

AI Governance and Compliance: A Practical Guide for General Counsel

A FDCC Continuing Legal Education White Paper

Authors: Hunter McMahon, Kiera McGroarty, Brent Messel, James Biese

Executive Summary

Artificial intelligence is no longer an emerging technology; it is embedded in how companies make decisions, how they litigate, and how they operate day to day. That integration has outpaced the legal frameworks meant to govern it. In 2026, three developments demand immediate attention from general counsel: AI-generated content (prompts, outputs, and interaction logs) is now discoverable, electronically stored information in litigation. Trade secret protections face new erosion risks as employees feed proprietary data into AI platforms and as AI-enabled reverse engineering expands what competitors can reconstruct. And the attorney-client privilege, long a bedrock of legal strategy, is being tested by courts grappling with whether communications with AI tools deserve confidentiality protections at all.

This white paper, developed as a continuing legal education resource for the Federation of Defense and Corporate Counsel, provides general counsel with:

- A practical framework for understanding when and how AI interactions become discoverable in litigation
- Analysis of the evolving risks AI poses to trade secret protection and attorney work product
- Guidance on the governance implications of custom-built versus generic AI tools
- A survey of major state legislation enacted in 2026
- Scalable governance program models tailored to legal departments of all sizes

The regulatory landscape remains fragmented; in 2025 alone, 1,208 AI-related bills were introduced across state legislatures. But the litigation landscape is moving faster than the statutes. Courts are already deciding whether AI prompts are discoverable, whether uploading trade secrets to AI platforms waives protection, and whether AI-assisted legal analysis qualifies for work product immunity. General counsel who wait for regulatory clarity before acting will find that the courtroom has already answered the questions their governance programs never thought to ask.

Commented [HM1]: AI DETECT [Structural — Strong]:
Em dash pair used as parenthetical where commas or parentheses would be standard. Em dashes are overrepresented in AI-generated text. Also flagged as a style violation (user preference: never use em dashes).

Key Takeaway: AI governance is no longer a compliance exercise in search of a regulation. The risks are here: in discovery motions, trade secret disputes, and privilege challenges that are being decided now. This paper provides the framework for understanding those risks and building governance programs that address them before they become someone else's exhibit.

Table of Contents

1. Introduction
2. The Current AI Regulatory Landscape
3. AI Prompts as Discoverable ESI
4. Trade Secrets in the Age of AI
5. Attorney Work Product and AI: An Emerging Framework
6. Custom AI vs. Generic: The Governance Implications
7. Major AI Governance Risks
8. Survey of Major 2026 State AI Legislation
9. Building an AI Governance Program
10. Conclusion and Best Practices

1. Introduction

The Awareness Gap

Artificial intelligence is embedded throughout modern business operations: in HR tools, customer service platforms, marketing analytics, financial systems, and product development. According to a 2024 survey by Littler, 82% of Chief Human Resources Officers reported that their companies use AI in human resources, yet only 48% of Chief Legal Officers and General Counsels at those same companies were aware of this AI use.¹

That gap between operational reality and legal awareness is the central challenge this paper addresses. Most organizations became AI-integrated companies before their legal departments realized it. The question facing general counsel is not whether to adopt AI governance; it is whether the governance program will be in place before the first discovery request arrives.

Why 2026 Is Different

- **AI-generated content is discoverable ESI:** Courts have ruled that prompts, outputs, and interaction logs are electronically stored information subject to preservation and production, with no special AI privilege shielding them from discovery
- **Trade secret protections are being tested:** Uploading proprietary information to public AI platforms may constitute a failure to maintain "reasonable measures" to protect secrecy, and AI-enabled reverse engineering is expanding what competitors can reconstruct from public data
- **Privilege doctrine is unsettled:** Federal courts have reached divergent conclusions on whether AI-generated legal analysis qualifies for attorney-client privilege or work product protection, with outcomes turning on whether counsel directed the AI use and whether the platform offered confidentiality
- **The regulatory patchwork continues to expand:** Multiple states enacted AI laws in 2026 with private rights of action and statutory damages, while federal enforcement agencies continue to apply existing authority to AI systems

The Role of General Counsel

In-house lawyers were once peripheral to AI decisions. Now they are expected to lead their organizations through them. A 2024 ACC and Everlaw survey found that fewer than half of in-house lawyers feel prepared for the impact of generative AI on their careers. That preparation deficit becomes a liability the moment opposing counsel requests production

¹ <https://www.littler.com/publication-press/publication/littlers-2024-ai-c-suite-survey-report>

of AI interaction logs, a departing employee's AI chat history reveals trade secret exposure, or a court must decide whether work product created with AI assistance retains its privileged character.

The sections that follow address each of these challenges with the specificity that CLE credit demands, grounded in cases being decided right now. The goal is not abstraction; it is practical guidance that scales from a three-attorney department to an enterprise legal team.

2. The Current AI Regulatory Landscape

Federal Framework: Executive Orders and Policy Shifts

Biden Administration (2023–2025)

Executive Order 14110 (October 30, 2023): “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence”² This order established extensive requirements for AI developers, including:

- Safety testing and reporting requirements
- Standards for privacy-preserving techniques
- Guidance on AI use in critical infrastructure
- Equity and civil rights protections

Trump Administration (2025–Present)

Executive Order 14179 (January 23, 2025): “Removing Barriers to American Leadership in Artificial Intelligence”³ This order marked a significant policy shift by:

- Revoking Executive Order 14110
- Directing development of an AI Action Plan focused on promoting “human flourishing, economic competitiveness, and national security”

² <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence>

³ <https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence>

- Emphasizing AI systems “free from ideological bias or engineered social agendas”

Executive Order (December 11, 2025): “Ensuring a National Policy Framework for Artificial Intelligence”⁴ This subsequent order established a framework for coordinated AI policy development across federal agencies.

Practical Impact for General Counsel: The federal landscape remains fluid. While the Trump administration has rolled back certain requirements, federal agencies (SEC, FTC, EEOC, CFPB) continue to enforce existing laws as they apply to AI systems.

NIST AI Risk Management Framework

The NIST AI Risk Management Framework (AI RMF 1.0),⁵ released in January 2023, remains the leading voluntary framework for organizations designing, developing, deploying, or using AI systems.

Core Functions:

1. GOVERN – Organizational structures and policies
2. MAP – Identify and document AI risks
3. MEASURE – Assess and track risks
4. MANAGE – Implement plans to mitigate risks

Risk Categories:

- Accuracy and reliability
- Fairness and bias
- Interpretability and transparency
- Accountability
- Privacy and security
- IP and confidentiality

Commented [HM2]: AI DETECT [Vocabulary — Moderate]: 'Foundational' and 'framework' in the same phrase. 'Framework' appears 11+ times in this document; 'foundational' is a formal inflation word. The clustering of these elevated-register terms is consistent with LLM output.

⁴ <https://www.federalregister.gov/documents/2025/12/16/2025-23092/ensuring-a-national-policy-framework-for-artificial-intelligence>

⁵ <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>

- Workforce and environmental impacts

Recent Updates:

- NIST AI 600-1 (July 2024): Generative AI Profile
- NIST Critical Infrastructure Profile (April 2026): Trustworthy AI in Critical Infrastructure

3. AI Prompts as Discoverable ESI

Every interaction with a generative AI tool creates electronically stored information: prompts, uploaded documents, outputs, and metadata about when and by whom. All of it is subject to the same preservation, collection, and production obligations that govern email and other digital records, and courts have made clear that no AI exemption is coming under the Federal Rules of Civil Procedure.

The Basic Framework: Familiar Rules, Unfamiliar Data

AI interactions fit comfortably within the existing ESI framework, which is precisely why the discoverability question is less novel than it appears. A prompt is text; an output is text, image, code, or another digital artifact; an activity log is metadata about who used a system, when, and how. Once those items are relevant to a claim or defense, within a party's possession or control, and not shielded by privilege, they are ordinary discoverable ESI. Courts in 2026 have been applying familiar discovery principles to AI data rather than inventing new ones.

The practical implication is that companies must begin treating AI interaction logs with the same seriousness they apply to email retention. This means understanding where AI data is stored (locally, in cloud provider logs, in enterprise platforms), who has access, and what retention policies apply. Many organizations have no litigation hold procedures that contemplate AI-generated content, a gap that creates spoliation risk.

The Defining Early Cases

United States v. Heppner (S.D.N.Y., February 2026). In what Judge Jed S. Rakoff characterized as "a question of first impression nationwide," the Southern District of New York ruled that a criminal defendant's written exchanges with Anthropic's consumer Claude platform were not protected by either attorney-client privilege or the work product doctrine. The defendant, a CEO charged with securities fraud, had used Claude to analyze his legal exposure and generate defense strategy documents on his own initiative, not at

counsel's direction. The FBI seized these materials during the execution of a search warrant.

The court's reasoning rested on three independent grounds: Claude is not an attorney and cannot form an attorney-client relationship; the consumer platform's terms of service provided no reasonable expectation of confidentiality; and the defendant's purpose was not to obtain legal advice from Claude, which expressly disclaims that capacity. Judge Rakoff emphasized that all recognized privileges require "a trusting human relationship" and that "no such relationship exists, or could exist, between an AI user and a platform such as Claude."

Critically, the court left the door open for a different outcome where counsel directs the AI use: "had counsel directed Heppner to use Claude, Claude might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege." This dictum, referencing the Kovel doctrine, signals that the analysis may shift significantly based on who initiates and supervises the AI interaction.

Warner v. Gilbarco, Inc. (E.D. Mich., February 2026). Decided the same week as Heppner, the Eastern District of Michigan reached a markedly different conclusion regarding work-product protection. A pro se plaintiff in an employment discrimination action used ChatGPT to prepare litigation materials. The defendants sought to compel the production of all documents concerning her AI use, arguing that disclosure to an AI platform would waive any applicable protections.

The court denied the motion. Its reasoning drew a critical distinction: a work product waiver requires disclosure "to an adversary or in a way likely to get into an adversary's hands." The court held that "generative AI programs are tools, not persons, even if they may have administrators somewhere in the background." Entering litigation materials into an AI platform did not constitute disclosure to an adversary sufficient to waive work-product protection.

In re OpenAI (S.D.N.Y.). In the consolidated multidistrict copyright litigation against OpenAI, the court ordered production of 20 million anonymized chat logs, including user prompts and model responses. While the context was copyright rather than trade secret litigation, the ruling established that AI platform providers may be compelled to produce user interaction data in litigation, creating a secondary discovery pathway that companies may not have anticipated when their employees began using AI tools.

Tremblay v. OpenAI (N.D. Cal.). The Northern District of California provided a counterpoint to the Heppner analysis in the copyright context. The court held that prompts crafted by

counsel could reflect counsel's mental impressions and therefore qualify for strong work-product protection. Importantly, the court limited any waiver to material actually relied upon in the complaint, rejecting the argument that producing some AI interactions waived protection for all related materials.

Practical Implications for General Counsel

The emerging case law points to concrete steps that general counsel should take now. AI interaction data must be included in litigation hold notices. An employee who used ChatGPT to brainstorm a response to a customer complaint, or who uploaded a contract draft to Claude for analysis, created discoverable records; organizations that fail to preserve this data face the same spoliation consequences they would face for destroying relevant emails.

Companies should also audit how their AI providers handle data retention, because the answers vary more than most legal departments expect. Consumer-tier AI products may retain prompts and outputs for model training, meaning a reservoir of potentially discoverable information persists even after the user deletes their local chat history. Enterprise-tier products typically include contractual commitments not to train on customer data and may offer shorter retention windows, but the specific terms vary across providers and must be reviewed at the contract level rather than assumed.

The distinction between consumer and enterprise AI platforms is emerging as a factor in both discoverability and privilege analysis, and it may prove to be the most consequential variable in this space. A company whose employees use personal ChatGPT accounts for work-related tasks has less control over data retention, less visibility into what was shared, and a weaker argument for confidentiality protections than one that has deployed an enterprise platform with documented security controls and attorney-supervised usage protocols.

4. Trade Secrets in the Age of AI

Trade secret law rests on a deceptively simple premise: the information must not be generally known, and the owner must take reasonable measures to keep it secret. Generative AI is pressuring both elements simultaneously, and the pressure is coming from directions that most trade secret protection programs were not designed to anticipate. Employees and contractors are feeding proprietary information into AI platforms that may share, retain, or train on the data they receive; at the same time, AI-enabled reverse engineering is expanding what competitors can reconstruct from publicly available

sources, which raises the question of whether certain trade secrets remain protectable at all.

Disclosure Risk: When Using AI Erodes Trade Secret Protection

The threshold question is whether inputting trade secret information into a generative AI platform constitutes a disclosure that undermines the "reasonable measures" requirement. In *Trinidad v. OpenAI*, a federal court in the Southern District of New York addressed this directly. The plaintiff alleged that OpenAI misappropriated proprietary AI development frameworks she had entered into ChatGPT. The court's analysis focused on the foundational requirement that trade secret owners take reasonable measures to maintain secrecy, and found that uploading proprietary information to a public AI tool that may use customer data for model training failed to meet that standard.

The reasoning has broad implications. When employees input proprietary formulas, customer lists, pricing strategies, or product development plans into consumer-tier AI tools, they may be disclosing that information to a third party whose terms of service permit data retention, model training, and in some circumstances, disclosure to other users or governmental authorities. A court applying *Trinidad's* logic could find that a company that permits uncontrolled employee use of AI has failed to take reasonable measures to protect its trade secrets, regardless of whether the information was actually incorporated into the AI model's outputs.

The Reverse Engineering Problem

The Defend Trade Secrets Act and the Uniform Trade Secrets Act both exclude reverse engineering from the definition of "improper means." If a competitor can divine otherwise secret information by analyzing a publicly available product, that does not constitute misappropriation. AI is rapidly expanding the scope of what reverse engineering can accomplish.

Competitors can now feed patents, publications, product data, and even redacted court documents into AI tools and ask them to synthesize the gaps. Trade secret defendants are expected to use AI-enabled reverse engineering to argue that a claimed trade secret was "readily ascertainable" all along. While no reported decision has directly addressed whether trade secret information surfaced through generative AI queries is thereby "readily ascertainable," courts are likely to treat AI-based discovery of information the same way they treat traditional publicly available sources. If so, trade secret information that generative AI can identify or reconstruct from publicly available data could lose its protected status.

For companies whose competitive advantage depends on proprietary know-how, the reasonable measures analysis must now account for the AI environment. Companies that have not yet updated their trade secret protection programs to address AI should, at minimum, move employees to enterprise AI tiers with appropriate contractual protections, adopt clear policies governing what information may be entered into AI tools, and audit existing trade secret inventories against the capabilities of current AI models.

Employee Departure and AI-Enabled Knowledge Transfer

Employee mobility disputes have long been a primary vehicle for trade secret litigation, and AI is complicating them in ways that existing forensic tools are not equipped to detect. Enterprise AI platforms maintain significant context about an employee's work through memory and project features. Many platforms now offer memory export functions that allow users to transfer their conversation history and context to other accounts. A departing employee can achieve wholesale knowledge transfer, in summarized or even disguised form, with a few simple prompts.

Critically, this transfer can occur incrementally. An employee need not attempt to obtain a trade secret in a single session. They can query an enterprise AI about different components across multiple sessions (each query seemingly innocuous in isolation) and later collate the results into a functional reproduction of the trade secret. This pattern defeats the bulk-download alerts that conventional data-loss-prevention tools are designed to catch. The relevant forensic question is not whether a single large extraction occurred, but whether a pattern of isolated queries adds up to one.

Prompt Injection and System Prompts as Trade Secrets

A related question is whether AI system prompts (the hidden instructions that define how an AI application behaves) qualify as protectable trade secrets. In *OpenEvidence v. Pathway Medical, Inc.* (D. Mass., 2025), the plaintiff alleged that the defendant used prompt injection attacks to trick an AI medical platform into revealing its proprietary system prompts. The case asked the court to decide whether manipulating an AI's interface to extract hidden instructions constitutes misappropriation under the DTSA.

Courts have previously held that scraping data using bots, impersonating users, or injecting malicious prompts into AI systems can constitute "improper means" under trade secret law. But the line between proper reverse engineering and improper prompt injection remains blurred, and this area of law is likely to develop rapidly as more companies build competitive products on proprietary AI configurations.

Discovery Battles over AI Activity

AI use is also expected to trigger new discovery battles in trade secret litigation. Because AI use may reveal or compromise purported secrets, parties will seek extensive discovery on each other's AI activity (policies, tools, prompts, outputs, logs, and historical usage) including employees' and contractors' usage at home, on personal devices. Courts developed frameworks for traditional e-discovery over the past two decades, but disputes over AI-related data are introducing new layers of complexity. Discovery may spread to third parties, including the AI platform providers themselves, as demonstrated by the In re OpenAI chat log production order.

5. Attorney Work Product and AI: An Emerging Framework

The question of whether AI-assisted legal analysis retains work product protection is developing rapidly, and the law is still finding its footing. Two federal court decisions issued within the same week in February 2026 reached conclusions that look contradictory at first glance: Heppner stripped protection from AI-generated legal strategy documents while Warner preserved it for AI-assisted litigation materials. The cases are more consistent than they appear, but reconciling them requires attention to the specific facts that drove each outcome.

The Heppner-Warner Split

As discussed in Section 3, the Heppner court held that a criminal defendant's unsupervised use of a consumer AI tool produced materials that were neither privileged nor protected work product. The Warner court, addressing a civil plaintiff's use of ChatGPT in litigation preparation, held that work-product protection survived because AI platforms are "tools, not persons" and that disclosure to them does not constitute disclosure to an adversary.

Read together, the cases are less contradictory than the headlines made them sound. Heppner turned on the absence of attorney involvement: a defendant acting unilaterally, without counsel's direction, on a consumer platform with no confidentiality guarantees. Warner involved a pro se plaintiff who was effectively her own attorney, using AI as a drafting tool in the ordinary course of litigation; the court applied the standard work-product framework and concluded that the AI tool was no different from any other technology a litigant might use to develop case materials.

What matters, then, is not whether AI was involved but the circumstances of its use. Courts appear to be converging on a handful of related factors: whether counsel directed or supervised the AI interaction, whether the platform provided reasonable confidentiality

protections, whether the resulting materials reflect the attorney's mental processes, and whether the user intended from the outset to use the outputs for legal purposes.

The Consumer-Enterprise Distinction

The Heppner court emphasized Anthropic's consumer terms of service, which permitted data sharing with third parties. By contrast, enterprise AI platforms typically offer contractual commitments to data isolation, no-training clauses, and compliance certifications. While no court has yet drawn a bright line between consumer and enterprise AI for privilege purposes, the Heppner reasoning strongly implies that the platform's confidentiality architecture will be relevant to any privilege analysis. General counsel should evaluate their organizations' AI platforms with this distinction in mind and consider whether current deployments would withstand a privilege challenge.

What General Counsel Should Do Now

The law in this area will continue to develop rapidly, but several practical steps are defensible regardless of how courts ultimately resolve the open questions. If AI is part of the litigation or investigation workflow, counsel should direct its use and maintain clear documentation that materials were created in anticipation of litigation. Organizations should not assume that sharing AI outputs with counsel after the fact creates privilege; the Heppner court explicitly rejected that argument. And companies should avoid disclosing confidential client information to public AI platforms where they cannot control downstream use or exposure. The time to protect information is before it enters the AI platform, not after.

6. Custom AI vs. Generic: The Governance Implications

AI use is not monolithic, and governance programs that treat every interaction the same way will either lock down low-risk tasks unnecessarily or leave sensitive ones underprotected. The distinction that matters most is between generic AI use (employees querying consumer chatbots for general assistance) and custom deployments where organizations have invested in purpose-built tools, curated knowledge bases, and domain-specific configurations.

This distinction matters for governance because it maps directly to the risk factors courts are weighing: confidentiality protections, attorney direction, data isolation, and the sophistication of the AI interaction.

The Spectrum of AI Deployment

At one end of the spectrum, an associate using a free ChatGPT account to brainstorm deposition questions creates a high-risk, low-governance interaction. The prompts may contain privileged strategy, the platform's terms permit data retention and training, and the organization has no visibility into what was shared. At the other end, an enterprise deployment, in which a law firm or legal department has configured an AI platform with custom instructions, domain-specific knowledge, defined workflows, and data isolation, presents meaningfully different risk characteristics.

This is a legal distinction as much as a technical one, and it carries real weight in litigation. A custom AI deployment with documented access controls, no-training clauses, and attorney-supervised usage protocols positions an organization far more favorably in a privilege challenge or trade secret dispute than an ad hoc collection of individual consumer accounts, because the custom deployment demonstrates the kind of intentionality courts look for: reasonable measures to protect information, directed use by counsel, and institutional control over the AI workflow.

Governance Implications of the Distinction

Effective governance programs should stratify AI usage by deployment type and enforce controls proportional to risk. For consumer-tier generic AI, the minimum viable governance includes: clear policies prohibiting the input of privileged, confidential, or trade secret information; employee training on the discoverability of AI interactions; and contractual review of the provider's data retention and training policies.

For enterprise and custom deployments, governance can be more permissive about what information enters the system while imposing structural controls, such as role-based access, audit logging, defined retention periods, contractual data isolation, and documented chain-of-custody protocols. The governance objective shifts from prohibition to managed use, acknowledging that AI can meaningfully enhance legal work while ensuring that the interaction is structured to preserve the protections the organization needs.

The Billable Hour and the Judgment Premium

The custom-versus-generic distinction also highlights a practical tension that in-house and outside counsel are navigating in real time: AI makes certain tasks significantly faster, but speed alone does not create value and may actually erode it if the efficiency gains compress the time available for the judgment and analysis that clients are paying for. The real differentiator for legal professionals working with AI is not how quickly they can produce a first draft, but whether they can identify when the AI's output is wrong, incomplete, or subtly misleading.

Custom AI deployments that encode institutional knowledge and domain expertise, rather than relying on a general-purpose model's default behavior, help preserve this judgment layer; practitioners start from a higher baseline, spend less time on mechanical drafting, and invest more of their attention in the strategic analysis, issue-spotting, and experienced judgment that no model can replicate. The governance question is ultimately about structure: how to deploy AI so that it amplifies expertise rather than quietly replacing it.

Governance programs should be designed with that distinction in mind. The goal is to deploy AI in a way that amplifies expertise; the risk is deploying it in a way that quietly replaces the judgment clients are paying for.

7. Major AI Governance Risks

1. Beyond the discovery, privilege, and trade secret issues addressed above, AI governance programs must also account for several additional categories of risk that span industries and legal functions.
2. Algorithmic bias remains one of the most heavily legislated areas. AI systems trained on historical data can perpetuate existing biases in high-stakes decisions: hiring, credit, insurance, and criminal justice. Multiple states enacted legislation in 2026 specifically targeting algorithmic discrimination, with some creating private rights of action and statutory damages. General counsel should confirm that AI systems used in consequential decision-making have undergone bias testing and that the organization can demonstrate compliance with applicable anti-discrimination requirements.
3. Privacy exposure follows a related pattern. AI systems require large volumes of data, often including personal information, to function effectively. That creates liability under state privacy laws, sector-specific regulations like HIPAA, and emerging AI-specific data governance requirements. The risk compounds when employees input personal data into third-party AI tools without contractual protections or data processing agreements in place.
4. Transparency and explainability present a different kind of challenge. Many AI models operate as functional black boxes, producing outputs without an auditable reasoning chain. When organizations must explain AI-driven decisions to regulators, courts, or affected individuals, the absence of an explanation may itself constitute a compliance failure. Several 2026 state laws now require disclosure when AI is used in certain decision-making contexts.
5. The IP and confidentiality risks discussed in Sections 3 through 5 extend further than discovery and privilege. AI raises open questions about the ownership of AI-

generated outputs, the copyright status of content produced with AI assistance, and the adequacy of existing IP frameworks to address a technology capable of generating novel works without a traditional human author.

- 6. Regulatory compliance is complicated by fragmentation. The state-level landscape is expanding rapidly amid limited federal coordination, and companies operating across multiple jurisdictions face requirements that may conflict. The introduction of private rights of action in several 2026 state laws creates litigation risk that did not exist a year ago.
- 7. Shadow AI may be the most immediate governance problem. Employees using personal AI accounts for work, departments procuring AI tools without legal review, contractors integrating AI into deliverables without disclosure: all of it creates governance blind spots. Shadow AI is especially dangerous because it generates discoverable records, risks exposing trade secrets, and undermines privilege protections, all without the organization's awareness.
- 8. Finally, AI system reliability is a concern that touches every use case. Generative AI systems produce probabilistic, non-repeatable outputs; the same prompt may yield different results on different occasions. When the system generates plausible but factually incorrect content (commonly called hallucination), the consequences in legal contexts range from inaccurate contract analysis to fabricated case citations, as the Mata v. Avianca incident made nationally visible. Verification protocols for AI-generated content are not optional, particularly when that content will appear in legal filings, regulatory submissions, or client-facing communications.

8. Survey of Major 2026 State AI Legislation

The following table summarizes major state AI legislation enacted or taking effect in 2026. This survey is not exhaustive; additional states have introduced AI-related bills that remain pending or have narrower scope. General counsel should monitor developments in their operating jurisdictions and consult current legislative trackers for the most recent status.

State	Bill Number	Effective Date	Description	Key Provisions
California	SB 243	Jan 1, 2026	AI Companion Chatbot Regulation	Disclosure, minor protections, suicide protocols, \$1,000 statutory damages

Oregon	SB 1546	Jan 1, 2027	AI Companion Platform Safety	Transparency, minor safeguards, engagement limits, private right of action
Washington	HB 2225	Jan 1, 2027	AI Companion Chatbot Consumer Protection	Hourly disclosures, self-harm detection, manipulative engagement prohibition
Idaho	S 1297	2026	Chatbot Regulation	Companion chatbot safety and consumer protection
Tennessee	SB 1580	July 1, 2026	AI Mental Health Services Prohibition	Bans AI from being marketed as mental health professionals
Colorado	SB 24-205	Feb 1, 2026	Anti-Discrimination in AI (ADAI)	High-risk AI regulation, impact assessments, AG enforcement

9. Building an AI Governance Program

Governance is only useful if it scales to the organization that has to implement it. A three-attorney legal department and a 50-person enterprise legal team face many of the same AI risks, but they do not have the same resources, infrastructure, or tolerance for process overhead. The sections below outline what governance looks like at three different scales, organized around the components that matter regardless of size: ownership, inventory, risk assessment, policy, and oversight.

Small Legal Departments (1–3 Attorneys)

In organizations with small legal departments, general counsel often wear multiple hats and must prioritize efficiency and strategic leverage over comprehensive documentation. AI governance in this context is typically characterized by lean, practical approaches that rely heavily on collaboration with other functions. Rather than building a standalone AI governance infrastructure, small departments often integrate AI risk management into existing compliance, privacy, or vendor management processes.

Ownership of AI governance may be assigned to the general counsel directly or delegated to a senior member of the compliance, privacy, or IT team with legal oversight. The priority is ensuring that someone has clear responsibility, even if AI governance is not their sole function. Inventorying AI systems may begin informally, through conversations with department heads or review of vendor contracts, rather than through enterprise-wide

software audits. The goal is to develop situational awareness of where AI is being used, particularly in high-risk areas such as hiring, performance evaluation, customer-facing decision-making, or automated pricing.

Risk assessment in small departments tends to be practical and issue-spotting-oriented, focusing on identifying obvious legal exposures (such as the use of AI in employment decisions without bias testing, or customer-facing AI that may trigger disclosure obligations under new state laws). Formal impact assessments may be reserved for the highest-risk use cases. Policy development often starts with a concise acceptable use policy that sets boundaries on employee use of generative AI tools and requires consultation with legal before deploying AI in regulated contexts. Small departments benefit from adopting model policies or frameworks published by industry groups or legal associations, customizing them to fit organizational realities.

Vendor management becomes a critical governance lever for small departments, as many AI systems are procured rather than built in-house. Legal teams can require vendors to provide representations regarding bias testing, data security, and compliance with applicable AI laws, and can build contractual protections such as indemnities, audit rights, and termination provisions tied to regulatory changes.

Medium Legal Departments (4–10 Attorneys)

Medium-sized legal departments have the capacity to implement more structured governance frameworks while still balancing resource constraints. At this level, AI governance often becomes a cross-functional effort, typically involving legal, compliance, privacy, HR, IT, and relevant business units. General counsel in medium-sized departments may designate a specific attorney or form a small working group to lead AI governance efforts, supported by subject-matter experts from other departments.

AI inventories become more systematic, often maintained in a shared repository or compliance management system. These inventories typically capture key attributes of each AI system, including vendor or developer, business purpose, data inputs, decision-making authority (automated vs. human-in-the-loop), and applicable regulatory considerations. Medium departments may use questionnaires or intake forms to gather this information as part of procurement or project approval processes.

Risk assessment becomes more formalized, with differentiated treatment based on risk level. Many organizations adopt a tiered approach, categorizing AI systems as low-, medium-, or high-risk based on factors such as impact on individuals, regulatory exposure, and potential for harm. High-risk systems—such as those used in hiring, credit decisions,

healthcare, or pricing—may be subject to heightened review, including bias and fairness testing, privacy impact assessments, and legal sign-off before deployment.

Policy development now includes not only acceptable use policies but also AI development standards, vendor selection criteria, and data governance requirements. Medium departments may establish approval workflows that require legal review and sign-off for new AI deployments in specified categories. Training and awareness programs become more feasible as legal teams educate business stakeholders about AI risks, regulatory developments, and company policies.

Monitoring and oversight mechanisms often include periodic reviews of high-risk AI systems, tracking of AI-related incidents or complaints, and participation in cross-functional governance committees. Medium departments may also establish basic incident response protocols to address AI failures, bias complaints, or regulatory inquiries.

Large Legal Departments (10+ Attorneys)

Large legal departments have the resources and scale to implement mature, enterprise-wide AI governance programs that integrate with broader compliance, risk management, and corporate governance structures. In these organizations, AI governance is typically led by a dedicated role (such as an AI Ethics Officer, AI Counsel, or Algorithmic Accountability Lead) or a formal cross-functional committee with executive sponsorship.

AI inventories in large organizations are centralized and maintained in enterprise governance, risk, and compliance (GRC) platforms. These inventories capture granular details about each AI system's lifecycle, from initial development or procurement through deployment, updates, and eventual decommissioning. Large departments often implement automated discovery tools to identify shadow AI use and ensure that the inventory remains current.

Risk assessment processes are formalized and standardized, often incorporating impact assessment templates aligned with frameworks such as the NIST AI Risk Management Framework or the EU AI Act risk classification system (even for U.S.-based companies operating globally). High-risk AI systems undergo rigorous pre-deployment reviews, including algorithmic audits, bias and fairness testing, explainability assessments, privacy reviews, and documentation of human oversight mechanisms. Large departments may engage external auditors or consultants to provide independent assessments.

Policy in large organizations covers AI development standards, ethical principles, procurement requirements, acceptable use, data governance, and algorithmic transparency, typically integrated into broader corporate governance documents and

updated as regulations change. Training programs at this scale are role-specific, with separate tracks for developers, business users, procurement teams, and executives.

Oversight and monitoring mechanisms are institutionalized, often including AI governance boards that review high-risk deployments, track performance metrics, investigate incidents, and ensure ongoing compliance. Large departments establish clear escalation pathways for AI-related concerns and maintain audit trails to demonstrate accountability. Incident response protocols are detailed and tested, covering scenarios such as algorithmic bias complaints, data breaches involving AI systems, regulatory investigations, and system failures.

Large legal departments also engage proactively with external stakeholders, including regulators, industry groups, and standard-setting bodies. At this scale, AI governance stops being a compliance obligation and becomes a competitive position.

10. Conclusion and Best Practices

AI governance in 2026 occupies an uncomfortable position: the technology is being adopted faster than the legal frameworks meant to govern it, but the consequences of ungoverned adoption are not waiting for the law to catch up. Discovery motions, privilege challenges, and trade secret disputes involving AI are being decided in courtrooms right now.

This paper has identified three developments that demand immediate attention from general counsel. AI-generated content is discoverable ESI, subject to the same preservation and production obligations as any other electronically stored information. Trade secret protections are vulnerable when employees input proprietary information into AI platforms that cannot guarantee confidentiality. And the attorney work-product doctrine is being tested by courts confronting whether AI-assisted legal analysis retains its protected character, with outcomes turning on the distinction between directed, enterprise-grade use and unsupervised, consumer-platform interactions.

Regardless of department size, the practical takeaways converge on the same core actions:

Conduct an AI inventory. Identify every AI system in use across the organization, including shadow AI adopted by individual employees. You cannot govern what you do not know about.

Include AI in litigation hold protocols. AI interaction logs (prompts, outputs, and metadata) must be addressed in preservation notices and ESI protocols. Failure to preserve AI-related records carries the same spoliation consequences as destroying relevant emails.

Stratify governance by deployment type. Consumer-tier generic AI and enterprise custom deployments create different risk profiles and warrant different controls. Apply restrictions proportional to the actual risk.

Audit AI vendor contracts. Understand how each provider handles data retention, model training, and confidentiality. The distinction between consumer and enterprise platform terms is legally significant and will likely be relevant in privilege and trade secret disputes.

Update trade secret protection programs. The reasonable measures analysis now includes the AI environment. Policies, training, and contractual protections must address what information employees may and may not input into AI tools.

Document the attorney's direction of AI use. If AI is part of the litigation or investigation workflow, ensure that counsel directs and supervises the AI interaction. The Heppner court left open the possibility that attorney-directed use of AI may qualify for privilege protection; unsupervised use almost certainly will not.

Invest in training and AI literacy. Legal departments must understand the AI tools their organizations use well enough to advise on governance, respond to discovery requests, and assess privilege claims. Less than half of in-house lawyers feel prepared for the impact of generative AI on their careers; that preparation gap is now a liability gap.

Monitor the caselaw, not just the statutes. The most consequential developments in AI law are happening in courtrooms, not legislatures. The Heppner, Warner, Trinidad, and In re OpenAI decisions discussed in this paper are early markers, not final answers. The law will continue to develop rapidly.

The question facing general counsel is not whether to allow AI; the workforce has already answered that. The question is whether governance will be in place before the next discovery request, the next trade secret dispute, or the next privilege challenge forces the issue. General counsel who build those programs now, rather than waiting for regulatory certainty that may never arrive, will be the ones who give their organizations a meaningful head start.

Resources and References

- NIST AI RMF: <https://www.nist.gov/itl/ai-risk-management-framework>
- EqualAI Governance Playbook: <https://www.equalai.org/>
- ACC AI Toolkit: <https://www.acc.com/resource-library/ai>
- State AI Law Tracker: <https://www.multistate.ai>
- FTC AI Guidance: <https://www.ftc.gov>
- EEOC AI Guidance: <https://www.eeoc.gov>

Key Cases Discussed in This Paper

United States v. Heppner, No. 25-cr-00503-JSR, 2026 U.S. Dist. LEXIS 32697 (S.D.N.Y. Feb. 17, 2026)

Warner v. Gilbarco, Inc., No. 2:24-cv-12333, 2026 U.S. Dist. LEXIS 27355 (E.D. Mich. Feb. 10, 2026)

Tremblay v. OpenAI, No. 23-cv-03223-AMO, 2025 U.S. Dist. LEXIS 38457 (N.D. Cal. Feb. 27, 2025)

In re OpenAI, No. 25-md-3143, 2026 U.S. Dist. LEXIS 64199 (S.D.N.Y. 2026)

Trinidad v. OpenAI, No. 25-cv-06328-JST, U.S. Dist. LEXIS 1129 (Jan. 5, 2026)

OpenEvidence v. Pathway Medical, Inc., No. 1:25-cv-10471 (D. Mass. 2025)

Additional Resources

DLA Piper, "Mark It: The Coming AI Shockwave in Trade Secret Disputes" (March 2026)

Quinn Emanuel, "Preventing, Detecting, and Litigating Trade Secret Theft in the Age of AI" (April 2026)

Hogan Lovells, "The Emerging Rules of the Road Governing AI Prompts and Outputs in Discovery" (February 2026)

Harvard Law Review Blog, "United States v. Heppner" (March 2026)