AI and Attorney-Client Privilege

How Private AI Deployment Protects Legal Privilege in the Wake of United States v. Heppner

A White Paper for Law Firms and Legal Departments
March 2026

SecureLLMs.org | michael@securellms.org

Executive Summary

On February 10, 2026, Judge Jed S. Rakoff of the United States District Court for the Southern District of New York issued a landmark ruling in United States v. Heppner that sent shockwaves through the legal profession. The court held that documents a criminal defendant generated using a public AI chatbot were not protected by attorney-client privilege or the work product doctrine, because the information had been disclosed to a third-party platform with no expectation of confidentiality.

For law firms and corporate legal departments, the implications are immediate and far-reaching. Every time an attorney or client enters case details, legal strategies, or privileged communications into a consumer AI platform like ChatGPT, Gemini, or a public instance of Claude, that information is transmitted to third-party servers, stored according to the provider's terms, and potentially used to train future AI models. Under the reasoning in Heppner, such disclosures may waive privilege entirely.

This white paper examines the Heppner ruling and the broader regulatory landscape surrounding AI and legal privilege. It then explains how deploying AI on private, self-hosted infrastructure—such as the platform provided by SecureLLMs.org—can allow law firms to harness the power of generative AI while maintaining the confidentiality that privilege requires.

The Heppner Decision: A Turning Point

Background

In October 2025, Eric Heppner was indicted on multiple federal charges, including securities fraud, arising from his time as an executive at several corporations. After receiving a grand jury subpoena and learning he was a target of the investigation, Heppner turned to a publicly available generative AI platform on his own initiative—not at the direction of his attorney—and used it to generate reports outlining potential defense strategies, legal arguments, and factual analyses.

When the FBI executed a search warrant and seized Heppner's electronic devices, they discovered these AI-generated materials along with the underlying chat logs. Heppner's defense team asserted attorney-client privilege and work product protection over the materials, arguing they were created for the purpose of communicating with counsel to obtain legal advice.

The Court's Ruling

Judge Rakoff rejected both privilege claims. The court's analysis focused on three critical elements of the attorney-client privilege, finding that at least two—and potentially all three—were not satisfied:

1. **Not a communication with counsel.** The AI tool was not an attorney, and Heppner was not consulting it at the direction of his lawyer. The communications were between a person and a software platform, not between a client and counsel.
2. **Not confidential.** Heppner used a public, consumer version of the platform. The provider's privacy policy explicitly stated that it collects user inputs and outputs, may use them for

model training, and reserves the right to disclose data to third parties—including government authorities. Under these terms, no reasonable expectation of confidentiality existed.

3. **Not for the purpose of obtaining legal advice from an attorney.** Heppner acted on his own initiative, not at his lawyer's direction. Even if the information he entered had once been privileged, sharing it with the public platform constituted a waiver.

The court also held that the work product doctrine did not apply, because the materials were not prepared by or at the direction of an attorney.

> **Key Takeaway from Heppner**
> *The court's reasoning hinged on the public, non-confidential nature of the AI tool and the absence of attorney involvement. This strongly implies that the outcome could differ if privileged information were processed within a private, confidential AI environment under the direction of counsel.*

Why This Matters for Every Law Firm

The Privilege Problem with Public AI

The Heppner decision did not create new law—it applied longstanding privilege principles to a new technology. The core issue is simple: attorney-client privilege requires that communications be kept confidential. When information is entered into a public AI platform, it is disclosed to a third party, and the confidentiality element is destroyed.

This is not a theoretical risk. Consider how law firms commonly use AI today:

- Associates paste deposition transcripts into ChatGPT to generate summaries and identify key admissions.
- Partners ask AI tools to analyze contract clauses and flag potential issues in pending transactions.
- Litigation teams use public AI to draft motions, outline legal arguments, or research case strategies.
- In-house counsel upload confidential corporate communications for AI-assisted review.

In each of these scenarios, privileged or work-product-protected information is leaving the firm's control and landing on servers owned and operated by a third-party AI provider. Under Heppner's reasoning, any privilege attached to that information may be waived the moment it is entered into the platform.

The Regulatory Environment

The Heppner ruling arrives in a regulatory environment that is already tightening around AI use in professional contexts:

- **ABA Formal Opinion 512 (July 2024):** The American Bar Association's first formal guidance on generative AI in legal practice emphasizes that attorneys must understand the security risks of AI tools they use. The opinion specifically addresses the duty of confidentiality under Model Rule 1.6, stating that lawyers must assess whether AI platforms could result in the disclosure of client information to unauthorized parties. It also confirms that client informed consent is required before entering confidential information into self-learning AI tools.
- **Model Rule 1.1 – Competence:** Attorneys must understand the technological tools they use, including the risks AI platforms pose to client data. Ignorance of how a public AI tool handles data is not a defense.
- **Model Rule 1.6 – Confidentiality:** Lawyers must make reasonable efforts to prevent the unauthorized disclosure of client information. Using a public AI tool whose terms of service permit data collection, retention, and sharing with third parties may fall short of this standard.
- **State-Level Requirements:** Multiple state bar associations—including California, Florida, New York, New Jersey, and Texas—have issued their own ethics opinions or guidance on AI use, with most emphasizing confidentiality protections and the need to vet the security of any AI tool used in legal practice.

The Competitive Intelligence Risk

Beyond privilege, public AI platforms present a competitive intelligence problem for law firms. When attorneys at different firms enter case strategies, legal analyses, and client details into the same public AI service, they are contributing to a shared pool of training data. There is no guarantee that insights from one firm's use of the platform will not influence the outputs delivered to opposing counsel at another firm.

As the Covington & Burling analysis of Heppner notes, the use of publicly available AI tools may be viewed as a disclosure to third parties that can undermine, weaken, or waive claims of privilege. Firms that fail to address this risk may find themselves in an increasingly untenable position—unable to use the most powerful tools available while competitors who have adopted private solutions pull ahead.

The Private AI Solution: How Self-Hosting Preserves Privilege

How It Works

Self-hosted AI means running large language models on infrastructure your firm controls, rather than sending data to public providers. It is the difference between storing confidential case files in your own secure filing system and handing them to a third party who reserves the right to read them, share them, and learn from them.

With a private AI deployment, every interaction—every prompt, every document uploaded, every response generated—stays within your firm's secure environment. No data is transmitted to external servers. No public AI provider has access to your information. No terms of service grant the provider rights to your inputs or outputs.

Why Private AI Preserves the Privilege Framework

The Heppner court's analysis provides a clear roadmap for understanding why private AI deployment addresses each of the privilege concerns the ruling identified:

Privilege Element	Public AI (Heppner)	Private AI (SecureLLMs)
Confidentiality	Data sent to third-party servers; privacy policy permits collection, training, and disclosure to others	Data never leaves firm-controlled infrastructure; single-tenant isolation
Third-Party Disclosure	Information shared with AI provider, a third party; provider may share with subcontractors and regulators	AI runs entirely within firm's private environment behind encrypted VPN tunnels
Training Data Risk	Provider may use inputs to train models; privileged information could surface in other users' outputs	Models do not learn from user inputs; no training on client data; complete data sovereignty
Data Sovereignty	No visibility into where data is stored, for how long, or who can access it	Full control over data location, retention, and access; firm maintains complete audit trail

Baker McKenzie's analysis of the privilege implications of AI confirms this approach, noting that lawyers can mitigate the risk of privilege waiver by ensuring AI models maintain strict confidentiality—specifically by deploying them in a closed, internal environment behind a firewall and implementing controls that prevent the model from learning from user inputs.

How SecureLLMs Deploys Private AI for Law Firms

SecureLLMs.org provides private AI infrastructure on demand, delivering the power of frontier language models—including offline versions of ChatGPT, Claude, and over 100 other models—in a secure, isolated environment where your data never touches the public internet.

Architecture Built for Confidentiality

- **Single-Tenant Isolation:** Every SecureLLMs deployment is a dedicated, isolated environment. Your firm's data is never co-mingled with other organizations' data on shared servers. This is fundamentally different from public AI platforms, where millions of users share the same infrastructure.
- **Encrypted VPN Tunnels:** SecureLLMs establishes dual-channel encrypted VPN connections between your firm's network and your private AI environment. Built on enterprise-grade IPsec VPN technology and backed by NIST SP 800-77 standards, these tunnels ensure that all model communications travel under the public internet through dedicated encrypted pathways that no outside party can access.

- **Zero Data Leakage:** Your prompts, documents, case files, and AI-generated outputs never leave your private environment. The models perform all inference computation entirely disconnected from the public internet. No data is transmitted to AI model providers.
- **No Training on Your Data:** Unlike public AI services that reserve the right to use your inputs to improve their models, SecureLLMs deployments do not train on your data. Your privileged communications, legal strategies, and client information remain yours alone.

Deployment Options for Law Firms

SecureLLMs offers three tiers of private AI infrastructure, each designed to meet different firm sizes and security requirements:

Feature	Professional	Business	Enterprise
Monthly Cost	$2,000 + tokens	$5,000 + tokens	$10,000+ + tokens
Users	Up to 10	Up to 50	Unlimited*
AI Models	3 models	All 100+ models	All 100+ models
Encryption	HTTPS/TLS	VPN Client	IPSec Site-to-Site
Authentication	Federated	Federated	OAuth, OIDC, Okta
Best For	Small firms, solo practitioners	Mid-size firms, practice groups	Large firms, corporate legal

**Enterprise Tier: $100 per user, per month over 100 users. All tiers include a one-time installation fee equivalent to one month of the selected tier.*

Capabilities That Enhance Legal Practice

Beyond security, SecureLLMs provides capabilities that make AI more useful for legal work than any consumer platform:

- **Secure Knowledge Bases (RAG):** Upload your firm's case files, precedents, internal memos, and legal research into a private knowledge base. The AI can then reference this material when generating responses, providing context-aware answers grounded in your firm's own work product—without any of that material ever leaving your environment.
- **Custom AI Agents:** Build specialized AI agents tailored to your firm's practice areas—a contract review agent trained on your firm's templates, a litigation research agent that knows your preferred arguments, or a regulatory compliance agent customized to your clients' industries.
- **Model Flexibility:** Access over 100 frontier language models from a single platform. Different practice areas can use different models optimized for their specific needs, all within the same secure infrastructure.
- **Role-Based Access Controls:** Assign granular permissions by attorney, practice group, or matter. Ensure that only authorized personnel can access specific case materials and AI interactions.

Practical Guidance: Implementing a Privilege-Safe AI Policy

The Heppner ruling, combined with ABA Formal Opinion 512 and emerging state-level guidance, provides a clear framework for what law firms should do now. The following steps represent a comprehensive approach to maintaining privilege in the age of AI:

1. Prohibit the Use of Public AI for Privileged or Sensitive Matters

Establish a firm-wide policy that no privileged communications, work product, client information, case strategies, or confidential business information may be entered into any public or consumer-grade AI tool. This is the most direct lesson from Heppner: public AI means third-party disclosure, and third-party disclosure means privilege waiver.

2. Deploy Approved, Private AI Infrastructure

Adopt a private AI platform—like SecureLLMs—that keeps all data within your firm's controlled environment. Ensure the deployment meets the confidentiality standards required to maintain privilege: single-tenant architecture, encrypted communications, no third-party data access, and no training on user inputs.

3. Ensure Attorney Involvement in AI-Assisted Work

One of the critical factors in Heppner was that the defendant acted without his attorney's direction. To maintain both privilege and work product protection, AI use for legal analysis should be conducted by attorneys or under the direction and supervision of counsel. Document this oversight as part of your firm's AI workflow.

4. Train All Personnel

ABA Formal Opinion 512 requires that managerial lawyers establish clear policies regarding AI use and that supervisory lawyers ensure all personnel are trained in the ethical and practical implications of AI tools. Training should cover the risks of public AI, proper use of private platforms, and the importance of maintaining confidentiality in all AI interactions.

5. Review and Update Engagement Letters

Consider updating client engagement letters and outside counsel guidelines to address AI use. Disclose to clients how AI tools are employed in their matters and confirm that only approved, private platforms with appropriate confidentiality protections will be used for work involving their information.

6. Maintain Audit Trails

Document your firm's AI usage practices, including which platform is used, what security controls are in place, and how attorney oversight is maintained. In the event of a privilege challenge, this documentation will be essential to demonstrating that reasonable steps were taken to preserve confidentiality.

The Broader Implications: AI Is Not Going Away

The legal profession cannot afford to ignore AI. Research indicates that professionals using AI within the boundaries of its capabilities can improve their performance by nearly 40 percent compared to those who do not use it. Law firms that adopt AI for legal research, contract analysis, document review, and case preparation gain meaningful competitive advantages in speed, cost, and quality of service.

But the Heppner ruling makes clear that how you adopt AI matters as much as whether you adopt it. Firms that continue using public AI tools for sensitive legal work are exposing themselves—and their clients—to privilege waiver, ethical violations, competitive intelligence leakage, and potential malpractice liability.

The answer is not to avoid AI. It is to deploy it correctly.

> **The Bottom Line**
>
> *A private AI deployment addresses every concern raised in Heppner: confidentiality is maintained because no AI provider accesses your data, privilege is preserved because communications remain within the firm's controlled environment, and work product protection is strengthened when AI is used under attorney direction. The question is no longer whether to use AI—it is whether you can afford to use it without proper safeguards.*

Conclusion

United States v. Heppner is likely the first in a long line of decisions addressing the intersection of artificial intelligence and legal privilege. The principles articulated by Judge Rakoff—that privilege requires confidentiality, that public AI platforms do not provide it, and that AI's novelty does not exempt it from established legal doctrines—will shape how courts evaluate privilege claims involving AI for years to come.

Law firms that act now to implement private AI infrastructure will be best positioned to benefit from AI's transformative capabilities while maintaining the ethical and legal standards their clients expect and the courts demand.

SecureLLMs.org was built for exactly this purpose: to give enterprises—including law firms, corporate legal departments, and other organizations handling sensitive information—access to the full power of frontier AI models without sacrificing security, confidentiality, or compliance.

To discuss how a private AI deployment could work for your firm—including approaches that do not involve us—we are happy to have that conversation.

SecureLLMs.org
michael@securellms.org
https://securellms.org

Sources and Further Reading

United States v. Heppner, No. 25-cr-00630 (S.D.N.Y. Feb. 17, 2026) (Rakoff, J.) – Written memorandum on attorney-client privilege and AI-generated documents.

Covington & Burling LLP, "AI and Legal Privilege: Key Takeaways from US v. Heppner," Inside Privacy (March 2, 2026).

Sidley Austin LLP, "Generative AI and Privilege: Practical Lessons from Two Early Decisions and What Comes Next" (February 2026).

Baker McKenzie, "Artificial Intelligence – United States," Global Privilege and Professional Secrecy Guide.

Norton Rose Fulbright, "Legal Professional Privilege in the Generative AI Era" (2025).

American Bar Association, Formal Opinion 512: Generative Artificial Intelligence Tools (July 29, 2024).

Morgan Lewis, "When AI Meets Privilege: Early Court Decisions" (February 2026).

Ogletree Deakins, "The Intersection of AI and Attorney-Client Privilege – A Cautionary Tale" (February 2026).

National Law Review, "Potential Privilege and Discovery of Generative AI Prompts" (2026).