Feel free to use as is or amend to suit your firm

AI Vendor Privacy & Security Vetting Template

Purpose: To be sent to any AI software provider (SaaS, Discovery Platform, or API) before uploading client data.

1. Data Usage & Model Training

- The "No-Training" Guarantee: Does the vendor contractually prohibit the use of our data (inputs, prompts, and outputs) to train, fine-tune, or improve their global models?
- Zero Data Retention (ZDR): Does the vendor offer a "Zero Retention" mode where data is deleted immediately after the API call is completed?
- Data Segregation: Is our data stored in a logically or physically segregated environment (Multi-tenant vs. Single-tenant)?
- Opt-Out vs. Opt-In: Is the "No-Training" setting the default, or does it require a manual configuration?

2. Security & Compliance Certifications

- Audit Reports: Can the vendor provide a current SOC 2 Type II or ISO 27001 report?
- Encryption Standards: Is data encrypted using AES-256 at rest and TLS 1.3 in transit?
- Key Management: Does the vendor offer Customer-Managed Keys (CMK) or "Bring Your Own Key" (BYOK) encryption?
- Data Residency: Can the vendor guarantee that data remains within a specific jurisdiction (e.g., US-only, EU-only) to comply with local privacy laws?

3. Technical Safeguards (The "Discovery Leak" Prevention)

- PII/PHI Masking: Does the platform have an automated "Sanitizer" that redacts sensitive info before it reaches the LLM?
- Human-in-the-Loop: Does the vendor’s staff have any "backdoor" access to view prompts or outputs for "quality assurance"? (This is a major privilege risk).
- Audit Logging: Does the tool provide a granular log of who accessed what data and what prompts were run? (Essential for Rule 37 defensibility).

4. Liability & Indemnity

- IP Indemnification: Does the vendor indemnify the firm against third-party claims of copyright infringement arising from AI-generated outputs?
- Breach Notification: What is the contractually mandated timeline for notifying the firm of a data breach? (Standard is 24–72 hours).

Feel free to use as is or amend to suit your firm

The "Cheat Sheet": Open-Loop vs. Closed-Loop AI

Add this simple table to the end of your template to help lawyers quickly categorize tools.

Feature	**Open-Loop AI (High Risk)**	**Closed-Loop AI (Preferred)**
Examples	Free ChatGPT, Public Gemini, Claude (Free)	Azure OpenAI (Enterprise), Discovery-specific AI
Data Usage	Your prompts train the model.	Your data is isolated; no training on your info.
Privilege	Likely Waived (Disclosure to 3rd party).	Preserved (Confidential service provider).
Security	Minimal / "Best Efforts."	SOC 2 Type II / Enterprise Grade.