



Apr 29, 2026

Categories:

[Affordable Housing](#)
[Artificial Intelligence at FBT](#)
[Data Centers](#)
[Publications](#)

Authors:

[Colby Roberts](#)
[Mason C. Clutter](#)

Generative AI in Business and Law: Protecting Privilege in the Age of AI

In today's world, artificial intelligence, and now generative artificial intelligence (GenAI), is everywhere. Many of us use it as a convenient tool — to draft emails, plan trips, or even choose a restaurant for date night. Generative AI tools like ChatGPT and Claude can be extremely useful, and companies should be thinking about how to adopt and responsibly use them in their business.

While these tools are powerful and convenient, they are not a substitute for legal advice. Generative AI is not a lawyer — and most importantly, it is not your lawyer.

A recent federal case highlights yet another example of the law failing to keep up with technology and how the use of GenAI must be in a purposeful and responsible manner.

In *United States v. Heppner*, No. 25 CR. 503 (JSR), 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026), the United States District Court for the Southern District of New York held that documents a defendant created using a public AI platform (in that case, Claude) were not protected by the attorney-client privilege or the work product doctrine.

Before his arrest, the defendant used a public GenAI tool to research issues related to a government investigation and generated a series of documents based on those prompts. In doing so, he input strategic information and facts that he later shared with his attorneys in privileged conversations, including his version of the facts and his understanding of the law. He ultimately provided the AI-generated materials to his attorneys.

After the government seized his devices, it discovered the AI conversations and documents and sought to use them in the case. The defendant argued that both his inputs and the AI's outputs were protected by the attorney-client and work product privileges. The court disagreed and allowed the government to access both the documents and the underlying AI communications (inputs and outputs).

The court's reasoning was straightforward:

- Communications with an AI tool are not communications with a lawyer.
- Information shared with a public third-party platform is generally not confidential.

The court specifically examined the public AI platform's user agreement, which allowed for the collection and potential disclosure of user inputs and outputs to third parties, including regulators. Because the defendant voluntarily shared information with a third party under those terms, the court concluded there was no reasonable expectation of confidentiality.

Notably, the defendant used a publicly available AI tool, rather than a closed or enterprise AI platform, where, through underlying terms with the provider, the organization may retain control over its data and the provider is contractually prohibited from accessing, sharing, or disclosing user inputs and outputs. Had the defendant used such an enterprise tool, the court's analysis regarding the expectation of confidentiality may well have reached a different conclusion.

The court also rejected the argument that the defendant was seeking "legal advice" in a way that would trigger privilege protections. The AI tool itself included disclaimers (e.g., "I'm not a lawyer..."), therefore, the court found that the defendant should not have expected "legal advice" from the tool.

Importantly, the court noted that the defendant could not "fix" the issue by later providing the materials to his attorneys. Privilege does not attach retroactively. And to the extent any otherwise privileged information was included in his AI prompts, the court found that privilege was waived when he disclosed the information to the AI platform – just as it would have been if he had shared the information with any other third party.

In contrast, one recent decision found that documents created with or by public generative AI by a pro se party in anticipation of litigation may qualify for work product protection. In part, the court noted that AI platforms are "tools, not "persons." See *Warner v. Gilbarco, Inc.*, No. 2:24-CV-12333, 2026 WL 373043, at *5 (E.D. Mich. Feb. 10, 2026). While the facts may differ, it remains unclear what

GenAI is from a privilege perspective: software used by a client or a third party.

So, what do these seemingly conflicting opinions mean? In short, do not assume that documents created using a public AI tool will be protected, even if prepared in anticipation of litigation. Enterprise tools with appropriate confidentiality safeguards are more likely than publicly available tools to preserve work product protection. And when in doubt, consult your human lawyer.

What does this mean for your business and its use of generative AI? Here are a few practical takeaways:

- **Adopt internal guidelines.** Implement clear policies governing employee use of GenAI tools, especially when dealing with privileged legal or other confidential information, including confidential information from or about your customers. Policies should address both publicly available tools and enterprise tools and clarify who may use GenAI and for what purposes. While it may be tempting to use GenAI for any reason, establish clear use cases and guardrails to ensure that you are using GenAI in a meaningful, responsible, and transparent manner. Until the law is further settled, we must rely on our own internal governance frameworks to ensure that we are safeguarding our confidential, trade secret, intellectual property, and privileged information. For example, see this [recent FBT Gibbons article](#) on the use of GenAI in the employment law context.
- **Treat publicly available AI tools as you would any outside communication.** Your existing policies and procedures regarding sensitive, confidential, trade secret, and privileged information should also apply to your use of a publicly available AI tool. If the information should not be disclosed outside of the organization, then do not input it into a publicly available AI tool.
- **Do not expect GenAI to provide privileged legal advice.** Even when using a closed, enterprise AI tool, asking the AI for legal advice does not create attorney-client privilege, because the AI is not a lawyer. This is true regardless of how

secure the platform is or what safeguards may attach. If you use GenAI to prepare documents in anticipation of litigation, those materials may qualify for work product protection, though work product protections are narrower than attorney-client privilege and vary by jurisdiction. Your attorneys can advise on whether and how to use GenAI safely or handle the AI task within protected channels.

As generative AI continues to evolve, so too will the legal landscape around privilege and confidentiality. Taking proactive steps now can help protect your business and preserve your legal options down the road.

At FBT Gibbons, we are leaning into the adoption and use of generative AI to help our clients realize its efficiencies while freeing our lawyers to focus on more complex and strategic legal issues. We also help clients evaluate appropriate tools to achieve their goals and craft approaches to support innovation and responsible use of AI that are designed to preserve privilege and confidentiality while building and maintaining trust with their own clients.

© 2026 FBT Gibbons LLP. May be considered attorney advertising in some jurisdictions.