WHITE PAPER

Human Decisions, Machine Research

Aligning AI-Assisted Title Research with the Texas Department of Insurance's Guidance on Artificial Intelligence

RecordsOnline

ChainOfTitleAI · RunsheetAI · CommitmentAI

June 2026

A position paper on responsible AI architecture in the title and settlement industry

Contents

Executive Summary

On June 12, 2026, the Texas Department of Insurance (TDI) issued Commissioner's Bulletin B-0003-26, reminding every regulated entity — and the third parties that serve them — that decisions affecting consumers must comply with existing insurance law whether or not artificial intelligence is involved.[1] The bulletin does not ban AI. It draws a line: AI may inform a decision, but a person must own it.

RecordsOnline built its products on that line before the bulletin was issued. ChainOfTitleAI, RunsheetAI, and CommitmentAI accelerate the research that underlies title work — assembling chains of title, compiling runsheets, and surfacing the documents and facts a title professional needs — while leaving every consequential determination to a licensed human. RecordsOnline also keeps customer data inside its own environment; it does not transmit client records or work product to third-party AI engines such as Claude, ChatGPT, or Gemini.

This paper explains the regulatory landscape behind the bulletin, what the bulletin asks of regulated entities and their vendors, and how a "machine researches, human decides" architecture maps onto those expectations. It also addresses, candidly, what architecture alone does not accomplish — governance documentation, bias and error testing, and examination readiness — and what an organization should put in place to close that gap.

1. A Regulatory Inflection Point for AI in Insurance

Regulated entities have always adopted new tools for efficiency. What distinguishes this moment is speed and scale: AI can shape large volumes of consumer-affecting work quickly, and regulators have responded by clarifying that the use of a new tool does not suspend old obligations. TDI's stated priority is consumer protection, and its message is direct — the standards that govern a decision apply regardless of the method used to reach it.

The three reference points

The bulletin sits atop a national framework. In 2020 the NAIC adopted its Principles on Artificial Intelligence, asking AI actors — expressly including third parties such as data and advisory providers — to keep AI fair and ethical, accountable, compliant, transparent, and secure.[2] In 2023 the NAIC followed with its Model Bulletin on the Use of Artificial Intelligence Systems by Insurers, which a majority of states have since adopted in full or substantially similar form.[3] Texas did not adopt that Model Bulletin; instead, TDI issued its own, lighter-touch guidance in B-0003-26, which expressly recognizes the NAIC's 2020 Principles as an appropriate guide. At the state level, the Texas Department of Information Resources adopted an AI Code of Ethics built

[1] Texas Department of Insurance, Commissioner's Bulletin No. B-0003-26, Use of Artificial Intelligence Systems (June 12, 2026).

[2] National Association of Insurance Commissioners, Principles on Artificial Intelligence (adopted Aug. 2020).

[3] National Association of Insurance Commissioners, Model Bulletin on the Use of Artificial Intelligence Systems by Insurers (adopted Dec. 2023).

on seven principles — human oversight, fairness, accuracy, redress, transparency, privacy, and security — which TDI expressly recommends companies review.[4]

Why a title-records company is in scope

The bulletin's reach is not limited to carriers. TDI states that its expectations "*extend to any third party working with a regulated entity*." Title insurance is insurance; title agents and underwriters are regulated entities; and the vendors whose tools shape their work product fall squarely within the bulletin's contemplated scope. A research platform that feeds a title commitment is precisely the kind of third-party system the guidance has in mind. For a company like RecordsOnline, that is not a liability to manage so much as a standard to meet — and, increasingly, a question its customers will ask before they buy.

2. What the Bulletin Requires

Existing law applies

The throughline of the bulletin is that AI changes the tooling, not the rules. Decisions reached with AI must satisfy the same legal standards as any other decision — at a minimum, that they are not inaccurate, arbitrary, capricious, or unfairly discriminatory. The relevant authority is broad; the table below summarizes the chapters TDI cites and why each bears on AI-assisted title work.

Authority	What it governs	Why it bears on AI-assisted title work
Title 11 (Texas Title Insurance Act)	Regulation of title insurance, agents, examiners, and promulgated forms.	Title work is insurance. The Act's standards govern the work product whether or not AI assists the underlying research.
Ch. 541 (Unfair Methods of Competition; Unfair or Deceptive Acts)	Prohibits unfair or deceptive acts and practices in the business of insurance.	An AI-influenced outcome that misleads or unfairly treats a consumer violates the same standard a manual one would.
Ch. 544 (Prohibited Discrimination)	Bars unfair discrimination, including discrimination tied to geographic and other protected bases.	Title research is inherently geographic. A research layer that treats records unevenly across areas can implicate this chapter.
Ch. 4201 (Utilization Review Agents)	Governs utilization review; prohibits AI from making an adverse determination.	Does not bind title agents directly, but signals how far regulators will go to keep humans in control of high-stakes adverse decisions.
Ch. 401 (Audits & Examinations) and Ch. 751	TDI's authority to examine regulated entities and require production of records.	TDI may examine AI use and require production of books, papers, procedures, and protections on request.

[4] Texas Department of Information Resources, AI System Code of Ethics, 1 Tex. Admin. Code § 219.11, adopted under Tex. Gov't Code ch. 2054, subch. S (SB 1964, 89th Leg.).

Authority	What it governs	Why it bears on AI-assisted title work
(Market Conduct Surveillance)		

A human must own consequential decisions

The single most operationally specific instruction in the bulletin is that, where AI is used to make a consequential decision, TDI expects a person to review and agree with all decisions before any action is taken. This is not a request for a human somewhere in the vicinity of the process; it is a requirement that a person stand behind the outcome.

One provision goes further. Chapter 4201, governing utilization review agents, prohibits using AI to make an adverse determination at all.[5] Title agents do not perform utilization review, so the prohibition does not apply to them directly — but it signals how far regulators are willing to go in the highest-stakes adverse-decision settings, and it underscores the wisdom of reserving determinations to humans by design rather than by policy.

Governance and testing are expected

TDI identifies governance, risk-management controls, and internal audit as core to keeping AI-driven outcomes lawful, and it encourages verification and testing methods to identify errors and bias. The clear implication is that human review is necessary but not sufficient: organizations are expected to demonstrate that the system feeding their people is itself sound.

You must be able to show your work

Finally, the bulletin reminds entities that TDI may examine AI use through market-conduct actions and product filings, and that books and papers relating to the business must be produced on request.[6] An organization using AI should expect inquiries into its governance frameworks, risk management, data and privacy protections, and internal controls — and should be ready to furnish its procedures and protections when asked.

3. The RecordsOnline Design Philosophy: AI Researches, Humans Decide

RecordsOnline's products rest on two deliberate choices that anticipate the bulletin's expectations.

Choice one: AI researches, humans decide

The products accelerate research, not judgment. ChainOfTitleAI assembles and orders the chain of title from the underlying instruments; RunsheetAI compiles the runsheet from the records it gathers; and CommitmentAI organizes the inputs a commitment draws on. In every case, the product delivers research for a professional examiner to review. The decisions that

[5] Tex. Ins. Code ch. 4201 (Utilization Review Agents).

[6] Tex. Ins. Code ch. 401 (Audits and Examinations); see also ch. 751 (Market Conduct Surveillance).

carry legal weight stay with the human. The consequential decisions in title work — what goes onto a commitment, what gets flagged as a defect or exception, what the chain ultimately shows — are exactly the ones the bulletin wants a person to own, and the architecture keeps them there.

As you can see from the picture below, the AI automation provides a research layer of processing. In this context, the research can identify anomalies and create audit insights into the data and documents themselves. But the automation does not “approve” the results. All research is presented for a human to review, modify as needed, and formally approve.

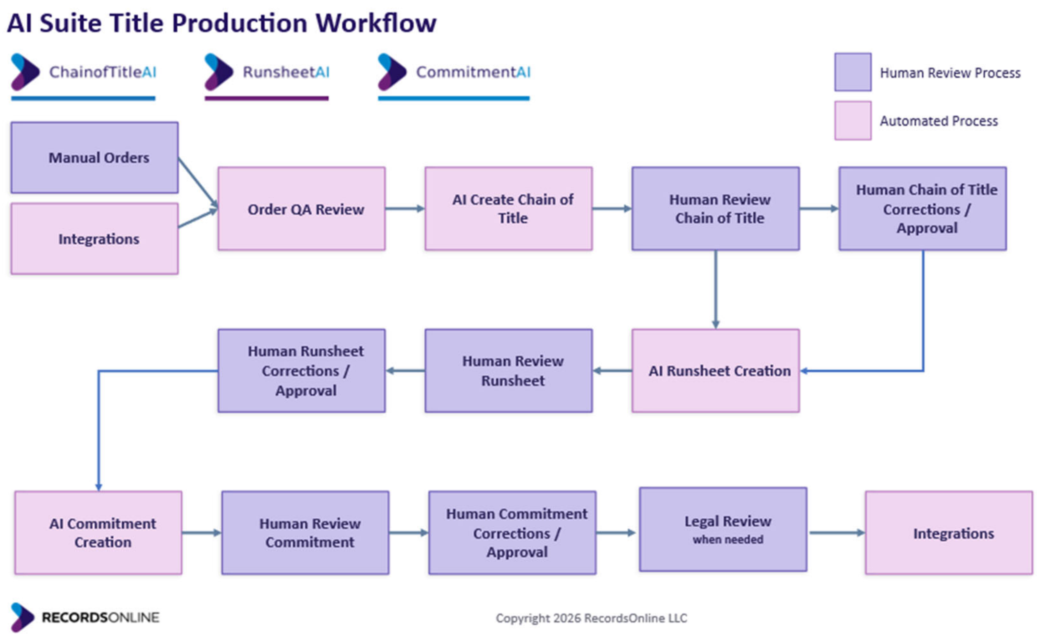

AI Suite Title Production Workflow

Choice two: data stays in-house

RecordsOnline does not transmit client records, document content, or work product to external, general-purpose AI engines such as Claude, ChatGPT, or Gemini. Processing stays within RecordsOnline’s controlled environment. This directly addresses the data-privacy and security expectations the bulletin names, and it removes an entire category of third-party exposure — the risk that sensitive consumer and property data leaves the organization’s control the moment an AI feature is used.

There are significant concerns when allowing an AI engine to have access to the actual data. Beyond the privacy issue is the very real concern that your data will be stored outside of a controlled environment and used to further the AI provider’s business model without your knowledge. In RecordsOnline’s use of AI, a key distinction from other solutions is the use of test data to train any AI solution — production data is never allowed outside of the RecordsOnline environment.

4. Mapping the Architecture to TDI’s Expectations

The table below maps the bulletin’s principal expectations to the way a research-versus-decision architecture responds to each.

TDI Expectation	How the Architecture Responds
A person must review and agree with consequential decisions before action is taken.	Determinations — what enters a commitment, which exceptions and requirements apply, whether a defect exists — are reserved to the professional examiner. AI prepares; the human decides.
Decisions must not be inaccurate, arbitrary, capricious, or unfairly discriminatory.	Human ownership keeps the professional examiner accountable for the judgment at the point of decision, rather than deferring to an opaque automated output. (Testing of the research layer remains necessary — see Section 5.)
Data and privacy protections; oversight of third-party AI.	Client records, document content, and work product stay within RecordsOnline’s controlled environment and are not transmitted to external, general-purpose AI services.
Robust governance, risk management, and internal controls.	The architecture supplies the central control — a human checkpoint on consequential decisions — around which a documented governance program is built.
Ability to furnish procedures and protections on request.	A bounded, auditable workflow with a defined human-review step is far easier to describe and evidence to an examiner than an end-to-end automated pipeline.

The pattern is consistent: the design directly satisfies the expectations that turn on who decides and where data goes, and it provides the foundation — but not the entirety — of the expectations that turn on governance and evidence.

5. Beyond Architecture: What Responsible Deployment Still Requires

Architecture is the strongest part of a compliance posture, but it is not the whole of one. An organization that wants to stand behind its AI use should pair the design with the following.

- **Documented governance.** A written program that names who owns AI governance, how systems are reviewed and approved, and how the human-review step is enforced. The architecture supplies the control; the document proves it exists.
- **Testing of the research layer.** A reviewer who trusts a skewed research output inherits the skew. Because title work is inherently tied to geographic location — a protected basis

under Chapter 544 — it is worth testing whether the research and surfacing steps treat records evenhandedly across areas, and validating accuracy against known-good results.

- **Meaningful, not nominal, review.** Human ownership only satisfies the bulletin if the human is genuinely positioned to catch and correct errors. Workflows should present enough context for an informed decision and should make overrides easy to record.
- **Examination readiness.** A short, current written description of how the system works, where the human checkpoints sit, and what controls exist — prepared before anyone asks for it.

None of these undercuts the architecture. They convert a sound design into a defensible, demonstrable program — which is what an examiner ultimately looks for.

6. What This Means for Title Agents and Underwriters

For the regulated entities RecordsOnline serves, the bulletin's expectations flow through to the vendors whose tools shape their work product. That makes a vendor's architecture part of the customer's own compliance story. A title agent who can show that the research feeding its commitments comes from a system with a human-owned decision point, isolated data handling, and a describable, auditable workflow is better positioned when an examiner asks how AI is used in the shop.

The practical takeaway for customers is to ask vendors three questions: who makes the consequential decision, where does the data go, and can the workflow be described and evidenced on request. A "machine researches, human decides" design answers the first two cleanly and gives a foundation for the third — provided the governance and testing described in Section 5 are in place.

Conclusion

AI is now part of how title work gets done, and regulators have made clear that the old standards still apply. The most durable response is not to slow the research but to keep the judgment human, keep the data controlled, and be able to show both. RecordsOnline built its products around that principle, and the principle is one any responsible deployment can adopt. In a field built on the reliability of the public record, that is exactly the kind of foundation worth building on.

About this paper

This document is provided for general informational purposes and describes how a human-in-the-loop AI architecture relates to TDI Commissioner's Bulletin B-0003-26. It is not legal advice. Organizations should confirm the application of the bulletin and the cited statutes to their specific operations with qualified counsel.