

CounselCore Partner Briefing Memo

A partner-level decision memo for evaluating in-house legal AI against privilege, discovery, ethics, and operational-risk requirements.

Decision Premise

Law firms are not deciding whether lawyers will use AI. They are deciding whether that use will happen inside a governed system or through fragmented, public, and hard-to-defend workflows.

CounselCore is built for the first path: controlled deployment, matter-aware retrieval, permissioned access, and evidence-ready logging.

Recommended Governance Position

Treat AI like a high-risk legal system rather than a productivity toy. Define approved use cases, restrict matter access, log activity, require verification of legal authorities, and avoid public systems for privileged or client-sensitive material.

PRIVILEGE	Keep confidential material inside the firm's own environment and document the boundary.
DISCOVERY	Avoid creating vendor-controlled repositories, logs, or unknown retention surfaces.
VERIFICATION	Require source-linked outputs and attorney review before legal conclusions are used.
PRESERVATION	Retain usage records where they may later be needed to explain or defend the work.

What the Case Law Is Signaling

The court examples cited below show a widening legal-risk surface around public AI: privilege disputes, fake case citations, sanctions, deleted AI records, appellate discipline referrals, and expert credibility failures.

Briefing Questions for Firm Leadership

- Where can attorneys safely use AI today, and where is public AI prohibited?
- Can the firm prove what source material supported an AI-assisted answer?
- Can the firm preserve AI usage records if a court later asks for them?
- Does the firm's AI system respect ethical walls, matter permissions, and client restrictions?

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>