Consilio

New Tools, Same Rules

Applying Established Best Practices to the Ever-Changing AI Landscape



Executive Summary

Artificial intelligence (AI) tools available to legal professionals are changing rapidly, but many of the responsibilities associated with using them are familiar. Understand the technology. Protect sensitive information. Preserve relevant evidence. Verify important conclusions. These established principles provide a foundation for navigating an AI landscape that continues to introduce new capabilities, new records, and new risks.

AI has moved from experimentation to operational reality. Leadership demands, client expectations, and AI features embedded in everyday applications are accelerating adoption. AI can increase legal teams' capacity across drafting, contract analysis, document review, and case assessment. Even organizations that have not adopted AI must be prepared for relevant evidence generated by employees, business partners, and opposing parties using it.

Legal departments and law firms must balance innovation with governance. Realizing AI's benefits requires training, integration, and oversight, supported by approved platforms, clear policies, confidentiality protections, and accountable professionals. Governance should enable useful applications while managing risk, and success should account for the effort required to verify outputs and correct errors.

Existing discovery principles still apply, but organizations must adapt them to AI-generated content and AI-assisted workflows. Chatbot conversations, meeting summaries, and Copilot interactions create potential evidence requiring attention to metadata, linked sources, and context. Court decisions involving hallucinations, protective orders, transparency, and prompts illustrate how established duties apply. AI use does not remove professional accountability or automatically determine discoverability or protection.

New tools, same rules. The landscape will continue to change, but legal professionals already have a foundation for navigating it.

Introduction

If you've ever returned to the neighborhood where you grew up after years away, chances are it has changed considerably. Restaurants that you used to frequent may be different or gone entirely. Your favorite grocery store may now be a gym. Previously undeveloped land may now be developed as an apartment building. Change is constant for urban development.

While many of the businesses and restaurants may have changed, the roads to navigate around the neighborhood haven't changed much, if at all. You still know the correct road to get from one place to another. What exists at the destination may have changed, but the road to get there stays the same.

Change is also constant for technology, and that's especially true for artificial intelligence (AI) technology. The tools we use are no longer just evolving – they are transforming rapidly. Generative AI chatbots like ChatGPT and Claude can be directed to generate all types of content – including text, images and videos – seemingly instantly. AI note-takers can create meeting summaries as soon as the meeting has ended. Tools like Microsoft Copilot automate everyday business tasks. AI agents are automating increasingly complicated workflows – including legal and eDiscovery workflows.

It's no wonder GenAI reached nearly 62% adoption among US adults in its first four years, considerably faster than previous transformational technologies like the personal computer and the internet. 71% of organizations reported regularly using generative AI in at least one business function.

Keeping up with the continual pace of change – especially given pressures to stay current with technology – can be extremely challenging. Given the enormous potential of GenAI, upper management is pressuring corporate legal departments to operationalize AI to manage and perform legal services faster, cheaper and better than they ever have before.

But the pressure to move so quickly has led to considerable challenges. The MIT State of AI in Business 2025 report found that approximately 95% of generative AI pilots fail to deliver measurable impact on the P&L.

The tools are constantly changing, but the best practices and professional responsibilities associated with using those tools never change. Understand the tools you use. Protect sensitive information. Preserve relevant evidence when required. Check the basis for important conclusions. Document decisions sufficiently to explain and defend them. These are principles that never change.

Formal Opinion 512 from the American Bar Association (ABA) is a prime example of how the same rules apply to new tools. ABA Formal Opinion 512 applies existing professional duties to generative AI, including competence, confidentiality, communication, supervision, candor, and reasonable fees – as discussed in the ABA Model Rules of Professional Conduct, which has been guiding US attorneys for over 40 years.

You've been down this road before. When predictive coding technologies first became popular in the early 2010s, questions abounded then as well. There were concerns about how these technologies could be used competently, safely and defensibly. There were even concerns expressed in publications like The New York Times in 2011 that these technological advances would replace "armies of expensive lawyers" in eDiscovery (five years later, that same author declared "not so fast" regarding technology replacing lawyers). Court rulings like Da Silva Moore approved the use of technology assisted review, and best practices for using the technologies were formalized and standardized.

So, how can legal professionals apply the same rules and best practices to newer AI technologies like generative AI? This paper will discuss the current landscape of the "new tools" of AI in terms of:

Why legal can no longer wait to embrace AI

Examples of legal use cases for AI that provide practical value today

Evidentiary considerations associated with AI-generated evidence

What the courts are saying about AI in case law rulings and how it impacts discovery

With the understanding of AI's "new tools," we will then discuss how the application of the "same rules" and established best practices can be used to manage the ever-changing AI landscape, including:

The Importance of a Strong AI Governance Program
How AI governance is an ongoing business process, regardless of how AI technology changes.

A Roadmap for Preparedness in the AI Era
A list of ten practical recommendations for your organization to achieve AI preparedness.

AI Has Arrived: Why Legal Can No Longer Wait

Given all the caveats and risks associated with generative AI, it can be tempting to wait "until the dust settles" to consider adopting GenAI to support legal processes and workflows. But legal professionals can no longer wait to embrace GenAI. Here's why.

Leadership Pressure is Accelerating AI Adoption

Legal teams are increasingly being asked to explain how they will operationalize AI. Understanding, choosing, and deploying new legal technologies was the leading strategic challenge in Consilio's 2026 Global Survey Report, selected by 54% of respondents. Workload followed at 52%, operational efficiency at 44%, and budgetary challenges at 36%.

Leadership is part of that pressure. In the same report, 41% of respondents identified leadership vision as a driver of change and innovation, while 51% selected emerging technology opportunities. These findings indicate executive direction and perceived practical opportunities are driving adoption of AI.

Another Consilio report, The AI Mandate in Legal, examines the leadership mandate more closely. It divides 678 respondents into a "mandated" cohort of 279 who selected leadership vision as an innovation driver and a "non-mandated" cohort of 399 who did not.

Within those cohorts, the differences are substantial:

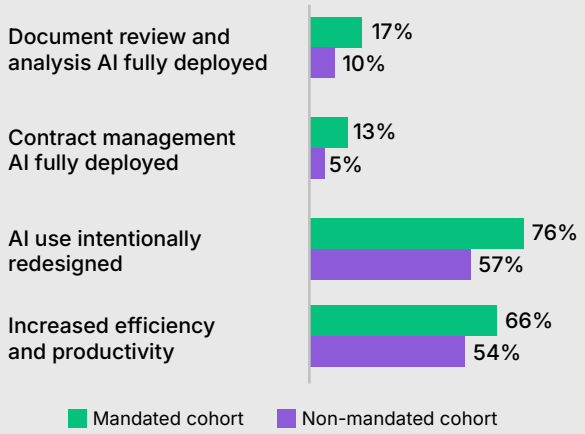

These results show an association between leadership involvement and greater deployment and reported benefits. They do not establish that leadership pressure alone caused those outcomes. Investment, organizational readiness, and other differences may also contribute.

Mandates for Legal Departments to Increase Efficiencies

For corporate legal departments, AI adoption sits within a broader expectation to support more business activity with disciplined spending and limited capacity.

The pressure extends beyond producing documents faster. Legal may need to shorten contract turnaround, investigate issues earlier, improve access to internal knowledge, and give executives timely assessments of business risk.

A useful response translates executive ambition into defined business outcomes. A mandate to "use AI" provides little guidance about priorities. A goal to reduce the time required to identify contractual change-of-control provisions, while maintaining an agreed quality standard, gives the team a task it can test.

The measure of success should include the entire workflow. Time saved generating an initial answer may be offset by correcting errors, transferring information between systems, or rebuilding the analysis in a usable format. In the report The AI Mandate in Legal, mandated respondents reported more manual workarounds between tools than non-mandated respondents **– 50%** vs. **32%** – and more fragmented systems, at **47%** vs. **37%.** Mandates alone don't guarantee success.

That's why corporate leadership should fund integration, training, and quality control alongside licenses. Otherwise, the department may increase its technological activity without achieving a corresponding improvement in service.

Setting AI Expectations for Outside Counsel

Corporate expectations flow down to outside counsel. In the 2026 Global Survey Report, in-house priorities were centered on internal productivity, while outside counsel firms placed greater emphasis on client delivery, litigation support, and professional risk.

Corporate legal departments can turn those expectations into practical outside counsel guidelines. Firms should be prepared to explain where AI adds value, how client information is protected, how outputs are checked, and how efficiency affects the proposed budget.

This creates opportunities for firms to consider fixed fees, phased fees, or other arrangements for defined services. Such arrangements should identify the work being performed, quality requirements, assumptions, and treatment of unexpected complexity.

Billing principles continue to apply. ABA Formal Opinion 512 explains that hourly billing must reflect actual time, including appropriate review of AI-assisted work. Flat fees and technology charges also remain subject to reasonableness requirements. AI-generated time savings do not justify billing hours that were never worked.

A productive client–firm discussion should connect price with verified service improvements: faster answers, better coverage, predictable delivery, or more useful analysis.

The Use of AI is Ubiquitous

While the "everyone is using it" argument isn't enough of a reason by itself to use it, AI is increasingly embedded in the legal applications professionals already use, making it part of everyday work across the legal department.

Legal research and drafting platforms integrate AI to answer research questions, analyze authorities, and help prepare and revise documents. Contract review and management tools incorporate AI through contracting workflows, supporting review, approvals, and access to contract information, while eDiscovery solutions embed capabilities for document review, privilege identification, and case strategy. These capabilities extend into familiar productivity and collaboration applications, where AI can assist with correspondence, document creation, and meeting summaries.

AI as a Force Multiplier for Legal

Generative AI can serve as a force multiplier for legal by accelerating drafting, summarization, contract analysis, and document review, giving professionals more time for judgment, strategy, and client needs. In eDiscovery, it can help identify significant documents, develop timelines, and flag potentially responsive, privileged, or personal information while enabling teams to explore questions and evidence that limited resources might otherwise leave unexamined.

Many eDiscovery teams are already applying GenAI and LLMs to multiple use cases. According to the 2026 State of the Industry report from eDiscovery Today, more than **30%** of **559** discovery professionals surveyed said they apply LLMs and generative AI to seven different eDiscovery use cases, with three of those use cases – Document Review Automation (at **66.2%**), Document Summarization (at **65.8%**) and Case Strategy/Early Case Assessment (at **52.4%**) – being applied by more than half of respondents.

eDiscovery Use Cases for LLMs and GenAI in 2026

"To which eDiscovery/litigation cases do you or your organization apply large language model/generative AI technologies today in actual projects and cases (not demos and testing)?

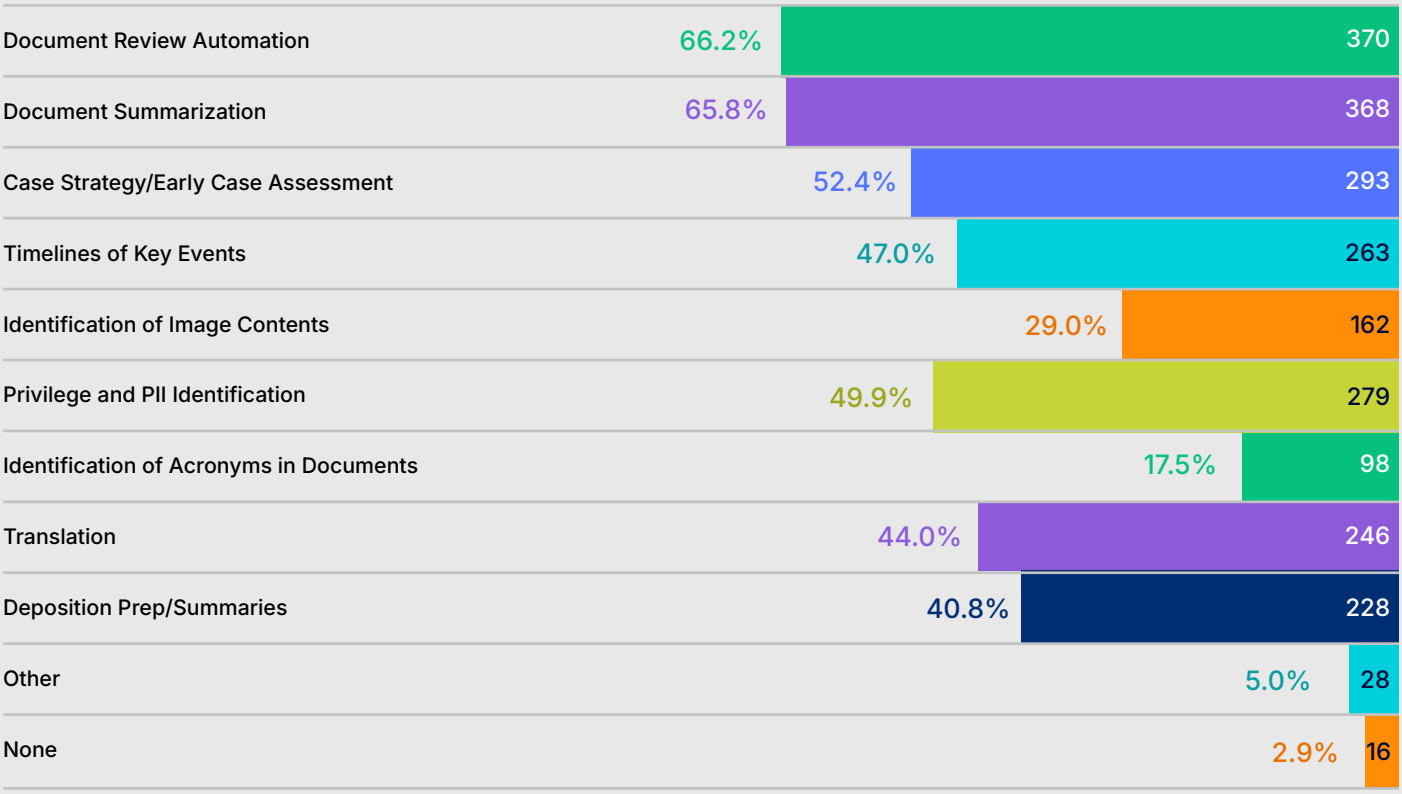

(Source: eDiscovery Today)

Legal Use Cases Providing Value Today

The following examples illustrate workflows and use cases that are already providing value to legal teams today and how legal professionals should evaluate these capabilities and use them defensibly.

Drafting Assistance

AI can prepare initial versions of correspondence, policies, discovery requests, interview outlines, and internal memoranda. Providing approved templates and clear factual inputs can make the first draft more useful. Legal professionals should verify factual assertions, commitments, legal positions, and consistency with the client's objectives before use.

Legal Research

AI can help identify issues, suggest search directions, and summarize authorities. While legal research platforms are purpose-built for accuracy, hallucinations can still occur. Verification should include whether an authority exists, whether it supports the proposition, and whether it remains applicable.

Contract Review

AI can extract clauses, compare terms against a playbook, identify missing provisions, and organize deviations for review. Corporate teams can use this approach to prioritize agreements needing attention. Evaluation should account for amendments, incorporated documents, exceptions, and commercial context.

Case Strategy & Early Case Assessment

Within eDiscovery platforms, GenAI can help organize allegations, locate relevant communications, identify potential witnesses, and construct preliminary timelines. A useful chronology should connect each event to supporting evidence and distinguish an event date from a document's creation or transmission date. Strategic conclusions should remain provisional until collection is complete.

Document Summarization

AI can condense lengthy records and, where supported, describe images or extract information from screenshots and scanned documents. Again, source references should allow reviewers to link back to the underlying material. Teams should distinguish visible content from interpretation, particularly where image quality, handwriting, or ambiguous objects affect the result.

Document Review Automation

In automating document review for discovery, AI can propose responsiveness and issue classifications, prioritize documents, and explain the basis for each call. The validation standards established for technology assisted review, including sampling, elusion testing, and quality control, apply to AI-driven review as well. As discussed later in this paper, at least one court has already treated a generative AI review tool as a form of TAR.

Privilege & PII Identification

Within eDiscovery platforms, AI can flag potentially privileged communications and personal information for further assessment. Privilege calls require context about participants, purpose, confidentiality, and governing law. PII handling also requires checking whether redactions remove the information from the produced file, underlying text, and relevant metadata. This is an area where more QC is typically expected to minimize the inadvertent disclosure of privileged, personal or confidential information.

Transcription & Translation

AI can make audio, video, and multilingual documents more accessible for investigation and review. Critical passages should receive additional verification where speaker identification, specialized terminology, accents, or translation choices may change their meaning.

Deposition Prep and Depo Summaries

AI can assemble witness-specific document sets, draft question outlines, and summarize testimony. Counsel should verify page-and-line references and confirm that any claimed inconsistency reflects the actual testimony of the witness, including qualifications and surrounding context.

Trial Preparation and Presentation

AI can assist with exhibit organization, timelines, demonstrative concepts, and proposed examinations. Any presentation should accurately distinguish original evidence from summaries, annotations, or generated illustrations.

AI's Value in Legal Workflows



Legal Research
AI identifies and summarizes legal authorities



Drafting Assistance
AI prepares initial legal documents



Contract Review
AI extracts and compares terms



Case Strategy
AI organizes case information and timelines



Document Summarization
AI condenses lengthy documents



Document Review
AI automates document classification and prioritization



Privilege Identification
AI flags privilege and personal information



Transcript & Translation
AI converts audio/video to text and translates languages



Deposition Prep
AI assembles witness documents and questions



Trial Preparation
AI organizes exhibits and timelines for trials

There will undoubtedly be many future legal use cases to come for the application of AI.

AI as a New Source of Discoverable ESI

Under Federal Rule of Civil Procedure 26(b)(1), discovery generally concerns nonprivileged information relevant to a claim or defense and proportional to the needs of the case. Whether or not legal professionals leverage GenAI to accomplish legal tasks, it is impacting them in a very important way – creating discoverable evidence under FRCP 26(b)(1).

Just as mobile device data, chat and collaboration app data (from apps like Slack and MS Teams) and data from enterprise solutions are modern sources of electronically stored information (ESI) to address in discovery, so is AI-generated data. Here are types of discoverable AI-generated data to be prepared to address.

AI Chatbot "Conversations"

As of June 2026, around 1.1 billion people worldwide used ChatGPT each month, up from about 358 million monthly users in January 2025. Chances are there are employees in your organization that use ChatGPT, Claude, Google Gemini or some other AI chatbot.

AI chatbot histories may reveal what employees knew, shared, considered, or used to develop documents. Interactions can be relevant even when a draft was never sent or an output was inaccurate, because they show information the user received or acted upon.

Custodian interviews should identify services and accounts used, including "shadow AI," along with uploaded files, shared conversations, and outputs copied elsewhere. Counsel should assess privilege and work-product protection under applicable law and circumstances; legal language in a prompt alone does not establish protection.

AI Meeting Transcripts and Summaries

AI note-takers now run in many work meetings. Each one can create several related records, including recordings, transcripts, summaries, action items, and follow-up messages. Each record may convey different information.

A summary may omit uncertainty or attribute an action item incorrectly. A transcript may contain speaker-identification errors. Where the underlying recording exists and falls within preservation scope, it may be necessary to evaluate those differences.

Even the slightest mistake can be hugely important. In U.S. v. Smith, law enforcement provided an AI-generated transcript of the defendant's interrogation to the court, which recorded him as saying: "But I need to have an attorney here." But, according to the court, from the video it was "clear that Smith asked the question, 'Do I need to have an attorney here?'" (emphasis added), confirming that the defendant did not request an attorney.

Configuration of the platform can also be important. Microsoft explains that Copilot can operate in a Teams meeting without creating a saved transcript, while prompts and responses may still be retained under the organization's Purview policies. So, disabling transcription doesn't necessarily eliminate AI interaction records.

Copilot Artifacts

Copilot documentation illustrates why AI evidence requires more than collecting the finished document.

Microsoft documents retention of user prompts and responses, including links and references, through compliance records stored in Exchange Online mailboxes. It also cautions that what users see in an AI application does not reliably indicate what remains retained for compliance purposes.

Audit logs provide another layer of ESI that could be discoverable. Microsoft identifies properties relating to the user, time, host application, accessed resources, and, where applicable, agents and plugins. Examples include resource identifiers, locations, sensitivity labels, and actions.

A discovery assessment for Copilot should consider prompt and response content, relevant audit events and available metadata, source files and the versions needed to understand the interaction, and documents, messages, or other records created or changed through the workflow. Where supported, the assessment should also address relevant stored memories or personalization records.

Microsoft separately identifies Copilot memory records and explains that deleting a conversation does not delete an associated memory. This reinforces the need to test each relevant record category rather than assume a single search or retention setting covers everything.

Available fields, collection methods, and retention behavior should be verified against the organization's actual licensing and configuration. By itself, a reference to a source file does not establish that the version used during the interaction has been preserved.

AI-Generated Drafts and Document Revisions

A final document may reveal only part of its history. Earlier AI-generated drafts, prompts, human edits, and approval communications may explain how a position developed or why the final version changed.

Retention and preservation should reflect the record's purpose, applicable obligations, and the matter's circumstances. Not every intermediate output requires indefinite retention, but collecting only the final document may omit relevant context.

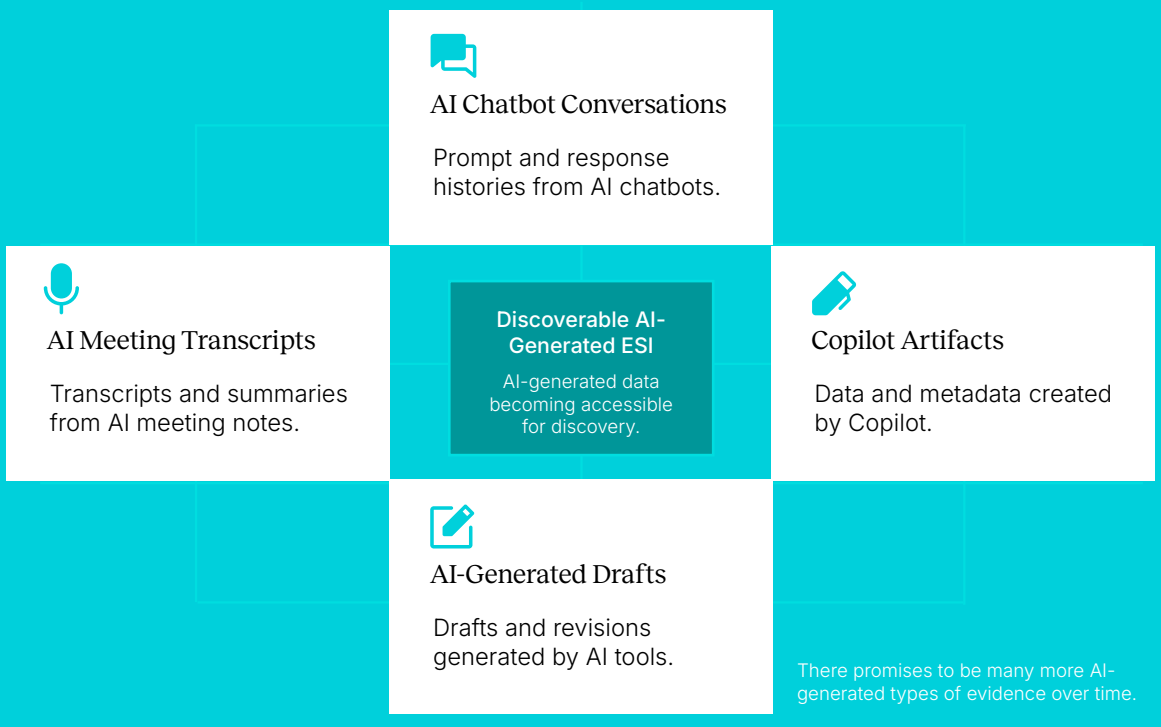

What the Courts Are Saying About AI

As GenAI adoption and AI-generated evidence expand, courts are confronting questions about hallucinations in filings, protecting confidential discovery materials from uploads to public LLMs, transparency in AI-assisted workflows, and whether discovery prompts qualify as attorney work product or must be produced.

Recent rulings, including many from 2026, are beginning to define how legal and eDiscovery teams should address these issues, applying established principles to emerging tools and evidence.



AI Hallucinations in Court Filings

When six submitted cases by an attorney in a filing were found to be AI hallucinations in Mata v. Avianca, Inc., in May of 2023, the attorney's mistake was covered by CNN and The New York Times and discussed in webinars and conferences for months.

Despite the widespread coverage of the failure to properly use AI in that case, the number of cases involving AI hallucinations has increased dramatically. According to Damien Charlotin's AI Hallucination Cases database, there have been 2,004 cases worldwide involving AI hallucinations since 2023. Here is the breakdown, year-over-year:

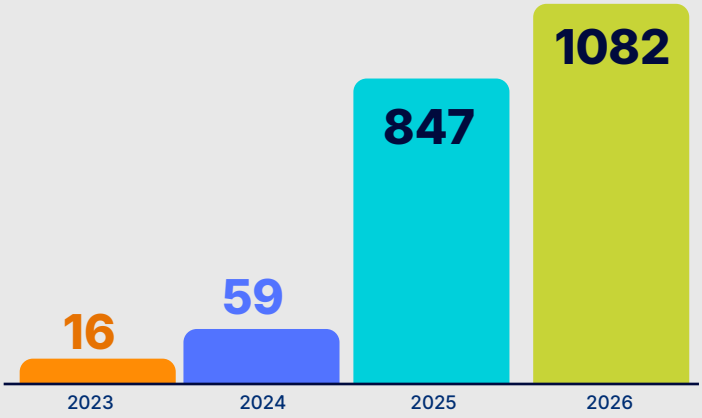

Reported Cases with AI Hallucinations 2023-2026
(Source: Damien Charlotin)

Even with some courts issuing severe sanctions to lawyers – including disqualification from the case, suspension of license to practice, referral to the State Bar – and other courts issuing orders and guidelines regarding the use of GenAI and the duty of lawyers to check and verify AI outputs, the numbers have continued to rise – considerably.

Pro se parties reflect more than half of the overall total referenced above – 1,149 cases total. Pro se filings have been spiking in the GenAI era as those parties have become more likely to use public AI models like ChatGPT to prepare their filings instead of using a lawyer.

Still, over 800 of the cases involve lawyer mistakes[2] – some of which are lawyers at Big Law firms. With ABA Formal Opinion 512 emphasizing competence when using GenAI, it's important to remember the same rules apply as they always have for these new tools. Nobody wants their mistake to be covered by The New York Times.

Protective Orders and Public AI

Organizations are increasingly seeking protective orders restricting opposing parties from uploading confidential discovery materials to public AI models. Recent rulings address the scope of these restrictions. Here are three of those rulings:

Jeffries v. Harcros Chems. Inc.

The court granted Defendant's proposal for protective order modifications to extend protection against uploading confidential documents into open generative AI tools to all documents and information produced in discovery, rejecting Plaintiffs' contention that the proposal would drive up Plaintiffs' costs by prohibiting their use of open AI Tools to analyze productions.

Morgan v. V2X, Inc.

The court wrote a balanced provision determining to what extent the protective order should restrict the use of AI which restricted uploading confidential documents only. The court also ruled that Plaintiff must disclose the name of any AI tool he had already used in connection with Confidential Information.

Litton v. Roblox Corp.

The court rejected Defendant's proposed changes from the court's Model Protective Order primarily addressing the use of Confidential/Highly Confidential Information with an AI tool, finding there was a set of provisions addressing AI tools in that Standing Order.

Protective orders that limit the ability of opposing parties to upload at least confidential produced documents into a public AI tool are a "must-have" today for organizations looking to protect sensitive materials.

2 There are also AI hallucination incidents from experts and even judges.

Discovery on Discovery

One big question that legal professionals have been asking: What level of transparency of the process will courts expect for parties using GenAI? At least one court has addressed that question.

In Schulte v. LinkedIn Corp., the court denied plaintiffs' request to prohibit LinkedIn from applying search strings or keywords to pre-cull the custodial documents reviewed by Relativity aiR and denied plaintiffs' request to compel LinkedIn to disclose various metrics about its use of Relativity aiR, finding "the number of documents in the target population alone does not warrant the requested discovery on discovery."

Equally important was the court's treatment of Relativity aiR as "a form of Technology Assisted Review", which establishes that courts can apply the **same rules** which many courts over the years have applied regarding process and discovery on discovery to **new tools**.

Discoverability of AI Prompts

Perhaps the most common question asked about the use of GenAI for eDiscovery is: Will AI prompts used during discovery be protected as attorney work product, or will they be discoverable by the opposing party?

Three recent case law rulings illustrate how courts are addressing this issue:

Tremblay v. OpenAI, Inc.

The court found that the order for plaintiffs to produce all prompts and outputs for their testing of ChatGPT was a "misapplication of law" because the ChatGPT prompts were crafted by counsel and reflected counsel's mental impressions and were attorney work product.

Concord Music Grp., Inc. v. Anthropic PBC

The court denied Anthropic's request to compel production of undisclosed prompts and outputs from Publishers' legal investigation, citing attorney work product privilege and limited waiver.

Conservation Law Foundation v. Shell Oil

A magistrate judge ordered production of AI prompts used by an expert as part of her methodology. The court later put the order on hold while the district judge considered an objection, illustrating the still-developing treatment of AI prompts in expert discovery. Based on initial case law, courts are applying the same rules to new tools. Prompts generated by attorneys as part of the discovery process are likely to be considered attorney work product because they typically reflect mental impressions of the attorney. Prompts generated by other parties (such as experts) may not benefit from work product protection, especially if they were not generated at the direction of counsel.

The Importance of a Strong AI Governance Program

Governance should operate throughout an AI system's use. Procurement approval is one checkpoint within that process.

Consilio's 2026 Global Survey Report identifies a significant concern regarding how companies organize governance around AI selection, implementation, and adoption: only 7% of respondents said they had a documented AI governance framework actively followed in practice as their governance arrangement. 32% of respondents indicated they use a cross-functional committee to do so, and 22% indicated that it's the responsibility of general counsel or legal leadership.

The voluntary Generative AI Profile (NIST AI 600-1) published by the National Institute of Standards and Technology (NIST) provides a useful reference for lifecycle risk management, including evaluation, oversight, documentation, and monitoring. Legal departments can use that framework to help structure a program, while tailoring implementation to their organization and legal obligations.



Components of an Effective AI Governance Program

The following elements constitute the framework of an effective enterprise AI governance strategy.

Enterprise AI Policies

An enterprise AI policy should define objectives, risk tolerance, decision ownership, and approval requirements, with safeguards tailored to each use case. Brainstorming with public information warrants different controls than analyzing confidential contracts, supporting employment decisions, or automating external communications.

Acceptable Use Guidelines

While most organizations have established initial AI policies, there is a significant gap between policy statements and practical execution. Many employees need specific guidance on what they can and cannot do in their daily tasks. Here are two considerations for acceptable use guidelines:

Simplified Guidelines: Instead of complex data classification schemes, simple "dos" and "don'ts" rules are typically more effective. For example, a guideline that says: "if you wouldn't email it, don't put it in AI" can cover many of the security concerns that organizations might need to address.

Clarity and Accessibility: Guidelines should be clear and visually accessible (e.g., a two-column list of permitted and prohibited actions) to ensure clarity for the workforce.

Confidentiality Protections

Trust in AI depends on protecting confidential information through vetted environments and controls over data access, processing, and storage. Maintaining equipment and data centers in specific jurisdictions (e.g., Canada, Australia, the EU) can support contractual, privacy, security, and data-sovereignty requirements. Organizations should assess those requirements separately from obligations under regulations such as the EU AI Act, which establishes requirements for certain AI systems based on their use, risk, and role in the AI value chain. Internal access controls are also important. AI can make sensitive information easier to locate, exposing it to unauthorized employees when permissions are too broad, even within an otherwise secure environment.

Approved AI Platforms and Policies Regarding Shadow AI

"Shadow AI" includes unauthorized AI tools and features introduced into existing platforms without security review. As vendors add AI assistants and new subprocessors, organizations must reassess how these changes affect data handling and risk.

An approved-tool register should document authorized products, accounts, uses, data, owners, and approval conditions. A tiered approval process can fast-track low-risk uses (such as public data or personal productivity tasks) while requiring greater scrutiny of use cases involving sensitive HR, compensation, or legal data, balancing productivity with appropriate safeguards.

Business leaders generally prefer a "freedom to use unless prohibited" approach rather than a rigid block list, so it's important to balance the freedom to use tools to boost productivity with restrictions to limit high-risk exposure.

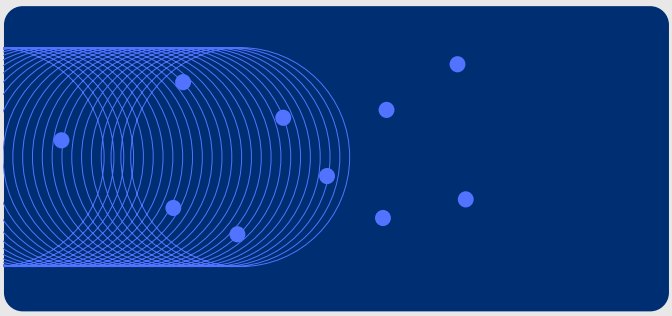

Human Oversight Requirements

Oversight should identify who reviews the work, what they check, and when the workflow must pause.

ABA Formal Opinion 512 recognizes that appropriate verification depends on the tool and task; it does not require an identical review method for every use. Lawyers nevertheless retain responsibility for exercising professional judgment.

The shift from deterministic systems to non-deterministic, probabilistic GenAI systems, which can produce varying or inaccurate outputs (even in closed, purpose-built systems) makes human review essential to identify errors. Agentic AI also requires oversight, including human authorization for consequential actions such as sending external communications.

Auditability and Vendor Due Diligence

Auditability is a fundamental component of a comprehensive enterprise AI governance structure.

When vetting third-party providers, client requests for audit rights have become increasingly complex and less straightforward than in traditional IT, due to how modern AI systems and architectures are set up:

Annual Risk Audits: Enterprise clients routinely audit their key vendors on an annual basis as part of ongoing third-party risk management.

The "Guided Tour" Approach: While vendors do grant audit rights to clients, many do not allow clients unrestricted access to navigate or touch the environment. Instead, they provide a structured "guided tour" that allows clients to inspect necessary security controls under supervision.

Data Retention

Federal Rule of Civil Procedure 37(e) addresses loss of ESI that should have been preserved, where reasonable preservation steps were not taken and the information cannot be restored or replaced. So, it's important to establish rules that apply to AI-generated evidence.

Organizations must align retention of AI-generated evidence – such as AI chatbot conversations and AI-generated transcripts – with their broader retention and destruction policies, understanding the importance of suspending those policies when appropriate to conform to legal hold requirements.

Governance Committees

Just as is the case for information governance – as illustrated by EDRM's Information Governance Reference Model (IGRM) – an effective AI governance program should connect legal, IT, information security, privacy, compliance, records management, procurement, and business stakeholders.

Balance is key between risk and opportunity. While responsible use of AI is important, governance committees risk becoming overly risk-averse echo chambers that could inadvertently block high-value AI use cases, without business unit leaders at the table.

The figure here illustrates the eight components of an effective AI governance program.

Legal's Role in Enterprise AI Governance

Legal helps translate professional, contractual, regulatory, and litigation obligations into requirements the business can implement. It should work as a part of the governance committee with technical and operational owners to determine whether those requirements function in practice. Within discovery, that includes reviewing protective orders and protocols before using AI with protected material.

Legal's contribution is strongest when it makes responsible adoption practical: clear rules, assigned ownership, tested controls, and an accessible process for resolving uncertainty.

COMPONENTS OF AN EFFECTIVE AI GOVERNANCE PROGRAM

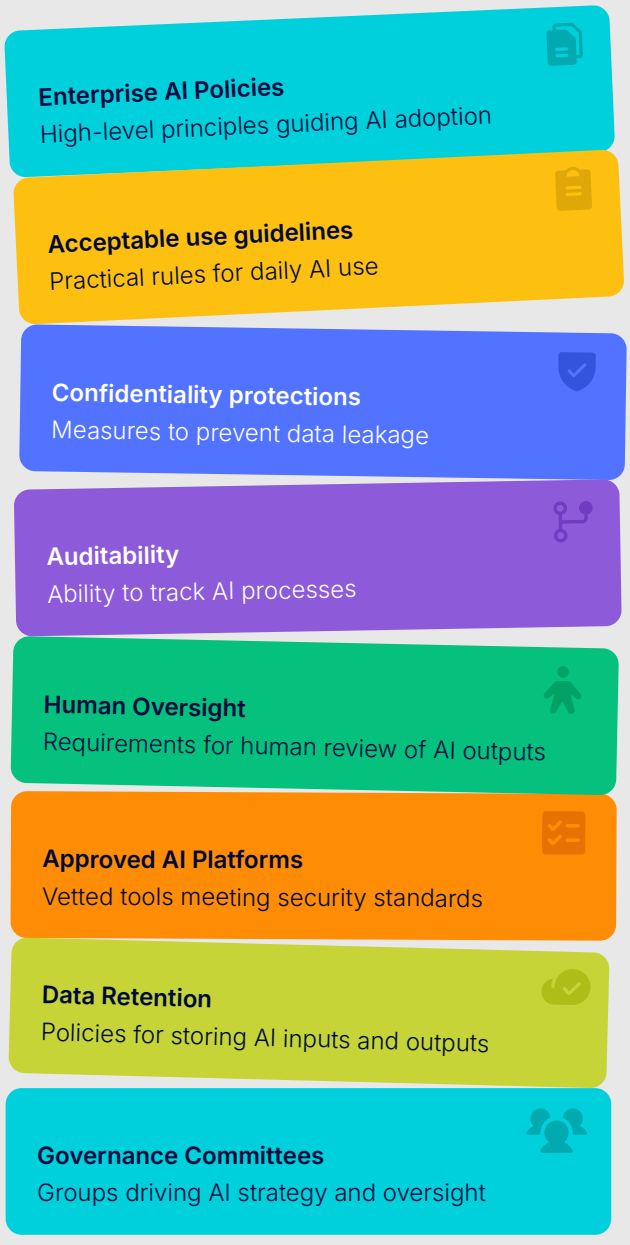

A Roadmap for Preparedness in the AI Era

You've been down this road before. Still, a roadmap is always important to enable your organization to achieve its goals – without one, you might find your team has gone off the road to successful AI preparedness.

Practical Recommendations for AI Preparedness

With that in mind, here are ten practical recommendations for your organization to achieve AI preparedness to face today's challenges, as well as those in the future.

1. Inventory AI Usage

An effective AI inventory should capture the full range of tools and capabilities employees use to conduct business, including standalone applications, AI features embedded in existing software, personal accounts used for work, meeting assistants, integrations, and agents that perform tasks across systems.

For each identified use, record the responsible business owner, users, purpose, connected data sources, information submitted, and outputs created. Document where those records reside, who can access them, how long they are retained, and what options exist to preserve and export them, including associated metadata.

Assign responsibility for maintaining the inventory and updating it as applications introduce new features, employees adopt new workflows and approved uses change.

2. Approve Enterprise AI Platforms

Evaluate AI platforms against business needs, confidentiality, integration, and preservation requirements, involving legal, technical, procurement, and user stakeholders. Define approved uses and permitted data, then pilot representative workflows and difficult examples to verify outputs and measure efficiency, including review and correction. Before broader deployment, confirm that relevant records and metadata can be preserved and exported, and document limitations and required controls.

3. Establish Governance

Establish AI governance by appointing an executive sponsor with the authority to secure resources, resolve competing priorities, and hold stakeholders accountable. Support that sponsor with a cross-functional operating group that includes legal, IT, information security, privacy, compliance, records management, and relevant business leaders. Give the group a clear operating framework that addresses:

Decision rights: Define who can approve tools and use cases, authorize access to sensitive information, accept identified risks, and suspend problematic workflows.

Risk categories: Classify uses according to information sensitivity, the consequences of incorrect outputs or unauthorized actions, and the human oversight required.

Exceptions: Document how exceptions are requested and approved, including conditions, expiration dates, and reassessment requirements.

Incident escalation: Establish reporting channels for suspected data exposure, unreliable results, and unauthorized AI use, with designated owners responsible for investigation and remediation.

Review frequency: Schedule regular reviews and additional assessments when new integrations, expanded uses, or significant incidents materially change the risk.

Resources and timetable: Provide a realistic budget, dedicated staff time, and measurable milestones to make governance an ongoing operational responsibility.

4. Develop Acceptable Use Policies

As discussed in the previous section about effective AI governance programs, AI policies need practical guidance that helps employees understand what they can and cannot do in daily work. Acceptable use guidelines should translate complex requirements into straightforward "dos" and "don'ts," with examples explaining which information employees may submit to approved AI tools and when additional review is required.

Present these guidelines in a clear, accessible format, such as a two-column list of permitted and prohibited actions. Employees should be able to find the guidance quickly, apply it to common tasks, and identify whom to contact when a situation is unclear.

5. Train Employees

Tailor AI training to employees' responsibilities and daily decisions. Custodians should recognize and properly handle AI-generated records under organizational policies and legal holds; reviewers should identify classification errors, omissions, and uncertain results; and lawyers should evaluate supporting authorities, factual accuracy, and AI-assisted conclusions. Administrators should know how to configure and verify access, retention, and other controls. Use realistic exercises involving incomplete information, misleading outputs, and escalation scenarios, refresh training as tools and workflows change, and provide clear contacts for questions and concerns.

6. Prepare for AI-Generated Evidence

Prepare for AI-generated evidence by testing the collection and review of conversations, summaries, audit records, source links, and document histories. Confirm that exports retain enough context to interpret the material, including available timestamps, user identities, and relationships between prompts, responses, and supporting documents.

Check whether linked sources and relevant document versions require separate preservation or collection and verify that the exported information remains usable in the review platform. Document missing fields, inaccessible records, and other limitations, along with practical workarounds and responsible owners. It's important to address these gaps before a live matter makes them urgent.

7. Update Legal Hold Procedures

Update legal hold procedures to address how employees use AI and where relevant records are stored. Add AI-specific questions to custodian interviews and clear instructions to hold notices, then coordinate preservation with IT, records management, and application owners. Address the following:

- **Tools and accounts**: Identify enterprise applications, embedded AI features, meeting assistants, and personal accounts used for relevant business activities.
- **Record categories**: Assess prompts, responses, uploaded files, summaries, audit records, and document revisions for inclusion within the hold's scope.
- **Connected information:** Determine whether linked sources, relevant document versions, and outputs copied into other systems require separate preservation.
- **Retention and deletion**: Identify settings or processes that could remove relevant records, including automatic deletion and account termination, and implement appropriate preservation measures.
- **Custodian instructions**: Explain which AI records are covered, what employees must preserve, and whom to contact about questions.
- **Verification**: Test whether holds preserve the intended data and document results, limitations, and corrective actions.

8. Validate AI Outputs

Define success before measuring AI performance, with clear quality criteria tied to the task and the consequences of errors. For document classification, evaluate both incorrect inclusions and missed responsive material, recognizing that a high overall accuracy rate can conceal important omissions. Where sampling supports statistical conclusions, document the population, selection method, estimates, and associated uncertainty so that results can be interpreted and defended.

Supplement representative sampling with a targeted review of consequential categories, such as potentially privileged communications or documents central to key issues. Establish thresholds for further review, correction, or escalation, and document how identified errors are resolved. Reassess performance after material changes to instructions, models, data, or workflows to confirm that earlier validation remains applicable.

9. Continuously Refine Governance

Review turnaround time, total cost, error patterns, rework, incidents, adoption, and preservation performance. Expand successful workflows and revise or retire those that fail to deliver. Give practitioners a regular opportunity to report what is working and where controls interfere with or improve the work.

A sensible starting sequence is to establish visibility and ownership, test a few useful workflows, and expand only after evaluating their results. Progress should be demonstrated through improved service and dependable execution.

Here's an illustration of ten practical recommendations to achieve AI preparedness:

10. Monitor Regulatory and Case Law Developments

Assign responsibility for tracking obligations in relevant jurisdictions, court-specific AI requirements, ethics guidance, and decisions affecting preservation or privilege. Assess each development against the organization's actual tools, workflows, and matters, distinguishing binding requirements from advisory guidance and emerging trends.

Translate applicable changes into updated policies, training, outside counsel instructions, and preservation or review procedures. Document who owns each update and confirm that affected personnel understand and implement the revised practices.

A Roadmap for Preparedness in the AI Era

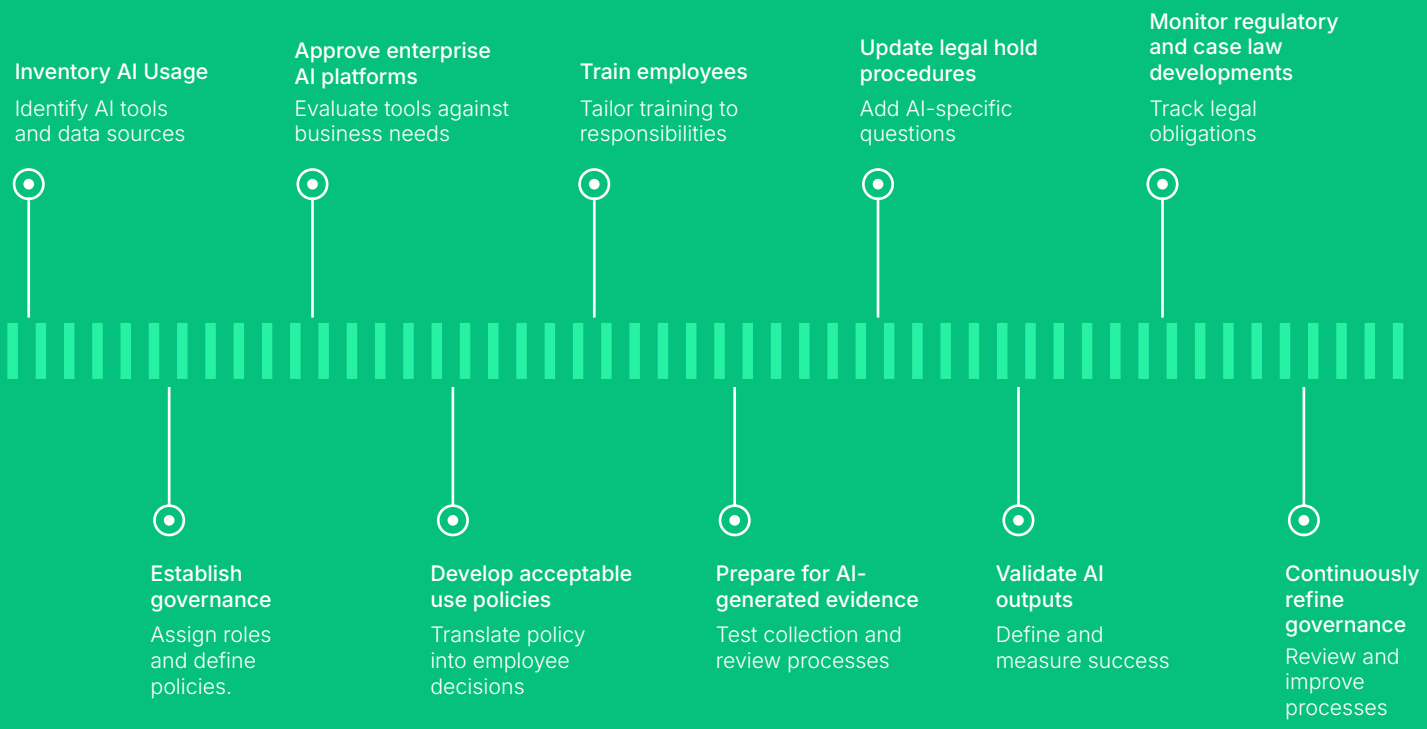

Conclusion

AI offers corporate legal departments and law firms an opportunity to increase capacity, understand information sooner, and devote more attention to consequential decisions. Achieving that value requires sustained professional involvement.

The Consilio report The AI Mandate in Legal emphasizes the importance of accompanying leadership direction with skills, workflow change, and governance. Its recommendations favor building lasting organizational capability alongside adoption.

AI can take over individual tasks and change staffing needs and service models. It does not eliminate the need for accountable professionals who understand legal obligations, evaluate evidence, advise clients, and exercise judgment.

Organizations that govern AI well will be better positioned to turn technical capability into a competitive advantage: reliable service, more informed decisions, stronger client relationships, and workflows they can explain.

The neighborhood will continue to change. New tools will appear, familiar applications will acquire new capabilities, and new forms of evidence will require attention. The sound judgment and discovery principles that have guided legal professionals for decades remain a foundation for navigating those changes.

New tools, same rules.
You've been down this road before.

© 2026 Consilio. All rights reserved.

Consilio