A.R. INTERNATIONAL CONSULTING · ARIC INSIGHT

VIENNA · DUBAI · WASHINGTON · MMXIV

ARIC INSIGHT · VOLUME I · ISSUE 07

The *Rawest* Form of Thinking.

A working note from the firm on the privilege-and-privacy crisis exposed by the recent federal ruling that admitted a defendant's pre-counsel artificial-intelligence chat history into the criminal record — a Kovel expansion the firm reads as overdue, a Riley contrast the firm reads as urgent, and the standard of care the firm now publishes for every client engagement.

FOREWORD · THE FIRM'S POSITION

On 10 February MMXXVI a federal judge in the Southern District of New York ruled that a defendant's pre-counsel conversations with a general-purpose AI assistant were **neither attorney-client privileged nor protected by any reasonable expectation of privacy**, and admitted thirty-one documents of those conversations into the criminal record. The decision is narrow on its facts and broad in its register. **AI chat history is the rawest form of human thinking ever recorded** — pre-counsel, pre-edit, candid, and exhaustive — and the consumer-tier surfaces on which most professionals, founders, and principals presently produce that record offer none of the protections the same act on a phone has already won under *Riley v. California*. This issue records the firm's working position on the doctrine, the legislative path, the platform two-tier gap, and the standard of care the firm now publishes for every engagement.

ISSUE
Vol I · Issue 07

DATED
13 May *MMXXVI*

VOICE
Firm · Reading Desk

SECTION I · A NEW CATEGORY OF EVIDENCE

AI chat history is the *rawest form of human thinking* ever recorded — and the courts have just begun to read it.

A short statement of what the February MMXXVI ruling actually held, why the category of evidence it admitted is constitutionally novel, and why the reading desk treats it as the load-bearing privacy event of the decade so far.

I.1 · *What was admitted into the record.*

On 10 February MMXXVI, sitting in the Southern District of New York, the Honourable **Judge Jed S. Rakoff** ruled on the discoverability of **pre-counsel conversations with a general-purpose AI assistant**. The court's rationale, in three lines: **the AI is not an attorney**, so attorney-client privilege does not attach; the consumer-tier terms-of-service authorise platform retention and use, so **no reasonable expectation of privacy** is preserved; and the use was **not under counsel's direction**, so no Kovel-style derivative privilege is in reach.

The decision is narrow on its facts. Its register is not. It is the first federal ruling of its kind on AI-conversation discovery, it sets the analytic framework on which every subsequent court will lean, and it forces a question the legal academy has been able to defer for three years: **what kind of evidence, exactly, is an AI chat session?**

I.2 · *Why this is constitutionally novel.*

The reading desk's position is that an AI chat history is a category of evidence the courts have never before had access to. A diary is edited; a notebook is selective; a private letter has an audience and edits itself for that audience. An AI session has none of those discipline gates. **The user thinks out loud, in writing, to a system that performs the dialogue of a counsellor without being one.** The session captures the half-formed hypothesis, the abandoned line of reasoning, the speculation that a human conversation would have absorbed and forgotten. It is the closest written record yet produced of *raw cognition* — and the court has just held, in the consumer register, that it is discoverable.

The discovery surface compounds. A user produces, in a year of ordinary work, more linear pages of pre-counsel reasoning across a consumer AI than across every diary, notebook, and personal letter of the prior decade. Multiply that by the working population of the country and the implication is direct: the criminal and civil discovery base, in five years, will include **more candid pre-counsel reasoning than the courts have seen in the prior fifty years combined**. The doctrine the courts adopt now will set the register for that surface for a generation.

> *A diary is edited. A letter has an audience. An AI session has neither.* **It is the rawest written record of human thinking the courts have ever had access to** *— and on the consumer tier, it is presently discoverable.*

SECTION II · THE RILEY CONTRAST

II.1 · *Riley v. California — what the cell phone bought, in 2014.*

In *Riley v. California*, 573 U.S. 373 (2014), the Supreme Court held unanimously that **a warrant is required before law enforcement may search the digital contents of a cell phone seized incident to arrest**. The reasoning the Chief Justice offered then is the reasoning the firm reads as load-bearing now: a modern smartphone, the Court wrote, holds "the privacies of life" — messages, photographs, location traces, search history, financial records, calendar, voice notes — in a volume and intimacy that the Founders could not have anticipated, and the Fourth Amendment's warrant requirement therefore attaches by structural necessity, not by analogy.

The Court's instinct in *Riley* was that the volume and intimacy of the record changed the constitutional analysis. **An AI chat history holds more, more intimately, than a phone** — the same messages, the same calendar, the same financial reasoning, but rendered through the explicit pre-counsel cognition the user produced in writing. The 2014 instinct, applied honestly to the 2026 surface, would require the same warrant register on AI chat history that *Riley* already requires on cell-phone contents. It does not, today.

II.2 · *Why terms-of-service "consent" cannot be the answer.*

The February ruling rests, in part, on the consumer-tier terms-of-service the defendant accepted on the platform. The court read those terms as a waiver of any reasonable expectation of privacy. The reading desk's position is that this analysis cannot hold the constitutional weight being placed on it. *Riley* itself was decided against a backdrop in which carrier and OS terms-of-service authorised broad data collection by the carrier and the OS vendor; the Court held the warrant requirement attached anyway, because the Fourth Amendment register does not bend to a click-through. **If the warrant requirement attaches to phone contents notwithstanding the terms accepted to use the phone, the same register must, in time, attach to AI chat history** notwithstanding the terms accepted to use the AI. The doctrinal work has not yet been done; this issue is the firm's reading of why it must be.

SECTION III · THE KOVEL DOCTRINE, AND THE EXPANSION THE FIRM PROPOSES

The doctrine that already shields the accountant in the room. *It must, now, be made to shield the agent on the screen.*

A short reading of United States v. Kovel, of why the doctrine does not yet fit agentic-AI assistance cleanly, and of the surgical expansion the firm proposes — one that protects principals, founders, and consultants without expanding privilege at large.

III.1 · *What Kovel covers, and what it does not.*

In *United States v. Kovel*, 296 F.2d 918 (2d Cir. 1961), Judge Friendly held that the work of a non-lawyer professional — in that matter, an accountant — sits inside the attorney-client privilege envelope **when the professional is retained by counsel for the purpose of rendering the legal advice**. The doctrine is narrow by design. It requires a counsel-of-record retention, a defined legal-advice purpose, and a working channel that runs through counsel. It does not protect the accountant the client hires before talking to a lawyer; it does not protect the consultant the client engages on the side; it does not, today, protect the AI assistant the client opens at midnight on a phone.

Kovel has been extended over six decades to translators, valuation experts, public-relations consultants in narrow fact patterns, and, in a small number of districts, to information-technology consultants. The *structural* condition is constant: the work moves through counsel, the retention is documented, and the channel is privileged from the first communication to the last.

III.2 · *The surgical expansion the firm proposes.*

The reading desk's working position is that the Kovel envelope must be made to cover, expressly, two categories of AI use that the present doctrine does not reach cleanly:

- **The agentic-AI consultant** — a non-lawyer professional whose working tool is an agentic AI system, retained by counsel for the purpose of the legal representation, operating under a documented engagement letter and an enterprise-grade AI subscription that contractually binds the platform to *zero data retention, no model training, and audit-log discipline*.
- **The AI-assisted advisor** — a Kovel-eligible professional (accountant, valuation expert, restructuring advisor, government-relations counsel) using an AI system as part of the engagement's working register, where the engagement letter binds the use to enterprise-grade infrastructure and the work-product moves only through counsel.

The expansion is surgical. It does not reach the consumer who opens an AI on a phone before talking to a lawyer; that user must be reached by the legislative path described in Section IV. It does not reach the consultant who runs unencrypted prompts through a free-tier platform; that conduct sits outside any privilege register and is, properly, discoverable. It reaches the structured channel: **counsel of record, enterprise SKU, written retention, audit-traceable use**. That is the channel the Kovel doctrine was built for; it is the channel the doctrine must, now, be made to recognise.

> *Kovel was the doctrine that brought the accountant inside the privileged room.* **It must, now, be the doctrine that brings the agent on the screen inside the same room** *— under the same discipline, on the same channel, with the same auditability.*

SECTION IV · THE FEDERAL LEGISLATIVE PATH

IV.1 · *Aligning AI-platform data handling with the Riley warrant register.*

The Kovel expansion reaches the structured engagement. It does not reach the citizen at large. The reading desk's second position is that **federal consumer-privacy legislation** is now overdue and must be sequenced behind a single substantive principle: *AI chat history is to the citizen of MMXXVI what the cell-phone's digital contents were to the citizen of MMXIV*, and the warrant register the Supreme Court applied in *Riley* is the doctrinal floor on which any honest federal statute must rest. In substance, the firm reads the necessary architecture as four lines:

- **Default no-training, no-retention.** Platform handling of consumer chat history defaults to ZDR — no model training, no operator review, no third-party access — with affirmative opt-in required for any other use.
- **Warrant requirement on access.** Government access to a user's chat history requires a probable-cause warrant on the *Riley* register, with statutory parity to the cell-phone contents standard already in force.
- **Audit-log discipline.** Platforms maintain a tamper-evident audit log of every operator and government access request, with annual transparency-report disclosure on the Stored Communications Act register.
- **Statutory Kovel-AI safe harbour.** Express recognition, in federal law, that AI use under a documented counsel retention and an enterprise ZDR contract sits inside the attorney-client privilege envelope. The expansion is statutory; it does not depend on a circuit-by-circuit doctrinal extension over the next decade.

SECTION V · THE TWO-TIER PRIVACY GAP

The enterprise SKU already protects. *The consumer SKU exposes.* The middle class is on the wrong side of a gap the platforms could close tomorrow.

A buyer-grade reading of the current AI platform stack against the four privacy rows that determine whether a chat history is discoverable on a click-through, or shielded inside an enterprise contract.

V.I · *Where the four major platforms presently sit.*

PLATFORM · TIER	ZERO DATA RETENTION	NO MODEL TRAINING	AUDIT LOG · ADMIN	BAA / SOC2 / ENTERPRISE POSTURE
ChatGPT *Free / Plus / Pro*	No (retained 30 days+, legal-hold extensions)	*Opt-out available, default on (Plus/Pro)*	No admin audit log	No BAA, no enterprise contract
ChatGPT *Team / Enterprise*	Yes (configurable ZDR)	Off by default, contractual	Admin console & audit log	BAA available, SOC 2 Type II
Claude *Free / Pro*	No (retained 30 days+, legal-hold extensions)	*No training on conversations by default; ToS-bound*	No admin audit log	No BAA, no enterprise contract
Claude *Team / Enterprise*	Yes (ZDR available, contractual)	Contractually no training on customer data	Admin console & audit log	SOC 2 Type II, HIPAA-eligible on Enterprise
Gemini *Free / Advanced (Pro)*	No (Activity retained, opt-out partial)	*Default on; opt-out via Activity controls*	No admin audit log on consumer	Consumer ToS only
Gemini for Workspace *Business / Enterprise*	Yes (no retention beyond session, Workspace contract)	No training on customer data, contractual	Workspace Admin audit log	BAA available, SOC 2 / ISO 27001
Microsoft Copilot *Free / Pro (consumer)*	No (Bing Chat history retained on consumer)	*Default on; partial opt-out*	No admin audit log	Consumer ToS only
Microsoft 365 Copilot *Enterprise (M365)*	Yes (tenant-bound, no retention beyond tenant policy)	Contractually no training on tenant data	Purview audit log & eDiscovery	BAA, FedRAMP High, SOC 2

The pattern is unmistakeable. **The four privacy rows the matrix tracks — ZDR, no-training, audit log, enterprise posture — are presently available on every major platform, but only on the enterprise SKU.** The consumer at the same provider, on the same model, signing the click-through ToS, is on the wrong side of the gap. The gap is contractual, not technical. The platforms could close it tomorrow as a default; today, they hold it open as a SKU upsell. The result is a two-tier privacy surface in which the corporate buyer is constitutionally protected and the middle-class user, on the same dialogue, is constitutionally exposed.

SECTION VI · THE FIRM'S STANDING POSTURE

VI.I · *Kovel-compliant AI use as the firm's standard of care.*

The firm's commitment to every client engagement, published here in writing on the open record, is a single line: **every ARIC engagement runs on enterprise-grade AI infrastructure under a Kovel-compliant working register.** Three operative disciplines stand behind the line:

- **Enterprise infrastructure only.** The firm's AI use, on every engagement and every surface, runs on Team / Enterprise tier subscriptions with contractual zero-data-retention, no-training, and admin-audit posture. No client work touches a consumer-tier surface.
- **Counsel-coordinated retention where the matter requires it.** Where the engagement touches active or anticipated litigation, regulatory enforcement, or pre-charge legal exposure, the firm is retained *under counsel direction* on the Kovel register, with the engagement letter co-signed and the work-product moving through the privileged channel.
- **Audit-log transparency to the client.** The firm's admin audit log on every engagement is, on request, made available to the client's counsel of record. The principal who chooses to work with the firm knows, in writing and in advance, the privacy register their work moves under for the life of the file.

> *The doctrine will, in time, catch up to the surface. Until it does, the discipline is contractual:* **enterprise infrastructure only, counsel-coordinated where the matter requires it, audit-log transparency to the client.** *Published here, in writing on the open record, as the firm's standard of care.*