CSAI Foundation

Cloud Security Alliance AI Safety Initiative

Federal AI Preemption: Enterprise Compliance and Security

What the White House National AI Policy Framework Means for Security and Compliance Leaders

Unofficial AI-assisted Research

2026-04-11

Key Takeaways

- On March 20, 2026, the White House released a National Policy Framework for Artificial Intelligence along with companion legislative recommendations, providing Congress with a blueprint to establish a unified federal AI regulatory regime and preempt conflicting state laws [1][2].
- The framework is not binding law and does not create new compliance obligations on its own; it is a legislative proposal aimed at Congress. State AI laws — including Colorado's SB 24-205, Texas's TRAIGA, and California's SB 53 — remain in effect unless and until Congress acts or courts intervene [3][4].
- Federal preemption is the framework's central pillar. The administration argues that a fragmented patchwork of state AI laws raises compliance costs, creates legal uncertainty, and undermines U.S. global competitiveness [2][5]. However, opposition has emerged across party lines: the Senate voted 99-1 to strip an earlier 10-year AI moratorium provision from unrelated budget legislation [13], and Democratic members of Congress simultaneously introduced the GUARDRAILS Act to repeal the December 2025 executive order entirely [6].
- The December 11, 2025 Executive Order that preceded the framework established an AI Litigation Task Force within the Department of Justice to challenge state AI laws on constitutional grounds, and directed the FTC and FCC to explore agency-level preemption mechanisms, creating immediate legal risk for states with active AI enforcement programs [7][8].
- Enterprises should interpret the current environment as a period of extended, structured uncertainty: state laws are enforceable today, legal challenges may reshape which laws survive over the next twelve to twenty-four months, and a federal floor standard is likely eventually, though its content and timeline remain unresolved [9][10].
- The framework, read alongside the administration's March 2026 Cyber Strategy for America, elevates AI security to a national security priority, with critical infrastructure operators and AI-enabled enterprises increasingly expected to align with federal cybersecurity baselines as de facto partners in national defense [11].

Background

The regulatory environment governing artificial intelligence in the United States entered a new and turbulent phase in late 2025. Throughout the year, states moved actively to establish AI regulatory frameworks in the absence of federal AI legislation: 1,208 AI-related bills were introduced across all 50 states in 2025, with 145 enacted into law [12]. The resulting landscape is heterogeneous in scope, focus, and legal theory. Colorado's AI Act (SB 24-205), delayed by subsequent legislation to an enforcement date of June 30, 2026, is the most comprehensive state-level regime, requiring high-risk AI developers and deployers to conduct algorithmic impact assessments, disclose AI use to affected consumers, and demonstrate reasonable care against algorithmic discrimination [4]. Texas enacted the Responsible Artificial Intelligence Governance Act (TRAIGA), effective January 1, 2026, which focuses primarily on government agency use of AI and takes an intent-based rather than impact-based approach to private sector liability [4]. California's SB 53, signed September 29, 2025, targets developers of frontier AI models with transparency and incident reporting requirements, rather than regulating AI applications broadly [4].

On December 11, 2025, President Trump signed an executive order titled "Ensuring a National Policy Framework for Artificial Intelligence," signaling a direct federal challenge to this growing body of state regulation [7]. The order directed the Attorney General to establish an AI Litigation Task Force within 30 days, charged with identifying and challenging state AI laws inconsistent with the administration's policy of a "minimally burdensome national policy framework" [7][8]. The order simultaneously directed the FTC to issue a policy statement on how existing federal authority could preempt certain state AI laws, and the FCC to initiate a proceeding on federal AI disclosure standards [8]. Separately, the order conditions access to certain federal grant funding on states' willingness to refrain from enacting laws the administration views as onerous [7].

The March 20, 2026 National Policy Framework translates the executive order's intent into a formal legislative roadmap for Congress. The document is not regulation and carries no enforcement authority, but it signals where the administration intends to push federal law and identifies which categories of state regulation it considers most problematic. On the same day, Democratic legislators introduced the GUARDRAILS Act – formally the Guaranteeing and Upholding Americans' Right to Decide Responsible AI Laws and Standards Act – to repeal the December 2025 executive order and block any federal moratorium on state AI regulation [6]. This legislative response underscores that the preemption debate is not resolved; it is at the beginning of a potentially multi-year legislative and judicial contest.

Security Analysis

The Regulatory Compliance Environment

A significant challenge for enterprise security and compliance leaders is that the federal preemption strategy has created legal uncertainty precisely where certainty is most needed. State AI laws remain on the books and enforceable until a court enjoins enforcement or Congress acts. The AI Litigation Task Force may secure injunctions against specific state laws, but litigation moves slowly: challenges will work through district courts and circuit courts, likely generating conflicting precedents across jurisdictions before any Supreme Court resolution. During this period, organizations operating in multiple states face the practical requirement of either complying with the most stringent applicable state standard or accepting jurisdiction-specific legal risk wherever they operate under less protective standards.

The framework's approach to preemption is also more bounded than early reporting suggested. The administration explicitly recommends that Congress not preempt state authority over children's safety, AI compute and data center infrastructure, or state government procurement and use of AI [2][5]. This means that state laws governing AI in educational settings, child welfare systems, or state contracting — a significant portion of high-risk AI deployment — would survive federal preemption even under the administration's own preferred legislative outcome. Enterprise compliance programs must account for this carve-out architecture, not a clean federal floor across all domains.

The framework's light-touch approach to private sector AI also has a structural implication for how security risk is allocated. Rather than establishing affirmative requirements for AI risk management, the framework relies on existing sectoral regulators — the FTC, financial regulators, healthcare agencies — to address AI-specific harms within their existing authority. This approach preserves regulatory diversity at the sector level even as it attempts to reduce it at the state level. A financial institution deploying AI in credit decisioning thus faces a different and potentially more demanding regulatory stack than a manufacturer deploying the same model in supply chain optimization.

The National Security Dimension

The March 2026 Cyber Strategy for America, released weeks before the National AI Policy Framework, provides critical context for how the administration frames AI security at an organizational level [11]. The strategy commits the federal government to adopting AI-powered cybersecurity solutions to defend federal networks and prioritizes securing the "full AI technology stack, from data centers to the models themselves" [11]. The strategy frames critical infrastructure operators — including energy, finance,

telecommunications, and healthcare entities — as essential to national cyber defense; in CSA's reading, this framing carries implicit expectations around information sharing, AI governance transparency, and supply chain security.

This creates a dual compliance imperative for enterprises in critical infrastructure sectors. Organizations in these sectors must simultaneously demonstrate that their AI governance aligns with emerging federal standards and that their AI-enabled systems meet rising cybersecurity baselines. The Cyber Strategy explicitly emphasizes zero-trust architecture, post-quantum cryptography readiness, and AI-powered threat detection as the modernization trajectory for both federal and aligned private sector networks [11]. For organizations already on federal contracting vehicles or operating in regulated critical infrastructure sectors, these are not aspirational signals — they are the direction of future procurement and regulatory expectations.

AI tools deployed in legally sensitive functions — document review, contract analysis, hiring screening, fraud detection — create a distinct category of cybersecurity risk under this framework. These functions are precisely those most likely to be subject to conflicting state-level AI obligations (Colorado and California specifically target algorithmic discrimination in high-stakes decisions) and most likely to generate regulatory enforcement exposure if AI systems produce outputs that cause consumer harm. Where a single AI system serves users across multiple states with different disclosure and impact assessment requirements, the exposure is not only technical but legal: a breach, hallucination incident, or discriminatory output may trigger regulatory investigation under multiple state regimes simultaneously, compounding incident response complexity.

The Preemption Risk to State-Level Protections

Security practitioners should understand that federal preemption, if achieved, would eliminate not only burdens but also protections that currently exist in state law. Several state AI laws include consumer notification requirements for consequential AI-driven decisions, mandatory human oversight provisions for high-risk systems, and data minimization standards for AI training pipelines [3][4]. An enterprise that has built compliance programs around these state-level requirements — accepting them as the operative standard — may face a period of regulatory ambiguity if those laws are enjoined or preempted, during which the applicable standard is unclear and legal exposure uncertain.

The absence of a federal floor containing comparable protective requirements means that preemption without replacement legislation would leave consumers with reduced protections and enterprises with less clarity about their obligations, not more. CSA's assessment is that this transitional ambiguity, rather than ultimate preemption outcomes, represents the near-term governance risk enterprises should prioritize planning for.

Recommendations

Immediate Actions

Enterprises should conduct a jurisdictional inventory of their AI deployments, mapping each system against the state laws currently applicable to its operation: Colorado's SB 24-205 enforcement (effective June 30, 2026), Texas TRAIGA (effective January 1, 2026), California's SB 53 frontier model requirements, and any other state laws enacted in their operating jurisdictions. This inventory should flag systems performing consequential decisions — credit, employment, housing, healthcare — as highest priority for compliance review given their prominence in both state and potential federal legislative attention.

Legal and compliance teams should engage outside counsel with AI regulatory practices to establish monitoring protocols for AI Litigation Task Force enforcement actions. Any court injunction against a state AI law applicable to the organization's operations would materially change that compliance obligation; monitoring systems for injunctions and stays should be active now, not reactive to press coverage. Similarly, any FTC or FCC rulemaking activity initiated under the December 2025 executive order should trigger internal legal review of potential preemptive effect.

Short-Term Mitigations

Compliance programs should be designed for jurisdictional modularity wherever operationally feasible. Rather than building to the standard of any single state regime or waiting for federal clarity, organizations should identify the core practices — impact assessments, consumer disclosure mechanisms, human oversight checkpoints for high-risk decisions — that satisfy the most protective applicable state requirements and would satisfy any plausible federal minimum. Building to this intersection maximizes regulatory durability across preemption outcomes, because it satisfies both state law as-is and would satisfy federal standards if and when enacted.

Security teams should assess whether AI systems deployed in legally sensitive functions are appropriately isolated and monitored. The compliance risk from AI errors in these domains is not merely reputational; it is direct regulatory exposure, and multi-state regulatory regimes compound that exposure. Investments in AI observability — logging inputs, outputs, and decision audit trails — should be treated as both a security control and a strong foundation for demonstrating compliance with impact assessment and disclosure obligations under applicable state law.

Organizations aligning with federal contracting or critical infrastructure frameworks should begin gap analysis against the Cyber Strategy's stated modernization priorities: zero-trust architecture, post-quantum cryptography readiness, and AI-powered security tooling. These are not immediate mandates, but they are the direction of evolving requirements, and gap closure requires time. Treating them as a 12-24 month roadmap item rather than a distant aspiration is appropriate given the pace of administration signaling.

Strategic Considerations

The federal preemption debate is likely to continue for at least the next two to three years. Neither a comprehensive federal AI law with preemption provisions nor a Supreme Court resolution of the constitutional questions is imminent. Enterprises should institutionalize a regulatory intelligence function — whether internal or through industry associations — that tracks state legislative sessions, DOJ task force litigation, and Congressional AI legislation on an ongoing basis rather than responding episodically. The compliance landscape will continue to shift, and organizations that develop durable monitoring capacity will be better positioned than those that attempt to stabilize their programs around any single predicted outcome.

Engagement with industry coalitions and standards bodies provides both intelligence and influence. The administration's framework explicitly relies on voluntary standards and industry self-regulation as alternatives to prescriptive regulation, citing the Frontier Model Forum and Coalition for Content Provenance and Authenticity (C2PA) as examples of useful technical inputs [1]. Organizations that participate in standards development have earlier visibility into emerging requirements and opportunities to shape frameworks that are operationally feasible. CSA's AICM is directly relevant here: it provides a structured, assessable control framework that organizations can use to demonstrate AI governance maturity regardless of which regulatory requirements ultimately emerge at the federal or state level.

CSA Resource Alignment

The regulatory environment described in this note intersects directly with several CSA frameworks and publications that enterprise security and compliance teams should be using to structure their AI governance programs.

The **AI Controls Matrix (AICM)** provides an 18-domain control framework that maps AI security governance requirements across model providers, application providers, orchestrated service providers, and AI customers. Because the AICM is a cross-jurisdictional framework grounded in security principles

rather than any specific legal regime, it provides durable compliance scaffolding that remains valid across the preemption scenarios described in this note. Organizations building AICM-aligned programs now will likely find their control evidence applicable to any federal standard the administration ultimately proposes, since the framework's governance, risk management, and transparency domains align with the categories CSA expects to appear in federal AI legislation.

The **MAESTRO** threat modeling framework for agentic AI is particularly relevant to organizations deploying AI in legally sensitive or high-stakes functions. MAESTRO's structured analysis of trust boundaries, data flows, and adversarial interaction points in agentic pipelines complements the compliance requirements in state AI laws that mandate risk assessments for high-risk AI systems. Organizations conducting algorithmic impact assessments under Colorado SB 24-205 or similar frameworks can use MAESTRO to ensure that their technical risk identification is thorough enough to satisfy regulatory review.

The **Zero Trust guidance** published by CSA aligns directly with the Cyber Strategy for America's emphasis on zero-trust architecture modernization. Critical infrastructure operators responding to federal cybersecurity expectations should be using CSA's Zero Trust Maturity Model as a baseline, particularly for networks and systems where AI tools are now integrated into monitoring, detection, or response workflows. The convergence of AI deployment and zero-trust architecture requirements means that these are not separate workstreams; they are the same modernization initiative viewed from different angles.

Finally, CSA's **AI Organizational Responsibilities** guidance addresses the governance accountability structures that both state AI laws and any plausible federal standard will require: designated responsibility for AI risk oversight, documented review processes for AI-driven decisions, and clear escalation paths when AI systems produce anomalous outputs. These structural requirements are jurisdiction-agnostic and provide a strong baseline governance posture for enterprises operating AI systems at scale, complementing peer frameworks such as the NIST AI RMF and ISO 42001.

References

[1] White House. "President Donald J. Trump Unveils National AI Legislative Framework." WhiteHouse.gov, March 20, 2026.

[2] White House. "National Policy Framework for Artificial Intelligence – Legislative Recommendations." WhiteHouse.gov, March 20, 2026.

[3] Crowell & Moring LLP. "White House National AI Policy Framework Calls for Preempting State Laws, Protecting Children." Crowell.com, March 2026.

[4] Swept AI. "State AI Regulations in 2026: Colorado, Texas, California, and What's Coming." Swept.ai, 2026.

[5] Ropes & Gray LLP. "The White House Legislative Recommendations: National Policy Framework for Artificial Intelligence and Federal Preemption of State AI Laws." RopesGray.com, March 2026.

[6] Rep. Don Beyer. "Beyer, Matsui, Lieu, Jacobs, McClain Delaney Introduce Legislation to Repeal White House AI Moratorium." Beyer.house.gov, March 20, 2026.

[7] White House. "Ensuring a National Policy Framework for Artificial Intelligence." WhiteHouse.gov, December 11, 2025.

[8] Sidley Austin LLP. "Unpacking the December 11, 2025 Executive Order: Ensuring a National Policy Framework for Artificial Intelligence." Sidley.com, December 2025.

[9] Akin Gump. "White House Releases Long-Awaited Artificial Intelligence Framework, Setting the Stage for Federal Preemption Debate and Further Legislative Action." AkinGump.com, March 2026.

[10] Ropes & Gray LLP. "Examining the Landscape and Limitations of the Federal Push to Override State AI Regulation." RopesGray.com, March 2026.

[11] White House. "President Trump's Cyber Strategy for America." WhiteHouse.gov, March 2026.

[12] National Conference of State Legislatures. "Artificial Intelligence 2025 Legislation." NCSL.org, 2025.

[13] U.S. Senate Commerce Committee. "Senate Strikes AI Moratorium from Budget Reconciliation Bill in Overwhelming 99-1 Vote." Commerce.Senate.gov, July 1, 2025.