CSAI Foundation | Cloud Security Alliance

# Texas's TRAIGA Complaint Portal Turns Compliance Into Enforcement

2026-09-13

AI-assisted Rapid Research

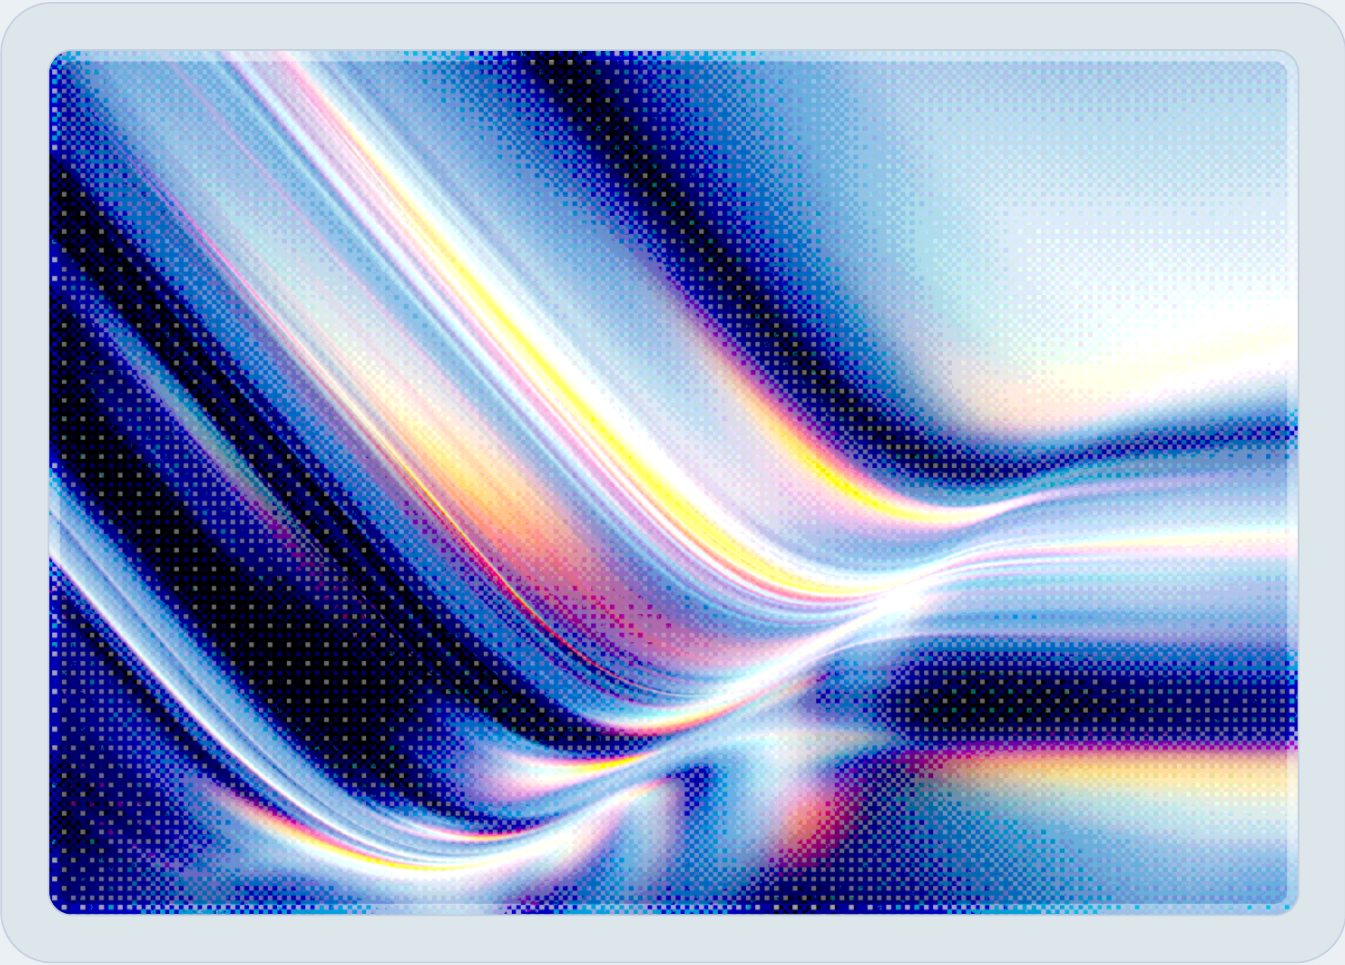

CSAI

CSA cloud security alliance®

## Key Takeaways

On September 1, 2026, the Texas Attorney General activated the online complaint mechanism required by the Texas Responsible Artificial Intelligence Governance Act (TRAIGA, House Bill 149), moving the state's AI law from a period with no operative enforcement channel into an active enforcement pipeline [1][3]. The portal is not administrative housekeeping; under Section 552.103 of the Texas Business and Commerce Code, a consumer complaint filed through the mechanism is the statutory precondition that permits, but does not automatically compel, the Attorney General to issue a civil investigative demand, meaning a single web form submission from a Texas resident can now be the event that opens the door to a formal inquiry, at the Attorney General's discretion [3][7]. TRAIGA took effect January 1, 2026, but the eight-month gap before the portal's launch meant that exposure through this complaint-driven pathway was largely theoretical until now; that gap has closed [2][3]. Organizations that built compliance programs anticipating a broad, EU-AI-Act-style risk-management mandate may find TRAIGA's actual reach substantially narrower: the statute prohibits five specific categories of AI use rather than mandating broad risk-assessment documentation, and violations premised on discrimination require proof of intent, with disparate impact alone insufficient to establish a violation [2]. Organizations that substantially comply with the NIST AI Risk Management Framework's Generative AI Profile, or an equivalent recognized framework, gain an affirmative defense under Section 552.105(e), which makes existing AI governance investment directly relevant to TRAIGA exposure even though the statute does not require it upfront [2][7]. Security, legal, and compliance teams whose AI systems reach Texas residents should treat this launch as the point at which documentation practices, incident handling, and complaint-response readiness move from aspirational to operationally necessary.

## Background

TRAIGA was enacted as House Bill 149 during the 89th Texas Legislature and signed into law in 2025, taking effect January 1, 2026 [2][4][7]. In CSA's assessment, the final statute represents a considerably narrower approach than the broad AI risk-management frameworks originally contemplated in earlier legislative drafts and in comparable statutes such as the EU AI Act. Rather than imposing comprehensive impact-assessment and documentation obligations on all AI developers and deployers, the enacted law functions closer to a targeted misuse statute: it prohibits four specific practices for private companies, separately restricts government use of AI for social scoring, and layers a set of disclosure duties onto

government agencies that deploy consumer-facing AI [2]. Private companies operating in or affecting Texas residents must avoid AI systems that intentionally incite self-harm, harm to others, or criminal conduct; that are designed with the sole intent of infringing a constitutionally protected right; that are designed to unlawfully discriminate against a protected class; or that produce child sexual abuse material under Texas Penal Code Section 43.26 or unlawful deepfakes under Texas Penal Code Section 21.165 [2]. TRAIGA separately bars government agencies specifically from using AI systems for social scoring that results in detrimental or unjustified treatment of individuals, a prohibition that does not extend to the private sector [2]. Government agencies bear a heavier compliance burden more broadly, too: any state agency offering a consumer-facing AI interaction must disclose, clearly and conspicuously and before or at the time of interaction, that the person is dealing with an AI system, with a narrow carve-out permitting healthcare providers to delay disclosure until treatment begins outside of emergencies [2].

TRAIGA's obligations attach differently depending on an organization's role in the AI supply chain. The statute defines a "developer" as an entity that develops an AI system offered, sold, leased, given, or otherwise provided in Texas, and a "deployer" as an entity that deploys an AI system for use in Texas, a distinction that TRAIGA applies to a narrower set of prohibited practices rather than a general risk-management regime [4]. Because the prohibited-use provisions apply regardless of whether an organization built the underlying model or simply integrated a third-party AI system into a Texas-facing product, deployers cannot rely on a vendor's representations alone; a company that licenses a general-purpose model and configures it in a way that produces a prohibited outcome bears its own exposure under the statute. This makes vendor due diligence and contractual allocation of TRAIGA-related risk a practical necessity for any organization embedding third-party AI into consumer-facing Texas products, even where the organization did not train the underlying model itself.

Enforcement authority under TRAIGA is exclusive to the Texas Attorney General; the statute creates no private right of action, so individuals harmed by a violation cannot sue directly and must instead route concerns through the state [2][3]. This design placed particular weight on the mechanism Section 8 of HB 149 directed the Attorney General to build: an online complaint intake that the statute required to be live no later than September 1, 2026 [3][7]. That deadline has now been met. The Attorney General's office maintains a "Consumer AI Rights" page within its Consumer Protection division that provides a plain-language summary of TRAIGA's prohibitions and a direct complaint-filing pathway, restricted to Texas residents acting in an individual or household capacity rather than businesses, employees, or competitors [1][3]. The practical effect is that Texas now has a channel with a low barrier to filing, through which any consumer frustration with a chatbot, hiring tool, or automated decision system can be converted into a documented complaint, and documented complaints are what the statute requires before the state's investigative machinery can engage [3].

# Security Analysis

The complaint-to-investigation pathway under TRAIGA follows a defined sequence that security and compliance teams should map against their own incident-response and documentation practices. A complaint filed through the portal permits, but does not require, the Attorney General to issue a civil investigative demand under Section 552.103, which can request a wide range of material: the AI system's stated purpose, the categories of training data used, input and output data types, performance metrics, known limitations, and the safeguards in place to monitor and correct problems [3][7]. If the Attorney General concludes a violation occurred, the office must issue a written notice identifying the specific statutory provisions at issue, after which the organization receives a mandatory 60-day cure period; violations remedied within that window incur no penalty [2][3][7]. Penalties for violations that are not cured, or that the Attorney General deems uncurable, follow a tiered structure: curable violations carry fines of $10,000 to $12,000 each, uncurable violations range from $80,000 to $200,000 each, and continuing violations accrue $2,000 to $40,000 per day, with licensed professionals facing additional suspension or revocation exposure [2][3][7].

Two features of this framework shift where organizational risk concentrates. First, because the discrimination prohibition is intent-based rather than outcomes-based, the evidentiary battleground is not primarily statistical disparate-impact analysis but internal documentation: design records, system prompts and their revision history, model cards, internal evaluation reports, and marketing or sales claims that could be read as evidence of discriminatory intent [3]. Organizations that maintain clean, contemporaneous records of why a system was designed the way it was, and that avoid overstated marketing language about what an AI system can determine about a person, are better positioned to rebut an intent-based claim than organizations relying solely on outcome statistics. Second, TRAIGA's Section 552.105(e) affirmative defenses reward organizations that already invest in recognized governance practices: substantial compliance with the current NIST AI Risk Management Framework's Generative AI Profile (NIST-AI-600-1) or another nationally or internationally recognized AI risk-management framework constitutes a defense, as does discovering and correcting a violation through internal testing, red-teaming, or adversarial testing before it is flagged externally, or discovering it through a documented internal review process [2][7]. In CSA's experience, much of this governance activity — internal testing, red-teaming, documented review — already occurs informally within security teams; TRAIGA's affirmative defense creates a direct incentive to formalize and log it, since the defense depends on being able to demonstrate the practice existed and was followed, not merely that it could have been.

The portal's low barrier to entry also changes the practical calculus around which complaints are likely to arrive first. Because any Texas resident acting in an individual capacity can file through the web mechanism without legal representation, and because the underlying statute is framed in consumer-

protection terms similar to existing deceptive-trade-practice law, the most likely early complaints are plain-language consumer grievances about a specific interaction with a chatbot, hiring tool, or automated support system rather than sophisticated legal filings [2][3]. Organizations should not assume that the technical narrowness of TRAIGA's five prohibited categories will limit the volume or noise of complaints reaching the Attorney General's office; a complaint alleging that a system "tried to talk someone into hurting themselves" or "discriminated against me" can be the event that opens a civil-investigative-demand inquiry at the Attorney General's discretion, regardless of whether the underlying conduct ultimately proves to be a statutory violation, meaning organizations should expect to respond to inquiries that require distinguishing genuine violations from complaints that do not meet the statute's intent-based threshold.

Texas is not acting in isolation, and its narrower, prohibition-based model is one point in a widening and inconsistent set of state approaches to AI regulation. Colorado's newly enacted Chatbot Safety Act imposes disclosure, age-estimation, and self-harm response obligations on conversational AI operators ahead of a January 1, 2027 effective date, backed by per-violation civil penalties of up to $20,000 with no cap on aggregate liability across multiple violations, and, like TRAIGA, exclusive attorney-general enforcement with no private right of action [6]. Illinois's SB 315 (the Artificial Intelligence Safety Measures Act) takes yet another approach, requiring the largest frontier AI developers to maintain a public safety framework, report critical safety incidents within 72 hours, and submit to annual third-party audits starting in 2028 [5]. The result is a compliance patchwork in which the same underlying AI system can face prohibition-based liability in Texas, disclosure-and-age-estimation obligations in Colorado, and audit-and-incident-reporting duties in Illinois, each enforced by a different state authority on a different timeline. Enterprises that build Texas-specific compliance in isolation risk both under-protecting against obligations that look similar across states and over-building documentation that TRAIGA itself does not require; organizations that modeled AI governance on earlier, broader draft versions of TRAIGA or on the EU AI Act may now find themselves maintaining documentation the enacted statute does not actually require [2].

# Recommendations

## Immediate Actions

Organizations whose AI systems are accessible to Texas residents should confirm, without delay, whether any current AI use case plausibly touches one of TRAIGA's five prohibited categories, particularly systems involved in employment, lending, or other decisions that could be characterized as designed to discriminate, and should inventory the design documentation, system prompts, and evaluation records

that would need to be produced in response to a civil investigative demand. Legal and compliance teams should also establish a clear internal owner for any TRAIGA complaint or Attorney General inquiry, given the mandatory 60-day cure window that begins running from formal notice.

### Short-Term Mitigations

Compliance and security teams should formalize whatever AI risk-management practice already exists into documented form aligned with the NIST AI Risk Management Framework's Generative AI Profile or an equivalent recognized framework, since substantial, demonstrable compliance is an explicit affirmative defense under Section 552.105(e) [2][7]. Internal testing, red-teaming, and adversarial evaluation of AI systems should be formalized and logged as an ongoing practice rather than an ad hoc exercise, both because it surfaces problems before a consumer complaint does and because documented internal discovery is itself a statutory defense. Marketing, sales, and product materials describing AI system capabilities should be reviewed for language that could later be read as evidence of intent to discriminate or to infringe protected rights.

### Strategic Considerations

Given the multi-state patchwork now forming around Texas, Colorado, and Illinois, organizations should design AI governance programs around portable controls that map to the strictest applicable obligation across jurisdictions rather than building statute-specific silos, an approach consistent with CSA's broader analysis of the state frontier AI regulatory landscape. Because TRAIGA's complaint portal materially lowers the friction for a complaint to reach the Attorney General's office, organizations should treat consumer-facing AI incident response and complaint-handling as a governance function with the same rigor applied to security incident response, including defined escalation paths, evidence preservation, and legal notification procedures the moment a Texas-originated complaint or inquiry is received.

## CSA Resource Alignment

CSA's *Colorado's Chatbot Safety Act: A New Compliance Floor* is the most directly applicable prior CSA analysis, since Colorado's statute shares TRAIGA's structural features of exclusive attorney-general enforcement, no private right of action, and reliance on organizational AI governance maturity as the strongest predictor of compliance readiness [6]. Organizations building the "absorptive" governance capability that report recommends for Colorado will find much of that same capability directly applicable to Texas. More broadly, CSA's ongoing analysis of the multi-state AI regulatory patchwork — spanning Texas's prohibition-based model, Colorado's disclosure-and-age-estimation regime, and Illinois's audit-

and-incident-reporting framework — recommends designing portable governance controls that map to the strictest applicable obligation across jurisdictions rather than building statute-specific silos. The portal launch analyzed in this note operationalizes exactly the kind of enforcement risk that patchwork analysis anticipated: a comparative regulatory landscape has become an active complaint intake with a defined investigative sequence in at least one state. The AI Controls Matrix (AICM v1.1) is the underlying control framework for operationalizing these obligations across all three states; AICM's transparency, accountability, and governance domains provide a structured way to translate TRAIGA's documentation and disclosure expectations, and its intent-focused discrimination provisions, into auditable internal controls rather than one-off legal responses [6].

# References


[1] Office of the Texas Attorney General. "Consumer AI Rights." Texas Attorney General, 2026.

[2] Secure Privacy. "Texas TRAIGA Compliance Requirements: What the 2026 AI Law Really Covers." Secure Privacy Blog, 2026.

[3] ComplianceHub.Wiki. "Texas TRAIGA Complaint Portal Opens September 2026." ComplianceHub.Wiki, September 2026.

[4] Norton Rose Fulbright. "The Texas Responsible AI Governance Act." Norton Rose Fulbright, 2026.

[5] Norton Rose Fulbright. "Illinois Enacts Artificial Intelligence Safety Measures Act." Norton Rose Fulbright, July 2026.

[6] Cloud Security Alliance. "Colorado's Chatbot Safety Act: A New Compliance Floor." Cloud Security Alliance, July 2026.

[7] Texas Senate Research Center. "Bill Analysis: C.S.H.B. 149, Texas Responsible Artificial Intelligence Governance Act." Texas Legislature, 89th Legislature, 2025.


 