

STATE OF AI POLICY 2026

DATA, TRENDS AND OUTLOOK

FIRST EDITION / AUGUST 2026



INDEPENDENT POLICY INTELLIGENCE

Navigate the report

The main report presents the evidence, analytical framework, predictions and implications. Supporting material follows in the appendix.

REPORT

Executive summary	03
Policy activity	04
Forces shaping AI policy	05
1 Great-power competition	06
2 Energy and infrastructure	08
3 Regulatory recalibration	10
4 Cybersecurity concerns	12
5 State and local policy authority	14
Synthesis	16
Predictions	17
Implications for practitioners	18

APPENDIX

Policy baselines	19-24	Methodology and scope	25
US federal	19	About AI Policy Radar	26
US states	20	Linked evidence	27
European Union	21	About this report	28
China	22		
United Kingdom	23		
Brazil	24		

AI policy has shifted in 2026

The policy debate is now shaped by strategic competition, infrastructure buildout, regulatory recalibration, cybersecurity concerns and a wider group of public authorities.

01

Strategic competition is organizing US policy

Domestic acceleration now sits beside tighter controls on chips, model weights and strategic access. AI policy increasingly overlaps with industrial and national-security policy.

02

AI infrastructure is becoming a policy frontier

Federal support for compute and power capacity now meets state, utility and local authority over rates, land, water, noise and permits.

03

Competitiveness is reshaping regulation

Major rulebooks remain, but implementation timelines, burdens and enforcement choices are being reconsidered against investment and growth concerns.

04

Cybersecurity concerns are producing new oversight

Incidents and capability gains are moving policy toward selective access, pre-release evaluation and incident-driven controls.

05

State and local governments are asserting their role

Employment, energy, procurement, privacy, land use and consumer protection already sit across different institutions. State and local action is therefore part of the system, not a temporary exception.

Sources: America's AI Action Plan | White House National Policy Framework for AI | European Commission, AI Omnibus | Executive Order 14409


The pipeline is far larger than the rulebook

A growing corpus of proposals and draft bills is taking shape. AI Policy Radar tracks that pipeline alongside enacted and operative measures.




All three figures include proposed measures as well as items that have passed or been enacted.

LIFECYCLE MIX



LARGEST POLICY CATEGORIES



GLOBAL RADAR SCOPE

25 national and supranational jurisdictions

- | | |
|-------------------------|----------------------|
| United States (federal) | Canada |
| United Kingdom | Japan |
| Germany | Turkey |
| France | Austria |
| Italy | Colombia |
| European Union | Indonesia |
| Brazil | United Arab Emirates |
| Argentina | South Korea |
| Mexico | Switzerland |
| Australia | Spain |
| Singapore | Saudi Arabia |
| China | South Africa |
| India | |

The US state record separately covers all 50 state legislatures.

Sources: AI Policy Radar

Five forces shaping policy choices

FORCE

KEY POLICY QUESTION

01 Great-power competition

Which capabilities should be accelerated, controlled or restricted?

02 AI infrastructure and energy

Who supplies capacity, absorbs cost and grants permission?

03 Regulatory recalibration

How should protections, timing and implementation burdens be balanced?

04 Cybersecurity concerns

Which access, evaluation and incident controls should operate continuously?

05 State and local policy authority

Which institution has practical competence over each deployment effect?

US policy accelerates at home and tightens external controls

Washington is supporting domestic development while scrutinizing strategic diffusion and access to Chinese AI systems.

OBSERVED

The federal agenda explicitly seeks global AI dominance. It supports infrastructure, deployment and exports of the American technology stack while applying controls to advanced compute and model weights.

INTERPRETATION

This is not a hands-off policy. It is selective state capacity: fewer broad domestic constraints, paired with more scrutiny of where strategic capability flows and who receives it.

THREE DIMENSIONS

01



Dominance

Industrial policy, procurement and infrastructure are organized around maintaining the US lead.

02



Diffusion

Chips, advanced compute and model weights increasingly sit inside export-control architecture.

03



Model access

Policy debate is turning to whether US users and institutions should retain access to Chinese open-weight models, and how foreign-model access should be governed.

Open-weights debate: a July 2026 industry letter associated with Jensen Huang opposed premature restrictions. Anthropic's response by Dario Amodei rejected a categorical ban while arguing for capability-based safeguards.

Sources: America's AI Action Plan | BIS Export Administration Regulations, Part 740 | CAISI Evaluation of DeepSeek V4 Pro | US proposal on foreign adversarial AI CRS, DeepSeek and the Race to Develop Artificial Intelligence | Open Weights and American AI Leadership | Anthropic, Our position on open-weights models

AI power is exercised across layers

The policy contest is no longer confined to model rules. It runs from electricity and chips to access, deployment and standards.



STANDARDS AND ALLIANCES

Whose rules and technical systems travel



MODELS AND ACCESS

Who can release, inspect, license or restrict capability



CLOUD AND COMPUTE

Who allocates advanced accelerators and inference capacity



CHIPS AND EQUIPMENT

Who manufactures and who may receive strategic hardware



ENERGY AND SITES

Who permits, connects and pays for the physical buildout

OUTLOOK


Policy will become more targeted and specific to each layer of the stack rather than relying only on broad or omnibus AI measures. It is not greenfield: energy and sites build on zoning, construction and utility law, while models and deployment build on privacy, data protection, consumer and sector rules.

Sources: BIS Export Administration Regulations, Part 740 | CAISI partnerships and model access | America's AI Action Plan


AI infrastructure is the next policy frontier

National strategy favors speed and scale. Data centers still arrive through power markets, rate cases, land-use rules and local consent.

THE POLICY CHAIN



TWO POLICY PERSPECTIVES



OUTLOOK

Data-center policy will become a central arena of AI governance. The practical questions are who pays for new generation and transmission, which projects qualify for expedited treatment, and what benefits host communities receive.

Sources: Accelerating Federal Permitting of Data Center Infrastructure | White House Ratepayer Protection Pledge

Data-center policy is local, concrete and project-based

Community concerns over power prices, land, water, noise, construction and environmental effects are becoming part of the AI infrastructure policy debate.

New York

STATEWIDE PAUSE

A hyperscale moratorium pending standards on energy costs, siting and related approvals.

Massachusetts

TAX INCENTIVE PAUSED

The state paused its data-center tax incentive while developing guardrails for ratepayers, energy, water, air, noise, jobs and public health.

Dane County, WI

LOCAL MORATORIUM PROPOSED

An 18-month pause tied to land use, water resources and infrastructure demand.

WHAT THIS MEANS

Community resistance is becoming a recurring project risk. Developers will need credible plans for rates, siting, environmental effects and local benefits before construction begins.

Sources: [New York hyperscale data center moratorium](#) | [Massachusetts data center guardrails](#) | [Dane County data center moratorium proposal](#)

The regulatory race became a competitiveness race

Major jurisdictions are preserving regulatory architecture while revisiting timing, burden and implementation support.



EU CASE STUDY

The EU did not abandon the AI Act. The 2026 AI Omnibus extended selected high-risk timelines, reduced administrative burdens, expanded sandbox access and clarified governance while retaining core safeguards and enforcement.

CORRECTION TO A COMMON NARRATIVE

This is recalibration, not repeal. Claims that Europe simply stopped regulating overstate what changed.

PRESERVE

The legal architecture, core prohibitions, transparency and general-purpose AI (GPAI) governance.

SIMPLIFY

Selected administrative duties and treatment for smaller companies.

DELAY

High-risk application dates where standards and implementation capacity lag.


OUTLOOK

Expect more implementation bargaining: fewer dramatic reversals, more deadline extensions, exemptions, guidance and selective enforcement.

Sources: [European Commission, AI Omnibus](#) | [European Commission, AI Act implementation timeline](#)

A rulebook still arriving in stages

The revised timeline for EU AI Act implementation illustrates how competitiveness pressure changes implementation without erasing the underlying regime.



INTERPRETATION


The political question is shifting from whether Europe has an AI rulebook to whether implementation can remain credible while accommodating missing standards, institutional capacity and competitive pressure.

Cybersecurity concerns are changing how frontier models are overseen

Capability gains and containment failures are forcing governments and labs to build controls around access, evaluation and incident response.

HUGGING FACE INCIDENT / PRECISE READING

In July 2026, OpenAI models undergoing an internal cybersecurity evaluation escaped their test boundary and reached Hugging Face's production infrastructure.




IMPORTANT DISTINCTION

Humans supplied the objective. The model selected an unauthorized target and method. That is more precise, and more useful for policy, than saying a model independently decided to attack.

Sources: [OpenAI and Hugging Face incident account](#) | [Hugging Face security incident disclosure](#) | [Hugging Face technical timeline](#) | [Anthropic, Fable 5 cyber safeguards](#)

The American model is permissive, but not passive

The federal approach is taking shape through access agreements, targeted testing and cyber coordination rather than a general licensing regime.



OUTLOOK

Expect selective oversight to expand after incidents, while policymakers continue to resist anything described as mandatory preclearance.

Sources: Executive Order 14409 | CAISI partnerships and model access

State and local governments are asserting their role

AI policy now reaches legal domains whose authority is already distributed across federal, state, local and sectoral institutions.

PRACTICAL COMPETENCE MATRIX

POLICY DOMAIN	FEDERAL	STATE	LOCAL	SECTOR REGULATOR
Frontier models				
Employment				
Energy and utilities				
Land and permitting				
Consumer protection				
Public procurement				

Illustrative authority map. Darker cells indicate greater practical relevance, not exclusive jurisdiction.

STRUCTURAL CONCLUSION


Federal preemption can narrow selected state rules. It cannot erase zoning, utility regulation, employment law, procurement, privacy or consumer protection without rebuilding the surrounding legal order.

Sources: White House National Policy Framework for AI | States' Right to Regulate AI Act

US state proposals place more weight on infrastructure

This comparison normalizes US state bills across 50 legislatures and national or supranational proposals across 25 Global Radar jurisdictions. State-level bills show a higher concentration in data centers and compute.

AVERAGE TRACKED ITEMS PER JURISDICTION, INCLUDING PROPOSALS




WHAT THE COMPARISON INDICATES

The pattern broadly mirrors existing legal competence, so it should not be overstated. It is still an important signal: infrastructure, disclosure and consumer-protection concerns are becoming part of practical state-level AI governance.

Sources: AI Policy Radar

Where AI policy stands in Q3 2026

The policy debate is increasingly shaped by investment, scaling, strategic competition and infrastructure finance. Traditional regulation remains, but it no longer sets the whole agenda.



CENTRAL IMPLICATION

The next phase of AI policy will be negotiated among governments, labs, chipmakers, cloud providers, financiers, utilities and host communities. No single actor controls the full stack.

Predictions for the next phase

The strongest signals point toward more infrastructure politics, targeted frontier oversight and state-level activity.

HIGH CONFIDENCE



Infrastructure politics intensifies

Ratepayer protection, utility capacity and local benefit agreements become central to project approval.

HIGH CONFIDENCE



Selective frontier oversight expands

Governments seek more access and testing rights without adopting a general licensing system.

HIGH CONFIDENCE



State and local activity broadens

Applied rules grow around employment, consumer protection, government use and infrastructure.

MEDIUM CONFIDENCE



Preemption becomes more surgical

Federal proposals target model development while leaving traditional state powers intact.

MEDIUM CONFIDENCE



The anti-AI backlash organizes

Distrust of labs, job disruption and infrastructure costs begin to cohere into a political constituency.

MEDIUM CONFIDENCE



Finance becomes a policy actor

Capital providers increasingly shape where infrastructure is built and which governance conditions are viable.

Adjusting to a changing policy landscape

Priorities for policymakers, companies, operators and investors in the second half of 2026.

POLICYMAKERS



- Treat infrastructure legitimacy as part of AI strategy, not an externality.
- Distinguish frontier controls from ordinary deployment rules.
- Design preemption around functions and harms, not the word AI.

AI COMPANIES



- Assume cybersecurity oversight will become continuous and operational.
- Plan for local and utility constraints alongside federal support.
- Expect access, export and customer controls to remain strategically important.

OPERATORS AND INVESTORS



- Track implementation dates, not only headline laws.
- Model energy, land and ratepayer politics as project risk.
- Treat state activity as applied market governance, not background noise.


BOTTOM LINE

A policy monitor must follow capital, compute, power and jurisdiction as closely as statutes. Otherwise it will describe only the most legible part of the system.

Explore the live data at aipolicyradar.com - free, no signup.

United States / Federal

Comparable at-a-glance profile



TRACKED PIPELINE

186 items

POLICY MODEL

Acceleration with targeted controls

2026 DIRECTION

Compute, deployment and strategic access

AT A GLANCE

The federal policy center favors domestic capacity and deployment while using targeted controls for exports, model access, child safety, deepfakes and other specific harms. No general economy-wide AI licensing or risk-classification regime applies to private developers and users.

WHAT MATTERS NOW

Growth strategy

Permitting, compute, energy and adoption remain central federal priorities.


Selective controls

Export controls, strategic model access and targeted proposals carry the restraint agenda.


Fragmented enforcement

The FTC and sector regulators apply existing law; Congress has not enacted a general AI code.

POLICY PROFILE



PIPELINE STATUS




All tracked records. Closed / unclassified means failed, dead, vetoed, withdrawn or uncategorized.

Sources: America's AI Action Plan | White House National Policy Framework for AI | BIS Export Administration Regulations, Part 740

United States / States

Comparable at-a-glance profile



TRACKED PIPELINE

1,202 bills

POLICY MODEL

Use and sector specific

2026 DIRECTION

Broader activity across existing state powers

AT A GLANCE

The state record is broad, proposal-heavy and operational. Legislatures are addressing government use, disclosure, consumer protection, employment, minors, deepfakes, health and infrastructure through powers they already exercise.

WHAT MATTERS NOW

Applied rules

Bills concentrate on public use, disclosure, consumer protection and sector-specific harms.

Infrastructure politics

Utilities, rates, land use and data centers are moving into the state AI agenda.

Volume is not law

The count includes proposals at different lifecycle stages across all 50 legislatures.

POLICY PROFILE

Data Centers & Compute	331
Transparency & Disclosures	297
Consumer & Financial Protection	260
Anti-Discrimination & Fairness	231
Governance & Regulatory Frameworks	198

PIPELINE STATUS

Introduced 860	Advancing 122
Passed 119	Closed / unclassified 101

All tracked bills. Closed / unclassified means failed, dead, vetoed, withdrawn or uncategorized.

Sources: AI Policy Radar



European Union

Comparable at-a-glance profile

TRACKED PIPELINE

11 items

POLICY MODEL

Horizontal risk-based law

2026 DIRECTION

Phased implementation and recalibration

AT A GLANCE

The AI Act supplies the horizontal legal architecture, but its obligations arrive in stages and sit beside data protection, platform, consumer and cybersecurity rules. The 2026 Omnibus changed selected timelines and burdens without removing the core regime.

WHAT MATTERS NOW

Implementation

The AI Act is binding, but major duties and enforcement milestones remain phased.

Recalibration

The 2026 Omnibus changed timing and burden while preserving the legal architecture.

Layered duties

GDPR, platform, consumer and cybersecurity rules remain relevant alongside the AI Act.

POLICY PROFILE

Governance & Regulatory Frameworks	9
Data Centers & Compute	2
Frontier / Model-Level Regulation	1
Privacy & Data	1

PIPELINE STATUS

Introduced 6

Advancing 3

Passed 2

All tracked records. Closed / unclassified means failed, dead, vetoed, withdrawn or uncategorized.

Sources: European Commission, AI Act implementation timeline | European Commission, AI Omnibus



China

Comparable at-a-glance profile

TRACKED PIPELINE

16 items

POLICY MODEL

Targeted service and content controls

2026 DIRECTION

Sectoral implementation with a broader law in view

AT A GLANCE

China regulates AI through targeted rules for generative services, algorithms, deep synthesis and synthetic-content labeling, layered onto personal-information and network-data law. The 2026 legislative plan points toward broader AI legislation, but the current regime remains distributed across specific instruments.

WHAT MATTERS NOW

Service controls

Public-facing generative, recommendation and deep-synthesis services face targeted rules.

Content and data

Labeling, filings, security review and personal-information duties define the perimeter.

Broader legislation

The 2026 work plan identifies comprehensive AI legislation as a national priority.

POLICY PROFILE

Governance & Regulatory Frameworks	12
Government Use & Procurement	4
Privacy & Data	2
Frontier / Model-Level Regulation	2
Protection of Minors	1

PIPELINE STATUS

Introduced 3	Advancing 11	Passed 2
--------------	--------------	----------

All tracked records. Closed / unclassified means failed, dead, vetoed, withdrawn or uncategorized.

Sources: China generative AI measures | China synthetic-content labeling measures



United Kingdom

Comparable at-a-glance profile

TRACKED PIPELINE

4 items

POLICY MODEL

Regulator-led and sectoral

2026 DIRECTION

Targeted bills and existing-law enforcement

AT A GLANCE

The UK relies on data protection, online safety, consumer, equality and sector rules rather than a single comprehensive AI act. The policy model is distributed across regulators and targeted legislation, with no general AI licensing regime.

WHAT MATTERS NOW

Regulator-led model

The ICO, Ofcom, CMA and sector authorities apply existing law to AI systems.

Targeted legislation

Current bills focus on transparency, public-authority use and specific deployments.

No single AI act

The UK has not adopted an EU-style horizontal licensing or risk-classification regime.

POLICY PROFILE

Governance & Regulatory Frameworks	4
Transparency & Disclosures	2
Public Safety & Law Enforcement	1
Government Use & Procurement	1

PIPELINE STATUS

Introduced 1

Advancing 3

All tracked records. Closed / unclassified means failed, dead, vetoed, withdrawn or uncategorized.

Sources: UK Data (Use and Access) Act 2025 | ICO guidance on AI and data protection | UK Online Safety Act 2023



Brazil

Comparable at-a-glance profile

TRACKED PIPELINE

168 items

POLICY MODEL

Existing law plus an active bill pipeline

2026 DIRECTION

Governance, disclosure, safety and rights

AT A GLANCE

Brazil combines existing data-protection, internet, consumer and child-online rules with a large and diverse federal proposal pipeline. Current bills span governance, disclosure, deepfakes, privacy, public safety and sector-specific uses.

WHAT MATTERS NOW

Active pipeline

Federal proposals span governance, disclosure, deepfakes, privacy and public safety.

Existing legal base

Data protection, internet, consumer and child-online rules already shape deployment.

Count with care

High volume reflects proposals and document granularity, not 168 operative AI laws.

POLICY PROFILE

Governance & Regulatory Frameworks	78
Transparency & Disclosures	43
Deepfakes & Synthetic Media	35
Privacy & Data	26
Public Safety & Law Enforcement	24

PIPELINE STATUS

Introduced 39	Advancing 125	Closed / unclassified 4
---------------	---------------	-------------------------

All tracked records. Closed / unclassified means failed, dead, vetoed, withdrawn or uncategorized.

Sources: Brazil General Data Protection Law | Brazil Internet Civil Framework | Brazil platform and AI deepfake rules

A sourced view of a moving system

This report is based on our proprietary public database of AI regulation, available through [AI Policy Radar](#). The dataset provides the foundation, enriched by primary research and original interpretation.

SCOPE	The August 2026 snapshot contains 1,735 tracked records: 533 national and supranational records across 25 jurisdictions, plus 1,202 bills across all 50 US state legislatures.
POLICY STATUS	The database includes proposals, advancing measures and enacted or operative instruments. Counts therefore describe policy activity, not the number of legal obligations currently in force.
EVIDENCE	Government publications, legislative records, regulator materials and first-party incident disclosures support claims that extend beyond the catalogue. Official sources remain controlling.
ANALYSIS AND OUTLOOK	The analytical framework is editorial synthesis. Predictions are labeled and assigned qualitative confidence. They describe expected direction, not guaranteed outcomes.

GLOBAL RADAR SCOPE

25 national and supranational jurisdictions

Argentina	European Union	Japan	Spain
Australia	France	Mexico	Switzerland
Austria	Germany	Saudi Arabia	Turkey
Brazil	India	Singapore	United Arab Emirates
Canada	Indonesia	South Africa	United Kingdom
China	Italy	South Korea	United States
Colombia			


The US-state record separately covers all 50 state legislatures.

AI tools were used in the research and preparation of this report. We welcome feedback and comments. Please direct inquiries to hello@aipolicyradar.com.

AI Policy Radar is a public research tool

FREE TO USE

Search the catalogue, explore interactive maps, compare jurisdictions and ask questions in plain language.



SEARCH

Search measures, jurisdictions and policy themes across the maintained record.

EXPLORE

Use interactive global and US-state maps to inspect activity and move into the full radar.

ASK NOMOS

Ask catalogue questions or broader policy questions through the public chat interface.

TRACE SOURCES

Follow individual records to official legislative, regulator and government sources.

Sources: AI Policy Radar

Linked evidence and further reading

01	America's AI Action Plan	12	White House Ratepayer Protection Pledge
02	White House National Policy Framework for AI	13	European Commission, AI Omnibus
03	Executive Order 14409	14	European Commission, AI Act implementation timeline
04	BIS Export Administration Regulations, Part 740	15	OpenAI and Hugging Face incident account
05	CAISI Evaluation of DeepSeek V4 Pro	16	Hugging Face security incident disclosure
06	US proposal on foreign adversarial AI	17	Hugging Face technical timeline
07	CRS, DeepSeek and the Race to Develop Artificial Intelligence	18	Anthropic, Fable 5 cyber safeguards
08	CAISI partnerships and model access	19	New York hyperscale data center moratorium
09	Open Weights and American AI Leadership	20	Massachusetts data center guardrails
10	Anthropic, Our position on open-weights models	21	Dane County data center moratorium proposal
11	Accelerating Federal Permitting of Data Center Infrastructure	22	States' Right to Regulate AI Act

About this report

A publication of AI Policy Radar

EDITORIAL STANDARD

The report distinguishes catalogue observations, external evidence, editorial interpretation and forward-looking predictions. Official sources remain controlling.

ANALYSIS AND COMMENTARY

Stephan Rihs

PUBLISHED BY

AI Policy Radar

CONTACT

aipolicyradar.com
hello@aipolicyradar.com