

BLOG 05 — THE EIGHT TIERS, PART 3

US State AI Laws: Where the Obligations Actually Are

Tier 3 of eight: over two thousand bills, a dozen that matter, and a jurisdiction question most trackers get wrong

By Robin Beetge

Co-founder, laibrary

Part 3 of nine

Tier 3 · ungated



Fifty jurisdictions, no federal statute, and reach that follows your users.

Short answer: there is no comprehensive US federal AI statute. The enforceable obligations sit at state level. As of September 2026 the material laws are **Texas TRAIGA**, **California SB 53**, **AB 2013**, the **AI Transparency Act** and the **CPPA ADMT regulations**, **Colorado SB 26-189**, **Illinois HB 3773**, **New York's RAISE Act**, and **NYC Local Law 144**.

More than two thousand AI-related bills have been introduced across the states. In the 2025 session alone every state introduced AI legislation and around thirty-eight enacted roughly a hundred measures. Most are narrow. A dozen are not.

Texas — TRAIGA

The Responsible AI Governance Act (HB 149) came into force **1 January 2026**.

It was narrowed substantially during passage. Most private-sector obligations were removed and the focus shifted to government use, with categorical bans that apply to everyone: AI designed to incite self-harm or violence, AI developed for unlawful discrimination, AI intended to infringe constitutional rights, and AI producing child sexual abuse material or certain deepfakes. State entities face additional restrictions on social scoring and non-consented biometric identification.

TRAIGA is intent-based, which is a meaningfully lower exposure than a disparate-impact standard.

Two features make it strategically important. It creates a regulatory sandbox and an AI Council. And **substantial compliance with the NIST AI Risk Management Framework operates as an enforcement safe harbour** — the clearest statutory endorsement of a voluntary framework anywhere in US law.

Enforcement sits with the Attorney General, with no private right of action.

\$10–12k curable violations	\$80–200k uncurable violations	\$40k/day continuing violations
---------------------------------------	--	---

California — four instruments running in parallel

SB 53 — Transparency in Frontier AI Act

Took effect **1 January 2026**. It applies to large frontier developers and requires published safety and security frameworks, catastrophic-risk assessment, incident reporting to the Office of Emergency Services, and whistleblower protections. It is the successor to the vetoed SB 1047, and deliberately transparency-first rather than prescriptive.

AB 2013 — training-data disclosure

Took effect the same day, requiring generative AI developers to disclose training-data information.

The AI Transparency Act

SB 942 as amended by AB 853 became operative **2 August 2026**. Covered generative AI providers must offer a free AI-detection tool, apply latent disclosures to output, and impose licence terms requiring revocation within 96 hours where a licensee strips the disclosure capability.

CPPA ADMT regulations

The automated decision-making technology regulations took effect 1 January 2026, with the significant-decision obligations — pre-use notice, opt-out, access rights and risk assessments — beginning **1 January 2027**.

Separately, amended FEHA regulations covering automated decision systems in employment took effect in October 2025.

Colorado — repealed and replaced

This one requires care, because a great many trackers still describe the old law.

SB 24-205, the Colorado AI Act, was the first comprehensive US state AI law and was widely expected to become the template. **It was repealed.** SB 26-189, the Automated Decision-Making Technology Act, was signed on 14 May 2026 and its obligations commence **1 January 2027**.

The successor is narrower: consumer notices, explanations of adverse outcomes within thirty days, and human review rights.

The consequential change is what did not carry over. The original law provided an **affirmative defence** for organisations using a recognised risk management framework such as the NIST AI RMF. That defence did not survive into SB 26-189. Adopt frameworks for the evidence they produce, not for statutory protection that may be withdrawn.

Illinois, New York and elsewhere

Illinois HB 3773

Took effect 1 January 2026, amending the Human Rights Act to address AI in employment decisions and prohibiting the use of zip code as a proxy for protected characteristics. Illinois also has BIPA, the most litigated AI-adjacent statute in America, and the WOPR Act restricting AI in psychotherapy.

New York's RAISE Act

Takes effect **1 January 2027**, requiring frontier developer safety protocols and incident reporting. It is the second state frontier regime after California's.

NYC Local Law 144

Has required annual independent bias audits and candidate notice for automated employment decision tools since July 2023.

Utah's AI Policy Act

Requires disclosure of generative AI use on request, and affirmative disclosure in regulated occupations.

The jurisdiction mistake


Most trackers model state jurisdiction as "where is your company." That is wrong, and it is why organisations discover obligations by accident.

Reach attaches through the residence of your consumers, applicants and data subjects. A company headquartered in one state crosses into another state's jurisdiction silently, through ordinary sales growth, usually finding out via a customer's security questionnaire rather than a regulator.

Your applicability question is not where you are.

It is where your users are, what role you play, and which thresholds you have crossed.

The preemption fight



Washington is pushing for one national rule. It does not have one yet.

Executive Order 14365, signed in December 2025, directs a Department of Justice AI Litigation Task Force to challenge state AI laws, requires a Commerce review of burdensome state laws, and contemplates conditioning federal funding. A White House legislative framework published in March 2026 recommends that Congress enact broad preemption.

Executive orders do not themselves preempt state law. Preemption normally flows from congressional enactment, and the directed actions face challenges on Tenth Amendment, Commerce Clause and delegated-authority grounds. The order expressly carves out state laws on child safety, AI compute and data centre infrastructure, and state government procurement.

The prudent posture is to keep complying with state law while tracking the litigation. A programme built on the assumption that preemption will arrive is a programme built on a contested executive action.

Questions people ask

Is the Colorado AI Act still in force?

No. SB 24-205 was repealed and replaced by SB 26-189, whose obligations begin 1 January 2027.

Does the NIST AI RMF provide legal protection?

In Texas, substantial compliance operates as an enforcement safe harbour under TRAIGA. Colorado's equivalent affirmative defence was removed in its 2026 replacement law.

Will federal law preempt state AI laws?

Unresolved. Executive Order 14365 directs agencies to pursue preemption, but executive orders do not preempt state law without congressional action.

THE EIGHT TIERS SERIES

Part 0: The eight kinds · Part 1: Comprehensive statutes · Part 2: The treaty nobody is watching · **Part 3: Where the obligations actually are** · Part 4: The regulators who were already there · Part 5: ISO/IEC 42001 · Part 6: NIST AI RMF · Part 7: Security taxonomies · Part 8: Principles · Part 9: Putting it together