CSAI Foundation

Cloud Security Alliance AI Safety Initiative

State AI Laws Take Hold as Federal Preemption Stalls

Enterprise Compliance Guidance for the US Multi-State AI Regulatory Landscape

Unofficial AI-assisted Research

2026-04-04

Enterprise Compliance Guidance for the US Multi-State AI Regulatory Landscape

Key Takeaways

- **No comprehensive federal AI law governs general AI use in the United States as of April 2026.** State AI laws enacted throughout 2025 and early 2026 remain fully in effect and legally enforceable. Enterprises must comply with them today.
- **The federal preemption campaign has not succeeded through legislation.** The US Senate voted 99–1 in July 2025 to strip a 10-year AI preemption moratorium from the One Big Beautiful Bill Act [1]. A December 2025 executive order directing the Department of Justice to challenge state AI laws does not itself preempt those laws [2].
- **145 state AI laws were enacted across 38 states in 2025 alone** [3], with major frameworks now in force in California (SB 53), Texas (TRAIGA), and Illinois (HB 3773) as of January 1, 2026, and Colorado's AI Act (SB 205) scheduled for June 30, 2026 [4][5][6].
- **Enterprises that delay compliance waiting for federal preemption assume a legal risk without legal basis.** No federal court has invalidated a state AI law; DOJ litigation is in its nascent phase, with no complaints filed as of this writing [7].
- **The NIST AI Risk Management Framework (AI RMF)** provides the most jurisdiction-agnostic governance baseline available. Colorado explicitly offers a safe harbor for NIST AI RMF compliance [8], and the framework aligns with ISO/IEC 42001 and the EU AI Act requirements, enabling a unified multi-jurisdictional program.

Background

The United States has entered a complex and fragmented regulatory moment: a proliferation of state AI laws without a federal counterpart, an executive branch actively campaigning to override those laws, and a Congress that has so far declined to act. The result is a legal landscape that enterprises must navigate

with urgency, because the operative compliance obligations are the state laws that are already in effect — not the federal framework that may or may not materialize.

The federal vacuum is not for lack of effort. The Trump administration made the curtailment of state AI regulation a stated policy priority in its first year. In December 2025, the White House issued an executive order declaring a national policy of "minimally burdensome" AI standards and directing the Department of Justice to establish an AI Litigation Task Force to challenge state laws on Dormant Commerce Clause and preemption grounds [2]. In March 2026, the White House released a National Policy Framework for Artificial Intelligence as a set of nonbinding legislative recommendations to Congress, proposing that federal law preempt state AI laws imposing "undue burdens" on frontier model developers [9]. Both documents establish executive branch priorities and direction, but neither constitutes statutory authority capable of preempting state law. Courts do not invalidate state statutes on the basis of executive orders, and as of April 2026, no federal legislation implementing the White House's recommendations has been formally enacted. The only legislative vehicle to emerge was a discussion draft — the TRUMP AMERICA AI Act — which was circulated beginning in late 2025 but had not been formally introduced as of this writing [20].

The most consequential federal preemption attempt came and failed in July 2025. A 450-word provision in the budget reconciliation package known as the One Big Beautiful Bill Act would have prohibited state enforcement of any AI regulation for 10 years. After a brief bipartisan negotiation over a modified five-year moratorium with child-safety carve-outs collapsed, the Senate voted 99–1 to strip the provision entirely [1]. President Trump signed the bill into law on July 4, 2025 without any AI preemption language. State laws emerged from that legislative cycle with their authority intact.

Meanwhile, state legislatures had been extraordinarily active. In 2025, state legislators introduced 1,208 AI-related bills across all 50 states; 145 were enacted into law, according to multistate.ai and Future of Privacy Forum tracking [3][4]. The National Conference of State Legislatures, using a narrower counting methodology, reports that 38 states adopted approximately 100 AI-related measures during the year [3]. California enacted the largest number of AI laws of any state, followed by Texas, Montana, Utah, and Arkansas among the leading jurisdictions for AI legislation [3]. The enterprise compliance challenge is not hypothetical. It is a present operational reality.

Analysis

The Current Legal Landscape

The operative state AI laws create compliance obligations that vary significantly in scope, mechanism, and enforcement model. Understanding their structure is the precondition for any rational enterprise compliance program.

California's **SB 53** — the Transparency in Frontier Artificial Intelligence Act — was signed by Governor Newsom on September 29, 2025, and its primary provisions took effect January 1, 2026 [5]. It targets "large frontier developers": companies with gross revenues exceeding $500 million developing AI models trained at more than 10^26 floating-point operations. Covered developers must publish a "frontier AI framework" explaining how they identify and mitigate catastrophic risks, report safety incidents to regulators within 15 days (24 hours for imminent harm situations), and maintain whistleblower protections. Maximum penalties are capped at $1 million per violation. SB 53 represents a narrower successor to the vetoed SB 1047, and its more limited scope is widely expected to make it more resistant to legal challenge than its predecessor, which Newsom rejected in September 2024 on the grounds that it would have created compliance burdens disproportionate to verifiable safety benefit [5].

Colorado's **SB 205** (the Colorado AI Act) was passed in 2024 and applies to developers and deployers of "high-risk AI systems" — those that make or substantially assist in consequential decisions affecting employment, credit, housing, healthcare, education, legal services, or insurance [8]. After a one-time delay from February 2026, it takes full effect on June 30, 2026 [11]. Developers must exercise "reasonable care" to prevent algorithmic discrimination, provide technical documentation, and issue disclosures to deployers and regulators. Deployers must adopt risk-management policies, conduct initial and annual algorithmic impact assessments, and issue pre-decision and adverse-decision notices to affected consumers. The Colorado law provides a meaningful compliance pathway: entities that demonstrably align to the NIST AI RMF or an equivalent standard receive a rebuttable presumption of having exercised reasonable care [8]. Enforcement is by the Attorney General only, with a cure period before any enforcement action is taken, and there is no private right of action [8]. Legal commentators have identified Colorado's SB 205 as a likely early target for DOJ challenge, given its broad scope and June 30, 2026 effective date.

Texas enacted the **Responsible AI Governance Act (TRAIGA)**, which took effect January 1, 2026 [6]. It applies broadly to any entity developing or deploying AI systems in Texas or offering AI-enabled products or services to Texas residents — a scope that, if applied broadly, could reach most nationally operating enterprises that offer AI-enabled services or products accessible to Texas residents. The law prohibits specific harmful AI applications (inciting self-harm, unlawful discrimination, unauthorized

deepfakes) and requires disclosures when government agencies and healthcare providers use AI systems that interact with consumers. A regulatory sandbox provision enables reduced-risk testing under state supervision. Enforcement is by the state AG without a private right of action. Notably, TRAIGA adopts an intent standard for discrimination claims: disparate impact alone is insufficient to establish a violation [6].

Illinois **HB 3773**, effective January 1, 2026, amends the Illinois Human Rights Act to cover discrimination resulting from employer use of AI in hiring, firing, discipline, tenure, and training decisions [12]. It prohibits using ZIP codes as a proxy variable in AI hiring models. This law is structurally distinct from others in one critical respect: it provides a **private right of action** for violations, substantially increasing litigation exposure for Illinois-operating employers relative to the AG-enforcement-only models in Colorado and Texas.

New York maintains active legislative initiatives in AI disclosure, training-data transparency, and employment discrimination but has not yet enacted a comprehensive statewide AI framework as of this writing. New York City's Local Law 144 on automated employment decision tools remains in effect since 2023. The New York RAISE Act, which would impose extensive frontier model safety reporting obligations, is targeted for effectiveness in 2027.

The Compliance Problem at Scale

The challenge enterprises face is not simply having multiple laws to comply with — it is the combination of different scope triggers, varying definitional standards, conflicting disclosure requirements, disparate enforcement models, and genuine legal uncertainty about which requirements will survive potential federal preemption challenges. Many enterprises operating at scale have not yet completed a systematic inventory of the AI systems deployed across their operations — an inventory that is the precondition for applying any of the above frameworks to their specific compliance exposure.

The definitional inconsistency problem is partially addressed but not resolved. California's AB 2885 (effective January 1, 2025) standardized the definition of "artificial intelligence" across California statutes [15]. But national enterprises must still reconcile scope triggers that differ across jurisdictions. What constitutes a "high-risk AI system" in Colorado, a "consequential decision" under that state's law, or "AI system" under TRAIGA involves legal interpretation that differs by jurisdiction and in many cases has not yet been tested in enforcement proceedings.

The enforcement model differences carry practical significance. Colorado and Texas provide AG-only enforcement with cure periods; Illinois provides a private right of action. This means that for employment-related AI applications, Illinois exposure includes class action litigation risk in addition to regulatory scrutiny — a materially different risk profile than similar operations in Colorado.

Enterprises must also contend with the temporal uncertainty created by the federal preemption campaign. Legal counsel broadly advises against delaying compliance in anticipation of federal action, given the timeline uncertainty inherent in constitutional litigation and legislative processes [7]. The DOJ AI Litigation Task Force was established by January 10, 2026, per the executive order's mandate [7], but as of late March 2026 had not yet filed its first complaint [7]. Even when it does, cases will proceed through district courts and appeals courts before any binding precedent is established — a process that constitutional preemption litigation typically unfolds over multiple years. State laws remain enforceable throughout that process.

State Law	Scope Trigger	Effective Date	Enforcement	Private Right of Action
California SB 53	Frontier developers (>$500M revenue, >10^26 FLOPs)	January 1, 2026	AG	No
Colorado SB 205	High-risk AI in consequential decisions	June 30, 2026	AG (cure period)	No
Texas TRAIGA	Any AI developer/deployer reaching TX residents	January 1, 2026	AG	No
Illinois HB 3773	Employer AI use in employment decisions	January 1, 2026	Civil litigation	**Yes**
New York RAISE Act	Frontier model developers	Targeted 2027	TBD	TBD

Recommendations

Immediate Actions

Enterprises should not wait for regulatory clarity that the federal process cannot credibly provide in the near term. The priority is to establish the factual foundation for compliance: a complete AI system inventory, accurate characterization of each system's use against state-law scope triggers, and clear assignment of compliance ownership.

For enterprises operating in California with revenues exceeding $500 million and using or developing large language models or other frontier AI: SB 53's requirements are in effect now. Verify whether any of your AI systems meet the "large frontier developer" threshold, and if so, confirm that the required frontier AI framework has been published and incident reporting procedures are established for your California-relevant operations.

For any enterprise that makes employment-related AI decisions affecting Illinois workers, the private right of action under HB 3773 should be treated as generating litigation exposure equivalent to established employment discrimination statutes. An internal review of any AI-assisted or AI-driven hiring, performance, disciplinary, or termination processes should be conducted, with particular attention to model variables (including geographic proxies such as ZIP code) and documented non-discrimination testing.

Short-Term Mitigations

With Colorado's SB 205 taking effect June 30, 2026, enterprises deploying "high-risk AI" as defined under that law have approximately 90 days from this writing to complete impact assessments, establish documentation, and implement the required consumer notice procedures. Colorado's safe harbor creates a rebuttable presumption of reasonable care for entities aligned to the NIST AI RMF — a material reduction in enforcement risk that makes NIST alignment an efficient compliance pathway [8].

Enterprises should also establish a monitoring function specifically for state AI legislation. The pace of state activity — 1,208 bills introduced in 2025 alone — means the landscape will continue shifting. Relevant monitoring should track both newly enacted laws and amendments to existing frameworks, particularly in states where an enterprise has material operations or customer relationships.

Vendor and supply chain review is a necessary component of near-term compliance. State laws in Colorado and Texas do not exempt enterprises from compliance because they use third-party AI rather than internally developed models; deployer obligations exist independently of developer obligations. Standard vendor security assessment and procurement processes should be updated to include AI governance due diligence, specifically the ability to obtain documentation necessary for algorithmic impact assessments and consumer disclosures.

Strategic Considerations

The highest-leverage long-term investment is alignment to a single, multi-jurisdictional governance baseline. The NIST AI RMF is particularly well-suited to the US multi-state context: it is regulation-agnostic by design, provides the explicit safe harbor in Colorado [8], aligns closely with AICM domains [16], and can be partially mapped to both the EU AI Act — whose primary obligations for most high-risk AI

systems apply from August 2, 2026, following GPAI rules that took effect in August 2025 — and ISO/IEC 42001:2023. Organizations with international operations will find that NIST AI RMF alignment, supplemented by ISO/IEC 42001 certification where appropriate, provides more comprehensive coverage than a jurisdiction-by-jurisdiction compliance approach.

The governance investment required to comply with state AI laws — AI system inventory, impact assessment, documentation, risk tiering, consumer notice procedures — substantially overlaps with what sound AI risk management demands independent of legal obligation. Enterprises that frame this work solely as compliance-driven may find that the value of governance investment extends well beyond any single regulatory obligation: the frameworks state legislators drew upon (NIST AI RMF, ISO/IEC 42001, OECD AI Principles) represent emerging professional consensus on what organizations operating AI systems owe to the people those systems affect. Building to those standards positions enterprises to meet evolving obligations and demonstrates to customers and partners a commitment to responsible AI development.

Finally, the legal situation counsels against allowing uncertainty to generate paralysis. The federal preemption campaign is real, but litigation and legislation are inherently slow-moving mechanisms for regulatory change, while state compliance obligations are immediate and enforceable. The prudent posture is to comply with existing law, build durable governance infrastructure, and monitor federal developments as they unfold — not to treat anticipated preemption as permission to defer.

CSA Resource Alignment

The compliance challenge described in this research note maps directly to CSA's existing AI governance frameworks and tools, which provide both conceptual structure and practical implementation guidance.

The **AI Controls Matrix (AICM)** provides the most directly applicable CSA framework for managing multi-jurisdictional AI compliance. As a superset of the Cloud Controls Matrix (CCM), the AICM provides control objectives spanning governance, risk management, security, and privacy that align to NIST AI RMF, ISO/IEC 42001, and OECD AI Principles simultaneously. Enterprises building a compliance program for Colorado's SB 205, California's SB 53, and Texas TRAIGA can use the AICM to map their governance activities to all three frameworks from a single control inventory, reducing duplicative effort and enabling consolidated auditor evidence [16].

The **AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects** publication provides direct implementation guidance for the enterprise compliance program structure described in this note's recommendations. It addresses the organizational structures,

policies, board-level accountability, and compliance reporting mechanisms that state AI laws implicitly or explicitly require, making it a practical companion for enterprises standing up new AI governance functions [17].

The **MAESTRO** (Multi-Agent Ecosystem Security Threat Reasoning and Oversight) framework is particularly relevant for enterprises deploying AI agents or automated decision systems in the consequential-decision domains targeted by Colorado, Illinois, and Texas law. MAESTRO's threat modeling structure for agentic AI systems enables systematic risk identification that, when documented, may contribute to the evidentiary basis for NIST AI RMF compliance and Colorado's safe harbor defense [18].

The **AI Risk Management: Thinking Beyond Regulatory Boundaries** publication, published in 2024, provides an audit methodology that explicitly bridges internal risk management and external regulatory compliance — precisely the dual purpose that a multi-state AI compliance program must serve. Its lifecycle-oriented approach aligns with the documentation, impact assessment, and incident reporting requirements appearing in the state laws reviewed here [19].

Enterprises engaged in the STAR (Security Trust Assurance and Risk) program should consider how their STAR assessment scope can be extended to AI system governance, enabling external attestation of AI controls maturity to regulators, customers, and partners — the kind of transparency increasingly demanded by both state AI laws and enterprise procurement requirements.

References

[1] NBC News. "Big Beautiful Bill passes without AI moratorium after 99-1 Senate vote." NBC News, July 2025.

[2] Sidley Austin. "Unpacking the December 11, 2025 Executive Order on AI Preemption." Sidley Austin Insights, December 2025.

[3] National Conference of State Legislatures. "Artificial Intelligence 2025 Legislation Summary." NCSL, 2025.

[4] Transparency Coalition. "2025 State AI Legislation Report." Transparency Coalition, 2025.

[5] Fisher Phillips. "California Lawmakers Pass Landmark AI Transparency Law for Frontier Models (SB 53 vs. SB 1047)." Fisher Phillips, September 2025.

[6] Norton Rose Fulbright. "The Texas Responsible AI Governance Act." Norton Rose Fulbright, 2025.

[7] BakerHostetler. "Navigating the Emerging Federal-State AI Showdown: DOJ Establishes AI Litigation Task Force." BakerHostetler Client Alert, January 2026.

[8] Baker Botts. "Colorado AI Act Implementation: NIST AI RMF Safe Harbor and Key Requirements." Baker Botts, September 2025.

[9] Holland & Knight. "White House Releases a National Policy Framework for Artificial Intelligence." Holland & Knight Insights, March 2026.

[10] CalMatters. "California AI Regulation: What Passed and What's Next." CalMatters, December 2025.

[11] Clark Hill. "Colorado's AI Law Delayed Until June 2026: What the Latest Setback Means for Businesses." Clark Hill Legal Update, 2025.

[12] National Law Review. "Several State AI Laws Set to Go into Effect in 2026 Despite Federal Government's Push to Eliminate." National Law Review, 2026.

[13] King & Spalding. "New State AI Laws Are Effective on January 1, 2026, But a New Executive Order Signals Disruption." King & Spalding Insights, January 2026.

[14] Credo AI. "Latest AI Regulations Update: What Enterprises Need to Know in 2026." Credo AI Blog, 2026.

[15] Pillsbury Law. "California AI Laws: AB 2885 and the Standardization of AI Definitions." Pillsbury Law, 2025.

[16] Cloud Security Alliance. "AI Controls Matrix (AICM)." CSA, 2024.

[17] Cloud Security Alliance. "AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects." CSA, 2024.

[18] Cloud Security Alliance. "MAESTRO: Multi-Agent Ecosystem Security Threat Reasoning and Oversight." CSA, 2025.

[19] Cloud Security Alliance. "AI Risk Management: Thinking Beyond Regulatory Boundaries." CSA, 2024.

[20] Jones Walker. "The TRUMP AMERICA AI Act: Federal Preemption Meets Comprehensive Regulation." Jones Walker AI Law Blog, 2026.