



Artificial Intelligence Toolkit for In-house Lawyers

AI Strategies to Support the Business and Elevate Your Career

Second Edition



WHAT IN-HOUSE LAWYERS SAY ABOUT THIS TOOLKIT:

"This second edition reflects the shift from understanding AI to operationalizing it. It equips in-house lawyers with the tools to move beyond issue-spotting and into true business leadership to help organizations adopt AI responsibly while managing real legal and operational risk."

Meganne Thaxton, *Principal Corporate Counsel, Red Ventures*

Thank you to ACC members

The ACC team warmly thanks the global group of ACC members and in-house legal professionals who shared their insight regarding this toolkit
(their input was personal and not on behalf of their organizations):

Aparna D. Williams

Chief Legal Officer and Corporate Secretary, Sophos

Douwe Groenevelt

Partner at Datacation, Data & AI for Professional Services

Malaika Roemer CIPP/US

Principal Paralegal and Compliance Program Manager, Legal & Compliance, Coalfire

Meganne Thaxton

Principal Corporate Counsel, Red Ventures

Schellie-Jayne (SJ) Price

Chief Legal Officer, Nooriam

Founder of ACC Australia's Legal Technology & Innovation Special Interest Group

ACC thanks Kilpatrick Townsend & Stockton LLP
for their generous sponsorship of this toolkit and
for the content they provided for this resource.



WELCOME

Artificial intelligence is no longer a future conversation—it is already reshaping how businesses operate and how legal departments deliver value. For in-house counsel, this moment represents both a challenge and an opportunity to lead.

ACC's latest research on generative AI shows just how quickly the landscape is changing. **GenAI adoption within legal departments more than doubled in the past year**, according to an ACC [global survey](#) of more than 650 in-house legal professionals across 30 countries. An overwhelming **91% of respondents say efficiency is the technology's greatest benefit**. Many legal teams are already using AI to streamline workflows and handle more high-value work internally.

This rapid adoption reflects a broader transformation in how legal departments operate. As businesses integrate AI into their strategies, legal teams are increasingly expected to help shape governance frameworks, manage risk, and guide responsible implementation. More than ever, organizations are looking to their legal teams not just for answers, but for leadership.

To support our members in this evolving environment, ACC developed the Artificial Intelligence Toolkit, including this updated Second Edition. This resource brings together practical guidance, frameworks, and insights to help in-house counsel understand emerging technologies, build effective AI governance, and support innovation on legal teams and across their organizations.

This toolkit reflects the strength of the ACC community. We are grateful to the many ACC members who contributed their expertise to this effort, and we extend our thanks to Kilpatrick Townsend & Stockton LLP for their continued partnership and thought leadership.

Artificial intelligence will continue to evolve—and so will the role of the in-house lawyer. At ACC, we are committed to equipping our members with the resources they need to navigate change, lead with confidence, and help their organizations harness the promise of AI responsibly.

We hope this toolkit serves as a practical guide as you move forward. As always, we welcome your feedback and ideas for future updates. Please share your thoughts with us at legalresources@acc.com.

Best regards,

Jason L. Brown

PRESIDENT AND CEO

ASSOCIATION OF CORPORATE COUNSEL



Kilpatrick is trusted counsel to the world's most innovative companies. To serve our groundbreaking clients, we are at the forefront of client service and on the cutting edge of law and technology.

Our multidisciplinary AI team partners with our clients to tackle their most pressing concerns and challenges while achieving the best results in line with overall business strategy and goals.

Learn more at ktslaw.com/generative-ai

Kilpatrick

© 2026 Kilpatrick Townsend & Stockton LLP
ktslaw.com | Attorney Advertising

CONTENTS

■ Introduction	7
Top of Mind: The Essentials	8
AI and In-house: Survey Data at a Glance	8
Four Immediate Steps for In-house Counsel	10
Key Data Insights for In-house Lawyers	10
■ On The Horizon: Ten AI Trends in 2026	11
■ Ten Checklists to Help You Succeed in Your AI Journey	15
Checklist 1: Seven AI Basics for In-house Lawyers	15
+ <i>Learn 15 Key AI Terms</i>	18
Checklist 2: 20 Tips to Be More AI Savvy and Boost Your Career	20
Checklist 3: Six Steps to Develop Governance and Compliance Strategies	23
+ <i>Generative AI Usage Policy Example</i>	27
+ <i>13 Principles of AI Regulatory Frameworks</i>	28
Checklist 4: Seven Steps to Integrate AI Governance Across the Business	31
+ <i>AI Maturity Roadmap – 2026 Edition</i>	34
Checklist 5: Incorporate AI into Legal Team's Work	35
NEW + <i>AI Use Cases for Legal Departments in 2026</i>	44
+ <i>ACC Member AI Use Cases: How In-house Lawyers Use AI</i>	45
NEW + <i>Track and Assess Outside Counsel's Use of GenAI</i>	46
+ <i>Top Seven Tips for Prompting Generative AI Tools</i>	48
Checklist 6: Meet Ethical Obligations While Using Generative AI	50
Checklist 7: Five Tips to Protect Intellectual Property When Using Generative AI	54
Checklist 8: Nine Steps to Evaluate AI Tools, Products, and Services	58
NEW + <i>Tips for Assessing EU AI Act High-Risk AI Systems</i>	63
+ <i>Evaluate These Four Key Risk Factors in AI Systems</i>	65
+ <i>Five Tips for Conducting an AI Impact Assessment</i>	68
Checklist 9: Top Eight Tips to Address AI Issues in Contracts with Third Parties	70
NEW Checklist 10: Five Tips to Develop Governance for Agentic AI	75
+ <i>Learn 10 Key Terms in Agentic AI Systems</i>	79
+ <i>Assess Agentic AI Capability and Oversight</i>	81
■ Additional Insights, Tools, and Resources	82

WHERE IN-HOUSE COUNSEL STAY AHEAD

*Built on knowledge.
Powered by connections.*

ACC ANNUAL MEETING

Continue the conversation
on AI—and so much more—
in person.
Explore how to apply insights
like these in a practical,
in-house context.




REAL-WORLD STRATEGIES
for AI, automation,
and legal tech

PRACTICAL GUIDANCE
on compliance,
governance, and risk

A POWERFUL NETWORK
of in-house peers

Be Part of the Conversation.

Visit acc.com/annualmeeting for details and dates.



Who is this toolkit for?

From **junior lawyers to Chief Legal Officers**, this toolkit is relevant for all in-house counsel who want to understand Artificial Intelligence (AI) and be AI-savvy leaders for their business.

- **AI presents major opportunities** – to automate tasks, make processes more efficient, analyze vast amounts of data, improve customer service, generate content, and more.
- **It also poses legal, ethical, and business challenges**, and requires fast adaptation and upskilling by business and legal teams.
- **A patchwork of laws and regulations** is unfolding around the world, as authorities seek to establish AI governance principles, foster innovation, and mitigate risks.
- **General counsel and their teams have a leading role** to play in helping the business take advantage of the AI wave in a way that optimizes benefits and mitigates risks.
- **This toolkit offers** practical checklists and insights to help you be an effective leader in your organization's AI preparedness.

Where is this insight from?

The insight in this toolkit is mainly based on:

- Resources from the **ACC Resource Library** that the ACC team has found relevant;
- Thought leadership from law firm **Kilpatrick Townsend & Stockton LLP**, and
- The perspectives of ACC members, experienced in-house counsel who kindly shared input.


The Association of Corporate Counsel

- The Association of Corporate Counsel (ACC) is a global bar association.
- ACC promotes the common professional and business interests of in-house counsel through information, education, networking opportunities, and advocacy initiatives.
- The ACC community includes more than 49,000 members from around the world.
- ACC is invested in your success as an in-house lawyer, a General Counsel, or a Chief Legal Officer. ACC advocates for you to have a “Seat at the Table” and provides value with resources for your team to succeed.
- Are you an in-house counsel? **Join us!** Learn, connect with peers, and boost your career.

Disclaimer: The information in this toolkit should not be construed as legal advice or legal opinion on specific facts. It should not be considered representative of the views of the Association of Corporate Counsel (ACC) or any of its lawyers, unless so stated. This toolkit is not intended as a definitive statement on the subject, but rather to serve as a resource providing practical information to the reader. AI is rapidly evolving. ACC makes no representations or warranties regarding the accuracy or completeness of this toolkit. This toolkit was developed mostly between December 2024 and March 2025, and the second edition was updated in March-April 2026.

AI and In-house Survey Data at a Glance


► Most in-house professionals are using or planning to use GenAI at work




► More than half use GenAI in their legal work



► Top benefits of using GenAI at work



► Legal teams not yet seeing cost savings from outside counsel's use of AI



Source: *Generative AI's Growing Strategic Value for Corporate Law Departments – Survey Results*,
Association of Corporate Counsel and Everlaw, October 14, 2025

AI and Attorney-client Privilege: Survey Insights

Risk perception, concern about AI's impact on privilege, and proposed actions to mitigate risks

56% believe the use of AI tools has the potential to compromise attorney-client privilege


27%

said that the risk of compromising privilege is dependent on the AI tool, its usage, or on the specific circumstance.


7%

of respondents globally say they have encountered specific instances where the use of AI has raised questions about the applicability of attorney-client privilege.

How concerned are you about the possibility of inadvertently disclosing privileged information when using AI-powered tools?



Measures currently taken to protect privileged information when using AI tools



Proposed Actions to Mitigate Risks to Attorney-client Privilege

- 80%** Create specific guidelines for AI use in legal practice
- 56%** Enhance data privacy and security standards
- 50%** Increase public awareness about the risk of AI and attorney-client privilege
- 49%** Develop laws and/or regulations for AI tool providers
- 49%** Use contract clauses or other means to establish liability for AI-related breaches

Source: *The Impact of AI on Attorney-Client Privilege*, Association of Corporate Counsel, September 11, 2024

Four Immediate Steps for In-house Counsel

1. **Inventory promising AI use cases** for your legal team and for the business.
2. **Identify top three high-risk AI deployments** that require immediate attention and governance.
3. **Implement an interim AI policy** to protect your organization and mitigate risk while you create a thoughtful, detailed AI governance policy.
4. **Launch an AI governance group** of key stakeholders across the business to begin the process of developing an AI strategy and policy for the organization.

Key Data Insights for In-house Lawyers

Learn how in-house lawyers view AI, through data from recent ACC surveys:

52% of in-house legal professionals are already using GenAI in their work.

From those, 91% report GenAI increased their task completion efficiency.

70% of respondents are concerned about non-lawyers relying on GenAI for legal guidance or decision making.

Over 35% of respondents say at least some of their outside counsel firms are using GenAI on their legal matters.

Nearly one in four CLOs list AI in the top three concerns for regulatory enforcement.

53% of CLOs list operational efficiency as their department's top strategic initiative for the year.

52% of CLOs say they plan to adopt new legal technology in their department to improve efficiency in the next year.

Over 60% of CLOs expect a similar or increased headcount in the legal department due to GenAI adoption, with legal roles evolving.

Learn more

- Gain more insights from the report on [*The on Role of Generative AI in Proving Corporate Law Department Value*](#) (November 2025, by ACC and Everlaw).
- Learn more about the growing importance of GenAI in Legal: [*Generative AI's Growing Strategic Value for Corporate Law Departments – Survey Results*](#) (October 2025, by ACC and Everlaw).
- Read the findings from the [*2026 ACC Chief Legal Officers Survey Report*](#) (January 2026).
- Read the findings from the ACC Foundation's [*2025 State of Cybersecurity Report: An In-house Perspective*](#).

Ten AI Trends in 2026

Generative AI (GenAI) is still new but hardly novel. Government regulation has begun to take effect, courts are no longer waiting for new developments, and this rapidly growing technology continues to present opportunities and challenges. Learn 10 key trends in the maturing AI landscape in 2026.

1. Defensible AI Procurement: Build a Litigation-Ready Record.

AI tool selection is no longer just a procurement or an IT decision. Plaintiffs, courts, and regulators are increasingly scrutinizing these decisions after an incident.

- **What record have you built** to demonstrate that your organization exercised reasonable care in selecting and deploying your AI tools **when – not if –** someone claims they were harmed?
- **Procurement Decisions:**
 - **Use a structured AI procurement review** that evaluates accuracy, bias risk, data sources, privacy exposure, and system integration before any tool is deployed.
 - **Require vendors to document their claims –** do not rely on marketing representations – and negotiate protective terms for data-use rights, audit rights, and liability allocation.
- **Pre- and Post-Deployment:**
 - **Keep a documented risk assessment and testing record** for intended use cases, known AI failure modes, mitigation steps, and human oversight controls.
 - **Maintain an auditable governance trail**, including approvals, testing results, oversight protocols, and aligned compliance controls.

2. Agentic AI and Autonomous Decision Making.

Semi-autonomous and fully autonomous AI agents introduce new opportunities and legal exposure.

- **Agents can initiate transactions, process payments, or trigger workflows**, raising issues of consent, authority, and attribution.
- **Liability risk emerges** when an agent acts outside its intended scope or causes financial or reputational harm.
- **Counsel must ensure guardrails**, human-override mechanisms, and clear allocation of responsibility in vendor contracts.

» WITH THANKS TO:

Kurtis G. Anderson,
Associate at
Kilpatrick,

for providing this
list of AI trends.

3. Use Enterprise RAG to Trade Opaque Training Risk for Governable Data Risk.

As businesses deploy internal AI tools, many are adopting retrieval-augmented generation (RAG) architectures. Doing so ensures that AI outputs are grounded in the organization's own data sets rather than uncertain AI model training data.

- **RAG architectures reduce reliance on public training data**, shifting your company's IP and privacy exposure from difficult-to-predict model behavior to internal data that organizations can actually oversee, audit, and control.
- Because RAG grounds outputs in the enterprise's own data, **it can improve accuracy and traceability. However, it also introduces new risks** if that data contains sensitive, privileged, outdated, or unlicensed materials.
- **Organizations must tightly govern RAG source repositories** to ensure they are permissioned, current, and auditable. They should also validate that RAG outputs continue to satisfy confidentiality, privilege, and data minimization obligations.

4. AI "Washing" Scrutiny.

Regulators are increasing scrutiny of exaggerated or misleading statements about AI capabilities, integration, and performance, using traditional consumer-protection and securities-disclosure theories.

- **This scrutiny is analogous to greenwashing concerns** but with heightened exposure in regulated or high-impact contexts due to safety, privacy, and consumer protection implications.
- **Companies should be able to substantiate AI-related claims** with retained support (testing results, limitations, and conditions of use).
- **Counsel should review** marketing, sales enablement, and investor-facing statements for accuracy, appropriate qualifiers, and consistency with contracts and disclosures.

5. Internal Use Risks and Privilege Protections.

GenAI platforms create new risks related to the waiver of the attorney-client privilege and attorney work product.

- **AI-generated internal documents may not qualify as privileged** or attorney work product, creating new risks.
- **GenAI platforms are not "lawyers,"** and, in general, no privilege attaches to AI-generated content.
- Employees using AI tools for drafting, summarizing, or analysis may **inadvertently create discoverable materials**.

6. Liability Under Existing Law (“Liability Now”).

Even as AI-specific regulations mature, traditional legal frameworks already apply to AI-driven harms, especially in the US, and courts are increasingly being asked to leverage them in novel AI contexts.

- **Product liability, negligence, employment discrimination, and privacy laws** remain enforceable regardless of AI novelty.
- **Plaintiffs are increasingly pleading AI harms as traditional negligence and consumer-protection failures**, including inadequate safeguards or warnings, insufficient oversight of delegated AI decisioning, and misleading or inadequate disclosures about AI use, capabilities, or data practices.
- **Counsel must map AI use cases to existing legal doctrines** and ensure risk mitigation measures are documented.

7. Mature Global and State AI Regulations.

The wave of 2024–2025 AI laws, including, among others, the [EU AI Act](#), [Vietnam AI Act](#), [Colorado AI Act](#), and various California AI laws, now enter enforcement, requiring operational compliance rather than planning.

- Companies must implement **risk classifications, documentation, and human oversight** requirements under new statutes.
- **State-level requirements in the US create a patchwork** that complicates nationwide deployment of AI tools.
- **Counsel must coordinate cross-functional compliance programs** and maintain jurisdiction-specific AI inventories regarding contract requirements, policies, claims substantiation, product disclosures, vendor governance, and escalation rules.

8. AI Supply Chain and Vendor Accountability.

Regulators expect companies to understand and manage the full lifecycle of AI components and data sources and not merely rely on vendors.

- **Vendor disclosures** on training data, model lineage, and evaluation methods are becoming mandatory.
- Regulators and contracting parties both increasingly require **audit rights, safety testing evidence, and shared liability frameworks**.
- **Counsel must build due diligence processes** that track AI dependencies and ensure compliance across the supply chain.

9. Explainability, Testing, and Safety Assurance.

High-risk AI systems face growing requirements for transparency, red team testing, and documented safety controls.

- **Explainability is becoming essential** for regulatory approval in employment, finance, and healthcare contexts.
- **Safety testing** must be repeatable, documented, and integrated into incident response frameworks.
- **Counsel must ensure** that technical teams maintain logs, testing records, and human review checkpoints.

10. Guardrails and Misuse: AI Safety Controls for Consumer-Facing Systems.

Governments, companies, and customers are increasingly focused on AI safeguards to prevent harm, deception, discrimination, and disinformation. Vendors are embedding safety controls, use restrictions, and policy enforcement mechanisms directly into foundation models.

- **As AI models reach broad and novel consumer use**, preventing AI-enabled fraud and misuse is becoming a growing risk area.
- **Expanding automation and consumer-facing use cases** present new oversight and guardrails questions, particularly with respect to vulnerable populations and minors.
- **Counsel should monitor** how AI vendors implement safety controls and misuse-prevention policies and ensure those safeguards align with the company's own compliance and risk-management practices.

Seven AI Basics for In-house Lawyers

Learn what artificial intelligence is (and is not), and get a grounding in the essential terminology.

1. **Artificial intelligence tools are not new.** You've likely been using them for years in your work, even if you did not know it.
 - Lawyers have been using AI tools to **assist with many tasks**, such as:
 - **Organizing and searching** through large volumes of documents;
 - **Identifying patterns** in legal documents and cases;
 - **Drafting** contracts and other documents; and
 - **Automating manual tasks** like approval workflows, data collection and entry, and scheduling.
 - For example, well-known **legal research platforms** (such as Westlaw® and LexisNexis®) have incorporated AI technology for several years to aid users in **searching and analyzing** large legal data sets like case law, statutes, and regulations.
2. **Hybrid AI models** combine multiple technologies, such as machine learning and **natural language processing (NLP)**, to create powerful tools like chatbots that can interact with users.
3. **Generative AI** is the newest category of AI, and is often what people now mean when they use the term "AI."
 - **GenAI** does not simply organize or analyze data. It is software that can **generate language, pictures, or even other computer programs**, typically from plain language text prompts.
 - This technology relies on **large language models (LLM)**, which analyze massive sets of human-created content to generate and classify text, answer questions in a conversational manner, and translate text from one language to another.
 - Essentially, LLMs use data analysis and NLP to **predict**, word by word, **what type of response a human would give** to a question or instruction. Because each new word is based on what the AI tool already generated, a single misstep by the AI can lead the entire output off track and cause the AI to produce an incorrect response.

CHECKLIST 1

« This checklist is based mainly on the following ACC resources and on insights from the sponsor Kilpatrick and from ACC members.

← [The AI-Empowered Counsel: A Practical Framework for In-house Legal Teams](#), by Hayden O. Creque, *Corporate Counsel Now*, October 18, 2025

← [Tech Toolbox: Implications of Generative AI for Law Departments, Part 1](#), by Greg Stern, *Corporate Counsel Now*, March 16, 2023

← [Tech Toolbox: Implications of Generative AI for Law Departments, Part 2](#), by Greg Stern, *Corporate Counsel Now*, June 8, 2023

← [Legal Tech: As AI Technology Advances so Must Our Legal System](#), by Olga V. Mack, *Corporate Counsel Now*, February 7, 2023

← Visit the [AI Center of Excellence for In-house Counsel](#).

4. GenAI does not “think” in the way that a human does.

However, newer AI models are being designed to engage in multi-step solutions to complex tasks and questions and even to analyze their own output for likely mistakes at each stage.

- **Human oversight of GenAI systems is essential.** These systems can enhance Legal’s productivity and outcomes, but they cannot replace human judgment. This is especially true for the new Agentic AI systems that are starting to handle more complex tasks and workflows with less human involvement.

5. GenAI models are only as good as the data they are trained on.

For a GenAI tool to provide quality output of text, images, or sounds, it must have a large volume of high-quality data to analyze.

- **“Unclean” data**, containing errors, contradictions, or irrelevant content, may cause a GenAI model to generate a response to the user’s prompt that looks correct but is inaccurate, sometimes wildly so. These **inaccurate** responses can often appear very convincing, so users must check the model’s output carefully.
- **“Hallucinations”** in AI outputs are usually caused by training data sets that are too small to ‘produce’ accurate information. A dataset can be fed with correct training data but may still make mistakes, especially if the set isn’t large enough or if its training hasn’t incorporated enough computing power.
- **Large public GenAI models** like ChatGPT use training data sets which may contain content scraped from the Internet, including copyrighted text, images, and audio. Even when they use large amounts of training data, GenAI models can yield inaccurate (or made-up) answers.

6. GenAI tools aimed at the legal industry often aim to be more precise and effective by using processes like “retrieval-augmented generation” (RAG, defined below) to enrich their AI models with additional legal information, such as case law and dockets, statutes and regulations, and regulatory filings.

7. Good practices for these legal-oriented AI tools:

- **Having the tools “sandboxed,”** i.e., tested in a secure environment isolated from the rest of your systems and data.
- **Having the tools “gated” or “walled off,”** i.e., operated as an instance that is either hosted solely within your organization’s systems or, if hosted in the vendor’s environment, an instance that is dedicated to you and not shared or comingled with other users’ data and instances (unlike large public models), to **ensure that your users’ private data is not publicly released** to the wider Internet (or shared with other users or with the vendor itself) when they interact with the tool.

» Also check out:

- ***Basics for Corporate Counsel to Consider About Generative AI***, by Meghan K. Farmer, Jon Neiditz, and John M. Brigagliano, Kilpatrick, August 24, 2023
- ***Start Small, but Start Somewhere: How Corporate Legal Departments Can Start Their AI Journey***, by Ethan M. Simon, *Corporate Counsel Now*, January 5, 2026

TAKE ACTION: Write one or two steps that you plan to implement.

Describe	By what date?
1.	
2.	
Notes:	

» WITH THANKS TO:

Kurtis G. Anderson,
Associate at *Kilpatrick,*

for his insights
in this glossary.



Learn 15 Key AI Terms

1. **Artificial Intelligence (“AI”).** A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.
2. **Machine Learning.** A set of techniques that can be used to train AI algorithms to improve performance at a task based on data.
3. **Large Language Model (“LLM”).** A language model using deep-learning algorithms and trained on a large data set to capture patterns and regularities present in natural language and make assumptions on previously unseen language fragments to perform a task or generate an output.
4. **Algorithmic AI.** Also known as traditional AI. Artificial intelligence that classifies data based on various attributes and then predicts outcomes bound by predetermined rules. It excels at pattern recognition. Examples include search engine results, auto-text generators, and digital content recommendations.
5. **Generative AI (“GenAI”).** A class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content. Examples of GenAI include OpenAI’s ChatGPT, Microsoft’s CoPilot, and Google’s Gemini.
6. **AI Agent.** An AI system that relies on a combination of machine learning, natural language processing, and automation technologies to interact with the world and act autonomously to achieve user goals without the need for constant human guidance. In contrast to GenAI, they are task followers, not predictive content generators.
7. **Input.** Data provided to or directly acquired by an AI system on the basis of which the system produces an output. Inputs are how most users will initially interact with GenAI to generate content.

8. **Output.** The data transmitted by the AI system in text, image, audio, or video format in response to an input. This is the product most end users are currently looking for from an AI program.
9. **Retrieval-Augmented Generation (“RAG”).** A method to enhance the accuracy of GenAI models. The RAG system first searches and retrieves relevant external information, such as from a legal-information database (or from the company’s internal database), and then adds this information to the user’s prompt. This provides more context that allows the GenAI tool to generate a more accurate and relevant response.
10. **Prompt.** The input provided by a user to interact with a GenAI model and generate a response, or output, from the model. Prompts normally include specific instructions for the model to follow or questions for it to answer. Skillful prompt engineering is key to successfully utilizing GenAI. **See [Top Seven Tips for Prompting Generative AI Tools](#).**
11. **Token.** These are the smallest units of data processed by an AI model, often words, characters, or sub-parts of words and punctuation.
12. **Token Limit.** The number of tokens the AI model is able to process between both the input prompt and the total generated output.
13. **Hallucination.** Incorrect, misleading, or entirely fabricated information generated by a GenAI model that appears to be true. This is a common caution for GenAI use: The risk of hallucinations must be managed.
14. **Deepfake.** AI-generated or manipulated image, audio or video content that resembles existing persons, objects, places, entities or events and would falsely appear to be authentic or truthful. While not all deepfakes are generated to mislead, they often enable fraud or misinformation.
15. **Bias.** An effect of training or preset instructions that results in an unrepresentative or distorted result, favoring certain outcomes and results. Unknowns around the training data and bias can lead to unintended or adverse outcomes.

CHECKLIST 2

» This checklist is mainly based on the following resources and on insight from ACC members:

- *The Human Edge: Lawyering Beyond AI*, by Laura Sulymosi, *Corporate Counsel Now*, December 15, 2025
- *Legal Tech: The Crucial Role of Legal Stewardship During the AI Revolution*, by Olga V. Mack, Kevin Keller, Kristina Podnar, *Corporate Counsel Now*, July 29, 2024
- *Prompt Engineering: Your Hidden Superpower for AI Mastery*, by Spiwe L. Jefferson, *Corporate Counsel Now*, July 23, 2024.
- *Legal Tech: AI, Legislation, and Your Career Story*, by Olga V. Mack, *Corporate Counsel Now*, February 7, 2024
- Visit the *AI Center of Excellence for In-house Counsel*.

20 Tips to Be More AI Savvy and Boost Your Career

The rise of AI presents opportunities for in-house lawyers to demonstrate leadership and advance their career. Sufficient familiarity with AI is also an increasing necessity for lawyers to serve their clients with competence.

1. In-house lawyers should become familiar with key AI concepts and their practical implications.

- This will allow you to better understand how new developments in the field of AI will impact your organization and its environment.
- Your business colleagues likely use or want to use AI tools. You need a minimum level of knowledge in order to be effective in helping them navigate these uses.
- Learn where you as a lawyer will continue to fit in with human abilities that AI cannot replace, such as people skills, emotional intelligence, judgment, and an understanding of business context.
- Also satisfy your ethical obligations to develop and maintain sufficient knowledge about new technology, for example under your duty of competence. Learn more at [Checklist 6](#).

2. Find opportunities to increase your knowledge.

- **Tip 1: Start somewhere.** Don't let the volume of AI-related news paralyze you. Explore key AI concepts such as those laid out in [Checklist 1](#) of this toolkit.
- **Tip 2: Gain industry-specific knowledge.** Read publications (if any) that are focused on AI-related developments in **your industry sector**.
- **Tip 3: Sign up for AI-related newsletters** relevant to **your practice**. For example, tailor your [ACC Newsstand](#) feed to include AI-focused legal updates.
- **Tip 4: Ask your outside counsel** if they can send you informative alerts at no cost regarding AI developments. They will probably be happy to showcase their knowledge.

- **Tip 5: Watch or attend educational programs** created for in-house lawyers on AI, such as on the [ACC Online Education](#) platform, or at [in-person conferences](#).
- **Tip 6: Join AI-focused discussion groups with peers, such as** via the [ACC IT Privacy and eCommerce Network](#).
- **Tip 7: Ask business colleagues** what tools they find the most useful, and what strategies they find the most effective when using these tools. Learn about those tools.
- **Tip 8: Use AI to educate yourself.** GenAI itself can help you learn more about how the technology works and how you can use it effectively in your practice.
- **Tip 9: Experiment with AI tools.** To truly become familiar with AI technology, you can't stick to mere theoretical knowledge. Experiment with AI tools, in a manner that is safe for your organization and that conforms with your ethical duties and professional obligations - e.g., don't compromise confidential or privileged client information when using AI tools.
- **Learn how to create effective prompts when you use generative AI tools (GenAI):**
 - **Tip 10: Give the tool context** for your query. For example, in your prompt, let the tool know what role/persona you want it to play, and which audience the output is intended for. Consider including background information (such as your organization's published mission statement).
 - **Tip 11: Also consider mentioning the tone and language complexity** that you want the tool to use in its output, depending on your audience. Are you planning to use the output for a presentation to fellow lawyers who are familiar with legal concepts, or to a business audience that expects simplified, plain-English language?
 - **Tip 12: Use follow-up prompts** to obtain more relevant results or output formats.
 - **Tip 13: Keep in mind that** you can't rely on the output from GenAI. You are responsible for verifying the accuracy of such output. Interactions with (and output from) GenAI tools feel human-like, but are typically based on a predictive algorithm.

» Also check out:

→ [AI Can Read the Law, But Can It Read the Room?](#), by Andrew Ting and Julie Xu, *Corporate Counsel Now*, July 8, 2025

→ [Career Advancement Toolkit for In-house Lawyers](#), an ACC Toolkit

→ Also watch: [Lex Machina- The Rise of Artificial Intelligence](#), an ACC webcast program with Melanie Laffin, Chief Innovation Architect, Artificial Intelligence, Booz Allen Hamilton, June 18, 2024



Don't just use AI for your work—use it to make yourself smarter. I asked ChatGPT to design a complete data science curriculum for me, with the goal of reaching a second-year bachelor's level of understanding. Then I asked it to break that curriculum into daily five-minute lessons—a short bit of theory and a quick quiz—and schedule them to arrive every morning using ChatGPT's 'Tasks' feature. It's become part of my morning routine, like checking the news.

Douwe Groenevelt
*Partner at Datacation,
 Data & AI for
 Professional Services*



3. Demonstrate leadership through AI-related projects.

- **Tip 14: If your company has a dedicated cross-departmental group** focused on AI, ask if you can get involved or invited to its meetings, and demonstrate a business mindset through your contributions to the discussion. If there is no such group, consider starting one.
- **Tip 15: Your legal department is probably already supporting or leading** your organization's AI governance or AI projects. If you're not already involved, identify relevant ways you can get involved.
- **Tip 16: Consider how AI intersects with practice areas in your existing portfolio**, and what you can do to further advance the company's goals or your department's goals by using AI. Identify how you can facilitate the use of AI and mitigate risks through your practice.
- **Tip 17: Help the business identify ways to measure outcomes** resulting from the use of AI tools.
- **Tip 18: Adopt a business mindset when approaching AI-related issues.** AI presents risks and challenges. When offering guidance to the business, show you balance risks and positive outcomes. Make sure your guidance is informed by the company's strategic goals.
- **Tip 19: Identify effective ways to bring your perspective** to the executive team or to the head of the legal department regarding how new developments in the AI landscape present new challenges and opportunities.
- **Tip 20: Show results.** Keep track of how your use (or your team's use) of AI tools yields measurable positive outcomes or improvements. Showcase results internally.

Six Steps to Develop Governance and Compliance Strategies

Learn six tips to develop a robust AI governance strategy for your company. Your business colleagues are probably using AI. They will experiment with use cases and will want to use new tools.

1. Audit your organization's current uses of AI.

An audit evaluates AI systems, algorithms, and data practices to identify risks, such as biases, privacy concerns, and ethical issues.

- It's likely that people within your business are already using generative AI tools for research or document drafting. If not, they may be experimenting with AI tools to evaluate their usefulness.
- **Assemble a cross-departmental team** to assess the current and potential uses of AI. Include these key stakeholders and any others you deem relevant:
 - Information Security
 - Information Technology (IT)
 - Developers and product teams
 - Human Resources (HR)
 - Finance
 - Investor relations
 - Business generators
 - Risk management
- **Engage people at all levels** to find out what is actually happening with AI and what work they intend to do with it. Conducting a **business-wide survey** may be the most practical way to get this information.
 - This survey may help you find patterns of usage and help determine which AI tools would be best to use across multiple departments.
- **Consider an immediate pause on the riskiest AI uses** while conducting this audit. Understand and communicate that "no AI usage ever" is not the end goal of this process.

CHECKLIST 3

« This checklist is based mainly on the following ACC resources and on insights from the sponsor Kilpatrick and ACC members.

- ← **Governing AI Inside the Enterprise: A Practical Framework for In-house Counsel**, by Alexander Lima, *Corporate Counsel Now*, February 25, 2026
- ← **Artificial Intelligence Laws (EU and US)**, by Meghan K. Farmer and Gregory P. Silberman, *Kilpatrick*, March 5, 2025
- ← **It's Go-Time – Creating an AI Governance Program**, by Adam Shedd and Mark Diamond, *Corporate Counsel Now*, April 1, 2024
- ← **No, Every Company Doesn't Need an AI Policy**, by Christopher Wlach, *Corporate Counsel Now*, April 1, 2024
- ← **Tips for Developing a Corporate AI Policy**, by Elizabeth Matecki, August 18, 2023
- ← **Employees Are Using Chatbots – Make Sure Your AI Use Policy Works**, by Robert Falk, Margo Lynn Hablutzel, Heather Peck, and Jonathan Yellin, *Corporate Counsel Now*, October 9, 2023

» Also check out:

- **Generative AI Risks: Legal and Compliance Insights - Part 1**, by Jeffrey Hunter and Colin Flood, Senior Managing Directors, Ankura Consulting Group, LLC, June 11, 2025
- **Generative AI Risks: Legal and Compliance Insights - Part 2**, by Jeffrey Hunter and Colin Flood, Senior Managing Directors, Ankura Consulting Group, LLC, June 11, 2025
- **5 Key Takeaways PTLF – Building an AI Governance Program**, by Amanda Witt, Ami Rodrigues, and Christina McCoy, Kilpatrick, March 7, 2025
- **AI: Virtual Patchwork**, by Stephen M. Anstey and Michael J. Breslin, Kilpatrick, November 8, 2024
- **Sample Use of Generative AI Policy Template**

2. Map out your regulatory obligations relating to AI.

- Jurisdictions worldwide are increasing their regulatory oversight of AI development and usage. Staying up to date is vital. Here are some useful resources that can help with this:
 - **ACC AI Center of Excellence for In-house Counsel.**
 - ACC newsletters such as *ACC Newsstand* and *In Brief*.
 - Legal research and news platforms.
 - Law firm blogs and alerts, like Kilpatrick's **Artificial Intelligence and Crypto blog**.
 - Relevant government alerts/guidelines, such as:
 - **National Institute of Standards and Technology** – NIST (US)
 - **Cybersecurity and Infrastructure Security Agency** – CISA (US)
 - **Regulators' Strategic Approaches to AI** (UK)
 - **European Parliament's page on AI** (EU)
 - **Office of the Australian Information Commissioner**
 - **Cyberspace Administration of China – Policies & Regulations**
 - **LinkedIn's Legal Generative AI Global Community**
 - **The Future of Privacy Forum**
 - The **AI Governance Center** of the International Association of Privacy Professionals (IAPP)
- Learn about key AI regulatory regimes, such as the EU Artificial Intelligence Act, and AI-related laws in California, Colorado and other US states. **Check out an overview of some of the key regulatory regimes: View the list in the ACC Resource Library.**
- **Develop a chart breaking down the jurisdictions** where your organization conducts business or has employees, then list regulatory obligations for AI users in each jurisdiction.
- It's possible that there will be **varying or even conflicting compliance obligations** across jurisdictions. An approach may be to choose the most stringent obligations as a company-wide compliance policy.

3. Develop business-friendly AI risk maps and governance structures.

- Map out AI **benefits, risks, and mitigation strategies**. Assemble a team to examine the organization's risk and needs. This could be the same group that conducts the AI audit discussed above. Including a broader group should make governance efforts more likely to succeed.
- Educate **top management and the board of directors** on the importance of AI governance and the specific risks if AI governance policies are not implemented.
- Put in place a formal **AI governance system**. An **AI governance policy** sets out the transparent and ethical use of AI. It also details how AI should be used by the organization. Having a sound AI governance policy helps to demonstrate to regulators and others that AI is being used responsibly.
- **AI Risk Management Tips:**
 - AI risks should not be considered in isolation.
 - AI-related risks should be integrated into the organization's larger risk management strategy.
 - Use a tiered risk approach; not all AI tools pose the same risk level.
 - Consider the tradeoffs between AI use for business operations versus risks such as privacy and cybersecurity concerns.
 - Consider risks to the company's intellectual property and confidential information, and risks of bias and discrimination.



It's important to deputize a team of AI ambassadors or have your Vice President level be the responsible layer for AI governance in the organization.

Aparna D. Williams
Chief Legal Officer and
Corporate Secretary
Sophos



4. Establish clear roles and responsibilities for the AI management team.

- Designate an **AI governance lead** or team. Consider appointing an **AI Steward** from the legal department to be the custodian of AI's ethical and strategic deployment.
- **Define roles** in AI governance for data scientists, engineers, legal, HR, and compliance.
- Clarify **accountability** for AI-related decisions. This isn't necessarily a single person, but more likely a designated leader from each team that incorporates AI into its workflows.
- **Define situations where it is necessary to involve senior management or the board:** regulatory inquiries, systemic errors, global compliance concerns, or whenever the AI system's usage conflicts with stated company values.



Using public AI tools to stay up to date can be a real life hack. For example, you can use the “Tasks” module of ChatGPT to define a daily task where you ask ChatGPT to search the web, including sites such as LinkedIn, for any news about AI regulations and create a curated overview for you each morning. [...] There are more tools that can do this, and it is not difficult to build your own.

Douwe Groenevelt
Partner at Datacation, Data & AI for Professional Services



5. Update non-AI policies to take AI into account.

- Many issues raised by GenAI **aren't new**. Businesses already deal with problems involving data security, confidentiality, bias, and intellectual property protection.
- If your organization already has written policies on these topics, **consider amending them to cover AI-specific issues** before you create a standalone AI policy to address them.

6. Create user-friendly AI policies for the entire organization.

- Your organization's AI policy should be **clear and easy** for employees to understand and follow.
- The policy should be **transparent** and explain the risks of AI misuse. An important caveat is that an organization that uses AI may not fully understand the risks (e.g. what's in the algorithms and training data).
- Your policy should also be **realistic** and reflect your company's culture and practice. The goal is to encourage responsible AI use, not to erect excessive barriers.
- Be clear that AI use requires employee **monitoring, training, and awareness**.
- Make clear that **non-compliance** may result in **disciplinary action**, which may include termination.
- Define whom the policy applies to. This will typically be **all employees, contractors, and other third parties** that use AI and/or its data on behalf of the organization.
- Set out the **terms of acceptable AI use** within the organization.
 - Explain which applications can be used, including their **permissible and prohibited use cases**.
 - Be clear that **security and confidentiality** are key considerations for employees, including **password-protected** AI accounts or VPN systems.



Generative AI Usage Policy Example


The following is an example of a simple, high-level Generative AI Usage Policy governing employees' use of commonly available GenAI tools. Customize to your own requirements. Please note this is only a sample policy, not a model policy, and it is provided below for informative purposes only (and not as legal advice).

- **Responsible Use:** Our organization encourages responsible use of AI in your work.
- **Verify Everything:** Don't trust the accuracy of the AI output. Check everything with reliable sources. You are responsible for the accuracy and completeness of your work. Beware of overreliance.
- **Don't include personal, sensitive, proprietary or confidential (or privileged) information** in your prompts to the AI (unless your organization has white-listed AI and clearly indicated that you can use it for such information.)
- **Good Judgment:** Generative AI use is not appropriate for all circumstances. Use good judgment and common sense about when to use it.
- **Acknowledgement:** Where AI-generated content is used in your work, **acknowledge** the use of AI as appropriate or as required by your organization.
- **Regular Review:** This policy will be reviewed regularly and will be updated to reflect our organization's experiences using AI, developing practice, and new technology.

» **WITH THANKS TO:**
Schellie-Jayne (SJ) Price, Chief Legal Officer, Nooriam, and Founder of ACC Australia's Legal Technology & Innovation Special Interest Group

for contributing this sample policy.

Visit the
ACC Resource Library
 to find more sample AI policies.



» WITH THANKS TO:

Meghan K. Farmer,
Partner, and
Gregory P. Silberman,
Counsel, at Kilpatrick,

for contributing
this list.



13 Principles of AI Regulatory Frameworks

While AI regulations differ, the principles below commonly appear across AI regulatory regimes and governance frameworks, together with practical implications for businesses:

1. Risk Classification and Role Mapping.

IMPLICATIONS:

- Classify AI use cases by risk, sensitivity, and business impact.
- Identify whether your organization is acting as developer, deployer, provider, customer, or integrator.
- Apply stronger controls to higher-risk or more autonomous use cases.

2. Transparency, Notice, and Provenance.

IMPLICATIONS:

- Disclose AI use where required or where doing so is necessary to avoid deception.
- Inform users when they are interacting with AI, where applicable.
- Consider labeling, provenance, or similar measures for AI-generated or AI-manipulated content.

3. Privacy and Data Governance.

IMPLICATIONS:

- Limit collection and use of data to what is needed for the use case.
- Govern prompts, outputs, logs, retained context, and connected data sources.
- Apply appropriate controls for personal, confidential, privileged, or regulated information.

4. Fairness, Non-Discrimination, and Use Restrictions.

IMPLICATIONS:

- Assess whether the system could create discriminatory or unfair outcomes.
- Test for bias in context, especially in employment, consumer, eligibility, pricing, or access decisions.
- Restrict or prohibit use cases that create unacceptable or prohibited risks.

5. Accountability, Documentation, and Governance.

IMPLICATIONS:

- Assign clear ownership for each AI use case.
- Maintain records of system purpose, design choices, limitations, approvals, and key decisions.
- Establish governance structures proportionate to the system's risk and reach.

6. Accuracy, Reliability, and Validation.

IMPLICATIONS:

- Test AI systems before deployment and periodically thereafter.
- Validate performance in the actual business context, not only in demos or pilot environments.
- Communicate material limitations and define fallback processes.

7. Safety, Security, and Misuse Prevention.

IMPLICATIONS:

- Protect systems against misuse, prompt manipulation, insecure integrations, and unauthorized access.
- Limit permissions and tool access based on least-privilege principles.
- Monitor for emerging vulnerabilities and misuse scenarios.

8. Human Oversight, Escalation, and Contestability.

IMPLICATIONS:

- Determine where human review, approval, or intervention is required.
- Avoid fully automated handling of high-risk or consequential matters without appropriate controls.
- Provide escalation paths and, where appropriate, a way to challenge or review AI-influenced outcomes.

9. Vendor, Procurement, and Contractual Controls.

IMPLICATIONS:

- Evaluate third-party providers for security, performance, transparency, and governance maturity.
- Address data use, confidentiality, subprocessors, audit rights, model updates, and incident support in contracts.
- Allocate responsibility clearly among the business, vendors, and implementation partners.

10. Monitoring, Incident Response, and Change Management.

IMPLICATIONS:

- Monitor deployed systems for drift, failures, misuse, and unexpected outputs or actions.
- Preserve relevant logs, inputs, outputs, and decision records for investigation and remediation.
- Reassess the use case when the model, tools, permissions, workflow, or business purpose changes.

11. IP, Publicity, and Content Rights.

IMPLICATIONS:

- Assess risks relating to copyright, trademark, trade secrets, and training or output provenance.
- Address ownership and use rights for AI-generated outputs in contracts.
- Consider voice, likeness, digital replica, and synthetic-content misuse risks where relevant.

12. AI Literacy and Internal Training.

IMPLICATIONS:

- Train personnel on how the relevant AI systems work, what they can and cannot do, and how to use them safely.
- Tailor training to role, use case, and risk level.
- Refresh training as systems, rules, and use cases evolve.

13. Truthfulness, Claims Substantiation, and Consumer Protection.

IMPLICATIONS:

- Substantiate claims about AI system performance, accuracy, reliability, safety, efficiency, or bias mitigation.
- Avoid misleading statements about whether outputs or interactions are AI-generated, AI-assisted, or human-generated.
- Ensure marketing, customer communications, and sales materials accurately describe the system's capabilities, limitations, and level of human involvement.

Seven Steps to Integrate AI Governance Across the Business

✦ This checklist was developed with the assistance of Google's Gemini AI tool.

Once your organization has developed an AI governance strategy and the necessary AI policies, it's time to implement them across the business.

1. Incorporate your AI strategy, governance, and compliance into your organization's business processes, services, and products.

- **Map AI use cases to specific business functions.**
Work with leaders across the business to make sure they apply the overall governance strategy as they implement AI into their workflows and business processes.
- **Embed governance checks into existing workflows** (e.g., product development, marketing, HR, procurement).
- **Monitor and reevaluate AI tools embedded in your organization's products and services.** Changes in both AI technology and law make product compliance an ongoing issue and not just a one-time process during development.
- **Ensure AI development aligns with the overall business strategy.** Many employees and leaders may be excited to try an AI tool as the "big new thing," but these tools should be analyzed based on how they serve business needs.

2. Prioritize explainability and transparency.

- **Strive for understandable AI decision-making processes.**
Document who is responsible for making AI-related decisions and how those decisions are made. Employees at all levels should be clear on whom to contact if they have questions or concerns about AI usage.
- **Document** how AI systems work and their limitations.
- **Communicate transparently about AI use to stakeholders.**
This can include both warnings about misuse and encouragement to make full use of AI tools within the bounds of AI policies.
- **Ensure defensible documentation** of AI procurement and usage decisions, including why an AI tool was adopted, what risks were identified, how those risks were mitigated, and who approved the deployment.

CHECKLIST 4

« This checklist is based mainly on the following ACC resources and on insights from the sponsor Kilpatrick and ACC members.

← **Governing AI Inside the Enterprise: A Practical Framework for In-house Counsel**, by Alexander Lima, *Corporate Counsel Now*, February 25, 2026

← **It's Go-Time – Creating an AI Governance Program**, by Adam Shedd and Mark Diamond, *Corporate Counsel Now*, April 1, 2024

← **Legal Tech: Integrate AI Compliance by Asking These Questions**, by Olga V. Mack, *Corporate Counsel Now*, April 10, 2023

← **Legal Tech: 5 First Steps Toward Responsible AI Oversight**, by Olga V. Mack, *Corporate Counsel Now*, September 12, 2023

← **Using AI in HR Decisions: Tips for In-house Counsel**, ACC Checklist, March 22, 2023

» Also check out:

- **Why Legal Should Be at the Center of Your Organization's AI Strategy**, by Michael T. Moore and Rebecca Hinds, *Corporate Counsel Now*, January 21, 2026
- **The Transformative Role of Artificial Intelligence (AI) in Cybersecurity**, by Ankura Consulting Group, LLC, July 17, 2025
- **5 Key Takeaways – Managing the Risks in Using Generative AI or How I Learned to Stop Worrying and Love AI**, by Steve R. Borgman and Jordan P. Glassman, *Kilpatrick*, February 20, 2025

3. Train employees to use AI responsibly.

- **Consider whether it would be helpful to partner with outside vendors for training.** This may be especially useful for issues involving cybersecurity and data privacy, which could be folded into existing cybersecurity training.
- **Educate everyone in the business** on AI risks, ethics, and compliance requirements. Employees may be aware of AI's potential but less aware of the business risks – or vice versa.
- **Provide practical training for all employees** on AI tools and responsible use. This will foster effective use of AI to improve productivity and ensure that AI is used safely.
- **Tailor this training to specific roles and departments.** AI use cases may vary across the business, so employees should receive training relevant to their work.

4. Implement robust data governance.

- **"Garbage in, garbage out!"** To use AI effectively, **you need quality data.**
- **Establish data quality standards and validation processes.** AI tools will not produce useful outcomes without clean, consistent business data to analyze. Relevant data sets should be validated and cleaned up before AI systems are installed or trained.
- **Implement data security measures and access controls.**
- **Ensure data privacy compliance**, e.g., as applicable, under the EU's General Data Protection Regulation (GDPR), or the California Consumer Protection Act (CCPA). Your organization should already have privacy compliance protocols in place for its business data. AI usage considerations can likely be added into these protocols.

5. Conduct regular technology audits.

- **Evaluate AI systems for bias, fairness, and accuracy.**
- **Consider engaging an independent auditor** to conduct your AI audits, as they will have the latest expertise on AI technology and may be able to suggest improvements to your AI governance practices.
- **Assess compliance with AI policies and regulations.** Government oversight of AI development and usage is expanding globally, making it necessary to review your systems' compliance regularly.
- **Document audit findings** and implement **corrective actions.**

6. Establish a feedback and monitoring system across the business.

- **Create channels** for employees to report AI-related concerns.
- **Monitor AI system performance** and identify potential issues.
- **Regularly review and update** the AI governance program.

7. Ensure that generative AI use in HR is monitored and augmented by human intelligence.

- **Anti-discrimination and other employment laws** apply to AI-enabled decisions as they do to human-made decisions.
- **AI is just a tool to enhance HR processes**, not a substitution for human judgment.
- **Consider implementing a full validation procedure** to ensure the tool meets business needs without introducing bias.
- **Keep human judgment in the loop** and ensure there is a mechanism to audit and override AI decisions as needed.
- **Ensure employees and job applicants are informed when AI tools are in use.** Consider asking them to confirm in writing they understand. Seek a signed consent/release from employees.
- **Consider privacy aspects** and how to set up the tool to mitigate privacy risks. For example, using AI tools could potentially expose employees' personal data to others, and AI systems should be designed to prevent this.



A big issue that we and many of our international peers talk about is the fact that AI tools can create risks when applied in combination with internal databases. When these databases suffer from suboptimal rights management, this can be hugely exacerbated by AI tools and can lead to people finding restricted information through internal tools.


Douwe Groenevelt
Partner at Datacation,
Data & AI for
Professional Services



» WITH THANKS TO:

*Joel D. Bush, Partner, and
Gregory P. Silberman,
Counsel, at Kilpatrick,*

for their input in this roadmap



AI Maturity Roadmap – 2026 Edition

Use this roadmap to assess current maturity and identify the next practical steps for scaling AI governance.

AI Roadmap	Ad-hoc (initial)	Defined	Integrated	Optimized	Transformative
Governance & Accountability	No clear owner, intake process, or approval path for AI use cases.	A basic owner and intake process exist, but application is inconsistent.	Cross-functional governance is documented, with defined roles, approvals, and escalation paths.	Governance is embedded into procurement, product, HR, privacy, security, and legal workflows.	Governance is strategic, measurable, and resilient enough to shape enterprise and industry practice.
Use-Case Inventory & Risk Tiering	AI uses are not inventoried, and systems are not meaningfully differentiated by risk.	The organization has a basic inventory and can generally identify high-risk uses.	Use cases are inventoried and tiered by impact, autonomy, data sensitivity, and jurisdiction.	The inventory is actively maintained and updated for new use cases, model changes, and connected tools.	The portfolio is managed dynamically, with clear triggers for enhancement, retirement, or escalation.
Policies, Training, & AI Literacy	Policies are limited, fragmented, or informal, and training is ad hoc.	Baseline AI guidance exists, and some teams receive introductory training.	Policies are standardized, and role-based training covers prohibited uses, disclosures, and human oversight.	AI literacy is refreshed regularly and tailored to function, risk, and the use of agentic or customer-facing systems.	AI literacy is embedded into onboarding, management expectations, and product or operational change processes.
Data, Model, & Vendor Controls	Little diligence is performed on data sources, model lineage, wrappers, or vendor dependencies.	Basic reviews address privacy, security, and core contract terms.	Inputs, outputs, retention, vendor dependencies, and model updates are documented and governed.	Structured controls address testing, provenance, labeling, third-party dependencies, and contingency planning.	The organization maintains a mature control framework for data, models, tools, agents, and vendors across the lifecycle.
Monitoring, Incident Response, & Change Management	Issues are handled case by case, with limited logging and little formal review after launch.	Basic monitoring and escalation exist, but updates and incidents are handled mostly manually.	KPIs, KRIs, incident playbooks, and review of material updates are defined and used in practice.	Continuous monitoring addresses drift, misuse, security events, synthetic content issues, and cross-system actions.	Monitoring and change controls are predictive, evidence-based, and continuously improve deployment decisions.

Incorporate AI into the Legal Team's Work

1. Treat the adoption of AI within Legal as a structured operational initiative. Recurring bottlenecks – such as NDA review, intake triage, first-pass contract review, or locating prior advice – can become focused AI pilots with a clear owner, limited rollout group, and defined success metrics.

- **Define clear goals** for how AI will improve legal operations, such as reducing contract cycle times, improving knowledge access, or reducing time spent on repetitive tasks.
- **Assign responsibility** for AI experimentation and implementation within the legal department, such as through legal operations personnel or a designated AI champion.
- **Establish timelines, testing procedures, and success metrics** for AI initiatives.
- **Train legal department team members** with hands-on learning based on real work scenarios. Training should demonstrate what is possible with AI in Legal's day-to-day work.
- **Monitor adoption and gather feedback** from attorneys and staff to refine the tools and workflows.
- **Track measurable outcomes** such as cycle time reduction, turnaround times, and attorney workload distribution.
- **Share successful use cases internally** to demonstrate value and encourage responsible adoption.
- **Recognize team members** who help improve legal department processes through thoughtful AI implementation.
- **Examples include:**
 - Launching an **AI-assisted intake portal** that categorizes legal requests and routes them to the right lawyer or team;
 - Piloting an **internal knowledge repository** that allows attorneys to quickly find prior templates, advice, investigation summaries, or playbook language;
 - Building a **self-service workflow for low-risk documents**, such as standard NDAs, with clear escalation points to Legal;
 - **Creating dashboards** that track contract cycle time, intake volume, and turnaround times before and after an AI pilot.

CHECKLIST 5

» WITH THANKS TO:

*Michael J. Breslin,
Partner, at Kilpatrick,
for his contribution
to this checklist.*

*Thanks also to ACC
members for their
valuable insights.*

2. Keep human legal judgment at the center of AI-supported work. AI tools do not replace legal judgment or professional responsibility. Human oversight is especially important where an AI output could influence a negotiation position, a compliance assessment, a litigation strategy decision, a reserve analysis, or advice to the business.

- **Treat AI output as a starting point** rather than a final product.
- **Carefully verify** the accuracy, completeness, and relevance of any AI-generated content.
- Remember that AI tools can generate **persuasive but incorrect answers** ("hallucinations").
- **Maintain human oversight** of legal analysis, advice, and decision-making.
- **Use AI to accelerate routine work** while reserving higher-judgment decisions for attorneys.
- **Ensure appropriate review processes** when AI outputs influence business decisions.
- **Examples include:**
 - **Asking AI to draft a first version** of a policy, training outline, or internal legal memo, then revising it based on business context and legal judgment;
 - **Using AI to summarize** a lengthy contract, pleading, or investigative record so the lawyer can focus faster on the key risk points;
 - **Generating an initial research summary**, while independently checking authorities, dates, and conclusions before relying on the work; and
 - **Using AI to suggest fallback language in a negotiation**, while reserving to counsel to determine whether the position is appropriate in light of commercial context and risk tolerance.

3. Identify practical use cases where AI can streamline legal workflows.

Focus on tasks that are repetitive, document-heavy, or time consuming. Common applications include:

- **Document drafting, review, and management.**
 - **Generating first drafts** of contracts, policies, memoranda, or internal communications.
 - **Summarizing lengthy documents** such as agreements, pleadings, or regulatory materials.
 - **Identifying key provisions** within agreements, such as indemnification clauses, termination rights, or governing law provisions.
- **Contract management and analysis**
 - **Comparing contract language** against internal playbooks or template standards.
 - **Flagging clauses** that deviate from approved standards.
 - **Creating searchable databases** of contractual provisions across large contract portfolios.
- **Legal research and regulatory analysis.**
 - **Producing summaries** of unfamiliar legal or technical concepts.
 - Creating preliminary **multi-jurisdictional legal surveys**.
 - **Monitoring regulatory developments** and generating summaries of legal updates.
- **Examples include:**
 - **Using retrieval-augmented generation (RAG) tools to compare contracts** against the department's approved fallback positions and identify deviations more quickly;
 - **Creating searchable repositories of negotiated clauses** across a large portfolio of agreements to support renewals, audits, diligence, or risk mapping;
 - **Using AI to assemble first-pass multi-jurisdictional survey charts** that lawyers can then validate and refine;
 - **Deploying AI tools during M&A diligence** to sort agreements into risk buckets, identify unusual provisions, and surface exposure trends for closer legal review;
 - **Using AI to monitor regulatory developments** across federal, state, EU, and sector-specific sources and generate summaries and implementation notes for the business.

« This checklist is based mainly on the following ACC resources and on insights from the sponsor Kilpatrick and ACC members.

← **Leveraging GenAI for Business Impact: A Playbook for Small Law Departments**, by ACC, March 17, 2026.

← **From Curiosity to Capability: What Responsible AI Adoption Really Looks Like for Legal Teams**, by ACC, *Corporate Counsel Now*, February 25, 2026

← **ChatGPT and AI Applications for In-house Lawyers**, by Spiwe L. Jefferson, *Corporate Counsel Now*, March 8, 2023

← **Prompt Engineering: Your Hidden Superpower for AI Mastery**, by Spiwe L. Jefferson, *Corporate Counsel Now*, July 23, 2024

← Visit the **AI Center of Excellence for In-house Counsel**

4. Use AI to improve internal legal service delivery.

AI tools can help legal departments respond more efficiently to internal business requests that are governed by established policies, templates, approval rules, or playbooks.

● **Examples include:**

- **An internal chatbot** trained on the company's policy library, code of conduct, signature authority rules, and template library, so employees can ask practical questions, such as whether a proposed contract can be signed at a given dollar threshold or whether a policy restriction applies in a given jurisdiction;
- **An intake tool** embedded in Microsoft Teams or another workflow platform that asks guided questions, then routes the request to commercial legal, employment counsel, privacy counsel, or compliance;
- **A self-service NDA process** that allows approved business users to generate a standard NDA automatically, while escalating non-standard requests to Legal;
- **AI-generated meeting summaries and action item lists** that reduce the administrative burden on in-house lawyers after recurring cross-functional meetings.

Douwe Groenewelt, *Partner at Datacation, Data & AI for Professional Services*, shared a shortlist of some of the compliance-related AI tools that European in-house legal professionals use, including

- EU Data Protection Impact Assessment (DPIA) Assessment Assistant
- Marketing Claim Analyzer, a type of tool that helps to check whether marketing claims are compliant with applicable regulations
- EU Corporate Sustainability Due Diligence Directive (CSDDD) Compliance Assistant
- Export Control Classification Assistant
- Sales Tender Review Assistant

5. Consider how emerging Agentic AI systems may automate portions of legal workflows. Agentic AI systems operate across multiple steps of a workflow rather than simply generating content. While these systems may significantly improve efficiency, they also **introduce new governance and oversight considerations**.

- **Ensure that appropriate safeguards are in place** before deploying systems that can access internal data sources or interact with enterprise systems. For more information on Agentic AI governance, see [Checklist 10: Five Tips to Develop Governance for Agentic AI](#).
- **Examples include:**
 - **A contract workflow** that receives intake, reviews the draft against playbooks, proposes fallback language, routes the issue for approval, and opens a ticket if business escalation is needed;
 - **A compliance questionnaire process** in which AI gathers information from approved internal sources, drafts responses, and routes the final output to a lawyer or compliance professional for review;
 - A workflow that takes a **routine legal request**, determines whether it fits a standard path, prepares a draft response, and escalates only if certain thresholds or exceptions are triggered; and
 - **An automated process** that checks signature authority, routes a document for approval and e-signature, updates the **CLM or matter system**, and notifies the relevant business stakeholders.
- In each of these scenarios, **the legal department should ensure there are clear controls** over what the system can access, what it can do, and when a human must remain in the loop.
- **Be aware that Agentic AI systems can make critical mistakes** in executing automated workflows and may even drift into inaccuracy over time. This can lead to catastrophic failure and liability when AI agents are supporting legal, financial, and HR decision making. Agentic systems must be tightly designed, carefully prompted, and continually overseen to avoid putting the business at risk.

» Also check out:

- [How Legal Ops Leaders Are Driving AI-powered Legal Departments](#), by Blake E. Garcia and Mauro Whiteman, *Corporate Counsel Now*, January 20, 2026
- [The AI-Empowered Counsel: A Practical Framework for In-house Legal Teams](#), by Hayden O. Creque, *Corporate Counsel Now*, October 28, 2025
- Watch: [Leading Through the Next Phase of GenAI Adoption](#) (ACC Webinar, January 20, 2026, available on-demand)

6. Leverage AI capabilities embedded within existing business tools.

Before adopting new software, consider whether existing tools already provide useful AI capabilities.

- **Do not assume that “existing” means “approved” or “low risk.”** Legal should still understand what features are enabled, what data those features can access, and whether the output is appropriate for the intended use case.
- **Examples include:**
 - **Using embedded AI features in document platforms** to summarize a long draft agreement, compare versions, or propose clearer language;
 - **Asking a CLM platform** whether it can automatically locate similar clauses, identify deviations from approved language, support bulk review across an agreement set, or help detect playbook drift over time;
 - **Using AI-enabled features in matter management or discovery tools** to categorize documents, cluster issues, or accelerate first-pass review;
 - **Using meeting transcription tools** to create internal summaries and task lists, where appropriate for the sensitivity of the discussion.

7. Measure and track efficiency improvements from AI adoption.

Many legal departments are interested in understanding whether AI tools deliver **measurable operational improvements**.

- **Consider tracking how outside counsel use AI** and whether those efficiencies are reflected in billing practices. For more on this topic, see Spotlight below on [tracking and assessing AI use by outside counsel](#).
- **Use GenAI tools to produce benchmarking reports, predictive analytics, and simple visualizations** that show increased efficiency and demonstrate the legal department's strategic value to the business.
- **Examples include:**
 - **Comparing average turnaround time** for routine commercial contracts before and after deploying an AI-assisted review workflow;
 - **Tracking how often an AI intake tool resolves or routes routine requests** without requiring individualized attorney triage;
 - Measuring whether **AI-assisted regulatory monitoring** reduces the time lawyers spend manually tracking legal developments;

- **Using analytics** to identify whether **outside counsel time billed** for routine research, document review, or drafting has declined after AI adoption;
- **Generating dashboards** that summarize legal department volume, contract cycle time, matter trends, litigation exposure, or AI-assisted workflow activity in a more board-ready format.

8. Continue to evaluate the risks and limitations of AI tools.

Even when used internally for efficiency purposes, AI tools raise legal, ethical, and operational risks.

- **Coordinate closely with IT, security, and compliance teams** when deploying AI tools that access internal data or interact with enterprise systems.
- **Examples include:**
 - Evaluating whether a **contract analysis tool** is drawing from the correct repositories and whether its retrieval logic could surface the wrong precedent or expose restricted information;
 - Confirming whether an **internal chatbot** is limited to approved policies and reference materials rather than drawing from uncontrolled sources;
 - Reviewing whether an **AI-based regulatory monitoring or litigation support tool** can be audited and whether its sources are sufficiently reliable for the intended use;
 - Testing whether a **workflow tool** properly escalates exceptions rather than treating every request as routine;
 - Confirming whether a **tool that generates analytics or executive summaries** is using data that is sufficiently current, complete, and reliable for the decision-makers who will receive it.



We work with our CISO and IT team to review every tool so we understand if it has an AI feature and whether or not our organization is actually using those features.

Aparna D. Williams
*Chief Legal Officer and
 Corporate Secretary,
 Sophos*



9. Encourage ideation and experimentation by the legal team.

AI implementation in Legal shouldn't feel like another compliance exercise. It's a **development process** that can be exciting, and even fun, for the entire team.

- **Set this tone at the top** to signal excitement, and not just permission.
- **Examples:**
 - One company's legal team launched an **internal competition** with the goal of building more than 350 AI agents across the department. Lawyers and legal professionals could submit their own agents, the best submissions were showcased, and the winner received a fully paid trip to Silicon Valley. People who had never built anything technical were prototyping workflows and sharing them with colleagues.

10. Learn from peers and share experiences.

AI adoption across legal departments is evolving rapidly. Many useful insights come from peer discussions and shared experiences.

- **This is also a useful way to benchmark what "good" looks like in the market** without treating every vendor pitch as a reliable indicator of maturity.
- **Consider:**
 - Participating in *ACC Community forums* and *practice networks*.
 - **Attending ACC programs and conferences** focused on legal technology.
 - **Sharing lessons learned internally** to improve future AI initiatives.
 - **Engaging with vendors and other legal departments** to understand emerging tools and best practices.
- **Examples include:**
 - **Asking peers** how they are measuring AI-related efficiency gains in contract review, intake, regulatory tracking, or research workflows;
 - **Comparing approaches** to outside counsel guidelines on disclosure of AI use, billing treatment, and confidentiality protections;
 - **Learning from other in-house teams** about which tools have actually reduced repetitive work and which proved harder to implement than expected; and
 - **Sharing internal lessons learned** after a pilot, including what worked, what failed, and what should be adjusted before scaling.

TAKE ACTION:
Write one or two steps that you plan to implement.

Describe	By what date?
1.	
2.	
Notes:	


EXPLORE NOW



» WITH THANKS TO:

Michael J. Breslin,
Partner, at *Kilpatrick,*

for contributing
this list.



AI Use Cases for Legal Departments in 2026

As legal departments become more sophisticated in their use of AI, **some are moving beyond basic drafting and summarization tools** to explore more integrated workflow uses.

- **RAG-powered contract intelligence platforms** that pull clauses from executed agreements, compare them against internal playbooks, identify deviation patterns, and help surface portfolio-level risk trends;
- **AI-assisted regulatory change radar tools** that monitor federal, state, EU, and sector-specific developments and generate summaries, impact assessments, and recommended next steps;
- **Litigation early case assessment tools** that help organize pleadings, prior similar matters, and internal documents to support reserve analysis, cost forecasting, and settlement strategy;
- **Contract playbook drift detection tools** that compare approved fallback positions against final negotiated language and help identify risk creep over time;
- **Knowledge memory engines** that allow attorneys to search prior legal advice, board materials, investigation summaries, and other internal work product to reduce duplicative analysis;
- **AI-assisted board and executive reporting tools** that help generate summaries of matter volume, risk themes, litigation trends, and operational metrics in a more digestible format;
- **AI-based deposition and discovery pattern analysis tools** that help litigation teams identify recurring themes, inconsistencies, or likely lines of inquiry across large data sets.

These tools may offer significant efficiency gains, but their usefulness will depend on the quality of the underlying data, the accuracy of retrieval and analysis, the appropriateness of the use case, and the degree of human oversight.



ACC Member AI Use Cases: How In-house Lawyers Use AI


Meganne Thaxton, Principal Corporate Counsel at Red Ventures, mentioned the following uses of AI tools:

- Review contracts using an add-on tool for the company's CLM software (Contract Lifecycle Management).
- Generate first drafts of:
 - New contract templates.
 - A contract clause more favorable to the company.
 - Email responses.
 - A legal memo.
- Personalize documents for tone and length.
- Write and rewrite annual performance reviews.
- Create a short sample form for legal disclaimers and disclosures.
- Produce summaries of:
 - Legal research.
 - Recent case decisions.
 - Regulation and its impact on the business.
- Generate insights and analysis across multiple sources to get up to speed quickly.
- Translate non-English agreements, emails, and case summaries.
- Brainstorm legal strategies, key considerations, and agendas.

» WITH THANKS TO:

Michael J. Breslin,
Partner, at *Kilpatrick,*

for contributing
this list.



Track and Assess Outside Counsel's Use of GenAI

As law firms increasingly deploy GenAI tools for research, drafting, and document review, many legal departments are evaluating how those tools affect billing practices, work quality, confidentiality, and overall value delivered to the client.

Some in-house teams are beginning to request greater transparency regarding outside counsel AI use and to track whether resulting efficiencies benefit the client.

1. Consider asking these questions of outside counsel, whether in outside counsel guidelines, engagement letters, or RFP processes:

- **Are GenAI tools used on this matter**, and for what tasks (e.g., research, drafting, document review, diligence)?
- **Which tools are used** (firm-developed systems, enterprise platforms, or third-party tools)?
- Is client **confidential or privileged information** entered into those tools?
- Is client information used for **model training or vendor development**?
- **What human review is applied** before AI-assisted work product is delivered?

2. Request AI usage reports. Some legal departments are asking firms to provide matter-level visibility into AI-assisted work, for example by:

- **Identifying tasks** where AI tools were used;
- **Logging AI-assisted tasks** within billing or matter management systems;
- **Estimating time or resource savings** associated with AI use; and
- Certifying that **client confidential information** was not entered into unapproved public models.

- 3. Track metrics.** While no uniform standard has emerged, some businesses are tracking outside counsel AI metrics such as:
- **AI transparency:**
 - **Percentage of matters** where outside counsel report AI use
 - **Types of tasks** where AI is used
 - **Efficiency and cost impact:**
 - **Reduction in hours billed** for routine research or drafting
 - **Reduction in document review hours** during discovery or diligence
 - **Changes in staffing patterns** or blended effective rates
 - **Quality indicators:**
 - **Number of revisions required** for AI-assisted work product
 - **Instances of hallucinated citations** or incorrect authorities
- 4. Expect continuing developments in assessment of AI use by outside counsel.** Many legal departments support responsible use of AI by law firms but increasingly expect transparency, appropriate safeguards, and fair sharing of efficiency gains as industry practices continue to evolve.

TAKE ACTION:
Write one or two steps that you plan to implement.

Describe	By what date?
1.	
2.	
Notes:	



Top Seven Tips for Prompting Generative AI Tools

Consider these strategies to increase the chances of obtaining relevant output when you query a Generative AI tool. Keep in mind that each tool has its own features, advantages, and limitations. Below are general prompt ideas. You will need to create your own prompts relevant to the context of what you are trying to accomplish.

1. **Use direct, simple and focused prompts.** Avoid including too many variables. Consider applying the **CO-STAR framework** or similar alternatives to design your prompt.
2. If possible, **include context** to help the tool offer a more tailored response. For example, ask the tool to assume a role, and to create output for a specific type of audience.
 - Depending on the tool, it may be helpful to upload relevant documents or website URLs it can reference to add context.
 - **Here is an example:** *"Assume you are the general counsel of a large global industrial group that needs to deliver a one-hour training session to sales staff regarding what constitutes a binding contract. The topic should include key points such as [...] Create a slides deck for a one-hour presentation. Don't include more than [...] slides. Use simple wording (fifth grade level)."*
3. To help the tool adopt the desired voice and tone, **consider providing examples of publicly available language or messaging that your organization has used.** For example, if you ask a tool to come up with the draft policy, consider pasting in the tool your organization's core values, for a result that aligns better with these values. For example: *"As context, assume that the trainees are familiar with industry x, and that core values of the company are x, y, and z."*
4. **Use iterative prompts.** Once you have a first answer from the tool, add follow-up questions or instructions to further tailor the output to your needs. For example, ask the tool to rephrase the output using a different tone; or ask the tool to present the content in a specific format (e.g., x number of bullet points of no more than x number of words in each bullet point).

5. **Bonus Tip 1: Adapt your prompts to the model you are using.** Newer AI models like Open AI's ChatGPT 5.5 and Anthropic's Claude 4.7 already incorporate iteration, backtracking, and multi-step reasoning. These models may actually work less well if "overprompted" to add those steps – a tip shared by **Douwe Groenevelt**, *Partner at Datacation, Data & AI for Professional Services*.
6. **Bonus Tip 2: Save successful prompts.** **Malaika Roemer**, *CIPP/US, Principal Paralegal and Compliance Program Manager at Coalfire*, shared that it may be helpful to save prompts that produce good responses, especially for AI uses likely to be often repeated. Some AI tools will let you store prompts for reuse.
7. **Bonus Tip 3: Think of GenAI like a helpful, but naïve first-year associate when crafting prompts.** Provide the tool with specific issues to spot and a clear format for the output.
 - **Poor prompt:** "Is this contract okay?"
 - **Better prompt:** "Review this vendor agreement against our company playbook, focusing on indemnity, limitation of liability, governing law, and data security. Flag any deviations from our standards in a bulleted list, explain the business risk of each deviation, and suggest our preferred fallback language."

Learn
more

Prompting strategies and tips:

- ***AI Use Case: Self-Service for Low-Complexity Business Agreements***, by Hayden O. Creque, January 21, 2026
- ***The AI-Empowered Counsel: A Practical Framework for In-house Legal Teams***, by Hayden O. Creque, *Corporate Counsel Now*, October 25, 2025
- ***Prompt Engineering: Your Hidden Superpower or AI Mastery***, by Spiwe L. Jefferson, *Corporate Counsel Now*, July 23, 2024)

CHECKLIST 6

» WITH THANKS TO:

Joel D. Bush, Partner,
at **Kilpatrick,**

for contributing
this checklist.

Meet Ethical Obligations While Using Generative AI

If you are an in-house lawyer in the US, your ethical obligations extend to your use of AI tools. Learn how to satisfy ethics requirements while advising the business.

1. Further your duty of competence by effectively using AI tools.

In the United States, under the American Bar Association's rules (**ABA Rule of Professional Conduct 1.1**), a lawyer is required to provide competent representation to their clients. **Comment 8** to this rule, adopted by 37 US states, indicates that this includes being up to date with "relevant technology."

- **Lawyers are thus ethically required to learn about new technology and develop skills for using new technologies**, such as generative AI, in ways that can serve their clients.
- Despite this ethical balance in favor of using GenAI, **it cannot replace independent analysis and judgment**. Use it wisely and responsibly on the client's behalf.
- **Uploading client information** into a public GenAI platform without client consent can violate both the duty of competence and the duty of confidentiality.
- **GenAI can hallucinate by producing fake results**, including making up summaries, quotes and citations for nonexistent court opinions.
 - As of March 2026, **over 1,000 documented court filings** have included hallucinated case citations and non-existent case law due to the misuse of GenAI in legal research.
 - In addressing this problem, **courts have shifted from money sanctions to more creative requirements for deterrence**.
 - One court disqualified the offending attorneys and required them to provide a copy of the court's order to all clients, opposing counsel, and judges with which the attorneys were involved (*Johnson v. Dunn* (N.D.Ala. July 23, 2025)).
 - Another court required the offending lawyer to share his experience at bar association events or with groups of law students (*Hodges v. Meridian Waste* (M.D.Fla. Jan. 16, 2026)).
 - The more obscure the topic, the more likely it is that the AI tool will hallucinate.

- **Verify the accuracy of your GenAI work product**, whether it is for a submission to a court, a memo to corporate leadership, or any other purpose.
- **Consider these parameters for making competent use of GenAI:**
 - Follow your company's GenAI use guidelines.
 - Be mindful before using any self-learning GenAI platform.
 - Depending on the type of information you want to enter in the AI tool, only use a private/enterprise GenAI platform with company information, unless you obtain informed written consent to use an open GenAI platform. Even if you obtain the client's informed consent to use a tool, ensure that your use of the tool will not contravene your duty of confidentiality or other ethics obligations.
 - Where productive, use specific and custom instructions. Provide inputs/instructions to use specific trusted information sources to generate work product.
 - Validate all outputs. Treat all GenAI outputs as a draft that must be reviewed and checked for accuracy.

2. **Maintain your duty of confidentiality when using AI tools.**

ABA Rule of Professional Conduct 1.6 requires lawyers not to reveal client information without the client's informed consent or under special circumstances.

- **Consider that public generative AI platforms are typically not secure (at least in their default settings).** Uploading confidential company information to these public platforms without consent could violate this duty of confidentiality, and may also potentially:
 - Waive any attorney-client privilege covering the information,
 - Violate privacy and data protection laws,
 - Destroy any trade-secret protection,
 - Violate export control laws,
 - Put patentability of inventions at risk, and
 - Risk liability under employment laws.

- **Use private/enterprise AI platforms whenever confidential company information will potentially be uploaded for legal work.** However, first ensure that these systems are, in fact, secure. This could include due diligence on vendor dependencies and data practices.
 - **Private AI platforms are becoming more common**, including downloading more compact GenAI models to in-house private servers walled off from the public Internet.
 - **“Private platforms” mean platforms that do not share information outside your enterprise.** Such private platforms may provide a greater degree of comfort that the user will not breach confidentiality, if adequate cybersecurity protocols are in place.

3. **Extend your duty of supervision to your oversight of Legal’s AI usage.** *ABA Rule of Professional Conduct 5.1*

requires supervisory attorneys to make reasonable efforts to ensure that their subordinates comply with the ABA Rules.

- This duty includes **deciding whether and how the legal department will use generative AI.**
- All supervising attorneys – not just the GC or CLO – need to be educated on how technology works, the risks of uploading company information, the risks of using downloaded content, etc.
- Supervising attorneys need to participate in determining parameters of AI use, setting expectations for verifying content, establishing adequate security protocols, etc.
- This duty to supervise extends to supervising GenAI vendors if their tools are used for legal work.

4. **Carry out your duty of communication to the business as your client.**

- *ABA Rule of Professional Conduct 1.4* requires lawyers to “reasonably consult” with their client and to keep them “reasonably informed” about the status of legal matters.
- **This duty includes communicating the risks, limitations, and potential uses of AI.**
- You should also discuss whether it may cost the business more not to use beneficial AI technology than to take on the costs and risks of implementing it.

5. Guard against loss of privilege protections with GenAI.

- Based on recent case law, if a non-lawyer poses legal questions to a GenAI platform, there is **no attorney-client relationship** created and **no privilege** attaches!
 - In *U.S. v. Heppner* (S.D.N.Y. Feb. 17, 2026), the court held that AI-generated documents created by a non-lawyer criminal defendant sent to, but not requested or directed by, his attorney were not protected by the attorney-client privilege or by the work product doctrine.
 - However, in *Warner v. Gilbarco* (E.D.Mich. Feb. 10, 2026), the court held that materials prepared by a *pro se* litigant using a public AI chatbot in anticipation of litigation **were** protected by the work product doctrine.
- **Lessons from these cases:** Attorney-client or work product privileges potentially apply if counsel uses (or directs the use) of a secure GenAI platform that safeguards confidentiality.

Learn more

- [*AI Counsel: When AI Addiction Becomes Professional Negligence*](#), by Spiwe L. Jefferson, *Corporate Counsel Now*, November 5, 2025
- [*Who's Really Advising the Client? The Blurred Lines of GenAI in Legal Practice*](#), by Susan Duarte and Jonathan S. Marashlian, *Corporate Counsel Now*, August 25, 2025
- [*Practical Lessons from the Attorney AI Missteps in Mata v. Avianca*](#), by William A. Ryan, Senior Vice President & Chief Compliance Officer (AT&T Services Inc.), Allen Garrett, Partner, and Brad Sears, Associate, Kilpatrick, August 8, 2023
- [**Watch:** *AI and the Legal Profession: A Responsible Approach to Adoption*](#) (ACC Webinar, August 14, 2025, available on-demand)
- [**Visit the** *AI Center of Excellence for In-house Counsel*](#)

CHECKLIST 7

» **WITH THANKS TO:**
Joel D. Bush, Partner,
and James A. Trigg,
Partner, at Kilpatrick,
 for contributing
 this checklist.

Five Tips to Protect Intellectual Property When Using Generative AI

Using GenAI comes with significant risks to intellectual property – both your organization’s IP and that of others. Use these tips to help avoid potential legal pitfalls.

1. Make sure your organization’s creative output includes human contributions to trigger copyright or patent protection, if ownership is a priority for the organization.

- **As of the date of production of this toolkit, creative works produced purely by GenAI** are not copyrightable in the view of the US Copyright Office, even if they result from hundreds of human prompts. The federal Court of Appeals for the District of Columbia **recently affirmed** the Copyright Office’s view. However, works containing a mixture of human-generated and machine-generated content can be protectable to the extent of human creation.
- **A significant judicial body in China** has held that AI-generated works can, theoretically, be protectable under copyright, in contrast to the approach taken in the United States.
- **Other major jurisdictions, such as the EU and UK, have yet to address this issue definitively.** The best practice for a global business is to ensure that public-facing creative works, such as marketing text and graphics, are made with significant human contributions.
- **GenAI tools cannot be “inventors”** without human contributions for purposes of patent protection in every jurisdiction, except South Africa, that has addressed the issue. Patent-generating businesses, from software to pharmaceuticals, should take this into account in their product development processes.

2. Be aware of the evolving legal landscape on fair use and GenAI if your organization builds or incorporates GenAI models.

- Many GenAI tools rely heavily upon the **ingestion of copyrighted works** in order to train and learn. This is an essential part of the process in developing a model that will generate meaningful content.
- **Copyright owners** across multiple industries and disciplines have taken **legal action** against the developers of GenAI products, arguing that the ingestion of their content without permission constitutes copyright infringement.
- **Developers** counter that training is **“fair use”** under copyright law.

- These legal battles are playing out globally in ongoing cases, for example, in the US, UK, India, and Canada.
- In June 2025, two separate judges in the Northern District of California issued rulings on the question of whether “training” large language models on copyrighted materials constitutes **fair use** under US law.
 - Both judges found, on the facts of the cases before them, that **ingestion of copyrighted materials for learning purposes was fair use**, but both decisions came with **caveats**:
 - In *Bartz v. Anthropic*, the presiding judge found that **using illegally downloaded “pirated” books for model training** was not fair use, even if the training process itself was a fair use.
 - In *Kadrey v. Meta*, issued two days after *Bartz*, a different Northern District of California judge similarly ruled that Meta’s use of copyrighted works to train its Llama large language model constituted fair use. However, that judge suggested that the plaintiffs had not developed a record that the defendant’s copying indirectly competed with the market for their works, and the opinion hinted further that plaintiffs who develop such a record in future cases might witness a different outcome.
 - Neither of these lower court rulings will be the final word on these issues. **Both could be appealed**, and there are **dozens of similar cases** currently proceeding in different federal jurisdictions throughout the United States. As of the time of publication, the *Bartz* case appears to be settled, with Anthropic agreeing to pay \$1.5 billion to rightsholders of some 500,000 works that were found to have been pirated – a figure that amounts to approximately \$3,000 per work. At present, this *Bartz* settlement awaits final court approval.
- **A court in China ruled in favor of a copyright owner** in a dispute over copyrighted training data.
- As the law on AI fair use evolves globally, companies seeking to develop AI models should **proceed with caution**:
 - A GenAI tool need not be built upon copyrightable material to be usable; factual data, government content, and public-domain material can be used as training data.
 - Developers can obtain permission or seek licenses for the training sets used to develop their AI model.
 - The nature and purpose of a GenAI tool can affect whether its training data will be covered by fair use. GenAI applications that are built for purely *internal* uses – such as for ideation or research and development – appear more likely to survive copyright scrutiny than products that are intended for external use, such as for the creation of public-facing output.

3. End users of GenAI tools may also face liability for using AI-generated output that infringes copyrights.

- For example, *cases over training data* brought in the US by Getty Images and The New York Times against AI developers allege that the **developers' products can be used to produce images and text** that are close enough to the originals used as training data as to be **copyright violations**.
- **Businesses should exercise caution when using "off the shelf" AI tools for content generation**, whether for text, graphics, music, or other materials. Without a special enterprise license, many AI programs are offered without representations and warranties and without guarantees that outputs will not infringe third parties' rights.
- **Use enterprise AI offerings when possible.** Such products may come with representations and warranties, along with undertakings to indemnify users, if a piece of content created with such a tool infringes a work on which the tool was trained.
- **Avoid prompts that invite AI tools to infringe.** Steer clear of prompts that ask a tool to replicate some specific piece of content. A prompt such as, "Provide a general recap of the outcome of Super Bowl LIX" is preferable to "Provide a copy of the *New York Times* article from February 10, 2025 regarding Super Bowl LIX."
- **As noted above, internal uses appear safer than external-facing uses.** Using GenAI tools for internal ideation and research seems to carry lower risk than using AI-generated content in public-facing settings.
- **Use programs to track which pieces of company collateral contain AI-generated materials.** This will assist the business in responding to potential copyright infringement claims and in assessing the viability of claims against third parties.

4. Use a closed GenAI tool to better protect your organization's outputs as trade secrets.

- Unlike with patent and copyright law, **AI-generated content can be protected as a trade secret, if confidentiality requirements are satisfied.** There is no requirement that the creator of a trade secret be a "natural person."
- **Ensure that your product development, IT, sales, and content teams use closed enterprise AI systems or subscriptions** rather than publicly available tools when engaging in work for the business. Using publicly available default versions of the tools will produce work that may not be confidential, and thus not a trade secret.

5. Avoid accidental public disclosure of your trade secrets when using GenAI.

- **Prohibit employees from discussing or uploading company trade secrets in conversations with public GenAI platforms.** Where needed, closed systems should be used for this purpose.
- GenAI platforms often warn users not to share sensitive information in conversations with the tool. These conversations are typically retained and used to further train the AI model.
- A *real-world example happened in 2023* when engineers at a prominent technology company unintentionally leaked confidential data – including source code and internal meeting notes – while using ChatGPT to fix errors in their source code. This data is now stored on OpenAI's servers and is potentially discoverable by competitors.

Learn more

- *Gen AI: The “Artificial” Threat to Trade Secrets*, by Joel D. Bush and Kurtis G. Anderson, Kilpatrick, April 4, 2025
- *Robots are Coming – But They Still Can’t Register Copyright*, by Gabriel M. Ross, James A. Trigg, Mehrnaz Boroumand Smith, and Joseph Petersen, Kilpatrick, March 24, 2025
- *AI and Fair Use: Two US Court Decisions Shaping the Landscape*, by ACC, July 18, 2025
- Visit the *ACC AI Center of Excellence for In-house Counsel*.
- Explore how other in-house lawyers use AI: **Connect with peers in ACC Community forums**
- Learn more about *Trade Secrets: Legal Framework and Best Practices for Enforcement (United States)*, ACC Guide, sponsored by Kilpatrick (2020).

CHECKLIST 8

» **WITH THANKS TO:**
Meghan K. Farmer,
Partner, and
Gregory P. Silberman,
Counsel, at Kilpatrick,
 for contributing
 this checklist.

Nine Steps to Evaluate AI Tools, Products, and Services

Learn nine steps to help the business vet the adoption of AI tools. Anticipate risks, propose practical safeguards, and shape AI deployments in ways that enable innovation while supporting compliant and successful launches.

1. Define the proposed AI use case. Confirm that the business has clearly defined the intended use of the AI system.

- Legal needs a clearly defined use case to determine whether the system may be **high-risk, regulated, or prohibited**, and to recommend safeguards early to avoid delays later in deployment.
- **Questions your team should ask:**
 - What specific **business problem** is the AI tool intended to address?
 - What **business functions or decisions** will the AI influence?
 - Will the system:
 - Inform human decision-making, or
 - Make **automated decisions**?
 - Could the proposed use case be considered **high-risk** under applicable AI regulations (e.g., the EU AI Act)?
 - Could the proposed use be **prohibited** or restricted under existing or emerging AI regulations?

2. Evaluate data sources and data rights. Confirm that the organization has the **right to use all relevant data**.

- **This review also helps identify risks** related to confidential data exposure, vendor data retention practices, and ownership of AI-generated outputs.
- **Questions your team should ask:**
 - **Training Data**
 - If the model is trained by the vendor, does the vendor have **legal rights to the training data**?
 - Could the training data include **copyrighted material, proprietary datasets, or personal data**?
 - Has the training data been reviewed for **bias or representativeness risks**?
 - Could training data raise **privacy or data protection concerns**?

- **Input Data (Prompts, Uploads, Integrations)**
 - Will employees input **confidential, proprietary, or regulated data**?
 - Does the vendor **retain prompts or input data**?
 - Is input data used for **model training or improvement**?
 - Are safeguards in place to prevent **unauthorized disclosure of sensitive data**?
- **Output**
 - Could outputs contain **hallucinations or inaccurate information**?
 - Could the outputs create **intellectual property infringement risk**?
 - Does the organization have rights to **use or commercialize AI-generated outputs**?

3. **Identify applicable laws and regulations.** AI systems may trigger obligations under **existing laws and emerging AI-specific regulations**.

- **Questions your team should ask.** Consider whether the proposed system may fall within:
 - **AI-specific laws.** Examples include:
 - *EU AI Act* (Check out Spotlight: *Tips for Assessing EU AI Act High-Risk AI Systems*)
 - *Colorado AI Act*
 - *NYC Automated Employment Decision Tool law*
 - Emerging US state AI governance laws
 - **Sector-specific regulations.** Depending on the use case, AI may trigger:
 - Employment discrimination laws
 - Fair lending regulations
 - Consumer protection laws
 - Financial services regulations
 - Healthcare regulations
 - **Data protection laws.** AI systems that process personal data may trigger obligations under:
 - *EU General Data Protection Regulation (GDPR)*
 - *California Consumer Privacy Act (CCPA)*
 - Other global privacy laws



We've seen vendors eager to jump on the AI bandwagon roll out **AI features with opt-out mechanisms rather than opt-in**. If you're not paying attention, your data could be feeding a model before you even know the feature exists. I try to get ahead of these surprises by **building in contractual requirements for advance notice and affirmative consent** before any new AI functionality is activated.

Malaika Roemer, CIPP/US
Principal Paralegal &
Compliance Program Mgr,
Legal & Compliance,
Coalfire



4. Review the model and technical documentation.

Request sufficient documentation to understand **how the system was developed and tested**.

- Documentation may be important to **demonstrate responsible AI governance and regulatory compliance**.
- **Questions your team should ask:**
 - What type of **AI model** is used?
 - Has the model undergone **bias testing, safety testing, and independent audits**?
 - Are **model cards, compliance reports, or technical documentation** available?
 - Has the model undergone any **conformity assessments** required under applicable AI regulations?

5. Evaluate vendor structure and dependencies.

Many AI tools rely on multiple vendors or underlying models, which can create legal and operational risk.

See also: [*Checklist 9: Top Eight Tips to Address AI Issues in Contracts with Third Parties*](#)

- **Questions your team should ask:**
 - Whether the vendor developed the AI model itself or relies on a **third-party model provider**.
 - Whether the product is a **wrapper** around another AI system.
 - The vendor's **security, compliance, and governance posture**.
 - The vendor's **financial stability** and **operational reliability**.
- **Contractual issues to consider include:**
 - Liability allocation
 - Indemnification
 - Data ownership
 - Service level commitments

6. Evaluate risks associated with AI outputs.

Assess how the organization intends to **use AI-generated outputs**. This will help determine whether **additional validation, oversight, or usage limits** are needed for higher-impact uses.

- **Questions your team should ask:**
 - What **business decisions** will rely on these outputs?

- Could **inaccurate outputs** create **regulatory risk, consumer harm, or reputational damage**?
- Could outputs create **biased or discriminatory outcomes**, particularly in HR or customer-facing systems?
- Are there **validation or review processes** before outputs are relied upon?

7. Confirm human oversight and governance. Make sure that responsibilities for monitoring, **review, and escalation** are clearly defined. This is particularly important for **Agentic AI systems**.

● **Steps your team should take:**

- Confirm that responsible teams have been identified to **monitor system performance**.
- Confirm that there are defined processes for **reviewing and escalating issues**.
- Ensure that humans will remain involved in **high-impact decisions**.
- Ensure that employees using the system will receive **appropriate training**.

8. Assess incident response and risk mitigation planning. Confirm that processes exist to **investigate incidents, mitigate harm, and coordinate responses** across teams.

● **Steps your team should take:**

- Determine whether the organization has plans to address:
 - Inaccurate or harmful outputs,
 - Biased decision-making,
 - System misuse, and
 - Operational disruption

● **Mitigation strategies may include:**

- Documented **AI risk management processes**;
- **AI impact assessments** prior to deployment (See *Spotlight: Five Tips for Conducting an AI Impact Assessment*); and
- **Cross-functional response teams** involving legal, IT, risk, and communications.


9. Establish ongoing monitoring and compliance processes.

AI systems and compliance requirements can change over time due to **model updates, new data inputs, performance drift, and evolving regulations**.

- **Steps your team should take:**
 - Ensure the organization has processes to:
 - Monitor system performance and accuracy,
 - Track **model drift**,
 - Periodically evaluate **bias and fairness**, and
 - Assess whether the system continues to comply with **evolving regulatory requirements**
- Organizations should also establish processes for **reviewing model updates or new system versions** before deployment.

Learn
more

- *Kilpatrick Connect: Your Trusted AI Partner*
- Visit the *ACC AI Center of Excellence for In-house Counsel*.



Tips for Assessing EU AI Act High-Risk AI Systems

*If an AI system may be used in the **European Union**, you should assess early whether it could qualify as a **high-risk AI system** under the **EU AI Act**.*

1. High-risk classification can arise because the AI is **embedded in certain regulated products** or because it is intended for **certain listed uses**:

- Critical infrastructure
- Education
- Employment
- Essential services
- Law enforcement
- Immigration
- Administration of justice-related contexts

2. Under the current European Commission timeline, the EU AI Act is generally in effect beginning **August 2, 2026**. The rules for high-risk AI systems **embedded into regulated products** will apply later, beginning **August 2, 2027**.

- **Note:** the Commission proposed amendments to the EU AI Act in November 2025 that could affect the timing of certain high-risk AI obligations.
- Unless and until those amendments are adopted, the current statutory timeline remains operative.

3. High-risk systems bring **added requirements** around:

- Conformity assessment,
- Risk management,
- Data governance,
- Technical documentation and traceability,
- Transparency and instructions for use,
- Human oversight,

» WITH THANKS TO:

*Gregory P. Silberman,
Counsel, at Kilpatrick,*

*for contributing
this checklist.*

- Accuracy,
- Robustness,
- Cybersecurity,
- Monitoring, and,
- In some cases, a fundamental rights impact assessment.


4. Evaluate a system’s high-risk status early, map its role in the value chain, and revisit the analysis when the system’s purpose, integrations, or deployment context changes.

5. Questions your team should ask:

- Is the system potentially high-risk because of its **product category** or its **intended use case**?
- **What is our role under the Act:** provider, deployer, importer, distributor, or another actor in the supply chain?
- **What documentation, testing, logging, human oversight, and governance measures would be required** if the system is high-risk?
- Could planned updates, retraining, new tools, or expanded use cases constitute a **substantial modification**?
- Do we need a **fundamental rights impact assessment**, and do we have a process for **post-deployment monitoring and incident escalation**?

TAKE ACTION:
Write one or two steps that you plan to implement.

Describe	By what date?
1.	
2.	
Notes:	



Evaluate These Four Key Risk Factors in AI Systems

When vetting an AI system and weighing the benefits vs. the risks, evaluate these four key areas for regulatory or compliance risks:

1. **Privacy and Data Governance.** AI systems that process personal data are likely subject to privacy laws that regulate data collection, usage, storage, and sharing.

- **Data Collection and Processing**

- Does the AI system process personal, sensitive, biometric, or other regulated data?
- Are prompts, uploads, outputs, logs, retained context, and retrieved data covered by your privacy and data governance policies?

- **Data Minimization and Purpose Limitation**

- Are data collection and processing limited to what is necessary for the AI system's intended function?
- Is there a clear legal basis and documented business justification for using personal data?

- **Transparency, Notice, and Transfers**

- Are profiling, AI-assisted decisions, or customer-facing AI interactions disclosed where required?
- Do cross-border transfers, retention periods, and deletion practices align with applicable requirements?

- **Vendor and Model Data Practices**

- Will the vendor retain inputs or outputs, use them for training or retraining, or share them with third parties?
- Does the organization understand where data goes throughout the model and tooling stack?

» WITH THANKS TO:

Gregory P. Silberman,
Counsel, at Kilpatrick,

for contributing
this checklist.

2. **Security and Resilience.** AI introduces new attack vectors and security risks, particularly related to data exposure, model manipulation, and adversarial threats.

● **Data Security and Encryption**

- How is confidential, trade secret, or personal data protected within the AI system and any connected tools?
- Does the system use encryption, least-privilege access, segmentation, and audit logs to safeguard data?

● **Model, Tool, and Integration Risk**

- Can the system be manipulated through prompt injection, adversarial inputs, model inversion, data poisoning, or insecure integrations?
- Are agentic features, APIs, and third-party dependencies properly constrained and monitored?

● **Incident Response and Recovery**

- Is there an AI-specific incident response process, including evidence preservation and escalation?
- Can the organization pause, roll back, or constrain the system if it fails or is misused?

3. **Trust, Safety, and Consumer Protection.** AI products and services must ensure user safety, prevent harm, and uphold ethical standards.

● **User Safety & Harm Prevention**

- Are there safeguards against generating harmful, deceptive, fraudulent or abusive outputs or actions?
- Does the AI system include limits on unsafe or unauthorized uses?

● **AI Transparency, Contestability & Disclosure**

- Can users understand when they are interacting with AI or when important content is AI-generated or AI-manipulated?
- Are AI-influenced outcomes contestable, reviewable, and documented?

● **Claims and Customer-Facing Risk**

- Are claims about the system's accuracy, reliability, and capabilities substantiated in writing? If so, by whom?
- Are customer communications, sales materials, and public claims aligned with how the system actually works?

4. Bias, Fairness, & High-Impact Use. Regulators are increasingly scrutinizing AI for bias, discrimination, and fairness. This is especially true in hiring, lending, health care, and applications directed towards children.

● **Bias Detection & Testing**

- Has the system been tested for bias, fairness, and representativeness in the relevant context?
- Are disparities in training data, outputs, or downstream outcomes assessed?

● **Automated Decision-making & Protected Risks**

- Could the AI use case trigger anti-discrimination, consumer protection, employment, lending, housing, education, or similar rules?
- Are human oversight and escalation mechanisms in place for consequential decisions?

● **Documentation & Remediation**

- Is there documentation showing how fairness risks were evaluated and mitigated?
- If an AI-influenced outcome is challenged, is there a process to review, explain, and correct it?

TAKE ACTION:
Write one or two steps that you plan to implement.

Describe	By what date?
1.	
2.	
Notes:	

» WITH THANKS TO:

*Gregory P. Silberman,
Counsel, at Kilpatrick,*

for contributing
this checklist.



Five Tips for Conducting an AI Impact Assessment

1. An **AI Impact Assessment ("AIIA")** is a **structured review** of an AI system's risks, benefits, and compliance obligations.
 - An AIIA helps organizations identify and mitigate legal, operational, security, and governance risks throughout the AI lifecycle.
2. The **level of documentation** should correspond to the system's potential impact and autonomy.
 - For example, an AI system analyzing surveillance footage to assess real-time threats would require a comprehensive AIIA.
3. **Low-risk AI tools** that do not process personal or sensitive data, make impactful decisions, or raise legal compliance concerns, may use a **streamlined review**, but the organization should still document the **rationale** and **approval** for the tool.
4. **Key Components of an AI Impact Assessment:**
 - **Use Case and Role Mapping:**
 - **Define** the purpose, intended users, affected stakeholders, outputs, and the business decisions the system will influence.
 - **Identify** the organization's role (for example, provider/developer, deployer, integrator) and the key third parties involved.
 - **System Boundaries and Dependencies:**
 - **Document** the model(s) used, connected tools, external APIs, retrieval sources, vendor dependencies, and any agentic tools (multi-step execution or cross-system actions).
 - **Data Inputs and Data Use:**
 - **Identify** the categories and sources of data processed (including personal, sensitive, confidential, and third-party data), where data is stored and retained, and whether any data will be used to train or fine-tune models.

- **Legal and Risk Review:**
 - **Assess** privacy, security, bias, fairness, intellectual property, consumer protection, sector-specific requirements, and any prohibited or restricted uses.
- **Human Oversight and Operational Controls:**
 - **Define** approvals, escalation paths, fallback procedures, contestability or appeals mechanisms where relevant, and identify accountable owners for monitoring and intervention.
- **Monitoring and Change Management:**
 - **Define** performance and misuse monitoring expectations, triggers for re-review, and a process to revisit the AIIA for new use cases or material changes (model version, data sources, tools, permissions, workflow, or autonomy).
- **Documentation and Approval:**
 - Record the key assumptions, known limitations, mitigations, owners, and final go/no-go decision so the organization can explain and defend the deployment if needed.

5. Jurisdiction-specific requirements may apply. Identify applicable jurisdiction and sector requirements and attach them to the AIIA as a supplement.

- **For example**, the use of **digital replicas and synthetic performers** in advertising and packaging would require additional consideration depending upon jurisdiction.

CHECKLIST 9

» **WITH THANKS TO:**
Joshua S. Ganz,
Partner, at Kilpatrick,
 for contributing
 this checklist.

Top Eight Tips to Address AI issues in Contracts with Third Parties

Your organization's vendors are most likely using AI. When entering into contracts with vendors, make sure their use of AI tools safeguards your business data and is legally compliant.

1. Ensure that your organization has proper due diligence mechanisms for vendor-provided or -used AI functionality.

- Have a **designated AI review process** to conduct due diligence on vendors' use or provision of AI functionality.
- Develop a **standard workflow** for seeking internal approvals for AI use cases.
- Create templates for **gathering and documenting vendor-provided information** about AI functionality.
- Depending on your organization's **risk tolerance**, consider the appropriate approach to your organization's contractual approval rights over vendors' use or provision of AI functionality. **For example:**
 - Some organizations include a **general prohibition** in their contracts that the vendor must obtain **prior written approval** for the use or provision of any AI functionality.
 - Other organizations take a **more permissive approach** and permit pre-approved "**white-listed**" uses tied to limited company data sets and subject to security and confidentiality parameters.

2. Understand your vendors' AI technology and usage.

- **Request a clear description** of the types of AI functionality involved and how the vendor and its subcontractors/third parties propose using any AI functionality to perform services for your company.
- This should include a description of **specific use cases** and the **data and information** that will be input into or generated from their AI systems.

3. Get documentation from vendors. Require vendors to provide detailed documentation about their AI systems, including details regarding:

- How their AI system is **architected and secured**;
- **How their AI system collects data and inputs and processes data** to generate outputs/results;
- **Assumptions, inputs, and decision paths** used by the AI to generate results;
- **Degree of human oversight** in managing AI-generated outputs or actions; and
- Circumstances under which **outputs, actions, or results can be overridden or modified** by the vendor or halted by a party.

4. Obtain assurances regarding the vendor's legal compliance.

- Check that the vendor has implemented **policies, procedures, and testing protocols** to
 - Monitor AI security and vulnerability risks,
 - Mitigate risks of illegal, unethical, biased, inaccurate, or infringing AI outputs, and
 - Comply with applicable laws and industry standards.
- Require the vendor to maintain and comply with these policies and procedures and **test and monitor the AI** on an ongoing basis to ensure compliance.
- Consult **your company's AI usage policies** to ensure that the vendor's policies and procedures are consistent and at least as protective.

5. Take subcontractors or third parties into account.

- If possible, confirm and require that any subcontractors of the vendor provide at least the same level of information as part of the diligence process and are **bound by the same AI-related contractual obligations** as the primary vendor.



It may be really tough for smaller organizations to get any traction on contractual changes, so another option is to make sure you obtain and review AI whitepapers and other technical information published by the vendor.

Aparna D. Williams
Chief Legal Officer and
Corporate Secretary,
Sophos





Something I'm constantly navigating is making sure **client expectations around AI in our offerings align with the third-party AI tools we're using internally and building into our products** or offerings. It's a **balancing act**, and I think other in-house folks at services or product companies may deal with this, too.

Malaika Roemer, CIPP/US
Principal Paralegal
and Compliance
Program Manager,
Coalfire



6. Develop standard AI-related terms to facilitate negotiations.

- More and more customers are using **form AI “addenda”** for procurement contracts, either as fixed schedules to their **standard procurement forms**, or as an addendum for use with **vendor forms**.
- Well-crafted AI “addenda” should contemplate that the vendor’s rights and obligations may differ depending on:
 - The type of AI technology and the attendant risks with potential use cases (for example, general GenAI versus Agentic AI),
 - Whether the vendor is the direct user of an AI system in delivering services or developing work product (for example, to create content that your company then uses in marketing campaigns), or
 - Whether the AI system is provided by the vendor for direct use by your company as part of the licensed technology or tool.

7. Tips for AI terms and conditions. Although the market standard for AI-related terms and conditions is still developing, consider the following:

- **Include clear definitions:**
 - Any AI-related contractual discussions will ultimately center on the **definition of “AI” technology, tool, or functionality**.
 - **Tip: Overbroad definitions** may result in protracted negotiations.
 - Include definitions for critical terms such as **AI Functionality, Inputs, Outputs and Training Data**.
 - **Tip: If customer-specific databases or workspaces** will be established — such as with a RAG AI solution — review or revise the “customer data” definition to ensure it includes such databases and workspaces (and any vectorized versions thereof).
- **Include clear limitations on:**
 - **The vendor’s ability to share inputs or outputs** with third parties, and the use of inputs or outputs by the vendor for its own purposes or for the benefit of anyone other than your company—including **prohibitions on the use of inputs or outputs for AI training purposes**;
 - **The vendor’s ability to use the AI system** for any purposes other than as approved;

- The **storage and purging** of input and outputs;
- **The use of AI systems other than non-public, enterprise-level AI instances**, with specifically articulated safeguards, to ensure privacy and security of your data and mitigate the risk of data leaks; and
- The use or implementation of **Agentic AI**. Clearly outline provisions for data ownership, confidentiality, and intellectual property rights.
- Ensure that all inputs and outputs are classified as **confidential and proprietary** to your organization.
- Where the vendor is providing **work product** that was developed using an AI functionality, **require that the vendor review and validate all AI-generated outputs** prior to providing it to your company.
- Seek **audit rights** and periodic information rights (such as responses to AI questionnaires) to provide ongoing visibility into the AI technology used or provided.
- Include requirements for monitoring for and remediating certain **“AI vulnerabilities”** and a process for receiving notification of certain **“AI incidents”**.
- **Include indemnification clauses** requiring the vendor to defend against claims arising from their use of AI functionality, including infringement or misuse of intellectual property.
- Include a provision acknowledging the need for **future changes** to the addendum based on changes in applicable law or the underlying technology.

8. Create a user-friendly playbook.

- Your standard **AI addendum** can be used as the basis for the development of a user-friendly negotiation playbook.
- **Develop approved fallbacks and approval mechanisms for deviations from the standard terms**, depending on the specific AI technology, the contemplated use cases, the types of data and information that will be input into the AI system, and/or the types of outputs or actions that will be generated or facilitated using the AI tool or functionality.

9. Stay ahead of potential contractual issues with your organization’s own customer-facing use of AI.

- Just as you ensure AI safeguards in your company’s procurement processes, make sure to account for **third-party AI features built into your organization’s own products.**



- Watch: *Navigating AI Partnerships: Essential Strategies for Success* (ACC Webinar, November 4, 2025, available on-demand)
- Visit the *ACC AI Center of Excellence for In-house Counsel*.

TAKE ACTION:
Write one or two steps that you plan to implement.

Describe	By what date?
1.	
2.	
Notes:	

Five Tips to Develop Governance for Agentic AI

Agentic AI is no longer on the horizon. It's here. This new category of AI tools has the potential to revolutionize the business, but it also makes many familiar AI risks more dynamic and harder to contain.

1. Learn the basics of Agentic AI.

- **Agentic AI** refers to an AI-enabled system that, can **select, sequence, execute, or help to coordinate actions** across software tools, data sources, and workflows. It is not limited to generating content in response to a prompt.
 - The more an AI system can access data, use tools, make or shape decisions, and operate across multiple steps or systems with limited human intervention, the more “agentic” it is.
- In **practical terms**, an agentic system may:
 - **Interpret** an objective or task;
 - **Plan** or dynamically select next steps;
 - **Use tools**, integrations, or connected systems;
 - **Evaluate results** and adjust its behavior; and
 - **Continue across multiple steps** until a stopping point is reached, such as task completion, escalation, human approval, a policy limit, or an error condition.
- **Agentic AI is not necessarily a wholly separate technology from GenAI.**
 - In many cases, it is best understood as **GenAI used within a broader system** that adds planning, tool use, workflow logic, memory, and action capability.
- **Agentic AI is not the same as an “AI agent”.** This is an important distinction.
 - An **AI agent** is a system that performs single, specific actions based on input from a user or another agent.
 - **Agentic AI** systems coordinate multiple agents to execute more complex, multi-step workflows through the process of orchestration.

CHECKLIST 10

» WITH THANKS TO:

*Gregory P. Silberman,
Counsel, at Kilpatrick,*

for contributing
this checklist.

2. Determine where and how Agentic AI is planned or already in use.

Internal stakeholders may say that a proposed system is “just a workflow,” “just automation,” or “not really an agent.” However, a system does not need to be called an “agent” to present agentic risk. Those labels can obscure real legal issues.

- Some **examples** of Agentic AI systems:
 - A **contract-management tool** for the legal department that reviews a contract, checks internal playbooks, suggests fallback language, routes the issue, and opens a ticket;
 - A **customer-operations system** that pulls information from multiple internal systems, determines the next best step, drafts a communication, and updates a CRM record; and
 - A **security system** that analyzes an alert, queries logs, opens an incident ticket, recommends containment steps, and notifies designated personnel.
- For **legal and governance purposes**, ask these questions:
 - What can the system access?
 - What can it do?
 - How much autonomy does it have?
 - What human oversight exists?
 - Can the company reconstruct and explain what happened?

3. Consider the possible legal risks posed by Agentic AI.

Agentic AI is not just a productivity tool, but a system of delegated action. As such, familiar legal risks can arise in a more operational form:

- **Privacy and confidentiality** issues may arise when the system can retrieve, combine, or transmit sensitive information across repositories, tools, or vendors.
- **Cybersecurity** issues may arise when the system has excessive permissions, relies on weak integrations, or can be manipulated through inputs, instructions, or connected tools.
- **Consumer protection, employment, and regulatory** issues may arise when the system influences consequential decisions, customer interactions, or operational outcomes.
- **Intellectual property and data rights** issues may arise when the system uses internal or third-party content in ways that are not well understood or contractually addressed.
- **Accountability and recordkeeping** issues may arise when the business later needs to explain, defend, investigate, or remediate the system’s conduct.

4. Address key questions in developing governance for Agentic AI systems. Ask these core questions not only during the initial review of a proposed system, but also when evaluating new use cases and as part of change management as the Agentic AI system evolves or is updated:

- **What can the system see?**
 - What data sources, repositories, messages, documents, or applications can it access?
 - Can it access confidential, privileged, personal, regulated, or commercially sensitive information?
 - What information is retained in prompts, traces, logs, or outputs?
- **What can the system do?**
 - Does it only recommend, or can it also act?
 - Can it send messages, update records, approve steps, trigger workflows, or interact with external tools?
 - Are any actions customer-facing, hard to reverse, or operationally significant?
- **How much autonomy does it have?**
 - Does it follow a fixed workflow, or can it choose its own path?
 - What decisions require human review or approval?
 - Are there limits on tool use, data access, or downstream actions?
- **What systems and tools are connected?**
 - What internal systems, APIs, databases, or third-party applications can it reach?
 - Are those connections limited by role, use case, or least-privilege principles?
 - Could an error in one system cascade into another?
- **How is it governed and monitored?**
 - Who owns the use case internally?
 - Is there a clearly identified business owner, technical owner, and escalation path?
 - How was the system tested before deployment?
 - What logging, traceability, and auditability exist?
 - Can the business reconstruct what happened if there is a complaint, failure, investigation, or regulator inquiry?

- **What happens if it is wrong?**
 - What is the fallback if the system hallucinates, acts on bad data, or fails to complete a task correctly?
 - Is there a clear escalation path?
 - Who is accountable for remediation, communications, and any necessary legal response?

5. **Develop practical and effective governance for Agentic AI systems.** Agentic AI should be treated as a governed business system, not merely a more capable chatbot.

- At a minimum, businesses creating **governance policies for Agentic AI systems** should address:
 - Approved and documented use cases;
 - Least-privilege access;
 - Limits on tool use and system permissions;
 - Human review for sensitive or high-impact actions;
 - Testing before deployment;
 - Logging, traceability, and monitoring;
 - Clear ownership and escalation paths; and
 - Periodic review as the system, its tools, or its connected data sources change.
- **Continue legal review and oversight of Agentic AI systems after the initial deployment.** Material changes to connected tools, permissions, use cases, or system autonomy may change the risk profile and warrant renewed review.
- **Address logging, traceability, and record retention policies early,** not after a problem arises
 - If the business later needs to explain, defend, investigate, or remediate the system's conduct, those capabilities will often determine whether the response is manageable.

Learn
more

- [Advising on Agentic AI: A Simple Framework for In-house Counsel](#), by Christopher Wlach, *Corporate Counsel Now*, February 4, 2026
- Visit the ACC [AI Center of Excellence for In-house Counsel](#).



Learn 10 Key Terms in Agentic AI Systems

1. **Guardrails.** Controls that limit what a system can access, decide, or do. They can include permissions, approval rules, policy checks, monitoring, and escalation triggers.
 - **Example:** A customer support agent can draft a refund email, but it cannot issue the refund without human approval.
2. **Human-in-the-loop / Human-on-the-loop.** Oversight models in which a human either reviews actions before they occur, or monitors the system and intervenes by exception.
 - **Example:** In one workflow, a lawyer must approve a contract escalation before it is sent. In another, the system sends routine escalations automatically and only flags unusual cases for review.
3. **MCP (Model Context Protocol).** An open protocol for connecting AI applications to external tools and data sources in a standardized way.
 - **Example:** An internal assistant uses MCP to connect to a document repository, ticketing system, and calendar through one common interface.
4. **Memory.** Information the system retains and reuses across steps or sessions, such as prior instructions, retrieved facts, or earlier actions.
 - **Example:** An HR assistant remembers that an employee's request concerns parental leave and continues using the relevant policy materials during later steps.
5. **Multi-agent system.** A system in which multiple specialized agents work together to complete a broader task.
 - **Example:** One agent gathers facts, a second drafts the response, and a third checks the result against company policy.

» WITH THANKS TO:

*Gregory P. Silberman,
Counsel, at Kilpatrick,*

for contributing
this checklist.

- 6. Orchestration.** The logic that coordinates steps, tools, handoffs, and approvals so the system can pursue an objective.

 - **Example:** The system checks the CRM, searches the knowledge base, drafts a response, and then requests approval if the issue meets an escalation threshold.
- 7. Planning.** The process of breaking an objective into steps and adjusting those steps based on intermediate results.

 - **Example:** To resolve a billing complaint, the system verifies account details, reviews prior tickets, drafts a response, and escalates if the amount exceeds a set threshold.
- 8. Scaffold/Scaffolding.** The surrounding software structure that makes an AI model work as part of an agentic system rather than as a one-step chatbot.

 - **Example:** A contract review tool uses prompts, playbooks, tool access, memory, and approval rules so the model can perform a repeatable review process.
- 9. Tool use/Tool calling.** The ability of a system to invoke external tools, APIs, databases, or software instead of only generating text.

 - **Example:** Rather than merely suggesting that a support ticket be opened, the system calls the ticketing API and opens the ticket directly.
- 10. Workflow.** A defined sequence of steps for completing a task, usually with predetermined rules, checkpoints, or handoffs.

 - **Example:** A compliance workflow collects intake details, checks policy, drafts a response, and routes certain matters to legal.



Assess Agentic AI Capability and Oversight

Use this Agentic AI Capability and Oversight Spectrum chart to evaluate AI systems based on what they can access, what they can do, and how much human oversight remains in place. It is designed to move the discussion away from vague labels and toward the operational features that most often matter for legal, compliance, privacy, security, and governance review.

» How to use the Agentic AI Capability and Oversight Spectrum:

- **Distinguish** low-action prompt/response tools from systems with broader delegated action.
- **Align** legal, technical, and business stakeholders on the level of oversight a proposed use case requires.
- **Trigger** Governance: as a system gains broader access, autonomy, or cross-system action capability, review should become more formal.

Level	Label	What the AI System Does	What the Human Does	Example (Illustrative)
0	Prompt/Response (Non-Agentic)	Generates content, analysis, or recommendations; does not directly act in business systems	Human initiates, reviews, and executes all actions	Internal policy Q&A chatbot
1	Tool-Using Copilot	Uses limited tools or data in a user-directed manner; assists with retrieval, drafting, or recommendation; does not independently complete a multi-step workflow	Human remains primary operator and approves any execution	Copilot that retrieves internal data and drafts a response
2	Workflow Agent (Bounded Multi-Step)	Executes a defined multi-step task within a bounded workflow or domain	Human sets objective and reviews check points or final output	Agent that triages tickets and prepares draft resolutions
3	Cross-System Agent (Delegated Action)	Takes or coordinates actions across multiple systems, such as updating records, opening tickets, or sending follow-up communications, within defined permissions and rules	Human sets authority, approval rules, and escalation thresholds; oversight may be approve-all or approve-exceptions	Customer success agent that updates CRM, opens tickets, and sends approved follow-up emails
4	High-Autonomy Orchestrator	Coordinates multiple tools, workflows, or specialized agents across more complex operations with broader delegated authority and exception-based oversight	Human monitors by exception, governs limits, and intervenes when needed	Operations assistant coordinating specialized systems across support, CRM, and internal workflows

AI Thought Leadership and Resources by Kilpatrick

1. *A Trademark Practitioner's Guide to Using AI: Guidelines, Use Cases, and Ethical Considerations*,
Takeaways by Kilpatrick, March 11, 2026
2. *Anthropic Challenges DoW's Supply Chain Risk Designation*,
Article by Kilpatrick, March 11, 2026
3. *The Governance of Innovation: Managing New AI, IP, and Regulatory Frameworks*,
Takeaways by Kilpatrick, February 17, 2026
4. *The Current AI Landscape in Trademarks and Copyright*,
Takeaways by Kilpatrick, February 5, 2026
5. *AI Data Centers and the Looming Energy Crisis in the United States*,
Article by Kilpatrick, January 15, 2026
6. *From Transparency to Oversight: New York's RAISE Act Raises the Bar for Frontier AI Developers*,
Article by Kilpatrick, January 12, 2026,
7. *New York Governor Signs Bill Requiring Disclosure of AI-Generated Performers in Advertising*,
Article by Kilpatrick, December 15, 2025
8. *Executive Order Seeks to Preempt Certain State AI Laws and Establish a National AI Framework - Initial Takeaways*,
Article by Kilpatrick, December 12, 2025
9. *USPTO Issues Revised Inventorship Guidance, Eliminating Separate Standard for AI-Assisted Inventions*,
Article by Kilpatrick, December 12, 2025
10. *Corporate Transactions: Deal Structuring, Governance Documents, and the Impact of AI*,
Takeaways by Kilpatrick, December 1, 2025
11. *Genesis Mission Artificial Intelligence Executive Order – President Trump Establishes National AI Platform and Objectives for Issues of National Significance*,
Article by Kilpatrick, November 25, 2025
12. *Trade Secrets Would Likely See Court Protection From GenAI*,
Article by Kilpatrick, May 22, 2025
13. *Kilpatrick's Generative AI Insights Hub*
14. *Kilpatrick Connect: Your Trusted AI Partner*

Additional Insights on AI Topics

1. *ACC AI Center of Excellence for In-house Counsel*
2. *Leveraging GenAI for Business Impact: A Playbook for Small Law Departments*,
by ACC, March 17, 2026
3. *Governing AI Inside the Enterprise: A Practical Framework for In-house Counsel*,
by Alexander Lima, *Corporate Counsel Now*, February 25, 2026
4. *Why Legal's Cross-functional Role Is the Key to Scalable Enterprise AI*,
by Michael T. Moore and Rebecca Hinds, *Corporate Counsel Now*,
February 2, 2026
5. *AI Use Case: Self-Service for Low-Complexity Business Agreements*,
by Hayden O. Creque, January 25, 2026
6. *Why Legal Should Be at the Center of Your Organization's AI Strategy*,
by Michael T. Moore and Rebecca Hinds, *Corporate Counsel Now*,
January 21, 2026
7. *Top 10 GenAI Transparency & Readiness Questions for Outside Counsel*,
by ACC, November 4, 2025
8. *The AI-Empowered Counsel: A Practical Framework for In-house Legal Teams*,
by Hayden O. Creque, *Corporate Counsel Now*, October 18, 2025
9. *AI-Augmented Contract Negotiation: Legal Strategy in the Age of Machine Collaboration*,
by Edward T. Paulis III, *Corporate Counsel Now*, September 30, 2025
10. *The Transformative Role of Artificial Intelligence (AI) in Cybersecurity*,
by Ankura Consulting Group, LLC, July 17, 2025

Sample Tools, Forms, and Checklists

1. [Sample Use of Generative AI Policy Template](#)
2. [Sample AI Guidelines for Outside Counsel](#)
3. [Using AI in HR Decisions: Tips for In-house Counsel](#)

More Resources

1. **Visit the [ACC Resource Library](#).** Find thousands of resources:
 - Sample contracts and policies
 - Checklists and articles
 - ACC in-house survey reports
2. **Explore [ACC Curated Collections](#):**
 - [ACC AI Center of Excellence for In-house Counsel](#)
 - [Privacy and Cybersecurity](#)
 - [ESG \(Environmental, Social, and Governance\)](#)
 - [Teaching Law School](#)
 - [New to In-house](#)
 - [Navigating the Current US Administration](#)

Connect with Peers and Boost Your AI Skills

» Connect with the ACC IT, Privacy & e-Commerce Network

- Monthly **live discussions** with other in-house counsel and outside speakers.
- A dedicated **e-forum** to connect with peers and seek their insights.
- [Join the Network!](#) For ACC members only.
- Not an ACC member? [Join ACC](#).

» Explore trends and tips at the ACC Annual Meeting

- **Enjoy** the largest annual gathering of the global in-house community.
- **Engage** in lively sessions and learn the latest in-house strategies, tips, and trends.
- **Bring** your in-house team. Get CLE/CPD credits.
- Enjoy the **unique experience** of our vibrant ACC in-house community.
- [Learn more](#)

» Learn strategies at the ACC Foundation Cybersecurity Summit

- Learn about the latest cyber threats and innovations for in-house.
- Participate in deep-dive sessions with leading cybersecurity and in-house professionals.
- Discuss emerging issues, and tips for preventing and responding to data breaches.
- [Learn More about ACC Foundation events](#)

Disclaimer: No information in this toolkit shall be construed as a recommendation or endorsement of any vendor or tool by ACC. Microsoft, Microsoft 365, Power Apps, Power Automate, Power BI, are trademarks of the Microsoft group of companies. Westlaw® is a trademark of Thomson Reuters or its affiliates. LexisNexis® is a trademark of RELX Inc. This toolkit is a publication of the Association of Corporate Counsel and is neither affiliated with, nor authorized, sponsored, or approved by, Microsoft Corporation, Thomson Reuters, or RELX Inc.