GOVERNING THE UNGOVERNED

A Framework for AI Defensibility and Operational Risk in Texas Regulated Industries

The 2026 Baseline on Shadow AI, PII Exposure, and the 3 Million Small Businesses Flying Blind

Matthew Bertram
President, ModalPoint | CEO, EWR Digital
Certified AI Professional (CAIP) | Goldman Sachs 10KSB

February 2026 | Version 2.0

Executive Summary 4
1. The Problem: AI Adoption Without Governance 5
1.1 The Scale of What's Already Happening 5
1.2 The 2026 Texas Small Business Baseline 5
2. The Three Risks Nobody Is Managing 7
2.1 Shadow AI: The Invisible Workforce 7
The Data Exfiltration Loop 7
2.2 Hallucinations: Confident, Plausible, Wrong 8
2.3 The Cost Barrier: Governance Priced Out of Reach 9
3. The GDPR Warning: What Happens When Regulation Ignores Operational Reality 11
4. What Texas Has Done So Far 12
4.1 TRAIGA and the AI Council 12
4.2 House Bill 3512 12
4.3 Two Mandates, One Gap 12
5. Where the Exposure Is Concentrated 13
5.1 Energy 13
5.2 Legal Services 13
5.3 Healthcare 13
5.4 Financial Services 14
5.5 The Small Business Multiplier 14
6. What Current Training Gets Wrong 14
6.1 The Prompt Engineering Fallacy 15
6.2 Awareness Is Not Competency 15
6.3 Enterprise-Only Design 15
7. A Governance Framework That Actually Works 16
7.1 Three-Layer Architecture 16
7.2 The Four Governance Competencies 17
7.3 Industry-Specific Modules 17
7.4 Framework Alignment 18
8. Making Governance Work for Small Business 18
8.1 The Micro-Governance Approach 18
8.2 The Cost Reality 19
8.3 Governance as Growth Infrastructure 19
9. Measuring What Matters: Competency, Not Completion 19
9.1 Competency Assessment 19
9.2 Organizational Readiness Score 20

9.3 Governance Readiness Scorecard 20
10. The Economic Argument 20
10.1 Liability Exposure Without Governance 21
11. What This Means for Executive Leadership 22
12. Conclusion 23
Next Steps 24
Appendix A: 2026 Workforce Pulse Baseline – Methodology and Data 25
Key Findings Summary 25
Investment Willingness 25
Demand Signal 26
About the Author 26

Executive Summary

Across Texas, AI is already making decisions in regulated industries. Legal documents are being drafted with AI-generated citations that no one verifies. Financial reports include AI-produced analysis that no one audits. Compliance filings incorporate AI output that no one traces. And in the overwhelming majority of cases, no policy governs the use, no training prepared the user, and no documentation exists to defend the decision if challenged.

This is not a future risk. It is the current operating condition for most Texas businesses.

The state's 3 million small businesses are adopting AI at scale—not because they have a strategy, but because competitive pressure left no alternative. They are doing it without governance guardrails, without awareness that regulations already apply, and without any affordable path to compliance. Shadow AI—employees using AI tools in regulated workflows without organizational knowledge or oversight—is the norm, not the exception. And the hallucination problem—AI generating confident, plausible, fabricated information—is embedding errors into regulated decisions at a rate no one is measuring because no one is looking.

Texas has recognized the problem. The legislature passed two AI training mandates: TRAIGA Section 554.102 and House Bill 3512. Both advance AI awareness. Neither delivers governance competency. And neither addresses the structural gap that will determine whether governance succeeds or fails: small businesses have no affordable way to comply.

If policymakers allow governance frameworks to develop without practitioner input, Texas risks replicating Europe's GDPR experience—well-intentioned regulation that imposed disproportionate administrative burden on the businesses least equipped to absorb it, while large enterprises simply added compliance departments. The result was not better governance. It was a compliance tax that punished small operators without reducing actual risk.

> *This paper argues that AI governance in Texas regulated industries requires a fundamentally different approach: affordable training at small business scale, standards calibrated to actual risk rather than theoretical completeness, and clear direction for business owners who want to do the right thing but have no roadmap. It provides an implementation framework designed for the real economy—not just for enterprises with compliance departments.*

1. The Problem: AI Adoption Without Governance

1.1 The Scale of What's Already Happening

The transformation is not coming. It has arrived. According to the Dallas Federal Reserve, Texas businesses using AI jumped from 20% in April 2024 to 36% by May 2025, with 59% of Texas Business Outlook Survey respondents using either generative or traditional AI.[1] This nearly 21-percentage-point increase in a single year represents one of the fastest adoption curves of any technology in Texas business history.

59% of Texas businesses now use generative or traditional AI (Dallas Fed, May 2025)	**21pt** increase in Texas AI business adoption in just one year
77% of workers expect AI to affect their career within 5 years; only 31% trained (JFF)	**68%** of US small businesses use AI regularly; 77% have no formal AI policy

Sources: JFF National AI Workforce Survey, 2024[2]; U.S. Chamber/Teneo Small Business AI Survey, 2024[3]

These numbers describe a workforce that has adopted AI faster than any governance infrastructure can keep pace. In regulated industries—energy, legal, healthcare, financial services—this gap is not merely a training deficiency. It is a liability engine running unmonitored.

1.2 The 2026 Texas Small Business Baseline

In February 2026, EWR Digital conducted a rapid workforce pulse survey—*AI in Your Business: A 90-Second Pulse Check*—targeting Texas small business operators to establish a governance readiness baseline.[4] While the sample (n=34) is directional rather than statistically representative, the pattern it reveals is consistent with national data and alarming in its specificity.

[1] Dallas Federal Reserve Bank, Texas Business Outlook Surveys, May 2025.
[2] Jobs for the Future (JFF), "AI and the Workforce," 2024 National Survey.
[3] U.S. Chamber of Commerce and Teneo, "Small Business AI Adoption Survey," 2024.
[4] EWR Digital, "AI in Your Business: A 90-Second Pulse Check," February 2026 (n=34).

94% of respondents use AI in their business	**79%** have no written AI use policy	**85%** have received zero AI training

82% unaware which regulations apply to their AI use	**88%** unaware of AI governance penalty structures

Among businesses actively using AI, 78% have no written policy governing that use. Among those who expressed concern about AI liability, 68% still have no policy in place.

When informed that states including Texas and Colorado have enacted AI governance laws with penalties ranging from $10,000 to $200,000 per violation, **88% said they were not aware.** Zero respondents reported being both aware and compliant.

> *Small businesses are not resisting AI governance. They have no idea it applies to them, no affordable path to compliance, and no one explaining what they need to do. They are not being defiant. They are flying blind.*

2. The Three Risks Nobody Is Managing

2.1 Shadow AI: The Invisible Workforce

Shadow AI is the use of artificial intelligence tools by employees without organizational knowledge, policy, or oversight. It is not a theoretical concern. It is the default operating condition in most Texas small businesses and many larger organizations.

Employees are using ChatGPT, Copilot, Gemini, and other generative tools in regulated workflows because the tools are free, fast, and available. No one told them to use AI. No one told them not to. No one is checking the output. In a five-person law firm, a paralegal uses AI to draft research memoranda. In a ten-person energy consulting shop, an analyst uses AI to generate compliance summaries. In a twenty-person financial advisory, a junior associate uses AI to produce client-facing reports. Each is acting rationally within their role. None has been given governance direction. All are generating regulated work product with zero documentation of AI involvement.

Shadow AI is not an employee behavior problem. It is a leadership infrastructure gap. When a business has no AI use policy, no training, and no documentation requirements, every employee's use of AI is shadow AI—including the owner's.

The Data Exfiltration Loop

Shadow AI creates a second category of risk that most businesses have not considered: data leaving the organization through the same tools employees use to generate output.

Prompt infiltration. When an employee pastes client PII, medical records, financial data, or trade secrets into a commercial AI tool to generate a summary or draft, that data has left the organization's control. There is no recall mechanism. The employee has not "leaked" data in the traditional cybersecurity sense—they have submitted it voluntarily to a third-party system in the course of doing their job. But the regulatory consequence is identical to a breach.

Unintended training exposure. Most commercial AI systems retain user inputs by default to improve their models. Unless an organization has negotiated enterprise licensing with explicit data retention exclusions, any information employees enter into free-tier or standard AI tools may be incorporated into the model's training data. A law firm's case strategy, an energy company's reserves analysis, a healthcare provider's patient records—each could theoretically surface in another user's query. The probability of exact reproduction is low. The regulatory exposure from the submission itself is not.

The Texas double-jeopardy. The Texas Data Privacy and Security Act (TDPSA), effective July 2024, regulates the collection, use, and processing of consumers' personal data, with civil penalties of up to $7,500 per violation enforced exclusively by the Attorney General.[5] TRAIGA governs AI decision-making in regulated environments. A business that leaks client data into an

[5]Texas Data Privacy and Security Act (TDPSA), Tex. Bus. & Com. Code § 541, effective July 1, 2024. Civil penalties up to $7,500 per violation. Texas Attorney General has exclusive enforcement authority.

AI system (TDPSA violation) and then uses that system's unverified output in a regulated filing (TRAIGA exposure) faces liability on both ends of the same workflow. The data going in is a privacy violation. The decisions coming out are a governance failure. Neither statute references the other, but both apply simultaneously to the same employee action.

Most businesses are not aware that either obligation exists. The 2026 pulse baseline found 88% of respondents unaware of AI governance penalty structures. The data exfiltration risk adds a privacy dimension that compounds the governance exposure this paper describes throughout.

2.2 Hallucinations: Confident, Plausible, Wrong

Large language models generate text that is fluent, structurally plausible, and confidently presented—even when the content is fabricated. Invented case citations. Nonexistent regulatory provisions. Fictitious statistics presented with precision. AI does not signal uncertainty the way a human researcher would. It produces wrong answers with the same tone and formatting as right ones. A 2024 Stanford RegLab analysis found that some AI models generate hallucinations in one out of three legal research queries.[6]

In unregulated contexts, hallucinations are an inconvenience. In regulated industries, they are a liability vector. And this is not theoretical—it is already generating enforcement consequences:

- **Mata v. Avianca (S.D.N.Y., 2023):** Two attorneys sanctioned $5,000 for citing six nonexistent cases generated by ChatGPT in a federal court filing. The case became the first widely reported AI hallucination sanction.[7]
- **Noland v. Topa Ins. Group (Cal. Ct. App., 2025):** Attorney fined $10,000—the largest California AI sanction to date—after 21 of 23 quoted case passages were fabricated. Referred to the State Bar.[8]
- **Johnson v. Dunn (N.D. Ala., 2025):** Attorneys at a Top 50 law firm disqualified from representing their client after filing hallucinated citations—despite the firm having a written AI use policy that prohibited unauthorized AI use.[9]
- **Morgan & Morgan sanctions (D. Wyo., 2025):** Three attorneys from the No. 42 U.S. law firm (by headcount) sanctioned after 8 of 9 cited cases were hallucinated by the firm's own internal AI platform.[10]

[6]Stanford RegLab, May 2024. Found some AI models generate hallucinations in one out of three legal research queries.
[7]Mata v. Avianca, Inc., No. 22-cv-1461 (S.D.N.Y. June 22, 2023). Attorneys sanctioned $5,000 for citing six nonexistent cases generated by ChatGPT.
[8]Noland v. Topa Ins. Group (Cal. Ct. App. 2025). Attorney fined $10,000 after 21 of 23 quoted case passages were fabricated by ChatGPT; referred to State Bar.
[9]Johnson v. Dunn, No. 2:21-cv-1701 (N.D. Ala. July 23, 2025). Attorneys at a Top 50 law firm disqualified from representing client after filing hallucinated citations, despite firm having an AI use policy.
[10]Morgan & Morgan sanctions, Case No. (D. Wyo. Feb. 24, 2025). Three attorneys from the No. 42 U.S. law firm (by headcount) sanctioned after 8 of 9 cited cases were hallucinated by internal AI platform.

- **Rochon-Eidsvig v. JGB Collateral (Tex. App.—Dallas, 2024):** A Texas Court of Appeals judge ordered an attorney to verify four case citations that neither the court nor opposing counsel could locate.[11]

These are legal services cases because courts create public records of sanctions. The equivalent failures in energy compliance filings, financial reports, and healthcare documentation are occurring without public visibility—not because they are less frequent, but because those regulatory enforcement mechanisms operate with less transparency. As of late 2025, AI citation incidents in courts alone have accelerated from approximately two per week to two to three per day.[12]

The more skilled a worker becomes at prompting AI to produce polished output, the harder hallucinations become to detect. Prompt engineering without governance literacy does not solve the problem. It makes the problem invisible.

2.3 The Cost Barrier: Governance Priced Out of Reach

Even where governance training exists, it is priced for enterprises. The gap between what governance costs and what small businesses can spend is not a market inefficiency. It is a structural barrier that guarantees non-compliance for the majority of Texas businesses.

Training Category	Market Price	Small Business Reality
Basic AI awareness	$500–$2,500/person	53% budget under $1,000 total
Intermediate governance	$2,000–$6,000/person	59% budget $100–$500/employee
Expert/executive AI governance	$5,000–$50,000/person	Only 12% budget $5,000–$10,000 total
Enterprise governance program	$12,000–$250,000 package	Exceeds annual revenue % for most SMBs
Small business governance setup	$2,000–$5,000 initial	Closest to accessible but still above median budget

Sources: Market pricing from Bizzuka, AI For Business Training Center, Liminal, Amra & Elma; small business budgets from EWR Digital 2026 Pulse Baseline.[13][14][15]

[11] Rochon-Eidsvig v. JGB Collateral, LLC, No. 5-24-00123-CV (Tex. App.—Dallas 2024). Texas Court of Appeals ordered attorney to verify four unlocatable case citations.
[12] Cronkite News/Arizona PBS, "As more lawyers fall for AI hallucinations, ChatGPT says: Check my work," October 28, 2025. Reports AI citation incidents accelerating from two per week to two-three per day.
[13] Medium/Bizzuka, "How Much Does AI Training for Businesses Cost?" 2024. Beginner: $500–$2,500/person; Intermediate: $2,000–$6,000/person; Expert: $5,000–$10,000/person.
[14] AI For Business Training Center, "AI Training Costs for Businesses," 2025. Executive-level AI training: €15,000–€50,000; corporate packages: €12,000–€250,000.
[15] Liminal, "Enterprise AI Governance: Complete Implementation Guide," 2025. Small business governance: initial cost $2,000–$5,000; ongoing $100–$500/month.

The 2026 pulse baseline quantifies this gap precisely. 59% of small business respondents indicated a per-employee training budget of $100–$500 annually. The lowest-cost governance training available on the market today exceeds that budget by a factor of four to twenty. This is not a willingness problem—the businesses want the training. It is an access problem. The training does not exist at a price point small businesses can reach.[16]

When governance training is unaffordable, the rational business response is predictable: skip it. Adopt AI for competitive advantage, hope the regulatory environment stays unclear long enough to figure it out later, and deal with enforcement if it comes. This is not negligence. It is a resource-constrained operator making the only decision the market allows them to make.

> *Any governance framework that does not solve the cost problem for small business is a governance framework that governs only enterprises.*

3. The GDPR Warning: What Happens When Regulation Ignores Operational Reality

Europe's General Data Protection Regulation was designed to protect consumers. It was implemented without meaningful consideration of small business operational reality. The result is instructive for Texas policymakers designing AI governance.

GDPR imposed the same compliance architecture on a 10-person marketing agency as on Google. Large enterprises absorbed the cost by adding compliance departments. Small businesses faced the same requirements with a fraction of the resources. An MIT Sloan study found GDPR effectively functions as a 25% tax on small firms,[17] with compliance costs ranging from $1.7 million for SMEs to $70 million for large enterprises.[18] Cookie consent banners replaced actual privacy governance. Compliance consulting became an industry unto itself. And the small operators GDPR was meant to protect became its primary casualties.

The pattern is predictable: regulation designed at the enterprise level, applied uniformly, creates a two-tier compliance landscape. Enterprises achieve certification and market it as competitive advantage. Small businesses either ignore the requirements or spend disproportionate revenue on compliance that does not meaningfully improve the thing the regulation was supposed to improve.

Texas is watching the same dynamic emerge with AI governance. The legislature has passed mandates. Training programs are being certified. Governance frameworks are being discussed. And the 3 million small businesses that employ the majority of Texans working in regulated industries have no seat at the table, no voice in the design, and no affordable path to compliance.

Texas's regulatory tradition—pragmatic, business-friendly, operationally grounded—should produce a different outcome. But only if the people designing governance frameworks understand what it actually looks like to run a small business in a regulated industry while trying to figure out AI at the same time.

[17] MIT Sloan Management Review, "GDPR Reduced Firms' Data and Computation Use," September 2024. Found GDPR functions as approximately a 25% tax on small firms.

[18] GDPR compliance costs range from $1.7 million for SMEs to $70 million for large enterprises. IT Pro, "GDPR costs are forcing firms to rethink data strategies," February 2024, citing prior survey evidence.

4. What Texas Has Done So Far

Texas has moved faster than most states on AI governance. The legislature has passed two relevant mandates, and the structural framework is in place. Understanding what exists—and what it does not cover—is essential to designing what comes next.

4.1 TRAIGA and the AI Council

The Texas Responsible Artificial Intelligence Governance Act (TRAIGA) establishes the Texas Artificial Intelligence Council as a seven-member body with authority to conduct training programs for state agencies and local governments on the use of artificial intelligence systems.[19] The Council may also issue reports to the legislature on AI compliance, ethics, data privacy, and legal risks. The statutory language is broad by design—the Council must deliver training, but the law does not prescribe content, format, duration, delivery mechanism, or competency standards.

4.2 House Bill 3512

HB 3512, effective September 1, 2025, mandates annual AI training for state and local government employees who use computers for at least 25% of their duties, as well as elected and appointed officials.[20] The Texas Department of Information Resources must certify at least five training programs annually. DIR has modeled this program on the existing cybersecurity awareness training framework.

The parallel is instructive—and reveals the gap. Cybersecurity awareness training teaches employees to recognize phishing emails and use strong passwords. It does not teach them to architect security systems. Similarly, AI awareness training will teach employees to recognize that they are using AI. It will not teach them to govern AI-assisted decisions in regulated environments.

4.3 Two Mandates, One Gap

Texas now has two separate mandates requiring AI training for government workers, administered by two different bodies, with no shared curriculum framework, no governance literacy requirements, and no mechanism for ensuring the training addresses the actual risk profile of regulated operations.

More critically, neither mandate addresses the private sector. The small businesses deploying AI in regulated workflows today—the law firms, the energy consultancies, the financial advisors, the healthcare practices—are entirely outside the scope of current training requirements. They are subject to the regulatory obligations that TRAIGA and existing industry regulations create, but receive no direction on how to meet them.

[19] Texas Responsible AI Governance Act (TRAIGA), Tex. Gov't Code § 554.102.
[20] House Bill 3512, 89th Texas Legislature (2025), effective September 1, 2025.

> *The risk is not that training will be absent. The risk is that training will be present for government employees and absent for the private sector businesses where the majority of regulated AI decisions are actually being made.*

5. Where the Exposure Is Concentrated

Texas's economy is disproportionately concentrated in industries where AI governance matters most. In each of the following sectors, AI is not merely a productivity tool. It is a decision-influencing system operating within a web of regulatory obligations that no current training program addresses.

5.1 Energy

Texas's energy sector employs hundreds of thousands of workers whose decisions involve regulatory compliance with the Railroad Commission, EPA, SEC, and FERC. AI is already embedded in workflows that carry direct regulatory exposure:

- AI-generated reserves commentary included in investor presentations and SEC filings, where hallucinated production estimates become disclosure violations
- AI summarizing environmental compliance language for permit applications, where fabricated or outdated regulatory references trigger enforcement
- AI assisting in FERC reporting, where unverified data in rate filings or market activity reports creates audit liability

Houston's transformation from energy capital to AI hardware manufacturing center is accelerating the need for governance-competent operators at enormous scale. Even sophisticated firms are operating like small businesses when it comes to AI governance—deploying tools into production workflows with little documentation and less oversight.

5.2 Legal Services

Tens of thousands of attorneys, paralegals, and legal staff are governed by the Texas Disciplinary Rules of Professional Conduct. AI-assisted legal work is generating specific, documented liability:

- AI-hallucinated case citations filed in court—now occurring at a rate of two to three incidents per day nationwide, with a Texas case already on record
- AI drafting settlement terms and contract language where fabricated or misapplied provisions create enforceability disputes
- AI assisting discovery summaries where incomplete or hallucinated document characterizations compromise litigation strategy and privilege review

CLE requirements create a natural delivery mechanism for governance training—but no governance-specific CLE program exists at scale in Texas.

5.3 Healthcare

Over 1.7 million Texans work in clinical and administrative healthcare roles regulated by CMS, HIPAA, and the Texas Medical Board. SB 1188, effective September 1, 2025, requires licensed practitioners to review AI-generated records[21]—but AI is already operating in workflows where unreviewed output creates direct exposure:

- AI-generated chart summaries affecting billing codes, where hallucinated or incomplete clinical details trigger CMS audit flags and reimbursement clawbacks
- AI-assisted prior authorization documentation where fabricated clinical justifications create fraud exposure
- AI drafting patient communications where inaccurate treatment descriptions generate informed consent liability

The gap between SB 1188's mandate and the availability of governance training to meet it is immediate and widening.

5.4 Financial Services

Banking, investment advisory, and insurance operations operate under SEC, FINRA, and state banking regulations. AI is already generating client-facing and compliance-critical output:

- AI-generated investment summaries included in client reports, where hallucinated performance data or risk characterizations create suitability violations
- AI-assisted suitability language in recommendation letters, where fabricated or misapplied regulatory standards expose advisors to FINRA enforcement
- AI producing credit decisioning narratives where undocumented algorithmic influence creates fair lending compliance exposure

Documentation requirements for these workflows are not optional—they are the basis of regulatory examination. Firms using AI without governance documentation are not merely underperforming. They are generating audit findings they do not yet know about.

5.5 The Small Business Multiplier

In each of these sectors, small businesses represent the operational backbone. The solo attorney. The boutique energy consultancy. The independent financial advisor. The small medical practice. These operators face the same regulatory obligations as their enterprise counterparts with a fraction of the resources. They adopted AI first because they had to—competitive pressure left no alternative. They will be governed last because no one designed governance for their operational reality.

6. What Current Training Gets Wrong

[21]SB 1188, 89th Texas Legislature (2025), effective September 1, 2025. Requires licensed practitioners to review AI-generated medical records.

The AI training market has responded quickly to demand — and the programs available today represent a necessary first step. But in regulated industries, the next step is the one that determines defensibility. Understanding what current programs were designed to do, and what they were not, is essential to building what the regulated workforce actually needs.

6.1 The Prompt Engineering Fallacy

The dominant AI reskilling paradigm teaches workers to be better users of AI tools—prompt engineering, workflow automation, AI-assisted analysis. Skills that increase throughput. In regulated environments, increased throughput without governance oversight means workers produce more AI-assisted output faster, with the same or greater regulatory risk per output. The result is not improved productivity. It is accelerated risk accumulation.

A paralegal trained in prompt engineering can produce twice as many legal research memoranda per day using AI. Without governance training, each memorandum carries the same risk of unverified citations, fabricated case references, and compliance violations. Doubling output doubles exposure.

The more skilled a worker becomes at prompting AI to produce polished output, the harder hallucinations become to detect. Prompt engineering without governance literacy does not solve the problem. It makes the problem invisible.

6.2 Awareness Is Not Competency

Awareness training creates recognition: employees learn that AI tools exist, that they have limitations, and that organizational policies govern their use. Governance competency creates operational capability: employees learn to evaluate AI output for accuracy, document AI's role in decisions, recognize when AI use crosses regulatory boundaries, and escalate situations that exceed their governance authority.

The distinction matters because compliance audits and enforcement actions do not test awareness. They test whether the organization maintained adequate governance controls. An employee who is "aware" that AI can be inaccurate but lacks the competency to verify AI output in their specific regulatory context has not reduced organizational risk. They have merely acknowledged it.

6.3 Enterprise-Only Design

Current governance programs are designed for organizations with compliance departments, training budgets, and the capacity to absorb multi-day curriculum. A Railroad Commission inspector, an Attorney General's office researcher, a Health and Human Services case worker—each operates in a specific regulatory environment that generic training cannot address. But at least these government employees will receive some form of mandated training.

The five-person law firm, the ten-person energy consulting shop, the solo financial advisor—these operators are entirely outside the system. No mandate reaches them. No affordable training serves them. And no one is explaining what they need to do.

6.4 The Foundation Is Built. The Next Layer Is Not.

None of this is an indictment of the organizations building AI training today. The managed service providers standing up AI platforms for municipal governments, the vendors earning DIR certification, the consultancies delivering AI awareness programs to state agencies—they are doing necessary work. They are building Layer 1: the awareness foundation that every governance system requires. Without that foundation, nothing else is possible.

The problem is not that the foundation is wrong. The problem is that the foundation is being mistaken for the finished structure. Awareness training, AI tool deployment, and platform adoption are prerequisites for governance—not substitutes for it. The gap this paper identifies is not between good providers and bad ones. It is between what Layer 1 was designed to accomplish and what regulated industries actually require to defend their decisions under audit, enforcement, and litigation.

The governance framework proposed in Section 7 is designed to build on top of these existing programs, not to replace them. Organizations that have already invested in AI awareness training, platform deployment, or workforce reskilling have a head start—they have a governance-ready workforce that needs governance competency, not a workforce that needs to start from zero. The framework treats current market offerings as the floor and provides the architecture for the ceiling.

The opportunity for the Texas AI ecosystem is not competition between awareness and governance. It is integration. The providers delivering AI platforms and awareness training today are the natural distribution channels for governance competency tomorrow. The question is whether that next layer gets built deliberately—or whether the market continues to treat awareness as the endpoint while regulated businesses accumulate unmanaged risk.

7. A Governance Framework That Actually Works

The following framework—the Digital Information Governance® Framework—is designed for the real economy. Not just for enterprises or government agencies, but for the regulated businesses of every size that are already using AI and need a practical path to governance competency.

7.1 Three-Layer Architecture

Layer	Purpose	Audience	Outcome
Layer 1: AI Awareness	Baseline AI literacy, tool recognition, policy awareness	All employees using AI	Recognition that AI is present and has limitations
Layer 2: Governance Literacy	Output evaluation, decision accountability, regulatory awareness, escalation judgment	Employees in regulated functions	Operational competency to govern AI-influenced decisions

Layer 3: Industry Governance	Sector-specific regulatory compliance, validation methodology, documentation systems	Professionals in energy, legal, healthcare, financial services	Audit-ready documentation and safe harbor capability

Layer 1 is already being addressed by DIR's vendor certification process for government employees. The framework does not replicate that effort. It builds on it.

Layer 2 is the critical gap—the governance competency that no current training program provides at scale or at a price point accessible to small business.

Layer 3 is the industry-specific application that makes governance operationally relevant rather than theoretically complete.

7.2 The Four Governance Competencies

Layer 2 transforms AI-aware workers into governance-competent professionals through four competency modules:

- **Output Evaluation:** Critically assess AI-generated content for accuracy, completeness, and fitness for regulatory use. Recognize hallucination patterns. Distinguish AI-retrieved information from AI-generated fabrication. Cross-reference against authoritative sources.
- **Decision Accountability:** Trace and document the role of AI in decision-making processes. Identify which decisions are AI-influenced. Maintain compliance audit trails. Recognize when human judgment must override AI output regardless of confidence scores.
- **Regulatory Mapping:** Map actual job functions to the regulatory requirements that govern them. Identify where AI involvement creates new compliance obligations. Maintain a personal Regulatory Impact Map.
- **Escalation Protocol:** Recognize when an AI-related situation exceeds governance authority and must be escalated. This prevents the most damaging failure mode: workers making AI-influenced decisions in regulatory contexts they are not qualified to assess.

7.3 Industry-Specific Modules

Layer 3 delivers sector-specific governance for the industries where AI decisions carry the highest regulatory stakes:

- **Energy:** Railroad Commission compliance, SEC disclosure obligations for AI-influenced reserves estimation, EPA compliance for AI-assisted environmental monitoring, FERC regulatory filing standards.
- **Legal:** State Bar Disciplinary Rules as applied to AI-assisted work, citation verification requirements, client communication standards, confidentiality obligations. Deliverable as CLE-accredited continuing education.

- **Healthcare:** HIPAA compliance for AI-processed PHI, Texas Medical Board standards per SB 1188, CMS compliance for AI-influenced billing and benefits, documentation standards for AI-assisted diagnostics.
- **Financial Services:** SEC and FINRA compliance for AI-influenced recommendations, fair lending compliance for AI credit decisioning, documentation for algorithmic trading and automated advisory systems.

7.4 Framework Alignment

The proposed governance architecture aligns structurally with the National Institute of Standards and Technology (NIST) AI Risk Management Framework, particularly the "Govern," "Map," "Measure," and "Manage" functions, while preserving Texas-specific implementation authority.[22]

8. Making Governance Work for Small Business

This is not an afterthought section. It is the core equity argument of this paper.

If the governance framework Texas builds only works for organizations with compliance departments, it will produce the same outcome GDPR produced in Europe: large enterprises achieve certification and market it as competitive advantage, while small businesses—which collectively employ far more Texans—operate outside the system entirely. The governance gap does not close. It institutionalizes.

8.1 The Micro-Governance Approach

For businesses with fewer than 50 employees, the framework delivers governance literacy through a compressed format designed around how small businesses actually learn: in short bursts, applied immediately, tied to actual business operations.

- **Compressed Delivery:** Core governance literacy delivered in two half-day sessions rather than multi-day curriculum. Focused on the highest-frequency governance decisions relevant to the business's industry.
- **Embedded Assessment:** Governance competency validated through scenario-based exercises using the business's actual AI tools and workflows, not generic case studies.
- **Owner-as-Governance-Officer:** In businesses without dedicated compliance staff, the owner or principal serves as the governance authority. Training equips them to set AI use policies, recognize escalation triggers, and maintain basic documentation.
- **Template-Driven Documentation:** Pre-built governance templates (AI use policies, decision logs, incident reports) that small businesses customize rather than create from scratch. These templates generate the documentation needed to invoke safe harbor provisions.

[22]National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0), January 2023.

8.2 The Cost Reality

A governance program that costs $5,000 per employee is not a solution for a business with ten employees and $800,000 in annual revenue. It is a cost that guarantees non-compliance.

The framework must be deliverable at the price point the market can absorb. The 2026 baseline found 59% of small businesses budgeting $100–$500 per employee annually for AI training. Governance training must meet them there—not at the enterprise price point that only reinforces the two-tier compliance landscape.

This does not mean diluting governance competency. It means designing for efficiency: compressed delivery, template-driven documentation, scenario-based assessment, and modular industry content that a small business owner can implement without hiring a consultant or a compliance officer.

8.3 Governance as Growth Infrastructure

The Goldman Sachs 10,000 Small Businesses (10KSB) program focuses on disciplined growth, financial management, and operational scaling. A recurring theme in growth-stage businesses is that operational expansion frequently outpaces governance infrastructure. AI adoption follows the same trajectory—businesses deploy AI to remain competitive before formalizing compliance systems. In regulated industries, this sequencing creates compounding exposure.

Governance literacy must be embedded into growth frameworks rather than treated as a post-growth compliance function. For small enterprises, governance is not a department. Governance is a leadership competency. The reskilling framework is compatible with 10KSB pedagogy: applied, practical, and directly tied to operational risk management.

> *Reskilling is not the budget alternative to governance infrastructure. For small businesses, reskilling is the governance infrastructure.*

9. Measuring What Matters: Competency, Not Completion

The existing training paradigm measures compliance through completion certificates. An employee watches a module, clicks through an assessment, and receives documentation of "training." In cybersecurity, this model has proven insufficient—phishing attacks succeed against organizations with 100% completion rates. In AI governance, the same model will produce the same result: documented compliance without actual competency.

9.1 Competency Assessment

The framework measures governance capability through scenario-based assessment, not knowledge-recall testing:

- **Output Evaluation:** Workers review actual AI-generated content relevant to their role and identify errors, hallucinations, and regulatory compliance issues.

- **Decision Accountability:** Workers trace AI involvement through realistic multi-step decision scenarios and produce audit-ready documentation.
- **Regulatory Mapping:** Workers map actual job functions to applicable regulations and identify AI-related compliance obligations.
- **Escalation Judgment:** Workers are presented with escalating scenarios and must identify correct escalation triggers and appropriate responses.

9.2 Organizational Readiness Score

Beyond individual competency, the framework measures organizational governance readiness: the percentage of AI-using employees who have achieved governance competency, the coverage of governance documentation across AI-influenced decisions, time-to-escalation for AI governance incidents, and safe harbor documentation completeness.

9.3 Governance Readiness Scorecard

Organizations can benchmark their maturity using a simple self-assessment:

Category	Assessment Question
AI Inventory	Do we know where AI is being used?
Policy	Is there a written AI use policy?
Training	Have AI-influencing roles received governance literacy training?
Escalation	Is there a documented escalation protocol?
Documentation	Are AI-influenced decisions logged and reviewable?
Oversight	Is responsibility for AI governance formally assigned?
Review	Are high-risk AI use cases periodically reviewed?

10. The Economic Argument

Workforce governance is not a cost center. It is infrastructure investment that determines whether Texas captures or loses the economic value of AI adoption in regulated industries.

Texas's $469 billion technology sector is growing at a rate that positions the state to capture 15–20% of North American AI investment by 2028. That investment will flow to jurisdictions where companies can deploy AI with confidence in both the regulatory environment and the workforce. A governance-literate workforce is as essential to attracting AI investment as energy infrastructure and innovation-friendly regulation.

The inverse is equally true. States with AI regulations but governance-illiterate workforces create an environment where regulation functions as friction rather than framework. Compliance

becomes a cost without a corresponding competency. Companies deploy elsewhere—not to avoid regulation, but to find workforces that can operate within it.

But the economic argument must be honest about cost. If governance compliance requires $50,000–$100,000 in training, documentation, and audit preparation for a small business, that is not a governance framework. It is a market-exit mechanism. GDPR drove an estimated 30% of small EU businesses to reduce their digital operations rather than absorb compliance costs. Texas cannot afford a governance regime that produces the same result in the regulated industries where small businesses represent the operational backbone.

Governance-literate workforces reduce regulatory uncertainty, which lowers perceived compliance risk and strengthens capital allocation confidence in regulated sectors.

10.1 Liability Exposure Without Governance

Where AI influences regulated decision-making without documentation or validation, exposure increases across several vectors:

- Administrative enforcement actions triggered by AI-generated errors in compliance filings
- Audit findings revealing undocumented AI involvement in regulated decisions
- Civil liability claims from clients or counterparties harmed by hallucinated output
- Contractual disputes where AI-generated deliverables contain fabricated information
- Public records challenges exposing AI-influenced government decisions without audit trails

The risk is not catastrophic system failure. The risk is routine, unverified AI-assisted decisions accumulating unnoticed within regulated workflows. A hallucinated citation in one legal brief. A fabricated regulatory reference in one compliance filing. A misrepresented data point in one financial report. Each individually minor. Collectively, a pattern that no insurer, no regulator, and no court will treat as excusable once the scale becomes visible.

> *The state that builds a governance-literate regulated workforce first becomes the default jurisdiction for responsible AI deployment in energy, legal, healthcare, and financial services. Texas has every structural advantage. What it lacks is the implementation guide—and the affordable delivery mechanism for the businesses that need it most.*

11. What This Means for Executive Leadership

AI governance is no longer an IT concern or a compliance department project. It is a board-level issue with direct implications for audit defensibility, regulatory standing, and competitive positioning.

AI governance is now a fiduciary obligation. In regulated industries, the decisions AI influences carry the same legal and regulatory weight as decisions made entirely by humans. The difference is documentation. When a regulator, auditor, or opposing counsel asks how a decision was made, organizations without governance infrastructure cannot demonstrate that AI involvement was identified, reviewed, or validated. The absence of documentation does not prove the decision was wrong. It proves the organization cannot defend it.

Failure to document AI involvement undermines audit defensibility. Every AI-influenced decision that lacks traceability is a potential audit finding. This includes not only the obvious cases—AI-drafted regulatory filings, AI-generated financial analysis—but also the routine decisions where AI played an undocumented supporting role. When enforcement catches up to adoption, the exposure will not be limited to the decisions AI got wrong. It will extend to every decision where AI involvement was not recorded.

AI risk is not technology risk. It is decision integrity risk. The relevant question is not whether an organization's AI tools are technically sound. It is whether the decisions those tools influence can withstand regulatory scrutiny, client challenge, and judicial review. Decision integrity requires governance infrastructure: documentation, validation, escalation protocols, and competency in the people who rely on AI output. Technology risk is managed by IT. Decision integrity risk is managed by leadership.

Governance maturity will increasingly differentiate firms. The firms that build governance infrastructure now will operate with lower regulatory friction, stronger audit positions, and greater client confidence than competitors who delay. In regulated industries, governance is not a cost of doing business. It is a competitive asset. The managing partner, the CEO, the board—these are the people who determine whether governance is treated as overhead or as strategic investment. The firms that treat it as investment will outperform.

> *Your competitors are already using AI. The question is not whether to adopt it. The question is whether your firm governs it well enough to defend every decision it touches—and whether your competitors can say the same.*

Executive Brief: 5 Questions for AI Defensibility

Before a regulatory inquiry, client audit, or enforcement action occurs, leadership should be able to answer the following:

1. The Inventory Check

Do we have a centralized registry of every AI tool—including shadow AI—currently processing company data?

2. The PII Perimeter

Can we prove that no Personally Identifiable Information is being used to train public AI models through employee prompts?

3. The Hallucination Filter

What is our documented process for human-in-the-loop verification of AI-generated regulatory filings, client deliverables, and compliance documents?

4. The Escalation Trigger

Does staff know exactly when an AI-influenced decision must be halted and escalated to legal or compliance?

5. The Audit Trail

If challenged today, could we produce a decision log showing the human's role in every AI-assisted output submitted to a regulator, court, or client?

If your organization cannot answer all five with confidence, the governance gap this paper describes is operating in your workflows now.

12. Conclusion

AI is already making decisions in Texas regulated industries. The question is not whether governance is needed. Two statutes confirm that it is. The question is whether governance will be designed for the real economy or for regulatory abstraction.

The risks are not theoretical. Shadow AI is operating in regulated workflows today. Hallucinations are generating fabricated content in legal, financial, and compliance documents today. Small businesses are deploying AI at scale with no governance infrastructure, no awareness of regulatory obligations, and no affordable path to compliance today.

If governance frameworks develop without addressing these realities, the result will mirror GDPR: administrative burden concentrated on small operators, compliance theater that satisfies documentation requirements without reducing actual risk, and a two-tier system where enterprises get certified and small businesses get left behind.

The framework presented in this paper provides a different path. Three layers of training, from awareness to governance literacy to industry-specific compliance. Four governance

competencies that transform AI-aware workers into governance-capable professionals. Sector-specific modules for the industries that drive Texas's economy. A small business model that makes governance accessible and affordable at every organizational scale. And a measurement system that tests competency, not just completion.

> **Governance makes trust defensible. Without it, AI scales risk.** *And without affordable, accessible governance infrastructure, small businesses do not scale governance. They scale exposure.*

Next Steps

The path to AI defensibility requires action at two levels. For the broad workforce—the small businesses, the growing firms, the teams that need governance literacy built into daily operations—the immediate priority is closing the competency gap through scalable, affordable reskilling. For high-stakes regulated environments—where a single unverified AI output can trigger enforcement, litigation, or audit failure—the priority is structural: comprehensive AI audits, PII exfiltration assessment, and the deployment of validation systems that ensure every AI-influenced decision can withstand scrutiny.

Both levels are necessary. Neither is sufficient alone. Training without audit leaves structural vulnerabilities undetected. Audit without training leaves the workforce unable to sustain governance after the auditors leave. The organizations that build both will outperform.

For policymakers and the Texas AI Council, this paper recommends three immediate actions. First, establish governance competency standards for DIR-certified AI training programs that go beyond awareness to include the output evaluation, decision accountability, regulatory mapping, and escalation competencies outlined in the Digital Information Governance® Framework. Second, extend AI training mandates to include private-sector businesses operating in regulated industries—not as punitive compliance requirements, but as tiered, affordable certification pathways calibrated to organizational size and risk profile. Third, create a small business governance pilot program through an existing state economic development channel to test compressed delivery models and validate that governance competency can be achieved at the $100–$500 per-employee price point the market requires.

For organizations ready to act now, the governance gap does not require waiting for legislative direction. The framework in this paper is operational. Firms in energy, legal, healthcare, and financial services can begin with a governance readiness assessment, implement the documentation templates and escalation protocols described in Section 8, and establish the audit trails that will determine defensibility when enforcement arrives. The infrastructure exists. The question is whether your organization builds it before or after the first audit finding.

Appendix A: 2026 Workforce Pulse Baseline – Methodology and Data

AI in Your Business: A 90-Second Pulse Check was administered by EWR Digital in February 2026 as a ten-question survey targeting Texas small business operators. The instrument was designed to establish a directional governance readiness baseline rather than a statistically representative sample. Responses (n=34) were collected February 22–26, 2026.

Key Findings Summary

Metric	Finding
Active AI usage	94% use AI (71% regularly, 24% occasionally)
Written AI use policy	Only 9% have formal policy; 79% have none
AI training received	85% have received zero AI training
Regulatory awareness	82% unaware of applicable AI regulations
Penalty awareness	88% unaware of AI governance penalty structures
Concern level	56% somewhat or very concerned about AI liability
Training demand	59% want both skills and governance training
AI users without policy	78% of active AI users have no written policy
Concerned but unprotected	68% concerned about liability but have no policy

Investment Willingness

53% of respondents indicated willingness to invest under $1,000 annually in AI training, 26% indicated $1,000–$5,000, and 12% indicated $5,000–$10,000. On a per-employee basis, 59% indicated $100–$500 annually, with an additional 15% indicating $500–$2,000. Only 9% indicated AI training was not a budget priority.

Demand Signal

When asked which resources would be most valuable, respondents selected (multiple selections permitted): AI skills training (82%), AI policy templates for immediate implementation (68%), AI risk audit of current operations (53%), and AI governance training for regulatory environments (38%).

The combination of high adoption, low governance, regulatory unawareness, and expressed willingness to invest confirms a market ready for governance training but not yet reached by existing programs.

About the Author

Matthew Bertram is the President of ModalPoint, an advisory firm specializing in Digital Information Governance® and operational risk for regulated enterprises in energy, legal, and financial services. He is also CEO of EWR Digital, a 26-year Houston-based firm delivering AI governance and digital strategy for complex B2B environments.

Matthew is a Certified AI Professional (CAIP) through OxEthica and a member of the International Association of Privacy Professionals (IAPP). He holds advanced AI certifications from Harvard Business School (AI Essentials for Business) and Oxford University (Artificial Intelligence Programme; AI Ethics, Regulation and Compliance Programme).

As a scholar in the Goldman Sachs 10,000 Small Businesses program and a FINRA Communications Compliance certified professional, Matthew bridges the gap between small business operational reality and enterprise-grade regulatory requirements. He is a Texas A&M University graduate and Corps of Cadets alumnus.

He holds two provisional patents (#63/948,546 and #63/948,571) covering authority-hierarchical validation and non-probabilistic governance override systems—the technical foundation for ModalPoint's AI audit methodology. His Digital Information Governance® framework underpins the governance architecture proposed in this paper.

Contact
Matthew Bertram | matthewbertram.com
EWR Digital | ewrdigital.com • ModalPoint | modalpoint.com
linkedin.com/in/matthewbertram