


AI, REGULATION, AND LEGAL ADVICE:

# Managing Risk and Governance After OpenAI's Policy Shift



# Executive Summary

On October 29, 2025, OpenAI updated its usage policy to prohibit “tailored advice that requires a licence, such as legal or medical advice, without appropriate involvement by a licensed professional.”

This is the first time a major AI vendor has formally set boundaries around regulated advice. For legal and compliance leaders, it marks a shift in how AI’s role in professional services is defined. Yet, ChatGPT still drafts contracts, analyses legal frameworks, and offers guidance that closely resembles legal work.

This gap between policy and practice creates both risk and opportunity. Law firms that implement structured AI governance will be best placed to use these tools responsibly while demonstrating the oversight regulators and clients expect.

This whitepaper examines the implications of OpenAI’s policy change and outlines how law firms can build the governance infrastructure needed to deploy AI confidently in regulated environments.



# Context: The Policy Shift and Its Significance

Since ChatGPT's launch in late 2022, both legal professionals and the general public rapidly adopted generative AI – often without clear oversight, as the tools freely offered detailed guidance accompanied only by brief disclaimers to “consult a qualified attorney.”

OpenAI's October 2025 policy update attempts to draw a line where previously there was none. Whether that line proves meaningful in practice remains an open question.

## What's Changed?

OpenAI's updated [usage policies](#), published on October 29, 2025, now explicitly prohibit:

*“...provision of tailored advice that requires a license, such as legal or medical advice, without appropriate involvement by a licensed professional.”*

The policy also states that OpenAI's services cannot be used for “automation of high-stakes decisions in sensitive areas without human review” including legal, medical, and essential government services.

This represents the first time OpenAI has formally restricted the use of its models for regulated professional advice. Previously, ChatGPT would offer disclaimers after providing extensive legal guidance. Now, the restriction appears in the terms themselves.

## Why This Matters Now

The timing of OpenAI's policy shift reflects a turning point in how legal work and professional accountability intersect with AI. Over the past two years, AI tools have slowly moved from experimental use to reliance in legal workflows. What was once a technical innovation has become a governance

concern, forcing law firms to consider not just what AI can do, but how it should be managed.

Clients, insurers, and regulators are now converging on the same expectation: **AI in legal practice must operate under the same standards of diligence and accountability as human work.** Corporate clients increasingly ask firms to disclose how AI is used in their processes, while insurers are beginning to assess exposure tied to AI-assisted outputs. These external pressures are reshaping how firms structure internal approval, risk, and review mechanisms.

OpenAI's decision to limit the use of its tools for “tailored advice” formalizes what many in the profession already recognized: the responsibility for compliance sits squarely with the user. For law firms, this shift means AI adoption can no longer rely on vendor assurances; it must be supported by firm-level governance frameworks, policies, and documentation that demonstrate professional control.

In short, the announcement arrives at a moment when AI's role in legal practice is being redefined – not by technology itself, but by the expectations of clients, regulators, and the profession.





### **The Reality: Policy Versus Practice**

However, the policy change may be more protective than preventative. When [Artificial Lawyer tested ChatGPT](#) immediately after the update, the model happily drafted an employment contract, explained relevant English law provisions, and created custom clauses for specific scenarios. Only when explicitly thanked for "legal advice" did ChatGPT clarify that it was providing "general legal information and example drafting guidance" rather than formal advice.

The distinction is telling. ChatGPT now frames its outputs as educational rather than advisory, but the substance remains functionally identical. [Testing by Legal IT Insider](#) confirmed similar results: the model still provides contractual templates, explains negligence claims with supporting case law, and walks users through legal frameworks step by step.

For law firms, this creates a challenging dynamic: the policy establishes a boundary, but enforcement relies entirely on user interpretation and internal governance.

# Implications for Enterprise Legal and Compliance Teams

## **The Blurred Boundary Between Legal Advice and Legal Information**

In common law jurisdictions, the distinction between legal advice and legal information turns on context and specificity. Legal information provides general explanations of law and process. Legal advice applies legal principles to specific facts to guide a particular course of action.

Generative AI operates precisely in this grey zone. When ChatGPT drafts a termination clause tailored to "an employee about to retire," is that information or advice? When it analyses whether workplace conduct meets the threshold for negligence, citing relevant precedent, where does explanation end and counsel begin?

The answer matters because unauthorized practice of law carries regulatory, civil, and reputational consequences. If internal teams rely on AI-generated text without appropriate legal review, they may inadvertently facilitate unlicensed practice. The fact that OpenAI now disclaims responsibility only heightens the importance of internal controls.





## Risks for Legal Teams

### Regulatory Risk

Misuse of AI-generated legal content could constitute unauthorized legal practice or violate consumer protection frameworks. As regulators clarify AI's role in professional services, expectations around human oversight will only increase.

### Operational Risk

AI outputs entering critical workflows without validation create exposure. Contract errors, compliance gaps, or misapplied legal standards can have material consequences that standard disclaimers will not mitigate.

### Reputational Risk

The perception that a law firm allows AI to provide legal advice without licensed oversight undermines trust with clients, investors, and regulators. In sectors where professional standards define competitive advantage, this matters.


### Vendor Risk

OpenAI's policy demonstrates that AI providers can change terms, capabilities, and access parameters with limited notice. Dependency on a single vendor whose policies are in flux creates strategic vulnerability.

## The Compliance Paradox

Law firms adopted AI to spend less time on routine tasks, and more time on strategy, client relationships, and meaningful legal work. OpenAI's policy reinforces that realizing these benefits depends on maintaining demonstrable human oversight and auditability.

The message is clear: AI should function as a tool supervised by licensed professionals, not as a replacement for professional judgement. The challenge for legal leaders is to operationalize that principle across dispersed teams, multiple use cases, and evolving technology.



# Building an AI Governance Framework for Legal Functions

## Governance Principles

*Effective AI governance in legal contexts rests on four foundations:*

### Transparency

Know which tools and models teams use, what data they access, and how outputs flow into decision-making processes.

### Accountability

Maintain clear records of who reviewed AI-generated content, when, and what changes were made before relying on it.

### Traceability

Document when and how AI outputs influenced legal or compliance decisions, creating an audit trail that satisfies internal and external scrutiny.

### Explainability

Require clarity on model limitations, training data sources, and known failure modes. Teams should understand what AI can and cannot reliably do.



## The Human-in-the-Loop Model

*Appropriate oversight varies by task and risk. For legal functions, a practical framework might look like:*

### Research and Analysis

AI generates preliminary summaries of case law or statutory provisions. A qualified lawyer reviews for accuracy and relevance before incorporating insights into advice or strategy.

### Drafting and Review

AI produces initial contract language or policy text. Licensed counsel reviews substantive provisions, confirms compliance with applicable law, and approves outputs before use.

### Process and Workflow

AI assists with document management, deadline tracking, or matter prioritization. Human users validate critical decisions and maintain supervisory control.



## Integration with Enterprise AI Policy

*Legal AI governance should align with broader responsible AI frameworks across the law firm.*

*This means:*

- Establishing tiered approval levels for different use cases (research, advice, execution)
- Requiring vendor proof of data security, model provenance, and policy compliance
- Defining roles and responsibilities for AI oversight within legal, compliance, and IT functions
- Creating escalation pathways for questions about appropriate AI use

## Metrics for Governance Maturity

*Measuring progress helps ensure governance operates in practice, not just policy:*

- Percentage of AI-assisted legal tasks with documented human validation
- Existence of an AI risk register updated periodically with new use cases and controls
- Inclusion of AI oversight requirements in external counsel guidelines and vendor agreements
- Regular training completion rates for teams using AI in legal or compliance work

# Strategic Response: What Legal Leaders Should Do Now

## Immediate Actions

### Audit Current AI Use

Map all AI tools in use across teams. Identify whether they access proprietary data or influence high-stakes decisions.

### Define Use Case Boundaries

For each tool and application, determine where it falls on the advice-versus-information spectrum. Flag areas where current use may cross into regulated territory without adequate oversight.

### Update Internal Policies

Incorporate OpenAI's restrictions and similar vendor policies into internal AI usage guidelines. Make clear that AI-generated legal content requires professional review before use.

### Educate Lawyers and Staff

Provide training on the distinction between permissible AI assistance and unauthorized legal practice. Create clear escalation procedures for ambiguous situations.





## Medium-Term Actions

### Formalize Oversight Structures

Implement an AI Governance Committee or expand existing legal risk committees to include AI oversight. Ensure representation from legal, compliance, IT, and business units.

### Develop Documentation Standards

Create template disclaimers, process notes, and approval workflows for AI-assisted outputs. Require that teams document AI's role in any work product that influences decisions.

### Partner with Compliance-Ready Providers

Prioritize AI platforms built for regulated contexts, offering audit trails, user-level controls, and transparent model behaviour. Platforms like [Alexi](#) that embed human oversight by design reduce governance burden.

### Prepare for Regulatory Evolution

Monitor emerging guidance from U.S. and Canadian regulators, including the White House AI Executive Order, the U.S. Department of Justice and Federal Trade Commission's AI enforcement priorities, and updates from the Canadian Department of Justice and provincial law societies. Expect professional regulators such as the ABA and Federation of Law Societies of Canada to issue clearer standards on the responsible use of AI in legal services.

## Strategic Positioning

Forward-looking law firms will treat AI governance as a differentiator rather than a burden. Clients and regulators increasingly expect demonstrable AI oversight. Enterprises that can articulate clear governance frameworks, supported by documentation and metrics, will be better positioned to:

- Win work from clients with AI risk concerns
- Navigate regulatory scrutiny with confidence
- Adopt new AI capabilities faster because governance infrastructure is already in place
- Attract and retain talent seeking responsible AI environments

The alternative is reactive compliance, where law firms scramble to retrofit governance after incidents or regulatory action.



# Regulation, Trust, and the Next Phase of Legal AI

## Anticipating Regulatory Developments

Regulators across North America are increasingly unified in their view that AI deployed in legal services must remain subject to human oversight, supported by transparent processes and accountable governance.

In the United States, the [White House Executive Order on Safe, Secure, and Trustworthy AI \(Executive Order 14110\)](#) establishes the federal framework for AI governance, emphasizing transparency, safety, and responsible deployment. Complementing this, the [NIST AI Risk Management Framework \(AI RMF\)](#) and its [Generative AI profile](#) outline standards for documentation, auditability, and continuous oversight. Agencies such as the Federal Trade Commission (FTC) and Department of Justice (DOJ) have also reinforced that existing laws on consumer protection, antitrust, and civil rights apply to AI tools. Several state bars (including California and Florida) are now evaluating the ethical implications of AI-assisted legal work, focusing on confidentiality, competence, and client communication.

In Canada, the forthcoming [Artificial Intelligence and Data Act \(AIDA\)](#) will introduce national standards for responsible AI use, supported by evolving guidance from the Federation of Law Societies of Canada and various provincial law societies.

As these frameworks mature through 2026, law firms should expect clearer professional obligations, potentially including certification or documented human review for AI systems used in client-facing legal tasks.

## The Shift from Experimentation to Compliance

*The arc of enterprise AI adoption is clear.*

**2023-2024:** Experimentation and pilots. Legal teams tested AI for research, drafting, and analysis, with few formal controls.

**2025:** Governance, audits, and accountability. Regulators, clients, and leadership now expect legal teams to demonstrate responsible use through policies, oversight, and documentation. Enterprises must now treat AI governance as an operational requirement, not a future consideration.

**2026 and beyond:** Integration and enforcement. AI governance will move from policy statements to measurable compliance. Law firms and corporate legal departments will be expected to maintain auditable AI inventories, certify model provenance, and demonstrate human review for regulated outputs. External audits and client due diligence will become standard practice.



# Why Legal Teams Choose Alexi Over Generic AI Tools

After OpenAI's policy shift, law firms need tools that support compliance, accountability, and control. Alexi provides a governed, secure AI environment built for legal work, ensuring firms can adopt AI responsibly while maintaining professional standards.

## Governance by Design

Generic AI tools leave firms exposed to shifting policies and unclear accountability. Alexi embeds governance into every layer – from data handling to model behavior – so firms can define how AI is used, reviewed, and approved within professional and regulatory boundaries.

## Secure and Auditable Architecture

Alexi's SOC 2–certified, AES-256–encrypted infrastructure ensures that every action is logged and every dataset is protected. Client information never leaves the firm's private environment, providing verifiable compliance and full traceability when audits or client reviews arise.

## Transparent, Controlled Intelligence

Firms need to know how AI reaches its conclusions. Alexi shows how outputs are generated, what sources are used, and where uncertainty exists. Firms can define citation rules, drafting standards, and permissions to ensure accuracy and accountability.

## Human Oversight

Alexi reinforces professional judgment by keeping lawyers in control. Review steps and approval records make oversight visible and provable, key to demonstrating responsible AI use to clients, regulators, and insurers.

## Future-Ready

Public AI tools can change terms overnight. Alexi evolves with the legal industry, aligning its roadmap with emerging regulations and professional guidance. Your AI governance remains stable, auditable, and adaptable to whatever comes next.

## The Foundation for Firm-Wide AI Governance

Alexi transforms AI governance into a strategic asset. By uniting security, oversight, and continuous learning, your firm builds private intelligence that compounds in value and trust over time.

### Ready to strengthen your legal AI governance?

Speak to us today to see how Alexi can support your firm's AI adoption with confidence.

*This resource reflects developments as of November 2025. AI regulation continues to evolve rapidly, and organizations should monitor updates from relevant regulators and professional bodies.*

