CSAI Foundation | Cloud Security Alliance

The Great American AI Act: A 2026 Compliance Problem

Federal Preemption Uncertainty Doesn't Pause State AI Compliance Obligations

2026-07-22

AI-assisted Rapid Research

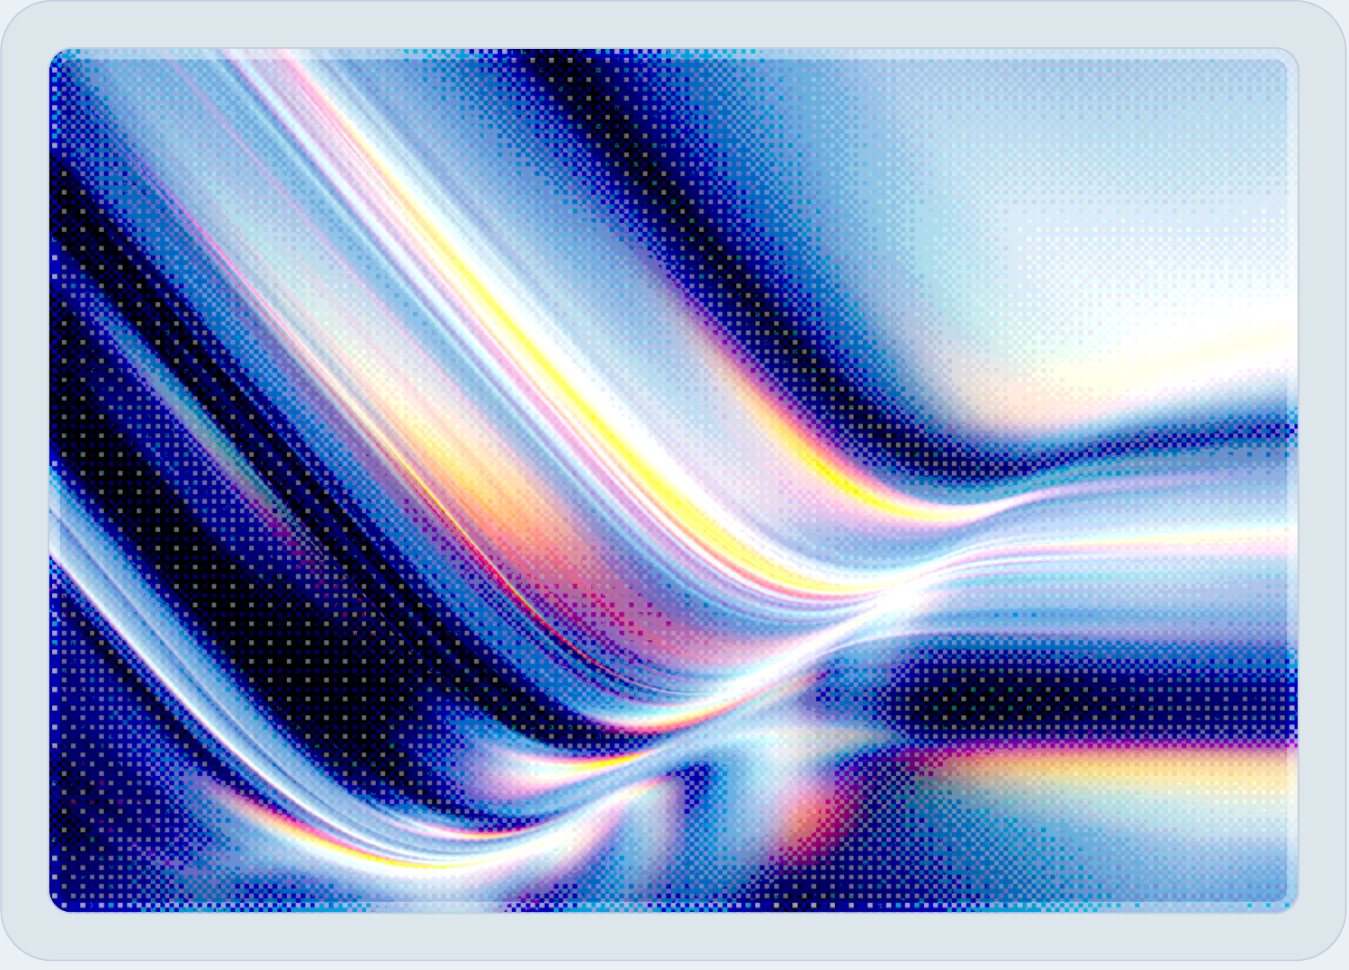

CSAI | CSA cloud security alliance®

Key Takeaways

On June 4, 2026, Representatives Jay Obernolte (R-Calif.) and Lori Trahan (D-Mass.) released a 269-page bipartisan discussion draft of the Great American Artificial Intelligence Act (GAAIA) — the first bipartisan discussion draft to pair binding frontier-developer obligations with a federal preemption of state AI law [1][5][6]. Its most contested provision — a three-year preemption of state laws "specifically regulating the development" of AI models, sunsetting around December 2029 — has drawn both praise as a needed federal floor and sharp criticism that its scope is legally undefined and could eliminate state frontier-safety statutes in California, New York, and Illinois before Congress has finished writing a durable replacement [1][2][3]. As of early July 2026, the bill remains a discussion draft that has not been formally introduced, and its sponsors are still soliciting stakeholder feedback with no set timeline for markup or a floor vote [10]. That legislative limbo is precisely the point security and compliance leaders should not misread: GAAIA's preemption clause, even if enacted exactly as drafted, would not touch the state deployment, privacy, consumer-protection, and sector-specific laws that already govern most enterprise AI use, and it would do nothing to stop states from continuing to amend their own frontier statutes in the meantime, as Colorado did in May 2026 when it replaced its original AI Act with a narrower successor [8]. The practical conclusion for enterprise risk programs is that the outcome of the Washington preemption fight is not a gating condition for compliance work; the obligations that matter for most organizations in 2026 arise from law that is already in force and continuing to change independent of what Congress ultimately does with GAAIA.

Background

GAAIA is organized into four titles. Title I, Frontier AI Governance, imposes binding obligations on "large frontier developers" — companies with more than $500 million in annual gross revenue that have trained qualifying frontier models, a group that captures OpenAI, Anthropic, Google, Meta, and xAI — while a lighter set of disclosure obligations applies to a broader tier of "frontier developers" with revenue above $50 million [1][3]. Covered developers must publish and annually update a frontier AI safety framework addressing risk thresholds, cybersecurity, and governance; issue transparency reports before or concurrent with model deployment; report "critical safety incidents" to a federal body within 15 days, or 24 hours where there is imminent risk of death or injury; and submit to semiannual third-party compliance verification by licensed Independent Verification Organizations [1][3]. Enforcement runs

through both federal and state attorneys general, with penalties reaching $1 million per violation per day, and the bill includes whistleblower protections offering reinstatement, double back pay with interest, and compensatory damages for employees who report violations [1]. Title I also funds the Center for AI Standards and Innovation (CAISI) at $100 million annually — what Rep. Trahan has characterized as a "tenfold" increase over its current $15 million budget, though the two figures work out to roughly a sevenfold increase — positioning CAISI as the primary federal body for frontier AI oversight [1][7]. Titles II through IV address AI's labor-market effects, cybersecurity information sharing, and international coordination, and are comparatively less contested than Title I [1].

The preemption language in Title I is what has generated the most substantive debate. As drafted, it displaces any state law "specifically regulating the development of an artificial intelligence model" for three years, but includes savings clauses preserving state authority over AI deployment and use, laws of general applicability, and existing privacy and consumer-protection statutes [1][3]. Rep. Trahan has described this as a considered trade: "I will preempt, but only if we're setting the strongest possible federal standard" [1]. Supporters, including the Business Software Alliance and the Information Technology Industry Council, along with a bipartisan group of cosponsors, argue that a single federal framework with real rulemaking authority is preferable to an accumulating patchwork of state statutes that no single company can practically track [1]. Opponents — including Public Citizen, Public Knowledge, the AFL-CIO, and the co-chairs of the House's bipartisan AI Commission, who stated the draft "cannot serve as the basis for productive dialogue" — argue the preemption sweeps far more broadly than the frontier-safety space it is nominally targeted at [1]. On June 18, 2026, the Consumer Federation of America joined a coalition of more than 130 organizations in a letter urging Congress to reject the bill outright, arguing that "a strong federal floor" should not become "a ceiling that ties [states'] hands through the most consequential years this technology will ever have" [4].

Security Analysis

The Preemption Clause's Legal Ambiguity Is the Real Compliance Risk

The specific phrase Congress chose to draw the preemption boundary — laws that "specifically regulate the development" of an AI model — has no settled meaning in existing case law, and that ambiguity is itself a governance risk rather than a drafting footnote. Writing in Lawfare, policy analyst Charlie Bullock argues that courts applying functional-interpretation precedents such as *National Meat Association v. Harris* are likely to preempt any state law that functions primarily to regulate development, "even if creatively drafted otherwise" [2]. Under that reading, the preemption clause would reach well beyond the frontier-safety statutes GAAIA's sponsors say they are targeting, potentially voiding Texas's TRAIGA

provisions on child-impersonation AI, California's AB 2013, and various state chatbot-safety laws that Congress has not otherwise addressed [2]. The Future of Privacy Forum's comparative analysis identifies California's SB 53 (the Transparency in Frontier Artificial Intelligence Act), New York's RAISE Act, and Illinois's SB 315 as the statutes most exposed to preemption, precisely because their frontier-development focus most closely mirrors GAAIA's own subject matter [3]. Bullock's proposed fix — narrowing preemption to a one-to-one basis that retires specific state frontier statutes only in exchange for an equivalent federal standard — has reportedly found some receptiveness in Trahan's office, but as of this writing it has not been incorporated into the public draft [2].

For enterprise security and compliance programs, the significance of this ambiguity is that it is unlikely to be resolved by closer reading of the bill text alone; based on Bullock's analysis, resolution would most plausibly come through litigation or a substantially revised draft, and neither appears likely on a predictable timeline [2]. A compliance program built on the assumption that a particular state obligation will or will not survive preemption is a program built on an outcome that will be decided by federal courts, not by Congress, on a timeline compliance teams cannot control.

GAAIA's Scope Is Narrower Than Its Political Profile Suggests

A separate risk is that GAAIA's high-profile preemption fight leads compliance teams to overestimate how much the bill would actually change for a typical enterprise. GAAIA's substantive obligations — the frontier safety framework, the transparency reports, the semiannual independent audits — apply only to the small population of developers that train frontier models above the bill's revenue thresholds. As one legal analysis of the draft observes, "for most companies using AI models in their daily operations, these requirements will not apply," because the risks a typical enterprise carries arise not from developing frontier models but from deploying third-party AI tools inside regulated business functions — a category of activity the preemption clause explicitly does not reach [9]. The bill's practical effect on that larger population of AI deployers is concentrated in two places that one legal analysis flags as underexamined in most GAAIA commentary [9]: enhanced fraud penalties for AI-assisted mail, wire, and bank fraud, which raise the stakes on unsubstantiated AI-capability claims in marketing and investor disclosures, and Title II's workforce-notification provisions, which would require 60 days' advance notice when AI is a "substantial factor" in mass layoffs [9][1].

State Law Is Not Standing Still While Congress Deliberates

The federal preemption debate is also not occurring against a static state backdrop. In May 2026, Colorado significantly narrowed its own AI Act, replacing the original SB 24-205 with SB 26-189, which eliminated the mandatory risk-management programs, annual impact assessments, and discrimination-prevention duties of the earlier law while retaining a narrower core of consumer notice, explanation rights

for adverse automated decisions, meaningful human review, and developer documentation obligations [8]. That amendment happened entirely independent of GAAIA and illustrates a dynamic compliance teams need to plan around regardless of what happens in Washington: state AI law in 2026 is not a fixed target that federal preemption will simply freeze in place. It is being actively rewritten by state legislatures on a timeline unrelated to the federal discussion draft, which means organizations anchoring their compliance architecture to the specific requirements of any one state statute — rather than to portable, framework-level controls — are exposed to a second axis of change even if the preemption question never resolves.

Recommendations

Immediate Actions

Security and compliance leaders should complete or update an inventory of every AI system in production, tagged by whether the organization is a developer or a deployer of the underlying model, by the model developer's approximate revenue tier, and by which state's residents or operations the system touches. This inventory should be built now, independent of GAAIA's outcome, because it is also the foundation for compliance with the state laws — California SB 53, New York's RAISE Act, Illinois SB 315, Colorado SB 26-189, Texas TRAIGA — that are already in force or on a near-term effective-date clock regardless of federal action. Teams should also flag any external-facing statement about AI capabilities, whether in marketing material or investor disclosure, for legal review given GAAIA's enhanced fraud penalties and the broader enforcement trend toward treating unsubstantiated AI claims as a securities and consumer-protection risk independent of the bill's passage [9].

Short-Term Mitigations

Organizations should strengthen AI vendor contracts to include audit rights, incident-notification requirements, and clear risk-allocation provisions, since responsibility for AI-assisted decisions in regulated functions stays with the deploying organization regardless of what a vendor's underlying model developer is or is not required to disclose under GAAIA [9]. Compliance documentation should be built for portability from the outset — mapped to jurisdiction-agnostic frameworks such as the NIST AI Risk Management Framework and ISO/IEC 42001 — rather than organized around the specific text of any single state or federal statute, precisely because both the state patchwork and the federal draft are still moving. Legal and policy teams should track the preemption clause's language changes across future

drafts specifically for whether it moves toward Bullock's proposed one-to-one narrowing, since that single drafting choice could determine whether existing state frontier-safety compliance work is preserved or discarded [2].

Strategic Considerations

At a program level, the durable lesson from GAAIA's preemption fight is that enterprise AI governance architecture should be designed to survive multiple divergent outcomes, not optimized for the outcome currently considered most likely. That means treating AI governance as an ongoing operational capability with recurring model inventory, documentation, and vendor-review cycles, rather than as a one-time compliance sprint tied to any single bill's enactment date. Security leaders briefing executives and boards on this topic should be explicit that "wait and see" is not a defensible posture: the compliance obligations that already apply to most AI deployments come from law already in force, the preemption clause's own scope will likely take years of litigation to settle even after enactment, and the states most affected by the preemption fight have shown in 2026 that they will continue amending their own AI statutes on an independent timeline.

CSA Resource Alignment

This analysis extends two CSA research notes published earlier in 2026 that anticipated much of the dynamic GAAIA has now made concrete. *AI Preemption Battleground: Federal Framework vs. State Regulation* (April 13, 2026) named what it called a "compliance uncertainty paradox" — organizations cannot build to a federal standard that does not yet exist, cannot ignore state enforcement deadlines that are already running, and cannot safely over-invest in state-specific programs that a future federal law might render moot [11]. That framing directly anticipates GAAIA's current posture: a discussion draft with no enactment timeline, sitting alongside state statutes with real effective dates. Its companion note, *State AI Laws Take Hold as Federal Preemption Stalls* (April 4, 2026), documented that federal preemption efforts had already failed once in 2025 — the Senate rejected a proposed ten-year preemption moratorium that July — while 145 state AI laws were enacted across 38 states in the same year, and concluded that enterprises should comply with existing state law now rather than defer action while awaiting an uncertain federal outcome [12]. GAAIA's discussion-draft status five months later is consistent with that earlier conclusion rather than superseding it.

CSA's *Federal AI Security Mandates: CISO Action Guide* (June 29, 2026) is also relevant, though its subject is the June 2026 executive order on AI security rather than GAAIA specifically. Its central argument — that nominally voluntary federal frameworks harden into de facto mandates through

procurement, insurance, and customer-contract pressure well before they become binding law — describes the same multi-jurisdiction compliance dynamic at issue here: enterprises should expect federal AI policy to shape enterprise obligations gradually and unevenly, through channels other than statutory enactment, long before any preemption question is settled in court [13].

Across all three notes, CSA's AI Controls Matrix (AICM) v1.1 is the consistent recommended anchor for organizations navigating this uncertainty. Its 247 control objectives across 18 security domains map to ISO/IEC 42001, the NIST AI Risk Management Framework, and the EU AI Act, giving compliance teams a single control set that remains useful under a federally preempted regime, an unpreempted state patchwork, or — as is most likely in the near term — an unresolved mixture of both [11][12][14].

References

[1] Ramsha Ejaz. "Unpacking the Great American Artificial Intelligence Act of 2026." Tech Policy Press, June 2026.

[2] Charlie Bullock. "Congress Should Do Something: The Case for (Fixing) the Great American AI Act." Lawfare, June 2026.

[3] Future of Privacy Forum. "Frontier AI Goes Federal: How the Great American AI Act Compares to State Laws." FPF Blog, June 2026.

[4] Consumer Federation of America. "CFA Joins Letter Calling on Congress to Reject the Great American Artificial Intelligence Act." CFA, June 18, 2026.

[5] Roll Call. "Bipartisan AI draft proposes three-year preemption of state laws." Roll Call, June 4, 2026.

[6] Rep. Jay Obernolte. "Obernolte, Trahan Release a Discussion Draft of the Great American AI Act." U.S. House of Representatives, June 4, 2026.

[7] Rep. Lori Trahan. "House Panel Approves Key Trahan AI Priorities." U.S. House of Representatives, June 2026.

[8] AI Accelerator Institute. "The US AI Rulebook Is Being Rewritten. Your Compliance Team Can't Wait.." AI Accelerator Institute, 2026.

[9] Subject to Inquiry. "The Great American AI Act: What It Means — and Doesn't Mean — for Companies Using AI." Subject to Inquiry, June 2026.

[10] Mintz. "AI: The Washington Report — July 2026 Edition." Mintz, July 8, 2026.

[11] Cloud Security Alliance. "AI Preemption Battleground: Federal Framework vs. State Regulation." CSA Lab Space, April 13, 2026.

[12] Cloud Security Alliance. "State AI Laws Take Hold as Federal Preemption Stalls." CSA Lab Space, April 4, 2026.

[13] Cloud Security Alliance. "Federal AI Security Mandates: CISO Action Guide." CSA Lab Space, June 29, 2026.

[14] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.