AMPLIFY IMPACT

What Healthcare Leaders Need to Know About Texas AI Regulations (TRAIGA and SB 1188)

A Decision Readiness Brief for Healthcare Organizations Operating in Texas

AUDIENCE	CEOs, COOs, CIOs, CFOs, and Chief Medical Officers
MARKET	Mid-market healthcare organizations operating in Texas

BOTTOM LINE

Texas AI governance laws are now operational facts. Leadership needs a repeatable Go / Hold / Defer judgment for every AI system already in use comply with TRAIGA and SB 1188's mandates.

Artificial Intelligence (AI) is moving through healthcare faster than the structures built to govern it. Automated clinical documentation, AI-assisted coding, and revenue cycle automation are already running inside Texas practices today. The question for leadership is no longer whether to adopt these tools - it is how to deploy them with clarity, ownership, and a decision you can stand behind with your board. Texas has now written that expectation into law.

The **Texas Responsible Artificial Intelligence Governance Act (TRAIGA / HB 149)** took effect January 1, 2026. Alongside **Senate Bill 1188 (SB 1188)**, effective September 1, 2025, it creates the most consequential set of AI obligations healthcare organizations operating in the state have yet faced. This brief sets out what both laws require, who they apply to, and the moves your leadership team can make now -as a protocol for board-defensible judgment.

In a market where every vendor sells certainty, the scarce asset is independent judgment – clear guidance on which tools to trust, which to pause, and what to tell your board. That judgment, not the technology, is what this brief is built to support.

For organizations that operate across state lines, Texas is rarely the only rulebook in play. But it is among the first states to put a comprehensive AI framework into force, and the architecture it establishes - disclosure, human review, data localization, and accountable ownership - is the same shape now appearing in other state markets. Get the decision right in Texas and you have a blueprint you can carry.

The Decision Framework

Read past the statutes and the practical task reduces to a single repeatable judgment. For each AI system already in your operation, leadership makes one of three calls:

Decision	What It Means	When It Applies
Go	Deploy or continue with documented controls, disclosure, and a named owner in place.	The system's purpose, risks, patient disclosures, and review process are mapped, and you can demonstrate reasonable care if the state asks.
Hold	Keep the system running but pause expansion until the controls catch up.	The tool is live and useful, but disclosure, practitioner review, or vendor documentation is still incomplete.
Defer	Do not deploy yet.	The system touches diagnosis, protected health information, or patient-facing decisions without the controls this brief describes.

Everything that follows gives you the factual basis to place each system - confidently - into one of those three columns.

WHERE THE LAW ALREADY STANDS

SB 1188 - in effect since September 1, 2025. Requires AI disclosure to patients, health care practitioner review of AI-generated records, and electronic health record (EHR) data localization (United States-only storage, effective January 1, 2026). Civil penalties run from $5,000 (negligent) to $250,000 (knowing or intentional use of protected health information for financial gain) per violation.

TRAIGA (HB 149) - in effect since January 1, 2026. Establishes prohibitions on harmful AI practices, healthcare-specific disclosure obligations, enforcement through the Attorney General, and a structured penalty framework. Civil penalties run from $10,000 to $200,000 per violation, plus $2,000 to $40,000 per day for continued violations.

ENFORCEMENT

Texas Attorney General (exclusive). No private right of action. An Attorney General online complaint mechanism is required by September 1, 2026. A 60-day written notice and cure period applies before the Attorney General may bring an action - a window that rewards organizations already prepared to respond.

1. What TRAIGA Requires

TRAIGA applies to any person who: (1) promotes, advertises, or conducts business in Texas; (2) produces a product or service used by Texas residents; or (3) develops or deploys an AI system in Texas (Sec. 551.002). The Act distinguishes between AI developers and AI deployers because your obligations differ depending on your role - and most physician practices are deployers. Knowing which you are is the first input to the Go / Hold / Defer call.

Role	Statutory Definition	If You Are a Physician Practice
Developer	A person who develops an AI system that is offered, sold, leased, given, or otherwise provided in this state (Sec. 552.001(2)).	Unlikely - unless you build custom AI tools in-house.
Deployer	A person who deploys an AI system for use in this state (Sec. 552.001(1)).	Almost certainly yes - if you use automated clinical documentation (ambient scribing), AI documentation, AI-assisted coding, or revenue cycle automation.

Statutory Prohibitions

Subchapter B of TRAIGA establishes six categories of prohibited or regulated AI conduct. These apply broadly - not only to a subset of "high-risk" systems:

Prohibition	What It Means	Statutory Ref
Healthcare AI Disclosure	If AI is used in relation to a health care service or treatment, the provider must disclose this to the patient or their representative no later than the date the service is first provided (except in emergencies).	Sec. 552.051(f)
Manipulation of Human Behavior	No person may develop or deploy AI that intentionally aims to incite self-harm, harm to others, or criminal activity.	Sec. 552.052
Social Scoring	Governmental entities may not use AI to evaluate or classify persons based on social behavior or personal characteristics to assign social scores.	Sec. 552.053
Biometric Data Capture	Governmental entities may not deploy AI for uniquely identifying individuals using biometric data without consent, where doing so would infringe constitutional rights.	Sec. 552.054
Constitutional Protection	No person may develop or deploy AI with the sole intent to infringe constitutional rights.	Sec. 552.055
Unlawful Discrimination	No person may develop or deploy AI with the intent to unlawfully discriminate against a protected class. Note: Disparate impact alone is not sufficient to demonstrate intent to discriminate (Sec. 552.056(c)).	Sec. 552.056

Where to Focus Governance Attention

While TRAIGA does not create a statutory category of "high-risk" AI systems, physician practices should concentrate their judgment on the systems most likely to trigger enforcement scrutiny or patient-facing obligations:

- Automated clinical documentation (ambient scribing, note generation)
- Revenue cycle management automation (coding, claims, denials)
- Clinical decision support and diagnostic recommendations
- Patient scheduling, triage, or prior authorization algorithms

2. What SB 1188 Adds for Healthcare

TRAIGA establishes the general AI governance framework for Texas. SB 1188 adds healthcare-specific requirements that apply directly to your clinical operations - and these are where most diagnostic and documentation decisions will land.

SB 1188 targets covered entities, defined as entities covered under Section 181.001 of the Texas Health and Safety Code (the Texas Medical Privacy Act). This definition explicitly includes health care practitioners - any individual licensed, certified, or otherwise authorized to provide health care services in Texas, including nurse practitioners, physician assistants, dentists, and therapists, not only physicians (Sec. 183.001(2), (4)).

Requirement	Details	Effective
AI Disclosure to Patients	A health care practitioner using AI for diagnostic purposes must disclose the use of that technology to patients (Sec. 183.005(b)).	Sept 1, 2025
Practitioner Review of AI Records	The practitioner must review all records created with AI in a manner consistent with medical records standards developed by the Texas Medical Board (Sec. 183.005(a)(3)). The practitioner must also be acting within their licensed scope (Sec. 183.005(a)(1)).	Sept 1, 2025
EHR Data Localization	All Texas patient electronic health record (EHR) data must be physically stored within the United States or a US territory - including data held by third-party or cloud providers (Sec. 183.002(a)). Note: This provision is retroactive - it applies to all EHR data stored on or after January 1, 2026, regardless of when the record was created.	Jan 1, 2026
Scope of Practice	AI diagnostic use is permitted only within the practitioner's licensed scope, and only if not otherwise restricted by state or federal law (Sec. 183.005(a)(1)-(2)).	Sept 1, 2025
Decision Tools & Biological Sex	Any algorithm or decision assistance tool in an EHR that assists with medical treatment decisions must include the patient's biological sex as a mandatory input (Sec. 183.007(a)(2)).	Sept 1, 2025

SB 1188 Entity Exclusions

SB 1188's definition of "covered entity" excludes seven categories of entities (Sec. 183.001(2)(A)-(G)): assisted living facilities, continuing care facilities, day activity and health services facilities, home and community support services agencies, HCS waiver program providers, intermediate care facilities, nursing facilities, and TxHmL waiver program providers. Physician practices and Federally Qualified Health Centers are not excluded - they are squarely in scope.

Important: These exclusions apply only to SB 1188 (Chapter 183). They do not exempt these entities from TRAIGA (HB 149), which has its own broader jurisdictional scope.

3. Protected Health Information, HIPAA, and AI Governance

TRAIGA intersects with the Health Insurance Portability and Accountability Act (HIPAA) in several ways. Section 552.054 explicitly exempts from its biometric data definition any "information collected, used, or stored for health care treatment, payment, or operations under the Health Insurance Portability and Accountability Act." While HIPAA permits use of protected health information (PHI) for treatment, payment, and operations, AI-enabled workflows require intentional governance to ensure compliance under both frameworks simultaneously.

To keep that judgment defensible, leadership should ensure:

- Clear definition of where protected health information is permitted or prohibited in AI system inputs and outputs
- Minimum-necessary standards applied to AI data processing, consistent with HIPAA
- Vendor contracts that restrict secondary data use or model training on patient data
- Documentation sufficient to demonstrate PHI governance under both HIPAA and TRAIGA simultaneously

4. Penalty Structure

The penalty ranges below are not a scare tactic - they are inputs to the decision. They tell you where the cost of a wrong call is highest, and therefore where documented judgment matters most.

SB 1188 - Civil Penalties

Violation Type	Penalty	Statutory Ref
Negligent violation	Up to $5,000 / year	183.011(b)(1)
Knowing or intentional violation	Up to $25,000 / year	183.011(b)(2)
Knowing/intentional use of PHI for financial gain	Up to $250,000	183.011(b)(3)

Disciplinary action: A regulatory agency may take disciplinary action (including license suspension or revocation) against a covered entity that violates Chapter 183 three or more times (Sec. 183.010).

No safe harbor in SB 1188: Unlike TRAIGA, SB 1188 does not provide a non-liability defense for organizations that adopt the National Institute of Standards and Technology (NIST) framework or any other standard. NIST

compliance protects you under TRAIGA but does not reduce your exposure under SB 1188. For SB 1188, you must meet each statutory requirement directly - there is no framework-based defense.

TRAIGA - Civil Penalties

Violation Type	Penalty Range	Statutory Ref
Curable violation, or breach of a cure statement submitted to the Attorney General	$10,000 - $12,000 per violation	552.105(a)(1)
Uncurable violation	$80,000 - $200,000 per violation	552.105(a)(2)
Continued violation	$2,000 - $40,000 per day	552.105(a)(3)
State agency sanctions (after Attorney General finding + recommendation)	Up to $100,000 + license suspension/revocation	552.106

Additional Attorney General remedies: Injunctive relief, recovery of attorney's fees, court costs, and investigative expenses (Sec. 552.105(b)).

Rebuttable presumption: There is a rebuttable presumption that the person used reasonable care as required under TRAIGA (Sec. 552.105(c)). The burden is on the Attorney General to overcome this presumption - a significant defendant protection.

Curability determination: Whether a violation is curable or uncurable is determined by the court, not by the organization (Sec. 552.105(a)(1)-(2)).

5. Defendant Protections and Safe Harbor

TRAIGA gives organizations that demonstrate good-faith compliance several layers of protection. Understanding them is what lets a leader choose Go with confidence rather than Defer out of doubt.

60-Day Cure Period (Sec. 552.104)

The Attorney General must provide written notice identifying specific violations before bringing an action. The Attorney General may not bring an action if, within 60 days of notice, the person:

- (a) cures the violation;
- (b) provides the Attorney General a written statement confirming the cure with supporting documentation; and
- (c) demonstrates that internal policies have been changed to prevent further violations.

Non-Liability Conditions (Sec. 552.105(e))

A defendant may not be found liable under TRAIGA if either of the following conditions is met:

- Third-party misuse: Another person uses the AI system affiliated with the defendant in a manner prohibited by TRAIGA.
- Self-discovery through a qualifying channel: The defendant discovers a violation through
- (a) feedback from a developer, deployer, or other person;

- (b) testing, including adversarial or red-team testing;
- (c) following guidelines set by applicable state agencies; or
- (d) an internal review process, if the defendant substantially complies with the National Institute of Standards and Technology (NIST) "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile" or another nationally or internationally recognized AI risk management framework.
-

WHAT THIS MEANS FOR YOUR PRACTICE

The NIST generative AI profile is not a standalone safe harbor. It is a qualifying condition for one of the non-liability channels - the internal review process.

To benefit from TRAIGA's liability protections, your organization needs both:

(1) a mechanism to discover violations (testing, feedback, agency guidance, or NIST-compliant internal review), and
(2) the ability to cure within 60 days if the Attorney General issues notice.

Adopting the NIST framework remains the single most strategically significant governance decision your organization can make - not because TRAIGA mandates it, but because it activates two of the statute's strongest protections at once: the non-liability defense for self-discovered violations and the evidentiary basis for demonstrating reasonable care.

Additional protections: The rebuttable presumption of reasonable care (Sec. 552.105(c)) applies to all defendants. A defendant may also seek an expedited hearing or declaratory judgment if they believe in good faith they have not violated TRAIGA (Sec. 552.105(d)). The Attorney General may not bring an action for an AI system that has not been deployed (Sec. 552.105(f)).

6. From Requirements to Readiness: Recommended Actions

TRAIGA does not mandate a formal governance program, impact assessments, or a written AI policy. But without these structures, your organization cannot cure a violation within the 60-day window, qualify for the non-liability defense through internal review, respond substantively to an Attorney General investigative demand (Sec. 552.103), or demonstrate the reasonable care the statute presumes.

The following actions are not statutory requirements - they are the practical inputs that let leadership make, and defend, the Go / Hold / Defer call.

Action	Why It Matters Under TRAIGA / SB 1188	Practical Implication
NIST Gen AI Profile Adoption	Activates the non-liability defense for self-discovered violations through internal review.	Evaluate and adopt the NIST framework or equivalent.
AI System Inventory	The Attorney General's notice includes a description of AI systems, their purpose, data types, inputs, outputs, and limitations.	Document every AI-enabled workflow, including vendor tools.
Impact Documentation	While not a statutory mandate, documented assessments of AI risks and mitigations form the evidentiary basis for demonstrating reasonable care and supporting a cure statement.	Write up records of what each AI system does, its risks, and your mitigations.
Patient Disclosure Protocol	Required by both TRAIGA for healthcare AI) and SB 1188 for diagnostic AI). Disclosure must be clear, conspicuous, and in plain language, and may not be designed to obscure, subvert, or impair the consumer's ability to understand it	A written disclosure policy, consistently followed, for all patient-facing AI use.
AI Governance Ownership	The 60-day cure period and the investigative demand process both require someone to coordinate the organizational response. Ad hoc oversight cannot produce a cure statement under time pressure.	Assign executive ownership for AI governance and board reporting.
Vendor Transparency Documentation	The Attorney General can demand descriptions of data types, performance metrics, limitations, and safeguards. Deployers who rely on vendor AI need this documentation to respond.	Request model cards, risk disclosures, and data processing documentation from all AI vendors.
Practitioner Review Process	Required by SB 1188. All records created with AI must be reviewed by a health care practitioner consistent with Texas Medical Board standards.	Implement a review workflow for all AI-generated clinical documentation.
EHR Data Localization Verification	Retroactively applies to all EHR data stored on or after Jan 1, 2026, regardless of creation date. The compliance deadline has already passed.	Verify all EHR data (including third-party/cloud) is stored in the United States.

7. The 10 Questions Your Board Will Ask

These are the questions a board raises when AI moves from pilot to production. Being able to answer them is what converts open-ended risk into a documented, defensible decision. Treat each as a readiness check that sharpens your judgment - not a verdict on your team:

1. What AI systems are currently deployed in our practice, and who authorized them?
2. Have we documented the purpose, risks, and mitigations for each AI system - sufficient to respond to an Attorney General investigative demand within the statutory timeline?
3. Do we have a written AI disclosure policy for patients - and is it being followed for both healthcare AI (TRAIGA Sec. 552.051(f)) and diagnostic AI (SB 1188 Sec. 183.005(b))?
4. Who in our organization is accountable for AI governance decisions and for coordinating a response if the Attorney General issues notice (see Sec. 552.104 and Sec. 552.103)?
5. Are our AI vendors providing the transparency documentation we would need to respond to an Attorney General civil investigative demand (Sec. 552.103)?
6. Have we evaluated our AI tools for potential disparate impact on patient populations - even though disparate impact alone does not demonstrate intent to discriminate under TRAIGA (Sec. 552.056)?
7. Are all AI-generated records reviewed by a health care practitioner consistent with Texas Medical Board standards (SB 1188 Sec. 183.005(a)(3))?
8. Are our EHR systems compliant with SB 1188's data localization requirement - United States-only storage for all records, including legacy data (Sec. 183.002, effective Jan 1, 2026)?
9. Are we substantially complying with the NIST Generative AI Profile or equivalent framework to activate TRAIGA's non-liability defense (Sec. 552.105(e)(2)(D))?
10. Can we produce a cure statement with supporting documentation and policy changes within 60 days if the Attorney General provides notice of a violation (Sec. 552.104)?

The Decision in Front of You

TRAIGA and SB 1188 did not arrive to slow Texas healthcare down. They arrived because AI is already inside the work - in the notes, the codes, the claims - and the state has decided that someone accountable must be able to explain it. That is the same standard leadership already holds for every other clinical and financial system it runs.

The organizations that move fastest will not be the ones that wait for certainty. They will be the ones that can look at each AI system in their operation and say, in plain language, why it is a Go, a Hold, or a Defer.

The law is settled. The deadlines have passed. The only open question is whether your next AI decision is one you can defend.

About Amplify Impact

Amplify Impact is a boutique advisory firm that helps mid-market healthcare leaders make defensible AI decisions with confidence, clarity, and accountability. We are an AI advisory firm that specializes in healthcare - working with physician practices, Federally Qualified Health Centers, and multi-state organizations operating in Texas.

Our flagship service, **AI Pulse - Healthcare Edition**, is a 10-business-day executive engagement that produces the **Execution Readiness Pack**: a board-ready set of decision artifacts including a Board Brief, an AI Risk Exposure Review (risk heatmap), a Vendor Claims Checklist, a Decision Roadmap, and an Executive Review.

The practice is led by a founder who has carried profit-and-loss accountability, scaled national programs, and managed corporate risk, the same seat the leaders reading this brief occupy. The work is not to lecture from the outside, but to help you make the call and stand behind it.

Chris Calitz, MPP | Founder & Managing Partner | chris@weareamplifyimpact.com

To work through these decisions with your leadership team, schedule an executive briefing. Learn more about the AI Pulse-Healthcare Edition at www.amplify-impact.ai.

Sources

1. Texas HB 149 (TRAIGA) - 89th Legislature, Enrolled Text, signed June 22, 2025, effective January 1, 2026
2. Texas SB 1188 - 89th Legislature, Enrolled Text, effective September 1, 2025 (data localization: January 1, 2026)
3. NIST Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1), July 2024
4. Holland & Knight, "Texas Enacts Comprehensive AI Governance Laws," June 2025
5. Baker Botts, "Texas Enacts Responsible AI Governance Act," July 2025
6. Texas Medical Association, "Physicians Must Disclose AI Use," 2026
7. Office of the Texas Attorney General, Consumer Protection Division enforcement guidelines

NEXT STEP

To pressure-test your readiness with your leadership team, schedule an executive briefing with Amplify Impact Consulting. Learn more about the 10-business-day **AI Pulse - Healthcare Edition** engagement at **www.amplify-impact.ai**.

INFORMATIONAL NOTE

This brief is for informational purposes only and does not constitute legal advice. Consult qualified legal counsel for compliance guidance specific to your organization.