preserving the same five-pillar framework, the same fifteen sub-criteria, and the same scoring rules documented in this edition, so that year-over-year comparisons remain meaningful rather than an artifact of a changed methodology.

A

ANNEX A

Why This Study Differs from Others in the Market

Positioning this Index against existing governance indices and compliance trackers.

AT A GLANCE

This annex positions the Index against existing governance-and-readiness indices and descriptive compliance trackers, and explains the specific gap it fills between them.

No publicly available AI-governance study, as of this Index's publication, combines four features at once: a roughly twenty-jurisdiction scope that scores individual US states and EU member states as separate units rather than folding them into national totals; a weighted quantitative composite score rather than a narrative or categorical rating; legal density and enforcement capacity scored as two explicitly separated axes; and a framing built for deployers and their compliance and legal teams rather than for policymakers or civil society.

The existing market splits into two camps that don't overlap. The first is scored governance and readiness indices — the OECD.AI Index, Oxford Insights' Government AI Readiness Index, the Tortoise Global AI Index, the Global Index on Responsible AI (GIRAI), and the AGILE Index. These produce genuine composite scores across many countries, but they're built for policymakers and measure national AI capacity, readiness, or responsible-AI implementation, not the regulatory burden a deploying organization actually experiences. None scores individual US states or EU member states as separate units.

The second camp is descriptive compliance trackers — the IAPP's Global AI Law and Policy Tracker, White & Case's AI Watch, DLA Piper's AI Laws of the World, and Baker McKenzie's AI Governance Tracker. These serve exactly the deployer audience this Index targets, and some of them (White & Case, DLA Piper) do cover individual EU member states and US states. But by design, none produces a weighted score or ranking. The IAPP's own framing is that its tracker "does not try to simplify this complexity into a single narrative."

GIRAI comes closest to this Index's second axis: it distinguishes framework existence from evidence of implementation, and finds that implementation evidence exists in only a minority of cases where a framework is nominally active. But GIRAI is built for civil-society and human-rights audiences, not deployer compliance, and it scores countries, not sub-national or sub-EU units.

THIS INDEX'S POSITION IN THE GAP

It rests on two features: scoring California, Texas, Colorado, New York, and Connecticut, and Italy, France, Spain, and Germany, as separate rows rather than folding them into "the United States" and "the European Union"; and scoring legal density and enforcement capacity as two distinct pillars, so that the gap between them, set out in Section 7, is itself a stated finding rather than left implicit in a narrative writeup.

B

ANNEX B

List of Abbreviations

An alphabetical glossary of every acronym and statute short-name used in this report.

AT A GLANCE

This annex is an alphabetical glossary of every abbreviation, statute short-name, and institutional acronym used throughout this report.

An alphabetical glossary of every abbreviation, statute short-name, and institutional acronym used throughout this report.

ACN — Agenzia per la Cybersicurezza Nazionale — Italy's national cybersecurity agency, designated AI Act supervisor

ADGM — Abu Dhabi Global Market — a UAE financial free zone with its own AI-relevant regulatory regime

AEPD — Agencia Española de Protección de Datos — Spain's data protection authority

AESIA — Agencia Española de Supervisión de Inteligencia Artificial — Spain's AI supervision agency

AG — Attorney General

AIGC — AI-Generated Content — the term used in China's AIGC Labelling Measures

ANPD — Autoridade Nacional de Proteção de Dados — Brazil's national data protection authority

APPI — Act on the Protection of Personal Information — Japan

CAC — Cyberspace Administration of China

CAD — Canadian dollar

CAI — Commission d'accès à l'information — Quebec's privacy and access-to-information regulator

CAIA — Colorado AI Act (2024), repealed before taking effect and superseded by SB 26-189

CCPA — California Consumer Privacy Act

CNIL — Commission Nationale de l'Informatique et des Libertés — France's data protection authority

CPPA — California Privacy Protection Agency

CRA — Cyber Resilience Act (EU)

CUTPA — Connecticut Unfair Trade Practices Act

DFS — New York Department of Financial Services

DIFC — Dubai International Financial Centre — a UAE financial free zone; Regulation 10 is its AI-specific instrument

DIGIT — New York's planned AI-oversight office, housed within DFS; not yet operational as of publication

DPA — Data Protection Authority

DPDP — Digital Personal Data Protection Act — India

DSA — Digital Services Act (EU)

EU — European Union

FCA — Financial Conduct Authority — United Kingdom

FTC — Federal Trade Commission — United States

GDPR — General Data Protection Regulation (EU)

HUMAIN — A Saudi state-directed AI investment vehicle; a proper name, not an acronym

IAPP — International Association of Privacy Professionals

ICO — Information Commissioner's Office — United Kingdom

IMDA — Infocomm Media Development Authority — Singapore

ISO — International Organization for Standardization

LGPD — Lei Geral de Proteção de Dados — Brazil's general data protection law

LLC — Limited Liability Company

METI — Ministry of Economy, Trade and Industry — Japan

MIC — Ministry of Internal Affairs and Communications — Japan

MIIT — Ministry of Industry and Information Technology — China

MPS — Ministry of Public Security — China

MSIT — Ministry of Science and ICT — South Korea

NDMO — National Data Management Office — Saudi Arabia

NDRC — National Development and Reform Commission — China

SAMR — State Administration for Market Regulation — China

OECD — Organisation for Economic Co-operation and Development

OPC — Office of the Privacy Commissioner of Canada

OSFI — Office of the Superintendent of Financial Institutions — Canada

PDPA — Personal Data Protection Act — Singapore

PDPC — Personal Data Protection Commission — Singapore

PDPL — Personal Data Protection Law — Saudi Arabia and the UAE

PIPC — Personal Information Protection Commission — South Korea

PL — Projeto de Lei — a Brazilian legislative bill, e.g. PL 2338

POPIA — Protection of Personal Information Act — South Africa

PPC — Personal Information Protection Commission — Japan

RAISE Act — Responsible AI Safety and Education Act — New York

RMB — Renminbi — the Chinese yuan

SB / AB — Senate Bill / Assembly Bill — US state legislative designations

SDAIA — Saudi Data and Artificial Intelligence Authority

SIA — The AI governance framework proposed under Brazil's PL 2338

SREN — Loi visant à sécuriser et réguler l'espace numérique — France's law securing and regulating digital space

TFAIA — Transparency in Frontier Artificial Intelligence Act — California (SB 53)

TRAIGA — Texas Responsible AI Governance Act

TSE — Tribunal Superior Eleitoral — Brazil's Superior Electoral Court

UAE — United Arab Emirates

UK — United Kingdom

US — United States

C

ANNEX C

Source Appendix

The primary legal instrument and enforcing authority behind every jurisdiction row.

AT A GLANCE

This annex lists the primary legal instrument and enforcing authority behind every jurisdiction row in this Index, plus the secondary cross-jurisdictional sources consulted.

Every jurisdiction row in this Index traces to the primary legal instrument and enforcing authority (or authorities) listed below, drawn from the jurisdiction-specific research briefs underlying this report.

JURISDICTION	PRIMARY INSTRUMENT(S)	AUTHORITY
Italy	AI Act (Reg. (EU) 2024/1689); Law No. 132/2025	ACN, AgID, Garante
France	AI Act; Loi Toubon	CNIL
Spain	AI Act; Royal Decree 729/2023	AESIA, AEPD
Germany	AI Act	BNetzA, BfDI, state DPAs
China	Cybersecurity Law; Data Security Law; PIPL; Algorithm Recommendation, Deep Synthesis, and Generative AI Provisions	CAC
EU (baseline)	AI Act; GDPR; DSA; CRA	AI Office (European Commission); national market-surveillance authorities
South Korea	AI Basic Act (Act No. 20676)	MSIT, PIPC
California	SB 53 (TFAIA); AB 2013; SB 942; CCPA	California Attorney General, CPPA
Canada	PIPEDA (federal); Quebec Law 25	OPC (federal); Quebec's CAI
Brazil	LGPD; PL 2338 (pending)	ANPD; TSE (electoral deepfake rules)
Saudi Arabia	PDPL; SDAIA AI Adoption Framework	SDAIA, NDMO

JURISDICTION	PRIMARY INSTRUMENT(S)	AUTHORITY
United Kingdom	UK GDPR; Online Safety Act 2023	ICO, FCA, Ofcom
South Africa	POPIA	Information Regulator
Japan	APPI; AI Guidelines for Business v1.1	PPC, METI, MIC
Texas	TRAIGA (HB 149)	Texas Attorney General
Singapore	PDPA; Model AI Governance Framework	PDPC, IMDA
UAE	PDPL; UAE Charter for the Development and Use of AI	Federal Authority for Artificial Intelligence and Data
US Federal	FTC Act Section 5	FTC; sectoral regulators
Colorado	SB 26-189 (successor to CAIA)	Colorado Attorney General
India	DPDP Act	MeitY; Data Protection Board (not yet operational)
New York	RAISE Act	New York Attorney General; DIGIT (DFS, not yet operational)
Connecticut	SB 5 (CUTPA-linked provisions)	Connecticut Attorney General

Secondary sources. Research for the underlying jurisdiction briefs also drew on cross-jurisdictional trackers and analysis, including the IAPP Global AI Law and Policy Tracker, the OECD.AI Policy Observatory, the Library of Congress Global Legal Monitor, the Council of Europe Treaty Office's CETS 225 signatory list, and Tier-1 law-firm AI-regulatory alerts, among them White & Case, DLA Piper, Latham & Watkins, Hogan Lovells, Bird & Bird, Baker McKenzie, Norton Rose Fulbright, Clifford Chance, and Linklaters.

D

ANNEX D

Country Scorecards

All 22 jurisdictions, ranked by composite score, with their full five-pillar breakdown.

AT A GLANCE

This annex presents all 22 jurisdictions as individual scorecards, each showing the full five-pillar breakdown behind its composite score, ordered by composite rank.

Each jurisdiction's composite score is the weighted sum of five pillar scores: legal density (25 percent), operational burden (25 percent), enforcement capacity (20 percent), penalty severity (15 percent), and extraterritorial reach (15 percent). Scorecards below are ordered by composite rank, matching Section 4.

LD Legal Density • OB Operational Burden • EC Enforcement Capacity • PS Penalty Severity • ER Extraterritorial Reach

01 **Italy** — *Comprehensive legislation* — **91.2** TIER 1

LD	95	OB	100
EC	88	PS	92
ER	73		

02 **France** — *Comprehensive legislation* — **90.6** TIER 1

LD	92	OB	100
EC	93	PS	90
ER	70		

03 **Spain** — *Comprehensive legislation* — **90.2** TIER 1

LD	96	OB	100
EC	93	PS	77
ER	73		

04 **Germany** — *Comprehensive legislation* — **87.5** TIER 1

LD	88	OB	100
EC	88	PS	78
ER	73		

05

China

Sovereign/centralized model

87.1
TIER 1

LD	92	OB	85
EC	91	PS	98
ER	67		

06

EU (baseline)

Comprehensive legislation

81.9
TIER 1

LD	92	OB	100
EC	65	PS	73
ER	67		

07

South Korea

Comprehensive legislation

77.6
TIER 1

LD	97	OB	83
EC	77	PS	53
ER	62		

08

California

Sectoral absorption

72.5
TIER 1

LD	92	OB	83
EC	65	PS	60
ER	45		

09

Canada

Sectoral absorption

66.8
TIER 2

LD	78	OB	72
EC	57	PS	53
ER	67		

10

Brazil

Comprehensive legislation

64.2
TIER 2

LD	75	OB	63
EC	77	PS	45
ER	50		

11 **Saudi Arabia** — 56.1 TIER 2

Sovereign/centralized model

LD	58	OB	50
EC	60	PS	47
ER	67		

12 **United Kingdom** — 55.0 TIER 2

Sectoral absorption

LD	68	OB	35
EC	68	PS	65
ER	38		

13 **South Africa** — 54.3 TIER 2

Sectoral absorption

LD	67	OB	50
EC	48	PS	53
ER	50		

14 **Japan** — 53.6 TIER 2

Sectoral absorption

LD	77	OB	33
EC	82	PS	27
ER	38		

15 **Texas** — 50.4 TIER 2

Sectoral absorption

LD	78	OB	30
EC	52	PS	45
ER	42		

16 **Singapore** — 50.0 TIER 2

Sectoral absorption

LD	67	OB	37
EC	78	PS	33
ER	23		

17 **UAE**

Sovereign/centralized model

46.8
TIER 3

LD	52		OB	35	
EC	37		PS	52	
ER	67				

18 **US Federal**

Sectoral absorption

46.3
TIER 3

LD	65		OB	27	
EC	53		PS	53	
ER	32				

19 **Colorado**

Sectoral absorption

39.6
TIER 3

LD	60		OB	53	
EC	25		PS	18	
ER	23				

20 **India**

Sectoral absorption

39.6
TIER 3

LD	48		OB	32	
EC	37		PS	43	
ER	38				

21 **New York**

Sectoral absorption

34.2
TIER 3

LD	40		OB	42	
EC	20		PS	33	
ER	32				

22 **Connecticut**

Sectoral absorption

31.4
TIER 3

LD	53		OB	28	
EC	25		PS	12	
ER	28				

END MATTER

About SenecAI Consulting

SenecAI is an on-demand EU AI Act compliance partner, helping EU startups and SMEs understand and implement the regulatory requirements introduced by the AI Act. Built exclusively around the AI Act and its converging EU law, GDPR, MDR, and NIS2, SenecAI translates regulatory obligations into operational checklists, playbooks, and workflows rather than abstract legal theory.

Mission

To help EU startups and SMEs comply with the EU AI Act and turn compliance into a competitive advantage, so founders can focus on building, innovating, and scaling without regulatory headaches or audit anxiety.

Services

SenecAI delivers end-to-end AI Act compliance across the full lifecycle: assessment, implementation, and post-deployment monitoring.

- **Assessment** — inventory and risk assessment of AI systems; AI Act readiness checks
- **Implementation** — practical compliance roadmaps; documentation preparation; assistance with conformity assessment
- **Ongoing support** — liaising with the EU AI Office and national authorities; policy monitoring; trainings and workshops

Connect

WEBSITE	senecai.eu
LINKEDIN	linkedin.com/company/senecai-consulting
EMAIL	matei.stefan@senecai.eu
PHONE	+40 770 875 267

ABOUT THE AUTHOR

Matei Ștefan

Founder & AI Act Compliance Consultant, SenecAI Consulting

Biography

Matei Ștefan is a legal professional operating at the intersection of AI, regulation, and innovation. He worked within the European Parliament. His background also includes work with the United Nations, as well as experience in corporate legal tech environments, where he supported compliance, risk management, and operational implementation, enabling businesses to innovate while remaining within regulatory boundaries.

As the Founder of SenecAI, he helps EU startups and SMEs adapt to life under the EU AI Act, providing structured compliance consulting, risk classification analysis, internal governance design, audit-ready documentation, and strategic regulatory guidance.

Areas of Expertise

EU AI Act compliance consulting, risk classification analysis, internal AI governance design, audit-ready documentation, and strategic regulatory guidance.

Selected Experience

European Parliament; the United Nations; and corporate legal tech, supporting compliance, risk management, and operational implementation.

Contact

LINKEDIN	linkedin.com/in/mateiștefan/
WEBSITE	senecai.eu
EMAIL	matei.stefan@senecai.eu

END MATTER

How to Cite This Report

When referencing this Index in academic work, legal filings, internal memoranda, or press coverage, please use the following suggested citation.

> **SUGGESTED CITATION**
>
> Ștefan, M. (2026). *Global AI Policy Report: What Awaits Organizations Deploying AI Globally* (Issue No. 01, First Edition). SenecAI Consulting.

Publication Metadata

TITLE	Global AI Policy Report: What Awaits Organizations Deploying AI Globally
AUTHOR	Matei Ștefan, Founder, SenecAI Consulting
PUBLISHER	SenecAI Consulting
EDITION	First Edition, Issue No. 01
PUBLICATION DATE	August 2026
SCOPE	22 jurisdictions, five weighted pillars, fifteen sub-criteria

This Index is published as the first issue of an annual series. Later editions should be cited by their own issue number and publication year.