



Why Legal Should Be at the Center of Your Organization's AI Strategy

Technology, Privacy, and eCommerce



Cheat Sheet:

- **AI is an organizational shift, not a tool rollout.** Successful AI adoption requires rethinking workflows, decision rights, governance, and risk — areas where Legal already operates by design.
- **Legal work is structured and repeatable.** From contract review to regulatory analysis, Legal follows consistent, pattern-based workflows that provide ideal training ground for reliable, scalable AI systems.
- **Legal decisions are verifiable and governed.** Unlike many business functions, legal outputs can be tested against statutes, contracts, and policies, making AI results auditable, explainable, and defensible.
- **Legal delivers outsized AI ROI.** Because legal work is both high-cost and highly structured, even modest AI automation can save thousands of hours, reduce outside counsel spend, and accelerate cycle times.

Every department feels mounting pressure to adopt AI and Legal is no exception. Pilots launch quickly, early demos look promising, and teams often get swept up in the appeal of the shiny new

tool. But the results rarely match the hype. A 2025 [Boston Consulting Group study](#) of more than 280 finance executives found the median ROI from AI initiatives was just 10 percent, far short of the 20 percent or more many companies were aiming for.

Often, the challenge isn't the technology itself. It's the assumption that AI is a simple "tool rollout" rather than a deep organizational rewiring. Implementing AI changes far more than software — it impacts workflows, accountability, incentives, data flows, decision rights, and risk exposure. It also requires legal professionals to adopt a new way of working, which naturally creates hesitation. In many ways, this mirrors the early days of search engines in the early 2000s: it took time and repetition for users to become comfortable searching, validating, and trusting the vast array of articles, opinions, and outright fiction that the search engines returned. But as search technology improved, and the quality and relevance of results increased, users grew more comfortable with this new way of searching and saw meaningful productivity gains as information became more accessible and reliable.

The misunderstanding that AI can be treated as just another tool rollout rather than a shift in how work gets done shows up most clearly in how companies involve Legal. In many organizations, Legal is brought into discussions late in the work cycle: to review and redline a contract, tighten a clause, approve a policy, or flag a risk. But that sequencing is backward and often misses the full potential of AI use for Legal.

AI success depends on building systems that are explainable, transparent, governed, and compliant with the full stack of constraints an organization operates under, including data-privacy regulations, industry standards, contractual commitments, and internal governance policies. These are all domains Legal already lives and breathes.

The past year marked a turning point, as in-house legal departments moved from planning to real-world implementation. In the US, [GenAI adoption more than doubled](#), rising from 23 percent of professionals in 2024 to 52 percent in 2025. This surge was driven largely by organizations removing restrictions on AI, as the share of professionals pointing to company policies banning its use fell sharply from 29 percent to 9 percent. Despite this progress, the most significant barrier to broader adoption continues to be concerns about the reliability and trustworthiness of GenAI.

That's no coincidence. The very qualities that make AI (generally) reliable and scalable, including clear reasoning, verifiable standards, documented decision trails, cross-functional collaboration, disciplined risk management, all are Legal's native language. Of course, any professional advice still requires a human in the loop and citations should be verified to catch potential [hallucinations](#) in legal documents.

The misunderstanding that AI can be treated as just another tool rollout rather than a shift in how work gets done shows up most clearly in how companies involve Legal.

That's why Legal's role doesn't just expand in the AI era; it becomes foundational to AI's long-term success. The organizations that scale AI effectively don't treat Legal as a gatekeeper. They treat it as the function that helps define the decision architecture within which AI must operate. Organizations such as ACC provide essential legal [guidance](#) and [resources](#), and privacy-focused

organizations like [IAPP](#) offer [certifications](#) in AI Governance. These are helpful to legal, privacy, and compliance professionals as they build this decision architecture.

This article explains why Legal (often in partnership with its Privacy and Security counterparts) is uniquely positioned to lead an organization's AI strategy — and offers concrete playbooks for in-house counsel to operationalize that leadership.

The structure AI needs already exists in legal

Legal work follows patterns that are unusually consistent across matters, legal professionals, and industries. A contract review in healthcare will typically look different from one in SaaS, but the underlying sequence is similar:

- Understand the overall posture of the contract and the business needs it is meant to address.
- Locate the relevant clauses, identify any deviations from the playbook or standard terms, and evaluate each clause against the preferred, fallback, and final positions (sometimes referred to as “bid, fallback, and walk” positions).
- Determine whether each clause introduces risk, attempt to quantify the risk in business terms, and document the rationale for accepting or rejecting the language.

Going further, AI can be used to explain a contract in “business language” tailored to each stakeholder. For example, AI is particularly effective at tasks such as:

- “Explain this contract in sales terms a junior salesperson would understand.”
- “Explain this contract in business terms an experienced C-level executive would understand.”
- “Explain this contract using logical structure and language a software engineer would understand.”

This same patterned logic often appears in incident response, employment matters, procurement, discovery, regulatory reporting, investigations, and policy interpretation. In fact, whether they realize it or not, lawyers have long been trained to generate useful prompts through the IRAC formula. Translated into AI prompting:

- **Issue:** the question the prompt is trying to solve;

-
- **Rule:** the constraints, standards, or parameters the AI should apply;
 - **Analysis:** the relevant facts and context needed to produce a useful, properly formatted response; and
 - **Conclusion:** the recommended action in legal terms, plus a plain-English explanation for business stakeholders and clients.

These workflows are designed to be stable and repeatable because inconsistency creates exposure, undermines credibility, and invites challenge.

That repeatability makes legal work especially conducive to AI. When AI observes hundreds of similar decisions handled in consistent ways, such as patterns of issue-spotting, interpreting, escalating, and documenting, it learns what "good" looks like. If Legal reliably flags uncapped indemnities, identifies missing DPAs, or treats certain security addenda terms as non-negotiable, those behaviors form a blueprint AI can follow. AI can also be used to transform your existing contracts into a structured playbook or preferred set of terms, and benchmark your contract templates against industry standards, competitor norms, or internal preferred criteria.

[Visit the ACC AI Center of Excellence for In-house Counsel today!](#)

Consider a legal team that has negotiated 300 sales contracts over three years. Each negotiation follows the same general framework: liability caps typically should not exceed contract value, data processing addenda must include specific provisions such as standard contractual clauses, and termination and refund clauses should be fair and unambiguous. When AI reviews the 301st contract, it doesn't need to guess what matters. The precedent is clear and documented.

Other departments rarely offer this level of structure. Marketing campaigns hinge on creative judgment. Product planning sessions depend on shifting priorities. Design critiques rely on aesthetic preference. While other functions follow patterns, those patterns shift far more frequently with context and changing objectives. Legal's predictability gives AI a more stable, disciplined scaffolding than perhaps any other function. This is particularly true in certain "structured document" use cases:

- Contracts, such as NDAs, DPAs, Terms of Service, and Purchasing Agreements;
- Questionnaires, such as privacy, security, compliance, and vendor onboarding forms;
- Requests for Information and Requests for Proposals; and
- Tracking legislative and administrative rules, and any changes to them.

Legal decisions are verifiable

Most corporate decisions can't be evaluated against a clear standard. There's no statute that determines whether a product roadmap was "good" or whether a design review was "effective."

Legal decisions are different. They trace back to source authority: contract language, statutes, regulatory rules, case law, or internal policy. An attorney's decision to classify data as personal, mark a clause as non-negotiable, or escalate a potential data incident or breach can often be evaluated against specific, objective criteria.

That verifiability makes Legal an ideal starting point for AI adoption. When an AI system proposes a redline, a reviewer can check it against fallback positions. When it flags an incident, the reviewer can compare it to contractual thresholds or reporting timelines. Legal often provides a rare "right or wrong" framework in an enterprise environment where most decisions are subjective or preference-driven.

And because legal outcomes often face external validation (Would this clause survive litigation? Does this language meet regulatory expectations?), AI outputs can be assessed not only internally but against real-world constraints. Few other functions offer that level of clarity.

Of course, a human should always review such AI-generated responses.

Legal is a data-rich environment for AI

Beyond contract review and traditional legal use cases, AI can already analyze large volumes of inputs (e.g., emails, chat messages, contract database fields, and related metadata) to produce quantifiable metrics for legal operations and spend reporting. For transaction and contracting functions, providing reliable data on the following metrics demonstrates operational rigor:

Executive view (monthly):

- SLA attainment by contract type and paper;
- Deal cycle time trends (sell-side) and procurement cycle time trends (buy-side);
- No touch rate (volume/value) and nonstandard rate (complexity/value); and
- Spend to budget reporting and outside counsel billing guidelines compliance.

Operations view (weekly):

- Task aging by stage and owner; top clause deviation hotspots;
- Workload and capacity by attorney and by issue type; and
- SLA misses by segment and their root cause.



The result is a domain where performance can be measured, audited, corrected, and improved with confidence, using AI both to perform core tasks, and to run the data analytics. Legal becomes a domain where organizations can clearly assess whether AI is performing as intended and build the systems needed to scale it across the enterprise.

The economic case: Legal Delivers Some of the Highest ROI

Legal work sits at a rare intersection in the enterprise: it is highly structured, highly repeatable, and highly expensive. Even routine matters often demand specialized judgment, detailed review, and rigorous documentation. A single commercial contract can take hours if reviewed “the old way” (i.e., manually). A regulatory inquiry can absorb weeks or even longer. Discovery can run into millions. Even a sudden spike in NDAs and DPAs at quarter end can drain capacity and stretch contracting teams. None of this is inefficiency. Rather, it’s the rigor the work requires.

That combination of repeatable workflows with high cost and high stakes makes Legal one of the most effective leverage points for AI.

AI excels where processes follow recognizable patterns. Legal generates those patterns every day. As a result, even modest automation unlocks meaningful returns. Automating portions of routine contract review frees attorney and paralegal hours immediately. Any questionnaire-type task is ripe for AI automation: privacy and security questionnaires, vendor onboarding checklists, requests for information (RFIs) and requests for proposals (RFPs), compliance and licensing responses, and many others. Long questionnaires can be completed in minutes using AI, with clear citations to source documents and supporting text, which greatly speeds up the response process. Of course, a human should always review such AI-generated responses, but using AI to do the heavy lifting can greatly improve the efficiency of the human reviewer, and help reduce the cognitive load and fatigue involved in filling out the responses.

AI-assisted triage reduces the volume sent to outside counsel and can often classify contracts by complexity and business risk. First-pass contract analysis shortens turnaround times without sacrificing accuracy. The net effect: legal professionals spend more time on the judgment-heavy work of negotiation, strategy, and interpretation and less time on repetitive manual tasks.

The impact on outside counsel spend is even more pronounced. When firms bill for every redline, memo, and review cycle, AI-generated first drafts and structured clause analysis can materially reduce costs, often offsetting the cost of the AI tools themselves. Discovery datasets that once required hundreds or even thousands of attorney hours can now be triaged in a fraction of that time, both reducing costs and accelerating turnaround. Regulatory deadlines that previously triggered emergency outside-counsel hours can be met with automated analysis and, in many cases, draft response generation.

Legal becomes a domain where organizations can clearly assess whether AI is performing as intended and build the systems needed to scale it across the enterprise.

AI also excels at tracking changes to online terms (from a webpage) and legislative drafts or enacted legislation. AI agents can be configured to repeatedly check any connected or public data source, flag changes for human review, and quantify the impact of those updates. Traditionally, this work required manual comparison on a recurring schedule — a task that is both repetitive and taxing. The math speaks for itself: reducing a routine contract review from 90 minutes to 20 minutes saves more than 1,000 hours a year for a team handling 100 contracts per month. At a blended BigLaw outside-counsel rate of often \$1,000–\$1,500 per hour, that single use case delivers \$1.5–\$2 million in savings.

These aren't hypothetical outcomes. Recent [real-world deployments](#) show the scale of results legal teams can achieve:

- A leading investment management company used AI-powered search and a Deal Status Agent to cut legal search time by more than 40 hours per month and reduce prep meetings by 33 percent.
- A publicly traded oil and gas company used AI-assisted policy tracking to save six hours per lawyer each week and reduce outside counsel costs by \$84,000 annually. By searching across policy documents, regulatory updates, and supporting evidence, AI can help lawyers

stay current with evolving requirements while dramatically reducing manual research time.

All of these gains, however, depend on the quality of the information the AI is retrieving. Legal workflows require the system to surface the correct clauses, policies, negotiation history, and governing documents with precision. If the retrieval layer pulls incomplete, outdated, or irrelevant materials, the downstream analysis will reflect those gaps — the familiar “garbage in, garbage out” problem. High-quality, context-relevant search ensures the AI operates on the same authoritative sources attorneys rely on, which is essential for producing outputs that are accurate, defensible, and aligned with established legal standards.

Conclusion: Why legal is poised to lead enterprise AI

AI is reshaping how work gets done, but its impact depends entirely on how well organizations design the systems around it. Legal is uniquely positioned to lead that effort. Its work is structured, repeatable, auditable, and grounded in standards — exactly the environment where AI performs reliably and where its outputs can be measured and improved with confidence. Legal already operates at the intersection of risk, governance, precedent, and decision architecture, making it an ideal starting point for building AI systems that are transparent, explainable, and compliant.

When Legal moves from being a late-stage reviewer to a strategic partner in AI adoption, organizations benefit twice: Legal teams gain powerful tools that reduce low-value work and improve cycle times, and the broader enterprise gains a disciplined framework for deploying AI safely and at scale, incorporating privacy and governance by design. AI doesn't supplant the role of Legal. It elevates it — reducing low-value work, improving cycle times, and enabling attorneys to focus on the decisions that drive key business outcomes. And because Legal's work is traceable, auditable, and grounded in standards, it offers one of the safest, clearest pathways for deploying AI at scale.

If enterprises want AI that is reliable, governed, and aligned to real-world constraints, Legal isn't just a stakeholder. It's the function that should lead the way.

[Join ACC for more AI insights!](#)

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

[Michael T. Moore](#)



Vice President, Head of Legal, and Company Secretary

Glean Technologies, Inc.

Michael T. Moore is Vice President, Head of Legal, and Company Secretary at Glean Technologies Inc., where he leads global legal, privacy, intellectual property, and AI governance. He has previously held senior in-house leadership roles at a cybersecurity startup, IQVIA, Pure Storage, and Rambus, spanning product counsel, patent strategy, privacy, cybersecurity, and regulatory compliance. Moore is an IAM-300 recognized patent strategist and holds multiple IAPP certifications, including CIPP/US, CIPP/E, CIPP/C, CIPM, CIPT, and AIGP. He is a frequent author and speaker for the Association of

[Rebecca Hinds](#)



Head of the Work AI Institute

Glean Technologies, Inc.

Rebecca Hinds is the Head of the Work AI Institute at Glean, where she leads research on how AI is reshaping the way people work. She earned her B.S., M.S., and Ph.D. from Stanford University. Her research and work focus on how emerging technologies — especially AI — change organizational behavior, collaboration, and performance. Her research and insights have appeared in outlets such as *Harvard Business Review*, *The New York Times*, *The Wall Street Journal*, *Forbes, Inc.*, and *Time*, as well as top academic outlets like *Organization Science* and *CSCW*. Hinds is a co-instructor for the CNBC *Make It* course *How to Use AI to Be More Successful at Work*, and is the author of *Your Best Meeting Ever*, named a Next Big Idea Club “must-read.”

