

CounselCore Leadership Overview

A concise leadership overview for law firms evaluating sovereign, in-house AI over privileged and prior-matter knowledge.

What CounselCore Is

CounselCore is an in-house AI infrastructure layer for law firms that want the benefit of modern AI without moving privileged material, prompts, logs, or institutional knowledge into public or vendor-controlled systems.

The platform is designed around firm custody: retrieval, reasoning, audit logs, and attorney-facing workflows remain inside infrastructure the firm can govern and explain.

Why This Matters Now

Recent court activity has made AI governance a legal-risk question, not just a technology preference. Courts are examining public AI prompts, hallucinated authorities, missing AI records, expert credibility, and counsel's duty to verify.

IN-HOUSE	Matter data, prompts, outputs, logs, and retrieval sources remain within the firm's controlled environment.
GROUNDING	Answers are tied back to internal documents, citations, and source material instead of unsupported public-model output.
GOVERNED	Usage can be permissioned, logged, reviewed, and explained to partners, clients, courts, and regulators.

Court Signals Driving the Business Case

The cases below do not all involve the same procedural posture, but together they show the same operational lesson: firms need AI systems that preserve confidentiality, verify sources, retain records, and make usage defensible.

Leadership Takeaways

- Public AI can create privilege, preservation, and confidentiality questions.
- AI-generated legal authorities must be verified before they enter any court-facing work.
- Logs and source records matter because they can become the evidence of how the work was produced.
- In-house AI gives firms a cleaner governance posture than ad hoc public-AI use.

Sources Cited

United States v. Heppner

Public generative AI use raised privilege and work-product concerns when prompts and outputs were sought in discovery.

<https://harvardlawreview.org/blog/2026/03/united-states-v-heppner/>

Miller v. Regions Bank

The court treated AI-generated fake citations and deleted ChatGPT records as a sanctions and preservation issue.

https://websitedc.s3.amazonaws.com/documents/Miller_v_Regions_Bank_USA_21_May_2026.pdf

Mata v. Avianca

The court sanctioned lawyers who submitted non-existent judicial opinions generated by ChatGPT.

<https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1%3A2022cv01461/575368/54/>

Park v. Kim

The Second Circuit referred counsel to its grievance panel after a brief cited a non-existent ChatGPT-generated case.

<https://law.justia.com/cases/federal/appellate-courts/ca2/22-2057/22-2057-2024-01-30.html>

Kruse v. Karlen

The Missouri Court of Appeals dismissed an appeal and awarded damages after fictitious AI-generated cases appeared in briefing.

<https://www.courts.mo.gov/file.jsp?id=205457>

Kohls v. Ellison

A federal court rejected reliance on an expert declaration after GPT-4o generated fake academic citations.

<https://www.law.berkeley.edu/wp-content/uploads/archive/2025/12/Kohls-v-Ellison.pdf>