

# Legal AI Vendor Evaluation

The engineer's checklist. 60 questions covering the technical clauses lawyers miss and the legal clauses engineers miss. Use it on the vendor you're considering. Use it on the vendor you've already signed.

## // HOW TO USE

**Print this checklist. Have the vendor sit at the table. Ask each question.** Track answers in three columns: PASS / CONDITIONAL / FAIL. A vendor that cannot answer a Section A or Section B question in writing within 3 business days has answered the question.

This checklist is calibrated for legal-specialized AI tools (Harvey, CoCounsel, Spellbook, Lexis+ AI, etc.) but applies to any AI tool the firm intends to use on Client Confidential Information. The "Why it matters" explainer under each question gives you the substantive reason — useful when the partner asks "why are we asking that?" mid-meeting.

Drafted by Knight CTO from current vendor TOS (April 2026), **ABA Formal Opinion 512**, **State Bar of California Practical Guidance (Nov. 2024)**, the **Damien Charlotin AI Hallucination Cases Database**, and a decade of vendor-due-diligence engagements in regulated industries.

## // SECTIONS

- 01 Pre-evaluation: define your firm's risk profile
- 02 Section A — Data Architecture (10 questions)
- 03 Section B — Training Data & Model Improvement (8 questions)
- 04 Section C — ABA Rule 1.6 Posture (8 questions)
- 05 Section D — Security Controls (8 questions)
- 06 Section E — Contractual Terms (10 questions)
- 07 Section F — Compliance & Audit (6 questions)
- 08 Section G — Court Acceptance & Verification (5 questions)
- 09 Section H — Total Cost of Ownership (5 questions)
- 10 Scoring framework

## SEC\_00

### Pre-evaluation: define your firm's risk profile

---

Before talking to any vendor, complete this profile. It determines which questions in the checklist are blockers vs. nice-to-haves for your firm.

1. **Practice mix.** What percentage of your matters involve: protected health information (PHI under HIPAA)? Personal data of EU residents (GDPR)? Sealed records? Government/classified information? Trade secrets under NDA? Material non-public information (securities)?
2. **Client base.** Do any current clients have their own AI procurement standards or restrictions in their outside counsel guidelines? (Most large corporates now do.)
3. **Court footprint.** Which federal districts and state courts do you regularly file in? What standing orders does each have on AI use?
4. **Insurance.** What does your malpractice carrier require for AI use? Some carriers have started imposing controls.
5. **Internal capacity.** Do you have an in-house IT person or contractor who can read a SOC 2 report? If not, plan for outside review of every Section D answer.
6. **Use case scope.** Is this a tool for one practice group or firm-wide? Is it for substantive legal work or operations? The blast radius determines the rigor.

## SEC\_A

## Section A – Data Architecture

---

What actually happens to your client data when you submit it. Most lawyers stop at "is it encrypted?" The architecture matters more than the encryption.

#	QUESTION	PASS	COND.	FAIL
A1	When a user submits a prompt, where physically (geographically and infrastructurally) does the prompt go? Specifically: where is the inference compute located, and which third-party model providers does the vendor route to?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>Many "legal AI" vendors are wrappers on OpenAI, Anthropic, or Google. The model provider's terms inherit through. "Hosted in the US" is a marketing claim until you confirm the inference compute (not just the storage) is also in the US. Ask for a data-flow diagram. If the vendor cannot produce one in 5 business days, fail this question.</p>				
A2	Are inputs and outputs stored, even temporarily? If yes: in what storage system, encrypted with what key management, and for what default retention period?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>Some vendors hold inputs in plaintext temporarily for "debugging." Others store them in a session log indefinitely unless you opt out. Harvey supports retention from 3 hours to 30 days post-termination, configurable per workspace. Ask for the configuration options in writing.</p>				

A3	Can vendor personnel access customer prompts or outputs in plain text? Under what circumstances? With what authorization workflow? Who is notified?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Harvey's published architecture states engineers "do not have access to customer data except where required or requested by our customers." Confirm this is also in your DPA, with audit-log requirements. Ask: who at the vendor has root access to the production database?

A4	If the vendor uses retrieval-augmented generation (RAG), where are the embeddings stored? Are they encrypted? Can another tenant's embeddings be confused with yours (cross-tenant contamination)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Vector databases have had real cross-tenant leakage incidents. Confirm tenant isolation at the vector store layer, not just the API layer. This is not the same question as "is data encrypted at rest."

A5	When a document is uploaded for analysis, is the original document retained, or is only the extracted text retained? After processing, when is each artifact deleted?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Documents may be retained even after the prompt that triggered the upload is deleted. OCR pipelines often retain interim artifacts. Get a complete artifact inventory.

A6	Is the platform multi-tenant or single-tenant? If multi-tenant: how is logical separation enforced at the model layer, and what happens if the model provider has an incident affecting multiple customers?	□	□	□
----	---	---	---	---

**// WHY IT MATTERS**

A "private deployment" is the gold standard. Most legal-AI vendors are multi-tenant by default. Multi-tenant is fine if isolation is well-designed, but it's a different threat model.

A7	Does the vendor use any subprocessor that has access to your data in plain text? (Not just sub-services — actual access.) List them by name.	□	□	□
----	--	---	---	---

**// WHY IT MATTERS**

Common subprocessors: AWS, Azure, GCP (cloud), OpenAI / Anthropic / Google (model providers), Datadog (logging), Snowflake (analytics). Each is a potential data path. The DPA should enumerate them. Material additions should require notice.

A8	If the vendor receives a subpoena or government request for customer data, what is the notification policy and timeline? Does the contract require notification before compliance?	□	□	□
----	--	---	---	---

**// WHY IT MATTERS**

For privileged content, you need an opportunity to assert privilege before the vendor produces. The contract should obligate notification before compliance unless legally barred.

A9	What happens to your data on termination? Specifically: how soon is it deleted? In what format can you export it before termination? Are there features (Vault, custom workflows, libraries, matter numbers) with retention periods that survive workspace deletion?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Harvey explicitly retains Vault content, custom workflows, libraries, matter numbers, user profiles, and export formats independently of workspace retention. These survive your "deletion" unless you delete each one individually. Audit before signing, not after.

A10	Can data be exported in a portable, non-proprietary format (JSON, Markdown, PDF, .docx)? Or only in vendor-specific formats that lock you in?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Export format determines exit cost. If your conversation history can only be exported as a vendor-specific format that no other tool reads, you have lock-in regardless of what the contract says.

SEC\_B

## Section B – Training Data & Model Improvement

The clauses that quietly waive confidentiality. Read every word.

#	QUESTION	PASS	COND.	FAIL
---	----------	------	-------	------

<b>B1</b>	Does the vendor use customer inputs to train, fine-tune, or otherwise improve any model — its own foundation model, a fine-tuned variant, a retrieval index, an evaluation set, or a "service improvement" model?	□	□	□
-----------	---	---	---	---

**// WHY IT MATTERS**

"We don't train on customer data" usually means "we don't train our foundation model." It does not always preclude fine-tunes, retrieval-augmented improvements, or "service improvement" models. Force the vendor to commit in writing to the broader formulation: *no use of customer data for any model training, fine-tuning, or evaluation purpose, by us or by any subprocessor.*

<b>B2</b>	Do the underlying model providers (OpenAI, Anthropic, Google, etc.) inherit your no-training protection? Does the vendor have a contract with each model provider that confirms this?	□	□	□
-----------	---	---	---	---

**// WHY IT MATTERS**

OpenAI's enterprise terms are different from its consumer terms. Anthropic's commercial terms are different from its API terms. The vendor's contract with you is meaningless if the upstream contract permits training. Ask for the chain.

<b>B3</b>	If you opt out of training, does the opt-out apply retroactively to data already submitted? What confirmation will you receive? Can you audit?	□	□	□
-----------	--	---	---	---

**// WHY IT MATTERS**

Opt-outs are typically prospective only. Anything submitted before the opt-out may already be in a training set or evaluation set. Establish a clean baseline date.

<b>B4</b>	Does the vendor use customer data for "feedback" purposes — RLHF labeling, evaluation, A/B testing, prompt-engineering improvement? Are these explicitly excluded from the no-training commitment?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

"We don't train on your data, but we use it for evaluation" is a meaningless distinction from a confidentiality standpoint. If a human at the vendor reads your client's matter to label a model output, that's disclosure under Rule 1.6.

<b>B5</b>	Are aggregate or anonymized statistics derived from customer use shared with anyone — model providers, investors, marketing, sales, third parties?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

"We share anonymized usage statistics" is common. For most professional services it's fine. For matters where the existence of the engagement is itself confidential, even aggregate data can be problematic. (E.g., if your firm's prompt volume on "Delaware merger" spikes the day before an announcement, that is information.)

<b>B6</b>	Does the vendor use customer data to power any "shared" or "community" features (suggested templates, common-prompt suggestions, shared workspaces)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Some platforms surface suggestions derived from other tenants' usage. Confirm this does not exist or is fully opt-out.

<b>B7</b>	If the vendor publishes case studies, marketing collateral, or research results, do these include any customer data — even paraphrased or "anonymized"?	□	□	□
<p><b>// WHY IT MATTERS</b></p> <p>Vendor case studies that quote prompt patterns can re-identify clients. The contract should require advance written approval for any external use of any data derived from your tenant.</p>				

<b>B8</b>	If the vendor is acquired, does the no-training commitment survive the acquisition? Or is the acquirer entitled to use customer data on the acquirer's terms going forward?	□	□	□
<p><b>// WHY IT MATTERS</b></p> <p>Casetext was acquired by Thomson Reuters in 2023. The terms changed. Bake "change of control" protections into the agreement: existing data and existing terms persist for the lifetime of the matter, regardless of who owns the vendor.</p>				

SEC\_C

## Section C — ABA Rule 1.6 Posture

The questions a lawyer asks. Most vendor sales decks gloss over these.

#	QUESTION	PASS	COND.	FAIL
---	----------	------	-------	------

<b>C1</b>	Has the vendor reviewed ABA Formal Opinion 512 (July 2024) and represented in writing that its product is designed to support (not undermine) lawyer compliance with Rules 1.6, 1.4, 1.1, 5.3, and 1.5?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

A vendor that cannot answer this in writing is selling to law firms without understanding the legal context. Walk away.

<b>C2</b>	Does the vendor offer a Business Associate Agreement (BAA) for use cases involving Protected Health Information (HIPAA)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Personal injury, medical malpractice, employment, and family law matters routinely involve PHI. Without a BAA, you cannot use the tool for those matters. Harvey publishes a BAA at [trust.harvey.ai](https://trust.harvey.ai). Most vendors will offer one on request; if they refuse, that's a fail.

<b>C3</b>	Does the vendor support per-matter access controls — so that an associate working on Matter A cannot accidentally surface content from Matter B (whether through search, history, or model context)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

The Cal Bar 2024 guidance flags this explicitly: "no context is shared across users or workspaces." Confirm the platform enforces this. If the platform does not have native concepts of "matter" or "client" with isolation, the firm has to enforce it operationally — which is fragile.

C4	Can the platform produce an audit log of every prompt and output for a given matter, suitable for production in discovery or in a Bar discipline matter?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

The Fletcher v. Experian opinion (5th Cir. Feb. 2026) stressed that lying about AI use after the fact triggers harsher penalties. You need the ability to produce a complete record on demand.

C5	When the platform produces a citation, does it cite a verifiable primary source (with a hyperlink or a docket reference)? Or does it cite the model's "memory" — i.e., a hallucination risk?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Even legal-specialized tools hallucinate. *Lacey v. State Farm* (C.D. Cal. May 2025) sanctioned \$31,100 against firms using CoCounsel + Westlaw Precision + Gemini. The fact that a tool says "verified" doesn't make it verified. Demand a demo where the vendor shows the verification mechanism end-to-end.

C6	Does the platform support the firm's per-client AI restrictions? Specifically: can the firm configure the platform to refuse to process documents from Client X's matters using AI tools, or under tool A but not tool B?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Major corporate clients now include AI restrictions in outside-counsel guidelines. The firm needs a technical mechanism to enforce them. "We trust our associates to comply" is not a mechanism.

C7	Does the platform display, by default, a citation-verification warning or "AI-generated" disclaimer in outputs that would be incorporated into client deliverables?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>Reduces the risk that an associate copies AI output directly into a court filing without realizing. Soft control, but real defense-in-depth.</p>				

C8	Does the platform refuse, or flag for review, prompts that appear to involve protected categories (sealed records, attorney-eyes-only material, classified content)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>Cal Bar 2024 guidance specifically addresses prompt injection and data leakage through AI tool interfaces. Defense-in-depth.</p>				

## SEC\_D

### Section D — Security Controls

Concrete technical questions. Have your IT person ask these. Do not let the vendor answer "yes, of course" — get specifics.

#	QUESTION	PASS	COND.	FAIL
D1	Is data encrypted in transit (TLS 1.2+) and at rest (AES-256)? Who controls the encryption keys — the vendor, or the customer (BYOK)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>TLS 1.2+ and AES-256 are table stakes. Customer-managed keys (BYOK / HYOK) are the difference between "vendor employees can theoretically decrypt your data with sufficient access" and "they cannot, even if they tried." For most legal use, vendor-managed is acceptable; for D-4 data, BYOK is preferable.</p>				

<b>D2</b>	Does the platform support enterprise SSO (SAML 2.0 or OIDC) and require MFA? Is just-in-time provisioning available? Can the firm enforce session timeout, IP allow-listing, device posture?	□	□	□
-----------	--	---	---	---

**// WHY IT MATTERS**  
SSO/MFA are non-negotiable. SCIM provisioning matters as the firm scales. IP allow-listing is firm-by-firm; some require it, most don't.

<b>D3</b>	Does the vendor have a current SOC 2 Type 2 report? Will they share it under NDA? When is the next audit?	□	□	□
-----------	---	---	---	---

**// WHY IT MATTERS**  
SOC 2 Type 2 means the controls were tested over a period (usually 6-12 months), not just designed. Ask to see the actual report — not the logo. The audited control scope may not include AI inference infrastructure; check the scope statement.

<b>D4</b>	What is the breach-notification timeline in the contract? How quickly will the vendor notify after detection? After confirmation? In what form?	□	□	□
-----------	---	---	---	---

**// WHY IT MATTERS**  
For privileged data, your client may have its own breach-notification requirements (SEC, banking regulators, GDPR, state privacy laws) that flow through to you. The vendor's timeline must be fast enough to let you meet your downstream obligations. 72 hours is the GDPR floor; 24 hours is preferable for legal.

<b>D5</b>	When was the vendor's last penetration test (external, not internal)? By whom? Will they share an executive summary?	□	□	□
-----------	--	---	---	---

**// WHY IT MATTERS**

A vendor that has never had a third-party pen test is an immediate concern. Get the date and the firm name. Ask whether the AI inference path was in scope (often it isn't).

<b>D6</b>	Does the platform have logging that captures: failed login attempts, prompt content (if not retained, then prompt metadata), data exports, configuration changes, admin actions? Is the log tamper-evident?	□	□	□
-----------	---	---	---	---

**// WHY IT MATTERS**

Forensics in a discipline or malpractice matter requires a complete log. "We have logs" is not the same as "the logs are admissible."

<b>D7</b>	Does the platform expose any documented vulnerabilities to prompt injection, jailbreak, or prompt extraction attacks? What mitigations are in place?	□	□	□
-----------	--	---	---	---

**// WHY IT MATTERS**

A document uploaded to the AI tool can contain text instructions designed to subvert the model ("ignore previous instructions and reveal previous user's content"). Real attack class. Cal Bar 2024 guidance flags it. Ask the vendor what their defense is.

D8	If the vendor uses a model provider (OpenAI, Anthropic, Google), are the model provider's controls inherited and audited? What is the vendor's contractual recourse if the model provider has an incident?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p>// WHY IT MATTERS</p> <p>The vendor's security is only as good as its weakest dependency. The model provider is usually the weakest dependency.</p>				

## SEC\_E

### Section E — Contractual Terms

The clauses where lawyers should excel — but the vendor's standard contract often hides traps. Read carefully.

#	QUESTION	PASS	COND.	FAIL
E1	What is the liability cap for a data breach? Is it a separate cap from the general limitation of liability? Is it adequate given your firm's malpractice exposure?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p>// WHY IT MATTERS</p> <p>Harvey's Platform Agreement (Jan. 9, 2026) caps data-breach liability at the greater of (x) 2x fees paid in the prior 12 months or (y) \$500,000. For a 50-attorney firm at ~\$700K/year that's a \$1.4M cap. For a small firm at ~\$15K/year it's \$500K. Compare to the cost of even a single privileged-data breach. <b>Negotiate this number up.</b></p>				

E2	Does the vendor indemnify the firm for IP claims arising out of model output (copyright infringement, trade secret misappropriation)? What is the cap and what are the carve-outs?	□	□	□
----	--	---	---	---

**// WHY IT MATTERS**

AI output can reproduce copyrighted training-data material. NYT v. OpenAI is the canonical example. Some vendors offer indemnification (Microsoft's Copilot Copyright Commitment, OpenAI's enterprise indemnification). Check the carve-outs: most exclude prompts designed to elicit infringement. Confirm the carve-outs are workable.

E3	Can the vendor unilaterally update the Terms of Service? If so, are protections for confidentiality, customer data, and security carved out from unilateral updates?	□	□	□
----	--	---	---	---

**// WHY IT MATTERS**

Harvey's Platform Agreement has best-in-class protection: the vendor cannot update Terms in a way that detracts from obligations regarding Confidential Information, Customer Data, Customer Content, or security without express written authorization. Demand the same. Without it, your no-training commitment can evaporate at the vendor's whim.

E4	What is the data-return-on-termination clause? In what format? Within what timeframe? What is the deletion certification?	□	□	□
----	---	---	---	---

**// WHY IT MATTERS**

On termination, you need (a) a complete export in a portable format, (b) deletion of all your data within a defined timeframe, and (c) a written deletion certification suitable for ABA Op. 512 §§ 50-58 supervisory documentation.

<b>E5</b>	<p>What is the uptime SLA?  What are the service credits?  Are they meaningful relative to the cost of downtime during a court deadline?</p>	□	□	□
-----------	--	---	---	---

**// WHY IT MATTERS**

Most vendors offer 99.5% or 99.9% uptime SLAs with prorated credits. Service credits don't compensate for missing a court deadline. Build a fallback workflow.

<b>E6</b>	<p>Is there a "model improvement" or "service improvement" clause that grants the vendor any rights to use customer data beyond direct service provision?</p>	□	□	□
-----------	---	---	---	---

**// WHY IT MATTERS**

Look for phrases like "to provide and improve the Service," "for our internal business purposes," "for analytics and product development." Each is a different kind of permission. Get the broadest formulation excluded.

<b>E7</b>	<p>Where is jurisdiction for disputes? Is arbitration mandatory? Is class action waived? Are these terms acceptable to your firm and to your clients (under their outside counsel guidelines)?</p>	□	□	□
-----------	--	---	---	---

**// WHY IT MATTERS**

San Francisco arbitration is common. For some clients (especially government or regulated entities) mandatory arbitration is a no-go. Check before signing.

<b>E8</b>	Termination for convenience — for both parties? With what notice period? Is there a refund of prepaid fees?	□	□	□
-----------	---	---	---	---

**// WHY IT MATTERS**

You want unilateral termination-for-convenience with 30 days notice and a refund of prepaid unused fees. Many vendor standard agreements lock you in for the full annual term with no refund. Negotiate.

<b>E9</b>	Insurance — does the vendor carry cyber liability insurance? Errors & omissions? In what amounts? Will they name the firm as additional insured?	□	□	□
-----------	--	---	---	---

**// WHY IT MATTERS**

\$5M cyber liability is industry minimum for SaaS handling sensitive data. \$10M+ for legal AI. Get a current Certificate of Insurance.

<b>E10</b>	"Most favored customer" clause — is the firm's pricing locked in if the vendor offers lower prices to comparable customers? Are renewals subject to a CPI or contractual cap?	□	□	□
------------	---	---	---	---

**// WHY IT MATTERS**

Vendor pricing in this market changes monthly. Lock the rate for the term and cap renewal increases at CPI + a defined ceiling. Otherwise your year-3 invoice triples.

SEC\_F

## Section F — Compliance & Audit

#	QUESTION	PASS	COND.	FAIL
---	----------	------	-------	------

<b>F1</b>	Compliance certifications: SOC 2 Type 2, ISO 27001, HIPAA-attested, FedRAMP (if government work), CCPA-compliant, GDPR-compliant. Which apply, and is the certification current?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Don't take logos as proof. Get the actual certificate, the audit period, and the scope statement. Many vendors are SOC 2 certified for the customer-facing application but not for the AI inference infrastructure.

<b>F2</b>	Right to audit — does the contract permit the firm or its third-party auditor to inspect the vendor's controls under NDA? Annually? On notice?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

For high-sensitivity engagements, the right to audit is what enables ABA Op. 512 supervisory compliance. Most vendors will not grant a free-text audit right; they will offer an annual share of the SOC 2 report. Negotiate.

<b>F3</b>	Sub-processor audit — does the vendor require its sub-processors (especially the model providers) to maintain compliance with the same standards?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
-----------	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**

Flow-through compliance. Without it, your no-training commitment is undermined at the model provider layer.

F4	Has the vendor been involved in any reported security incident, lawsuit, or regulatory action in the past 24 months? Disclose with detail.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>Public records help. Search the vendor name + "breach," + "lawsuit," + "regulator," + "FTC," + "complaint" in news and legal databases.</p>				

F5	EU AI Act, Colorado AI Act, NYC Local Law 144, California's SB-1047 (or its successor) — which jurisdictions does the vendor's product fall under? Is the vendor compliant with each?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>AI regulation is fragmenting fast. EU AI Act high-risk classification triggers documentation and oversight obligations. Colorado AI Act applies to consequential decisions. The vendor's compliance posture flows through to the firm.</p>				

F6	Government surveillance — does the vendor maintain a transparency report? Has it received national security letters? CLOUD Act requests? FISA orders? What was the disposition?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>For matters with US-government adverse parties, foreign government clients, or national-security implications, the vendor's posture toward government data requests is material. Most vendors do publish transparency reports; a vendor that does not is a question mark.</p>				

SEC\_G

## Section G — Court Acceptance & Verification



#	QUESTION	PASS	COND.	FAIL
G1	Can the vendor produce, on demand, an export of every prompt and output associated with a specific matter, suitable for production in discovery?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>AI work product is increasingly subject to discovery requests. Get the export mechanism in writing. Confirm the format is admissible (not vendor-proprietary).</p>				

G2	Has any court issued a sanctions order against an attorney specifically for using this tool? (Search the Damien Charlotin database — <a href="http://damiencharlotin.com/hallucinations">damiencharlotin.com/hallucinations.</a> )	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p><i>Lacey v. State Farm</i> involved CoCounsel + Westlaw Precision + Gemini. <i>Coomer v. Lindell</i> involved Copilot + Gemini + Grok. <i>Wadsworth v. Walmart</i> involved an in-house tool. The fact that a tool has been involved in a sanctions order is not disqualifying — the lawyer's verification protocol was the actual failure — but it shapes how you'll be questioned by opposing counsel.</p>				

G3	Does the vendor provide guidance, training, or templates for compliance with federal court AI standing orders (Castel S.D.N.Y., Wang D. Colo., Starr N.D. Tex., and the 300+ others)?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p><b>// WHY IT MATTERS</b></p> <p>Vendors that know the standing orders are partners. Vendors that say "compliance is your responsibility" are commodity providers — fine, but you'll bear more of the operational burden.</p>				

G4	If the firm receives an order to show cause or sanctions motion arising from AI use of this tool, will the vendor provide technical declarations or witness support?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p>// WHY IT MATTERS</p> <p>A vendor that will explain how its tool works, under oath, in your defense, is a different vendor than one that disappears.</p>				

G5	Does the platform present hallucination risk transparently — for example, by labeling outputs with confidence scores, source-quotation requirements, or "verified vs. inferred" tags?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p>// WHY IT MATTERS</p> <p>Stanford research found RAG-based legal AI tools still hallucinate 17-34% of the time. The honest tools surface this. The misleading tools claim "no hallucinations." Your \$9,800/citation lookup is not a hallucination-free guarantee.</p>				

SEC\_H

## Section H — Total Cost of Ownership

#	QUESTION	PASS	COND.	FAIL
H1	What is the per-seat or per-user list price? Is there a minimum seat count? Are there volume discounts?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<p>// WHY IT MATTERS</p> <p>Industry rough orders of magnitude (2026): Claude Pro / ChatGPT Plus = \$20-25/mo per user. Microsoft 365 Copilot = \$30/mo per user. Spellbook = \$99-200/mo per user. CoCounsel = \$100-200/mo per user. Harvey = \$1,200+/mo per user. A 50-attorney firm on Harvey is ~\$720K/year before discounts.</p>				

H2	Are there usage-based charges on top of seat fees? (Token usage, document processing, API calls, storage overage?) What are the rates, and what is the historical mean usage per attorney?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	--	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**  
 "Pay for what you use" sounds good and ends badly. Get a usage-based estimate from the vendor based on a comparable firm. Build a usage cap into the contract.

H3	What are the implementation costs? (Onboarding, training, integration, custom workflows.) Are these one-time, or recurring?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**  
 Implementation can be 50-100% of year-1 software cost. Build it into the budget.

H4	What is the cost of the integrations the firm needs? (Document management — iManage, NetDocuments, SharePoint. Practice management — Clio, MyCase, PracticePanther. Microsoft 365 / Google Workspace. Time entry. Billing.)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
----	---	--------------------------	--------------------------	--------------------------

**// WHY IT MATTERS**  
 Integration cost is rarely on the price page. A tool that's \$200/seat on the page becomes \$400/seat when you add the iManage connector. Confirm in writing.

H5	<p>What is the realistic 3-year total cost of ownership, including: software, implementation, integrations, training, internal IT support, vendor diligence (this checklist), insurance increment, and exit cost?</p>	□	□	□
<p><b>// WHY IT MATTERS</b></p> <p>The list price is rarely the real number. Force a 3-year TCO conversation. Use the result to compare against alternatives — including local-first architectures where the marginal cost per query is approximately zero.</p>				

SEC\_99

## Scoring framework

---

After completing all 60 questions, score each section and the overall posture.

SECTION	PASS = STRONG YES; CONDITIONAL = YES WITH CAVEATS; FAIL = NO, EVASIVE, OR VENDOR CANNOT ANSWER IN 5 BUSINESS DAYS	THRESHOLD
<b>A. Data Architecture</b>	10 questions. Any "Fail" on A1, A2, A3, A8, or A9 is a deal-breaker for D-3 use.	≥8/10 Pass; zero Fails on the listed critical questions.
<b>B. Training &amp; Improvement</b>	8 questions. Any "Fail" on B1 or B2 is a deal-breaker.	≥7/8 Pass; zero Fails on B1 or B2.
<b>C. ABA Rule 1.6</b>	8 questions. C1 and C4 are gating.	≥6/8 Pass; zero Fails on C1 or C4.
<b>D. Security</b>	8 questions. D1, D3, D4 are gating.	≥6/8 Pass; zero Fails on D1, D3, D4.
<b>E. Contractual</b>	10 questions. E1, E3, E4 are gating; the rest are negotiable.	≥6/10 Pass after negotiation; zero Fails on E1, E3, E4.
<b>F. Compliance</b>	6 questions. Practice-area dependent.	≥4/6 Pass.
<b>G. Court Acceptance</b>	5 questions. G1 is gating.	≥3/5 Pass; zero Fails on G1.
<b>H. Total Cost</b>	5 questions. Informational; influences ROI not safety.	All 5 answered; the answers fit the firm's budget.

## Decision matrix

- **All sections meet threshold:** Proceed to contract negotiation. Use this checklist as the basis for the schedule of representations.
- **One section below threshold, no critical fails:** Conditional. Negotiate the gaps. Re-score after vendor responses. Pilot before full rollout.
- **Multiple sections below threshold, or any critical fail:** Do not sign. The vendor either has architectural issues or contractual issues that no amount of marketing can fix. Look at alternatives — including local-first deployments.

// FINAL HONEST NOTE

No vendor will pass every question with a clean Pass on the first try. Some failures are negotiable; some are architectural. The point of this checklist is not to find a perfect vendor — it is to **know exactly where the gaps are before you sign**, so you can build operational controls around them, negotiate them out of the contract, or walk away.

If the vendor refuses to put answers in writing, that is itself an answer.

---

Published by **Knight CTO** – Fractional CTO for the practice of law.

This checklist is provided as a working template, free for any law firm or in-house team to use, customize, and distribute internally. It is not legal advice. Verify against your jurisdiction and your insurance carrier's requirements before relying on it. Vendor terms change continuously; the analysis reflects publicly available terms as of April 27, 2026.

Source: <https://knightcto.com/resources/vendor-evaluation-checklist.html> · © 2026 Knight CTO · Walnut Creek, California