CSAI Foundation | Cloud Security Alliance

US AI Governance Fragmentation: The State Patchwork Burden

2026-05-30

AI-assisted Rapid Research

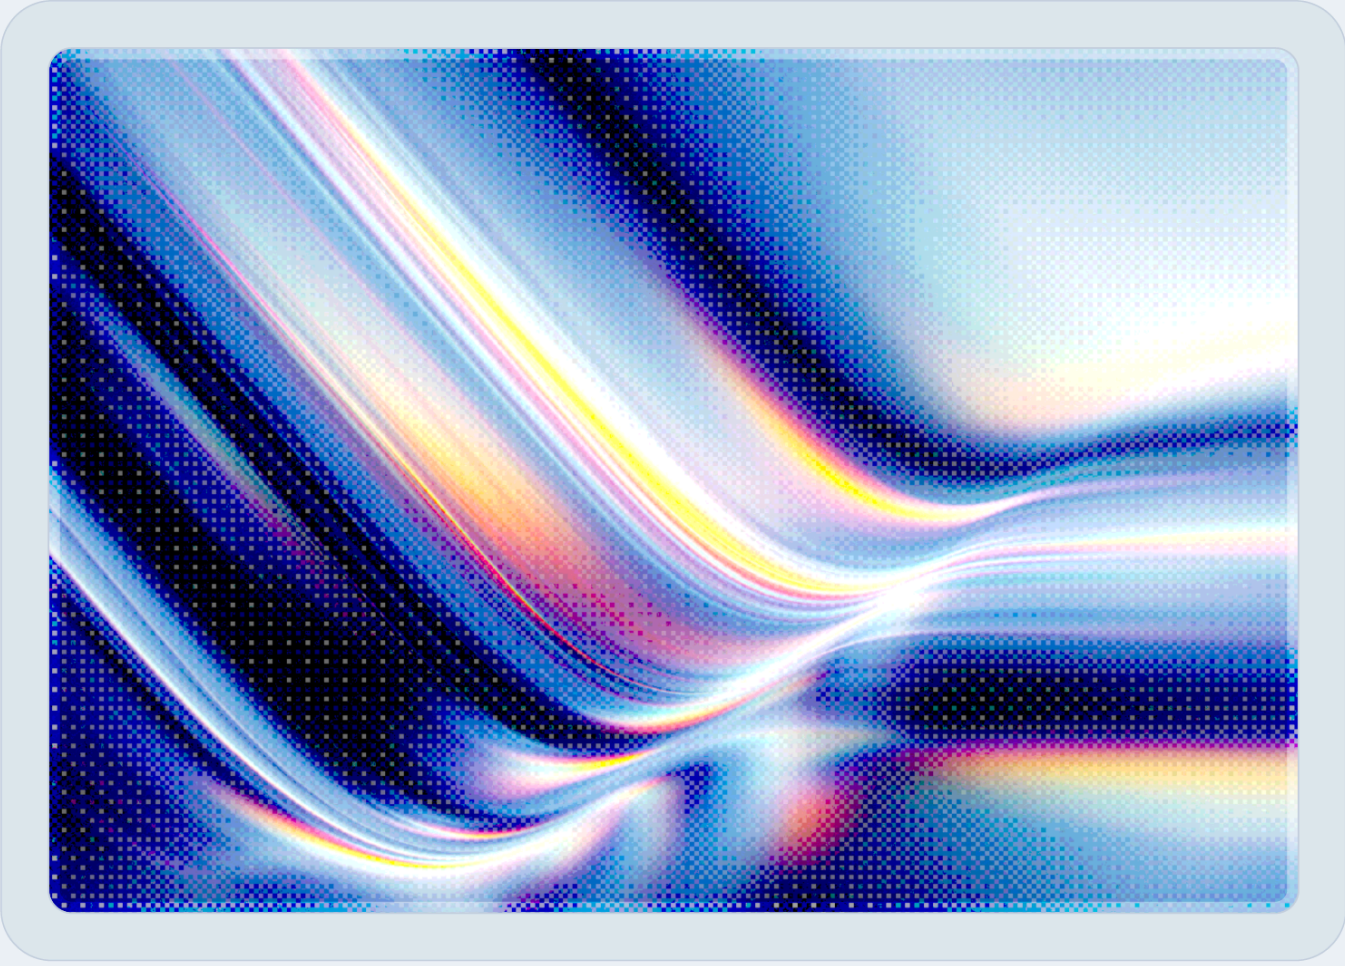

CSAI

CSA cloud security alliance®

Key Takeaways

- All 50 states have introduced AI-related legislation. As of May 2026, legislators across the country have introduced over 2,182 AI-related bills [26], with inconsistent definitions, audit timelines, and enforcement mechanisms that create conflicting compliance obligations for multi-state operators.
- Colorado's repeal and replacement of the nation's first comprehensive AI law (SB 24-205) in May 2026 illustrates the planning risk inherent in state-level AI governance: compliance programs built around specific requirements can be rendered moot mid-implementation by political pressure, industry litigation, and executive intervention.
- The Trump administration has used executive power to challenge state AI laws through an AI Litigation Task Force, federal funding leverage, and a non-binding legislative framework calling for federal preemption — but Congress has not enacted preemption, and a proposed 10-year moratorium on state AI enforcement was defeated 99-1 in the Senate.
- Per-company annual AI compliance spend now ranges from $50,000 to $500,000, with initial audit cycles running $25,000–$150,000 and enterprise-scale monitoring systems adding $1.5–3 million annually — burdens that disproportionately affect mid-market and small organizations relative to hyperscalers [18].
- NIST AI RMF alignment offers a strong and extensible compliance foundation under current conditions: statutory safe harbors referencing the framework exist in Colorado's replacement law and Texas RAIGA, and the framework provides a jurisdiction-neutral governance structure that can accommodate additional state requirements as they evolve.

Background

For the first decade of modern AI deployment, regulatory governance of AI in the United States operated through a diffuse set of agency guidance documents, voluntary frameworks, and the piecemeal application of existing law — antidiscrimination statutes, consumer protection authority, and financial services regulation — to AI-specific fact patterns. This changed materially between 2024 and 2026. Colorado's SB 24-205, signed in May 2024, became the nation's first comprehensive state AI law, establishing algorithmic discrimination obligations, mandatory risk management programs, impact

assessment requirements, and consumer disclosure duties for deployers of "high-risk AI" systems in consequential decision contexts. Within months of its signing, Governor Polis — who had signed the bill reluctantly — publicly called for its revision, describing the law's compliance architecture as an impediment to economic growth. That political dynamic, combined with industry opposition and a federal lawsuit by xAI joined by the Department of Justice, ultimately produced SB 26-189, signed on May 14, 2026, which formally repealed and replaced the original law with a substantially narrower transparency-and-notice framework effective January 1, 2027 [1][2].

Colorado's experience unfolded against a backdrop of accelerating legislative activity in all fifty states. By March 2026, legislators had introduced over 1,561 AI-related bills in 45 states, and more than 300 AI bills were introduced in the first six weeks of 2026 alone [26]. Several major laws took effect on January 1, 2026, including California's Transparency in Frontier AI Act (SB 53), targeting frontier model developers operating above 10^26 FLOPS training compute with revenue exceeding $500 million; California AB 2013, requiring generative AI developers to publish training dataset summaries; and Texas's Responsible Artificial Intelligence Governance Act (RAIGA/HB 149), which applies broadly to both developers and deployers serving Texas residents and establishes a prohibited-use framework covering social scoring, behavioral manipulation, and specific discriminatory applications [4][5]. Connecticut followed with SB 5, passed on May 1, 2026, covering automated employment decision technology, AI companions, synthetic media, and frontier AI whistleblower protections in what analysts described as the most extensive state AI legislation enacted to date [6][7].

The federal government has contributed to uncertainty rather than resolution. The Trump administration revoked President Biden's Executive Order 14110 on its first day in office and replaced it with EO 14179, signaling a shift from precautionary oversight to deregulation and innovation primacy [8]. A December 2025 executive order (EO 14365) established a DOJ AI Litigation Task Force to challenge state AI laws on constitutional and preemption grounds and authorized the withholding of $42 billion in BEAD broadband infrastructure funding from states enforcing laws the administration characterizes as "onerous" [9]. In March 2026, the White House released a non-binding National Policy Framework for AI calling on Congress to preempt state laws that impose "undue burdens" while preserving narrow state authority over child safety and government procurement [10]. Congress has not acted on these recommendations. A proposed 10-year moratorium on state AI enforcement, inserted into the budget reconciliation process by House Republicans, was stripped from the bill by a 99-1 Senate vote, and comprehensive federal preemption language was also rejected in the National Defense Authorization Act [11][12].

Security Analysis

The primary compliance risk created by AI governance fragmentation is not regulatory complexity per se — enterprises routinely manage multi-jurisdictional legal regimes — but rather the structural incompatibility of obligations across state laws, combined with rapid and unpredictable change that defeats long-horizon compliance planning. The definitional problem alone is significant: Colorado defines high-risk AI by reference to specific consequential decision use cases; California anchors regulatory scope to training data characteristics and model compute scale; Illinois applies no "high-risk" framing at all; and a proposed Nebraska statute at one point would have swept basic spreadsheet software within its AI definition [26][13]. A single AI tool may carry mandatory bias audit obligations in one state, developer safety reporting obligations in another, and no specific regulatory obligations in a third — simultaneously, for the same organization. The five main compliance obligation categories — training data transparency, AI-generated content provenance, frontier model safety assessments, child-safety disclosures, and human oversight mandates in healthcare and employment — overlap imperfectly across these frameworks, with distinct definitions, timelines, and enforcement mechanisms in each [26][14].

The employment AI context illustrates the conflict most concretely. New York City's Local Law 144 requires annual independent bias audits of automated employment decision tools and mandates ten-day advance notice to job applicants before use, with civil penalties of $500 to $1,500 per violation. Illinois requires notification and consent before AI analyzes video interviews. Colorado's original SB 24-205 imposed broad risk assessments across eight consequential decision domains, which would have applied independently of the NYC or Illinois requirements. A company hiring remotely across all three jurisdictions must comply with all three frameworks concurrently, despite their incompatible definitions of covered AI, different audit frequencies, and different notice content requirements [15][13]. Connecticut's SB 5, which explicitly codifies that use of automated decision technology is not a defense against discrimination claims and phases in deployer obligations through October 2027, now adds a fourth distinct framework for companies with employees or applicants in that state [6][7]. Leading legal practitioners advising on AI compliance generally recommend that enterprises not await resolution of ongoing legal challenges but instead identify the most stringent applicable requirements and design compliance programs around those — a strategy that by definition escalates cost and complexity for every jurisdiction added to an organization's operational footprint [3][5]. Organizations subject to laws currently under active federal court stays may apply a different calculus, but this represents the prevailing guidance for those operating under currently enforceable requirements.

Colorado's rollback from SB 24-205 to SB 26-189 deserves examination not only as a regulatory policy event but as a case study in compliance planning risk. Organizations that had built compliance programs around the original law's requirements — duty of care standards, risk management program

documentation, impact assessment workflows, and attorney general notification procedures for high-risk system deployments — found those programs substantially moot when the replacement law stripped all of those obligations and replaced them with a disclosure-only model [1][2][16]. The replacement law also eliminates the algorithmic discrimination framework, removes the private right of action that had created litigation exposure, and restricts enforcement exclusively to the attorney general, with a 60-day cure period sunsetting in 2030 [2]. Colorado's experience illustrates a governance risk that high-profile state AI laws may face: political and legal opposition can effectively prevent implementation even after enactment, leaving organizations that invested in compliance with stranded costs and limited legal clarity about what any replacement regime actually requires in practice. The Colorado outcome was shaped by an unusual confluence of conditions — a governor who had always opposed the original law's architecture, a federal court stay of enforcement obtained on April 27, 2026 following an xAI lawsuit joined by the DOJ, and an industry coalition that had spent two years opposing implementation [1][17]. Connecticut's SB 5, equally high-profile, passed in May 2026 without the same political reversal, illustrating that different state contexts produce different outcomes and that no single state's trajectory should be treated as the inevitable pattern.

The fragmentation dynamic creates structural incentives toward deferring governance investments that do not produce near-term competitive returns — though this is an anticipated market effect rather than a documented behavioral outcome. Safety evaluations, bias testing, impact assessments, and incident response documentation are cost centers whose value is realized primarily through risk reduction and regulatory compliance; when regulatory floors are absent or unstable, the business case for these investments weakens. The compliance cost asymmetry exacerbates this dynamic: organizations subject to California's SB 53 or Connecticut's SB 5 bear substantial ongoing obligations that competitors operating exclusively in less-regulated contexts do not face, creating a form of regulatory arbitrage in which organizations that invest in governance compliance face cost burdens that competitors in less-regulated jurisdictions avoid, disadvantaging early adopters [3][12]. The global AI governance compliance market is projected at $2.54 billion in 2026, growing to $8.23 billion by 2034, yet only 23 percent of organizations report confidence in their AI governance frameworks and 70 percent of IT leaders cite AI compliance as a major deployment challenge [18]. The gap between rising compliance investment and low governance confidence is consistent with fragmentation degrading governance quality rather than building it — though causation cannot be established from available data. A plausible mechanism is that distributing organizational attention across conflicting regulatory targets redirects resources away from coherent risk management, but alternative explanations — the inherent difficulty of AI governance, reactive rather than maturity-building spend patterns — cannot be ruled out.

The federal-state tension also creates second-order risks for organizations attempting to engage constructively with regulators. The DOJ AI Litigation Task Force, authorized under EO 14365, creates ongoing legal uncertainty about which state laws will remain enforceable — a question that cannot be answered quickly, given that constitutional preemption challenges typically resolve over years of

litigation and appeals [9][19][20]. The administration's threat to condition BEAD broadband funding on state AI law compliance represents what legal analysts characterize as a novel use of federal spending power, one that has not been tested in court against existing Dole doctrine precedents, and whose ultimate effect on state legislative behavior is unknown. Organizations planning multi-year AI governance programs must therefore budget not only for the current landscape of state obligations but for the possibility that laws they have invested in complying with will be challenged, stayed, or repealed; that laws they are not yet subject to will be enacted in new jurisdictions; and that a federal preemption framework, if it ultimately arrives, may not include a substantive safety baseline to replace the state protections it displaces. Unlike the GDPR transition in the European Union, which at minimum established a unified definitional framework and substantive baseline across member states, any U.S. federal AI preemption legislation under current political conditions may arrive without a replacement federal baseline, leaving enterprises in a governance vacuum rather than a governance regime [12][20].

Recommendations

Immediate Actions

Organizations with material AI deployments across multiple states should take the following steps to establish a defensible governance baseline before the October 2026 effective dates for Connecticut's automated employment decision provisions. Because state law applicability turns on the location of employees, customers, and applicants rather than corporate headquarters, the geographic scope of compliance obligations is often broader than legal teams initially estimate.

- Conduct a jurisdiction-specific AI system inventory that maps each deployed AI tool to the state laws applicable based on the locations of employees, applicants, customers, and service recipients — not just the organization's headquarters state.
- Identify the most stringent applicable disclosure, bias testing, and human oversight obligations currently in force across the organization's operational jurisdictions, and verify that deployed automated decision tools meet those requirements by their current effective dates. Connecticut's AEDT provisions take effect October 1, 2026 for employers; California's SB 53 and AB 2013 are already in effect.
- Designate a cross-functional AI Governance Group — including legal, privacy, security, product, and human resources — as the accountable body for tracking regulatory developments and coordinating compliance responses, rather than treating AI compliance as a legal department function alone.

Short-Term Mitigations

Organizations should adopt NIST AI Risk Management Framework alignment as the default governance architecture for all material AI deployments. This recommendation rests on practical grounds: Colorado's replacement law and Texas RAIGA both provide statutory safe harbors for organizations that can demonstrate NIST AI RMF adoption, making framework alignment directly relevant to legal exposure in two of the most active state regulatory jurisdictions [21][22]. More broadly, the AI RMF's structure — Govern, Map, Measure, Manage — provides a documented, auditable governance process that can absorb additional state-specific requirements as incremental obligations rather than requiring wholesale program rebuilds. Organizations with significant international compliance exposure may also consider ISO/IEC 42001 certification as a complementary or alternative framework. The NIST Generative AI Profile (AI 600-1, published July 2024) and the preliminary Cybersecurity Framework Profile for AI (NISTIR 8596, December 2025) provide sector-specific extensions that address concerns specific to generative AI and cybersecurity-adjacent deployments [23][24].

Vendor contract terms should be reviewed and updated to reflect the principle that organizations cannot contractually transfer AI compliance obligations to third-party vendors. This is particularly relevant for automated employment decision tools, where both NYC Local Law 144 and Connecticut SB 5 hold deployers liable for discrimination resulting from third-party AI tools regardless of vendor representations. Contracts should require vendors to provide documentation of training data categories, known limitations, bias evaluation methodologies, and change-notification procedures — the same documentation that constitutes the core developer obligation under Colorado's replacement law and California AB 2013 [2][25][5]. Representations that tools are "fair" or "legally compliant" should be replaced with documented contractual obligations tied to specific testing protocols and audit rights.

Strategic Considerations

The longer-term strategic question for AI-deploying organizations is whether to design compliance programs as state-law tracking functions or as substantive governance programs that happen to satisfy legal requirements. These produce different organizational outcomes. A tracking-based approach treats each new state law as an additive obligation that expands the compliance footprint incrementally, requiring continuous legal monitoring and jurisdiction-specific program modifications. A governance-based approach invests in the underlying safety and transparency capabilities — impact assessments, bias evaluations, human oversight mechanisms, incident response documentation — that more comprehensive state AI laws, such as Connecticut's SB 5 and the pre-repeal Colorado SB 24-205, are attempting to require, and treats legal compliance as a secondary output of those capabilities rather than their primary purpose. The governance-based approach is more expensive to establish but more

durable across a changing regulatory landscape, and it produces the kind of auditable evidence that regulators and enterprise customers require — and that insurers may increasingly require as AI risk underwriting matures.

The federal preemption debate, regardless of how it ultimately resolves, does not reduce the organizational need for substantive AI governance. If federal preemption passes without a meaningful safety baseline, organizations in regulated sectors — financial services, healthcare, critical infrastructure — will remain subject to sector-specific federal agency guidance under existing authority, including FFIEC examination guidelines, ONC health IT certification requirements, and CISA sector risk management obligations, none of which are displaced by preemption of general-purpose state AI laws. If preemption does not pass, the state patchwork will continue expanding. Either way, an organization that has built genuine AI governance capabilities is better positioned than one whose compliance program is optimized for a specific set of current state obligations that may change before they are enforced.

CSA Resource Alignment

The governance fragmentation described in this note creates direct demand for the frameworks and tools that CSA has developed to address AI risk systematically. The AI Controls Matrix (AICM), as a superset of the Cloud Controls Matrix, provides the control inventory most relevant to organizations attempting to build multi-jurisdictional compliance programs: its domains covering data governance, identity and access management, audit logging, and transparency disclosures align with the documentation and audit trail categories that appear across California's SB 53, Colorado's replacement law, and Connecticut's SB 5. Organizations using the AICM as a baseline can map individual state requirements to specific controls rather than building separate control architectures for each jurisdiction, which reduces the compliance cost differential between single-state and multi-state operators.

The STAR (Security, Trust, Assurance, and Risk) Registry provides a mechanism for organizations to demonstrate AI governance posture to enterprise customers and regulators through third-party-validated attestations. As state laws increasingly require deployers to obtain and retain documentation of AI system characteristics from developers and vendors, STAR Registry attestations can serve as a standardized evidence base that deployers can draw on when responding to such documentation requests. Regulatory acceptance of third-party attestations varies by jurisdiction and has not been formally recognized by California or Connecticut enforcement bodies, but the structured and auditable nature of STAR attestations positions them as a useful complement to jurisdiction-specific compliance documentation. The MAESTRO (Multi-layer AI Threat and Risk Ontology) framework provides the threat

modeling substrate applicable to the impact assessments and risk evaluation requirements that several state laws require — and is directly relevant to the systematic identification of harms in consequential decision contexts that regulators across all active state frameworks are attempting to address.

CSA's Zero Trust principles also bear on the governance fragmentation problem in a dimension that often goes unexamined. The automated decision systems most heavily targeted by state AI laws — employment screening tools, credit evaluation models, healthcare triage systems — frequently operate at the intersection of AI inference and privileged access to personal data. Zero Trust architecture, with its emphasis on continuous verification, least-privilege access, and detailed audit logging at every decision point, provides an operational framework that both reduces the attack surface of AI-integrated systems and produces the audit trail that bias investigations, impact assessments, and regulatory examinations require. Organizations implementing Zero Trust for AI-adjacent data flows are not simply improving their security posture — they are building the infrastructure that makes substantive AI governance verifiable.

References

[1] The Colorado Sun. "Colorado's fierce two-year fight over AI regulation ends with watered-down law, little fanfare." The Colorado Sun, 2026-05-12.

[2] Troutman Pepper. "Colorado Legislature Passes Bill to Repeal and Replace Colorado AI Act." Troutman Pepper Privacy + Cyber + AI Blog, 2026-05.

[3] Cooley LLP. "State AI Laws – Where Are They Now?" Cooley LLP, 2026-04-24.

[4] King & Spalding. "New State AI Laws are Effective on January 1, 2026, But a New Executive Order Signals Disruption." King & Spalding, 2026-01-01.

[5] Nelson Mullins. "Texas Responsible Artificial Intelligence Governance Act (HB 149)." Nelson Mullins, 2025-07-01.

[6] CT Mirror. "Connecticut Passes AI Regulations After Years in Development." CT Mirror, 2026-05-01.

[7] DLA Piper. "Unpacking SB5: Connecticut's New AI Law." DLA Piper, 2026-05-15.

[8] National Law Review. "Trump Revokes Biden's AI Executive Order." National Law Review, 2025-01-23.

[9] White House. "Ensuring a National Policy Framework for Artificial Intelligence (EO 14365)." WhiteHouse.gov, 2025-12-11.

[10] Morgan Lewis. "White House AI Framework Puts Federal Preemption at the Center of the Debate." Morgan Lewis, 2026-03-20.

[11] Reboot Democracy. "House Republicans Include AI Regulation Preemption in Budget Reconciliation Bill." Reboot Democracy, 2026-05-01.

[12] Center for American Progress. "Moratoriums and Federal Preemption of State AI Laws Pose Serious Risks." Center for American Progress, 2026-01-01.

[13] Akin Gump. "The Growing Patchwork of State AI Laws: What It Means for Employers." Akin Gump, 2026-01-01.

[14] Jones Walker LLP. "The Fragmentation Problem: Why Your AI Governance Can't Stop at State Lines." Jones Walker LLP, 2026-01-01.

[15] Warden AI. "NYC Local Law 144 Compliance Guide 2026." Warden AI, 2026-01-01.

[16] Davis Wright Tremaine. "Colorado AI Act Repealed and Replaced by Narrower Statute Focused on Transparency Requirements and Enhanced Consumer Rights." Davis Wright Tremaine Privacy & Security Law Blog, 2026-05.

[17] Law and the Workplace. "Major Developments Put Colorado's AI Law on Ice Ahead of Implementation." Law and the Workplace, 2026-05.

[18] SQ Magazine. "AI Compliance Cost Statistics 2026." SQ Magazine, 2026-01-01.

[19] Ropes & Gray. "Examining the Landscape and Limitations of the Federal Push to Override State AI Regulation." Ropes & Gray LLP, 2026-03-01.

[20] NPR. "Trump is trying to preempt state AI laws via an executive order. It may not be legal." NPR, 2025-12-11.

[21] Morrison Foerster. "Colorado Hits Reset on AI Regulation With a New AI Act: What Developers and Deployers Need to Know." Morrison Foerster, 2026-05-15.

[22] SecureWorld. "Texas Passes Most Comprehensive AI Governance Bill." SecureWorld, 2025-07-01.

[23] NIST. "NIST AI Risk Management Framework." NIST, 2023-01-26.

[24] NIST CSRC. "NISTIR 8596 (Draft) — Cybersecurity Framework Profile for Artificial Intelligence." NIST CSRC, 2025-12-16.

[25] Future of Privacy Forum. "SB 5 in Five: What to Know About Connecticut's New AI Law." Future of Privacy Forum, 2026-05-01.

[26] Multistate.ai. "AI Legislation Tracker." Multistate.ai, retrieved 2026-04-24.