

Texas governor signs Responsible AI Governance Act

June 26, 2025 | [Client Update](#) | 4-minute read

The Texas Responsible AI Governance Act that will go into effect in 2026 is a significant departure from the comprehensive legislation first introduced in 2024, highlighting continued challenges of state-level AI regulation.

Background

On June 22, 2025, Texas Governor Greg Abbott signed the Texas Responsible Artificial Intelligence Governance Act (TRAIGA) into law. TRAIGA prohibits intentional development or deployment of AI to discriminate, impair constitutional rights, or incite harmful or criminal acts; establishes a penalty structure for violations as well as a 36-month regulatory “sandbox” with legal protections for AI developers, and requires companies to provide the Texas Attorney General, upon request, “high-level” information regarding their AI systems, including descriptions of training data, performance metrics, and post-deployment monitoring and user safeguards. Although this update focuses on private sector implications, TRAIGA also establishes various limits on the Texas state government’s use of AI.

When first introduced in December 2024, the draft TRAIGA reflected a comprehensive, risk-based approach to prohibiting discrimination through the development and deployment of AI systems – similar to the Colorado AI Act and the EU AI Act. Following President Trump’s January 2025 Executive Order (described in our prior [client update](#)), which focused on fostering “America’s global AI dominance” and removing “unnecessarily burdensome requirements for companies developing AI[,]” the Texas legislature significantly pared back the bill’s private sector requirements to focus on intentional acts.

TRAIGA will go into effect on January 1, 2026, though its future remains subject to the outcome of negotiations over the federal budget reconciliation bill, which in its current form includes a 10-year moratorium on enforcement of state or local AI-specific laws or regulations.

Applicability and prohibitions

TRAIGA applies to “developers” and “deployers” of AI systems that promote, advertise or conduct business in Texas, produce a product or service used by Texas residents, or develop or deploy AI systems in Texas. “Developers” include those whose AI systems are offered, sold, leased, given or otherwise provided in Texas. Among other things, TRAIGA prohibits:

- **Manipulation of human behavior.** TRAIGA prohibits developing or deploying an AI system “in a manner that intentionally aims to incite or encourage a person to (i) commit physical self-harm, including suicide, (ii) harm another person; or (iii) engage in criminal activity.”
- **Impairment of constitutional rights.** The Act forbids the development or deployment of an AI system “with the

sole intent to infringe, restrict or otherwise impair an individual's rights under the United States Constitution.”

- **Unlawful discrimination.** TRAIGA prohibits AI systems developed or deployed with the intent to unlawfully discriminate against a protected class in violation of federal or state law. A “protected class” means a group or class of persons with a characteristic, quality, belief, or status protected from discrimination by state or federal civil rights laws. This provision carves out insurance entities and federally insured financial institutions and notes that a disparate impact is not sufficient by itself to demonstrate an intent to discriminate.

Regulatory sandbox

To encourage AI development, TRAIGA establishes a regulatory sandbox program that allows testing of AI systems in Texas for up to 36 months, subject to preliminary approval by the Texas Department of Information Resources and compliance with quarterly reporting requirements. TRAIGA also establishes a Texas AI Council to, among other things, ensure AI systems are “ethical and developed in the public’s best interest” and identify existing laws and regulations that impede AI innovation. These two provisions effectively mirror aspects of Utah’s AI Policy Act law (see our [prior client update](#)).

Enforcement

The Texas Attorney General (AG) has exclusive authority to enforce TRAIGA, although a Texas state agency may also impose sanctions against a person licensed, registered or certified by that agency. TRAIGA also requires the Texas AG to create an intake mechanism on its website for consumer complaints. If the Texas Attorney General determines a violation has occurred, it notifies the putative defendant, who is then permitted 60 days to cure the violation, provide supporting documentation, and make any necessary policy changes to prevent further violations. TRAIGA does not provide a private right of action.

In addition to injunctive relief, the Texas AG may seek penalties from statutory tiers that vary depending on whether the violation is curable or continuing. Curable violations may be penalized between \$10,000 and \$12,000; incurable violations between \$80,000 and \$200,000; and for a continuing violation, between \$2,000 and \$40,000 for each day the violation continues.

The Texas AG may also request, pursuant to a civil investigative demand, documentation regarding AI systems, including: (i) a high-level description of the purpose, intended use, deployment context and associated benefits of the AI system, (ii) a description of the type of data used for training, (iii) a high-level description of the categories of data processed as inputs, (iv) a high-level description of the outputs produced, (v) any performance metrics, (vi) any known limitations, (vii) a high-level description of the post-deployment monitoring and user safeguards, and (viii) any other relevant documentation.

TRAIGA also provides affirmative defenses to liability where (i) another person uses the AI system affiliated with the defendant in a manner prohibited by the law, or (ii) the defendant discovers a violation through feedback, testing, following state agency guidelines, or if the defendant substantially complies with the NIST AI Risk Management framework. The Act also exempts certain AI systems such as those developed for purposes including investigating or preventing fraud, preventing or responding to security incidents, preserving the integrity or security of a system.

If you have any questions regarding the matters covered in this publication, please reach out to any of the lawyers listed below or your usual Davis Polk contact.

Matthew J. Bacal

+1 212 450 4790

matthew.bacal@davispolk.com

James W. Haldin

+1 212 450 4059

james.haldin@davispolk.com

David Lisson

+1 650 752 2013

david.lisson@davispolk.com

Howard Shelanski

+1 202 962 7060

howard.shelanski@davispolk.com

This communication, which we believe may be of interest to our clients and friends of the firm, is for general information only. It is not a full analysis of the matters presented and should not be relied upon as legal advice. This may be considered attorney advertising in some jurisdictions. Please refer to the firm's privacy notice for further details.