Drafting for TRAIGA

Contractual safeguards for private sector AI technology agreements.

WRITTEN BY **IRINA BESCHIERIU**

The Texas Responsible Artificial Intelligence Governance Act (TRAIGA), effective January 1, 2026, reshapes private sector AI contracting in ways that most practitioners haven't yet recognized. While initial analysis focused on government applications, TRAIGA's broad jurisdictional reach extends to any entity conducting business in Texas, producing services for Texas residents, or deploying AI systems within the state.[1] This expansive scope can potentially transform every AI-related commercial agreement into a regulatory compliance document.

For Texas practitioners, TRAIGA isn't merely another compliance requirement. It's a reconceptualization of how AI contracts must allocate risk, document intent, and structure ongoing obligations. The act's intent-based liability framework creates both unprecedented risks and significant opportunities for sophisticated contract drafting.

THE INTENT DOCUMENTATION IMPERATIVE

TRAIGA's core innovation lies in its focus on prohibited intent rather than discriminatory outcomes. This approach impacts contract drafting because every agreement must now create an evidentiary record of lawful purpose. The act's primary prohibition—developing or deploying AI systems with intent to unlawfully discriminate—requires contracts to document not just what parties will deliver, but why they're building it and how that purpose will be proven.

STATEMENTS OF WORK AS LEGAL EVIDENCE

Under TRAIGA, Statements of Work (SOW) transform from project descriptions into critical legal documents. Every SOW must meticulously document the AI system's intended lawful purpose, specific business objectives driving development decisions, bias mitigation measures incorporated into system design, training data selection criteria and quality assessments, and testing protocols implemented to identify discriminatory potential.

This documentation serves as primary evidence against intent-based violations. A well-crafted SOW demonstrating comprehensive bias awareness and mitigation efforts provides powerful defense against discrimination allegations. Conversely, ambiguous or incomplete SOWs create significant liability exposure by failing to establish lawful intent clearly.

Practitioners should require SOWs to include detailed bias assessment protocols, documented training data provenance, explicit non-discriminatory purpose statements, and comprehensive testing methodologies. These elements must be specific enough to demonstrate genuine compliance efforts while flexible enough to accommodate evolving technical requirements.[2]

WARRANTY STRUCTURES FOR INTENT-BASED LIABILITY

Traditional technology warranties prove inadequate for TRAIGA's intent-focused approach. New warranty frameworks must address the mental state behind AI development, requiring vendors to represent that systems were designed without discriminatory intent, training data was selected and reviewed for bias potential, development teams received

Leo Lintang - stock.adobe.com

appropriate bias awareness training, and comprehensive testing was conducted throughout the development process.[3] These warranties should extend beyond final deliverables to encompass the entire development methodology. Vendors must warrant that algorithmic choices were made for legitimate business reasons, data selection decisions incorporated bias considerations, and quality assurance processes included discrimination testing protocols.

Customers, meanwhile, should warrant their intended use cases comply with TRAIGA prohibitions and provide vendors with accurate information about deployment contexts to enable compliant system design.

RISK ALLOCATION AND COMPLIANCE IN THE NIST ERA

TRAIGA's safe harbor provision for substantial compliance with the National Institute of Standards and Technology (NIST) AI Risk Management Framework creates powerful incentives for comprehensive compliance programs while fundamentally altering contractual risk allocation. The NIST framework becomes the de facto industry standard, making compliance demonstration a critical competitive advantage.[4]

Forward-thinking agreements should require vendors to implement and maintain NIST AI RMF compliance throughout the contract term. This requirement should include documented implementation of framework principles, regular compliance assessments and third-party audits, ongoing monitoring and risk management procedures, and immediate notification of any compliance gaps or violations.

Customers should secure audit rights to verify NIST compliance, including access to compliance documentation, rights to engage independent auditors, and review of vendor risk management procedures. These audit rights must be sufficiently broad to enable customers to demonstrate their own due diligence in vendor selection and oversight.

ENHANCED INDEMNIFICATION FRAMEWORKS

TRAIGA's enforcement structure creates complex indemnification scenarios that traditional technology contracts don't address. The distinction between development intent and deployment intent requires sophisticated liability allocation between vendors and customers.

Vendors should seek indemnification limitations covering claims arising from customer deployment decisions, modifications to AI systems contrary to specifications, use cases not disclosed during development, and integration with third-party systems not contemplated in the original design. However, vendors cannot escape liability for discriminatory design intent or failure to implement warranted bias mitigation measures.

Customers should demand broad indemnification covering Office of the Texas Attorney General enforcement actions related to system design, discrimination claims arising from warranted system capabilities, and penalties resulting from vendor compliance failures. The indemnification should extend to regulatory defense costs and cure period remediation expenses ongoing monitoring requirements imposed by enforcement actions.[5]

SPECIALIZED AGREEMENT TYPES AND COMPLIANCE CHALLENGES

The varied landscape of AI technology agreements requires tailored approaches to TRAIGA compliance, with each contract type presenting unique risk allocation and documentation challenges. Understanding how TRAIGA impacts different commercial relationships, from platform services to custom development, enables practitioners to craft agreements that provide genuine protection rather than mere regulatory gestures.

Software-as-a-Service (SaaS) providers face perhaps the most complex challenges because they provide AI capabilities to different user bases with varying intentions and use cases. Unlike traditional software licensing, AI platform agreements must address the fundamental challenge of monitoring user intent, a mental state rather than observable behavior. Traditional acceptable use policies focused on preventing abuse and illegal content prove inadequate for TRAIGA's intent-based prohibitions.

Modern SaaS agreements must create sophisticated monitoring and response frameworks that balance platform protection with user functionality. These frameworks should establish technical controls limiting AI system capabilities to prevent obvious violations while requiring users to warrant their intended use cases comply with TRAIGA.[6] The challenge lies in creating policies specific enough to provide legal protection while flexible enough to accommodate legitimate AI use cases, focusing on prohibited intentions rather than specific applications.

Platforms must also implement comprehensive monitoring capabilities that can be contractually documented to demonstrate good faith compliance efforts during regulatory proceedings. This documentation should include technical architecture descriptions, monitoring protocol specifications, response procedure timelines, and ongoing compliance assessment methodologies. The key is creating systems that can detect potential violations without overly restricting legitimate uses or creating unmanageable compliance burdens.

Application Programming Interface (API) providers face particularly complex liability allocation challenges because they may be several degrees removed from end-use applications. Consider a machine learning API that gets incorporated into a hiring system later found to discriminate intentionally—the resulting liability questions require sophisticated contractual solutions that traditional API agreements simply don't address.

Effective API licensing must create compliance chains extending throughout the technology supply chain, requiring licensees to impose substantially similar TRAIGA obligations

on their downstream users while implementing technical controls preventing prohibited system integration. These agreements must balance API provider risk limitation with sufficient downstream visibility to demonstrate compliance efforts, including rights to audit downstream applications and require disclosure of high-risk use cases.

The technical integration requirements become equally critical, with API agreements needing to address mandatory bias testing protocols for downstream applications, logging and audit trail capabilities for regulatory proceedings, and integration standards ensuring compliance documentation remains available throughout the technology stack. Licensees must warrant that their technical integration supports rather than undermines TRAIGA compliance, creating a contractual framework that extends regulatory responsibility across complex technology relationships.

Custom AI development contracts demand the most comprehensive TRAIGA compliance frameworks because they govern system creation from inception through deployment and beyond. These agreements must establish exhaustive documentation requirements covering every aspect of system creation, including detailed algorithmic design rationale, comprehensive training data assessment criteria, bias testing methodologies and results, and ongoing monitoring protocols.

The documentation must be sufficiently detailed to support both NIST safe harbor protection and intent-based violation defense, with each design decision documented with explicit rationale demonstrating compliance considerations and bias mitigation efforts. Unlike traditional software development, AI systems require continuous compliance monitoring throughout their operational lifecycle, creating ongoing vendor obligations that extend well beyond traditional warranty periods.

For healthcare providers, TRAIGA creates specific disclosure requirements that generate unique contracting challenges. Healthcare AI contracts must ensure vendor systems support patient disclosure requirements, specifying responsibility for disclosure language development, technical integration with clinical workflows, and warranties ensuring disclosure functionality. These contracts must address disclosure timing requirements, before or during patient interaction except in emergencies, while providing technical capabilities enabling consistent and verifiable compliance.

The enforcement landscape adds another layer of contractual complexity. TRAIGA's exclusive enforcement through the Office of the Texas Attorney General creates procedural advantages compared to private litigation but requires specialized response capabilities that traditional technology contracts don't contemplate. The 60-day cure period provides significant remediation opportunities that must be contractually leveraged through comprehensive regulatory response frameworks.

These frameworks should establish immediate notification requirements between contracting parties, evidence preservation procedures, joint defense coordination mechanisms, and equitable cost allocation for regulatory proceedings. The protocols must designate responsible personnel with appropriate expertise, establish secure communication channels for sensitive regulatory matters, and create realistic timelines enabling effective cure period utilization while protecting both parties' interests throughout enforcement proceedings.[7]

THE NEW CONTRACTING REALITY

TRAIGA represents more than regulatory compliance. It's a notable shift toward intent-based AI governance that requires sophisticated contractual responses. The act's combination of intent-focused liability, NIST safe harbor protection, and Office of the Texas Attorney General enforcement creates both significant risks and powerful defensive tools for well-prepared practitioners.[8]

Success in this new regulatory environment requires contracts that document lawful intent comprehensively, implement NIST compliance frameworks proactively, and allocate risks appropriately between parties with different capabilities and exposures. The practitioners who master these requirements will provide genuine value to clients navigating an increasingly complex AI regulatory landscape. **TBJ**

NOTES

1. H.B. 2060, 88th Leg., R.S. (Tex. 2023) (codified as Tex. Bus. & Com. Code ch. 2001) (effective Jan. 1, 2026), https://capitol.texas.gov/tlodocs/88R/billtext/html/HB02060F.htm.
2. Pauline T. Kim, *Resetting Antidiscrimination Law in the Age of AI,* 138 Harv. L. Rev. 1138 (2025), https://harvardlawreview.org/print/vol-138/resetting-antidiscrimination-law-in-the-age-of-ai/.
3. Emilio Ferrara, *Fairness and Bias in Artificial Intelligence: A Brief Survey of Sources, Impacts, and Mitigation Strategies,* Thomas Lord Department of Computer Science, USC Viterbi School of Engineering, University of Southern California (Apr. 2023), https://arxiv.org/pdf/2304.07683.
4. Artificial Intelligence Risk Management Framework (AI RMF 1.0), National Institute of Standards and Technology (Jan. 2023), https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf.
5. For general principles on SaaS indemnification provisions, *see* H. Ward Classen, *SaaS Agreements: Key Contractual Provisions,* Business Law Today (Nov. 2021), https://businesslawtoday.org/2021/11/saas-agreements-key-contractual-provisions/.
6. Shahed Kader, *Important Provisions of API License Agreements,* Contract Nerds (August 3, 2021), https://contractnerds.com/api-license-agreements/.
7. The Office of the Texas Attorney General has broad civil enforcement authority, including the power to seek injunctive relief and civil penalties. *See* Tex. Bus. & Com. Code § 17.47.
8. Steve Charkoudian and Omer Tene, *Contracting around AI: Reading the fine print,* IAPP (Nov. 27, 2024), https://iapp.org/news/a/contracting-around-ai-reading-the-fine-print (discussing how AI lawyers negotiate finely crafted agreements to allocate risks related to copyright, trade secrets, privacy, fairness, and cybersecurity in the absence of comprehensive AI-specific regulations).

IRINA BESCHIERIU
is a Texas- and New York-licensed attorney with 18 years of international legal experience, focusing on outsourcing and technology transactions. Educated in both civil and common law systems and certified in advanced contract and commercial management (WCC CMAP), data privacy (CIPP/EU), and AI governance (AIGP from IAPP), she frequently publishes thought leadership on outsourcing, technology contracting, data privacy, and AI in legal practice. In Beschieriu's daily practice, she negotiates complex, multimillion-dollar outsourcing and technology engagements.