IS YOUR FIRM READY FOR AI?

The 10-Minute AI Readiness Diagnostic for Law Firms

10 questions in 10 minutes to tell you
whether your AI governance is protecting your clients
or exposing your firm.

From the TosiBriefs Legal AI Deployment Evaluation Kit
Distilled from NIST AI RMF · ISO/IEC 42001 · OWASP Top 10 for LLMs

TosiBriefs
tosibriefs.com

Why this diagnostic exists

Your firm might already use AI tools or might be considering using them. Perhaps it's a contract review platform, a research assistant, a document summarizer, or something a vendor pitched as "transformative." Most firms haven't asked whether the security certifications those vendors wave around, like SOC 2 or ISO 27001, cover the risks introduced by AI, or whether they effectively handle the unpredictability of AI systems.

Traditional security audit wasn't built to catch several risks that come up with AI: a contract review tool that hallucinates a clause that doesn't exist, a research assistant that fabricates case citations, or a due diligence platform that routes client data through a third-party model provider's servers without your knowledge.

Although this diagnostic won't fix those problems, it will, in about ten minutes, tell you whether they exist at your firm and how serious they are.

These ten questions are drawn from the TosiBriefs Legal AI Deployment Evaluation Kit, a 78-item framework built on NIST AI RMF, ISO/IEC 42001, and the OWASP Top 10 for LLMs.

How to score

Rate each item on a 0–5 scale. The scores are for your private diagnosis, so be brutally honest.

Score	What it means
0	No capability or documentation exists
1	Ad hoc or informal processes only; no documentation
2	Some documentation exists but is incomplete or inconsistently applied
3	Documented processes are in place and mostly followed
4	Documented, consistently applied, and periodically reviewed
5	Mature practice with continuous monitoring, evidence, and third-party validation

The ten questions that matter most

1. DATA FLOW

Can you diagram exactly where client data goes when it enters each AI tool—every server, every subprocessor, every cache?

If you can't draw the map, you can't secure the territory. Most firms discover third-party data processors they didn't know existed. Your SOC 2 certification doesn't cover what happens after client data leaves your environment and enters a vendor's AI pipeline.

Your score (0–5): ____ Notes: _____

2. PRIVILEGE

Has your AI vendor tested for data extraction attacks, where an adversary queries the model to reconstruct your client's confidential information?

Attorney-client privilege doesn't survive a training data pipeline. If privileged information can be extracted from a model through targeted queries, you may have an ongoing privilege exposure you don't know about. If the vendor has never tested for this, they cannot tell you it isn't possible.

Your score (0–5): ____ Notes: _____

3. SECURITY

Does the vendor test for prompt injection, where carefully designed input overrides the AI's instructions to ignore adverse contract clauses or fabricate case citations?

Prompt injection is the SQL injection of the AI era. A document submitted for contract review could contain hidden instructions that cause the AI to skip adverse clauses. A research query could trick the system into fabricating case law. Traditional penetration tests don't check for this.

Your score (0–5): ____ Notes: _____

4. SUPPLY CHAIN

Can the vendor identify every model and component in their AI supply chain, including the foundation model provider?

Your AI vendor probably didn't build their model from scratch. They're layering their product on top of a foundation model from OpenAI, Anthropic, Google, or Meta. A vulnerability in that upstream model affects every application built on it, including the one your associates use every day.

Your score (0–5): ____ Notes: _____

5. GOVERNANCE

Is there a designated individual or committee at your firm with defined authority and budget for AI governance—not just "the IT department"?

If nobody owns AI risk, nobody's watching it. "IT handles it" is not a governance structure. A named owner with budget authority and reporting lines to firm leadership is the minimum viable governance structure.

Your score (0–5): ____ Notes: _____

6. SHADOW AI

Do you have technical controls—not just a written policy—preventing attorneys from pasting privileged documents into personal ChatGPT or Claude accounts?

Right now, associates at most firms can access consumer AI tools from firm devices with no detection. A written policy without technical enforcement is insufficient. Consumer AI tools often lack the enterprise privacy protections required for privileged legal work.

Your score (0–5): ____ Notes: _____

7. ACCURACY

Does the vendor publish hallucination rates or accuracy benchmarks specific to legal use cases?

A model that scores 95% on general knowledge benchmarks might hallucinate case citations 30% of the time. General accuracy numbers are meaningless for legal work. If the vendor can't provide legal-specific benchmarks, they either haven't tested or don't like what they found.

Your score (0–5): ____ Notes: _____

8. OVERSIGHT

Are there documented, enforced requirements for human review of all AI-generated work product before it reaches a client or a court?

ABA Model Rule 1.1 requires competent representation. Supervisory obligations under Rule 5.1 don't disappear because the work was generated by a machine. The attorney of record is always responsible, but your governance framework needs to make that operationally clear.

Your score (0–5): ____ Notes: _____

9. INCIDENT RESPONSE

Does your firm have an incident response plan that specifically addresses AI-related incidents?

An AI tool that fabricates case law in a filed brief requires a different response playbook than a phishing attack. A privilege breach through an AI pipeline requires different notification procedures than a laptop theft. If your plan doesn't distinguish between these, you're improvising during a crisis.

Your score (0–5): ____ Notes: _____

10. INSURANCE

Does your malpractice insurance explicitly cover errors originating from AI-generated work product? Have you asked your carrier?

Most firms have never asked their malpractice carrier whether AI-generated errors are covered. Most cyber policies contain exclusions for automated decision-making or third-party AI processing. If your answer to this question is "I don't know," your firm may be carrying uncapped financial exposure from every AI tool in production.

Your score (0–5): ____ Notes: _____

Your score

TOTAL SCORE: _____ / 50

What your score means

40–50 — AHEAD OF THE CURVE

Your firm has stronger AI governance than most. You've addressed the basics and have controls in place across multiple risk areas. The full Legal AI Deployment Evaluation Kit will help you identify remaining blind spots and create documentation to prove your governance maturity to clients and insurers.

25–39 — COMMON GROUND

This is where most firms land. You have some controls in place but still have meaningful gaps that don't cause problems until they do. The difference between a score of 25 and a score of 40 is often the difference between "we're working on it" and "we can prove it to a client who asks." The full evaluation breaks each of these ten areas into 78 detailed assessment items to surface exactly where your gaps are.

10–24 — MATERIAL GAPS

Your firm has significant unaddressed AI risks across multiple domains. Several of the gaps surfaced by this diagnostic likely represent immediate exposure, particularly around privilege protection, insurance coverage, and shadow AI. This is not a score to file away and revisit next quarter. The full evaluation provides a structured framework for prioritizing remediation and building the governance your firm needs before an incident forces the issue.

0–9 — IMMEDIATE ACTION NEEDED

Your firm is operating AI tools with minimal safeguards. The risks identified in this diagnostic—privilege exposure, potential insurance gaps, absent governance, undetected shadow AI—represent current vulnerabilities. Consider pausing AI use for client-facing matters until a minimum governance framework is in place. The full evaluation kit provides a 90-day action plan to move from where you are to a place where you can more confidently deploy AI systems while mitigating the risks.

Next Steps.

This diagnostic sampled ten questions from a much larger evaluation framework. It is designed to tell you whether you have a problem, not exactly where it is or how to fix it.

The full Legal AI Deployment Evaluation Kit contains 78 scored assessment items organized into five modules, each with detailed scoring rubrics, red flag indicators, evidence checklists, and write-in sections for documenting your findings:

Module 1: Data governance and client confidentiality (18 items)
Data flow mapping, privilege and confidentiality controls, data retention and deletion, including the model inversion and data extraction tests that most vendors have never been asked about.

Module 2: AI-specific security risks (15 items)
Prompt injection, sensitive information disclosure, supply chain and model integrity, mapped directly to the OWASP Top 10 for LLM Applications (2025), filtered for the risks most relevant to legal practice.

Module 3: Governance and accountability (20 items)
AI governance structure, human oversight and professional responsibility, vendor accountability, and a full shadow AI audit covering technical controls, sanctioned tool lists, and no-fault disclosure mechanisms.

Module 4: Output reliability and professional competence (10 items)
Hallucination management, accuracy benchmarks, bias testing to protect you from the malpractice scenario where an AI tool fabricates case law and nobody catches it before filing.

Module 5: Incident response and business continuity (15 items)
AI-specific incident response plans, business continuity and vendor dependence, and a full insurance alignment assessment covering malpractice, cyber coverage, and financial exposure analysis.

The full kit also includes a 390-point scoring system with four readiness levels, a red flag tracker with owner assignments and target dates, and a 90-day action plan template that breaks remediation into three phases: Stabilize, Build, and Operationalize.

Disclaimer

This diagnostic is an educational self-assessment tool. It does not constitute legal advice, compliance certification, or a guarantee of security. Organizations should consult qualified legal and cybersecurity professionals for guidance specific to their circumstances. Framework references (NIST AI RMF, ISO/IEC 42001, OWASP Top 10 for LLM Applications) are used for educational purposes; this tool is not endorsed by or affiliated with NIST, ISO, or OWASP.