Wolters Kluwer

Strategic Perspectives

Ethical guidelines for use of AI explored at ACC 2026 Annual Meeting

By Anne H. Gibson, J.D., LL.M.

AI use in the legal profession is quickly expanding, to the point that AI is no longer seen as a bonus but rather as the status quo. However, adoption of ethical guidelines and oversight continues to lag behind the technology. These trends and how to address them were explored in the session "AI, Ethics, and the In-house Lawyer: Guardrails, Governance, and What Comes Next," at the Association of Corporate Counsel 2026 Annual Meeting on October 6. The panelists explored how firms can safely utilize AI tools and the pitfalls they should watch out for.

Existing ethical rules apply to AI

The panelists, Autumn George, Director, Galloway Johnson Tompkins Burr & Smith; Katherine Currie Fallas, Director, Galloway Johnson Tompkins Burr & Smith; Aileen Schwartz, General Counsel & Chief Legal Officer, QC Servion, LLC; and Anthony Begon, Vice President, Legal - Enterprise Risk & Compliance, Trellix, started their discussion by pointing out that lawyers are not operating with an entirely blank slate in this area. Existing ethical rules, such as the ABA Model Rules, apply to the use of AI as they do to other resources.

In particular, they pointed to the following ABA Model Rules:

- Rule 1.1: Competence
- Rule 1.6: Confidentiality of Information
- Rule 3.3: Candor Toward the Tribunal
- Rule 5.1 & 5.3: Supervision & Responsibility
- Rule 8.4: Misconduct (misrepresentation)

The panelists note that the rule on competence requires technical competence with the tools the attorney is using. In the case of AI, while an individual does not need to be a computer scientist or know how to code, he or she does need to know enough to understand the benefits, limits, and risks of the tool being used. The panelists noted that adequate competency can be key to avoiding some of the other problem issues with AI. Better understanding of how the tools work and the way they access and store data can help a lawyer avoid misusing them.

Confidentiality is another crucial concern for lawyers using AI. The same rules that require lawyers to keep client and customer data private apply when information is provided to an AI tool, just as they would in other cases of sharing information with a third party. Lawyers must be sure that confidentiality is maintained. The panelists encouraged attendees to complete a quick poll during the session, and the results showed that the attendees' top concern when it came to AI was confidentiality, with 60% voting for that, while concerns about hallucinations, such as fake case citations, came second at 38%. This shows a shift from the outset of AI usage several years ago.

Concerns about hallucinations remain important, however. The ethical rules around candor towards the tribunal relate directly to inaccurate information provided by AI. When attorneys rely on what an AI tool provides without confirming its accuracy, they may find themselves inadvertently lying to the court. This also ties into the rules on supervision, which apply to AI tools as they would to other attorneys or staff members, and to misconduct, which includes misrepresentation.

AI pitfalls for companies

The panel next discussed the major areas where companies can run into trouble when using AI, areas that the guardrails outlined later can protect against. The first they discussed was using AI indiscriminately in human resources functions, for instance in screening resumes for open positions. The panelists pointed out that AI tools are notorious in these circumstances for biased thinking. If an AI tool disproportionately screens out members of protected classes, for example, it not only can cause the company to miss out on qualified candidates, it can also open the company up to discrimination claims.

Privacy issues and data leaks, as well as hallucinated content, were also key concerns. The panelists discussed the Canadian case of *Moffatt v. Air Canada*, 2024 BCCRT 149 (Feb. 14, 2024). That case involved a chatbot on the Air Canada website that gave a customer inaccurate information about bereavement fares. The customer relied on the information and later sued when the promised discounted fare was not available. Although Air Canada argued that the chatbot was

© 2026 CCH Incorporated and its affiliates. All rights reserved. Reprinted with permission from CCH Incorporated.

"responsible for its own actions," the tribunal member in the case did not agree and found that Air Canada "is responsible for all the information on its website. It makes no difference whether the information comes from a static page or a chatbot." If a company offers and relies on an AI tool in its customer interactions, it will be responsible for the tool's actions, just as it would be with a human employee.

Another issue that the panel discussed also had to do with the use of AI tools created by third parties. IP infringement and ownership confusion can often occur when an AI vendor is used. If a unique tool is developed for a company and then undergoes continuing improvements and learning, who does it belong to? Who do the outputs of the tool belong to? Could the developer retain access to the outputs and use them for training other tools, unrelated to the company? These questions would need to be answered in the contracts, but that too can prove tricky, as the panelists discussed later.

Finally, a big pitfall for companies using AI is a lack of documentation for their AI governance plans. The panelists noted that a surprising number of companies are afraid to have something in place at all. However, the panel stressed that creating a concrete, straightforward policy that employees can follow is crucial, and at this point in the adoption of AI, it is irresponsible not to have one. They pointed out that the ACC offers sample plans to members.

AI pitfalls for lawyers

Lawyers themselves also face issues in using AI. The top issue for individuals, according to the panelists, continues to be blind reliance – accepting AI output without question. This is often followed by a refusal to correct such mistakes when found. As one panelist said, the worst thing an attorney can do if outside counsel challenges a case citation, for example, is to double down without actually checking it themselves. The panel noted that AI outputs in signed filings in particular need to be scrutinized carefully before they are submitted. Signing certifies that there was reasonable inquiry and review. If there are mistakes, the lawyer could be personally liable.

Another issue lawyers face, whether in-house or outside counsel, is getting pushed by the business side or by clients to take positions that are not supported by the law. While this has always been an issue, the use of easily accessible AI-search tools seems to have empowered nearly everyone to feel that they can do their own legal research. "Everyone thinks they're a lawyer now," one member of the panel quipped. Deflecting these requests while maintaining good relationships can sometimes be difficult.

Six Questions to Ask Before Using AI

If you can't answer one of these, pause before proceeding.

1. What is the task?
2. What data will I provide?
3. What tool am I using?
4. What rules or client limits apply (e.g., contractual)?
5. How will I verify the output?
6. Should I disclose or document the use?

Practical guardrails for in-house counsel

The panel then presented its six guardrails for AI use for in-house counsel, which can also apply to any case of using AI in a legal setting.

Guardrail 1 – AI governance

The panel pointed out that having a robust AI governance plan in place is one of the most important steps a company can take in reducing its risks related to AI usage. The first step, the panelists emphasized, is ensuring that the board of directors and the leadership of the company are aware of the risks inherent in using AI. Having buy-in from leadership on the project of setting up a system to govern AI use is essential.

The next step is to put together a Governing Committee to oversee and take responsibility for the AI plan. Its first duty would be creating an "AI inventory" of all the AI products that the company is using, who has access to them, what the approved uses are for each, and so forth. The Committee would draft the company's AI policies and standards, discussed further below. This would include defining high-risk scenarios that would require additional oversight. Companies should also consider how long they are retaining the information they get from AI. Many companies do not have a policy on this, the panel noted, but it has big downstream implications for data privacy and preventing breaches. Finally, the plan should provide for periodic audits.

A common AI governance challenge is ensuring that contracts for AI products are consistent with a company's internal AI policies. The panelists noted that with AI in high demand and every company worried about being behind the curve, it can be difficult to negotiate with the providers. Typically, the contracts are presented in a take-it-or-leave-it fashion, and there is not a lot of ability to change the terms.

Another important task is ongoing management of AI usage, after the policies are in place. This includes creating concrete plans for monitoring, mitigation of issues, and incident response. If a breach in data

© 2026 CCH Incorporated and its affiliates. All rights reserved. Reprinted with permission from CCH Incorporated.

privacy occurs, for example, the company needs to have a plan in place to remediate it.

Guardrail 2 - AI policy

The AI policy is an important guardrail in using AI. The panel laid out some basic provisions such a policy should contain, including: disallowing any confidential data to be used in unsupported tools; requiring human verification of AI outputs; and requiring approval from the legal department for high-risk AI uses. High-risk triggers should include any circumstances in which AI plays a role that could negatively impact people's jobs, finances, housing, safety, and so forth.

The AI policy should also include provisions for meaningful supervision. A good explanation of this, the panel stated, is ABA Formal Opinion 512, which describes how supervisory principles that apply to other humans also apply to AI tools and emphasizes the responsibility of the supervising lawyer.

What To Do After an AI incident

- Stop using the output & preserve it
- Report to supervisor/AI Governance Committee immediately
- Identify the tool, prompt, date used, etc.
- Assess whether there was any breach of confidentiality, privacy, etc.
- Correct the record promptly, if applicable

Guardrail 3 - Working with AI vendors

The next guardrail relates to working with AI vendors. Whether a company is using an existing AI tool or working with a developer to create a bespoke solution for the firm, the company will be dealing with a third-party vendor. The information that the company puts into the tool, as well as its outputs, will potentially be accessible by the vendor. Consequently, it is important that the company has policies in place to protect itself and maintain its responsibilities to clients and customers. These include policies around due diligence and vetting when choosing the vendor, and required disclosures for the vendor regarding access to or use of the relevant data. That said, the panel emphasized that the contract controls in these situations, and contract negotiations can be difficult.

Guardrail 4 - Bias & fairness controls

As mentioned above, bias in AI systems can be a problem, especially in human-resources matters. To protect against this, the panelists suggested a "Bias Control" checklist when deploying AI tools to handle human resources tasks. The checklist includes: pre-deployment testing of the tool; post-deployment monitoring by responsible persons; a mechanism to flag complaints and investigate them; retesting of the tool at regular intervals; and, finally, consistent documentation of mitigation steps that are taken. At the end of the day, the panel said, the key is human oversight. There should be people keeping an eye on the system and taking ultimate responsibility.

Guardrail 5 – IP rights & ownership

One question the panel posed to the audience was, "Who owns AI-generated work product?" Is it the company? The vendor? The panel's answer was, "No one knows." There is not a single answer here, but rather, the answer should be in the relevant contract. Despite the difficulty of negotiating these contracts in the current context, companies should work to ensure that anything containing confidential client information or other personal data should belong to the company, and not to the vendor. For example, does the vendor own the outputs of the tool? Could the vendor use those outputs, which include client data, to train future models, potentially passing the sensitive data on? Ultimately, the contract controls, and this should be something companies scrutinize carefully.

Guardrail 6 - Insurance & incident response

The final guardrail centered on protecting against the eventuality of problems occurring. The panel stressed the importance of ensuring that the company's insurance was updated to cover any issues arising from AI tool use. It also referred back to the incident response component of the AI governance plan. Having a system in place and responsible individuals to contact in the event of a data breach or a hallucination being submitted to the court, for example, is very important to keep the problem from growing larger. Ensuring that employees understand the incident response plan and what to do is crucial.

What's coming in 2027?

The panelists closed by discussing what they thought was coming in the world of legal AI in 2027. They predicted the development of much more targeted systems and processes. Rather than off-the-shelf products, they predicted that everyone would be using tailored solutions for their workflows.

Additionally, they thought that the question that clients would be asking would no longer be, "Are you using AI?" but rather, "Can you explain how you're using AI?" A greater depth of knowledge about the systems, and how and why they are used will be expected.

5 Step Action Plan

What every company should do as best practices for AI use.

1. Inventory AI tools
2. Approve AI policy
3. Update vendor templates
4. Enforce human verification
5. Train teams

© 2026 CCH Incorporated and its affiliates. All rights reserved. Reprinted with permission from CCH Incorporated.