Advisor's AI vendor audit checklist

What to ask an AI vendor before letting them touch client data

2026-06-16

A 20-question framework for evaluating any AI-enabled vendor against the obligations an SEC-registered RIA actually carries: Reg S-P, Rule 17a-4, Marketing Rule §206(4)-1, Rule 204-2, and Rule 206(4)-7. Each question maps to a structural concern, not a vendor's marketing language.

Use it to compare vendors on a level field. Anything that cannot be answered structurally — "trust us" answers, NDA-walled answers, "it's on the roadmap" answers — is a signal in itself.

Section 1 · Data handling (Reg S-P §248.30(b))

1. Where physically does my client data reside, and in which jurisdiction is it processed? Cloud region, sub-processor regions, and any cross-border processing. A vendor that cannot answer this in one sentence is not Reg S-P-ready.

2. Is my data used to train any AI model — yours, your vendor's, or your vendor's vendor? The honest answer is structural: contractual Zero Data Retention with the LLM provider, plus evidence of how the no-training contract is enforced (e.g., workspace ID, dedicated tenancy).

3. What is the inference-region commitment? US-only? Multi-region failover? If multi-region, what client-disclosure language covers it?

4. What is the retention horizon for inputs, outputs, prompts, and logs — separately? "Zero retention" only applies to one of these unless explicitly stated for all four.

5. Can you provide a written attestation of current AI posture on letterhead, on request, to a qualified reviewer under NDA? The answer should be yes, with a named accountable person and a typical turnaround time.

Section 2 · Immutability + recordkeeping (Rule 17a-4)

6. How does your platform satisfy SEC Rule 17a-4(f)(2)(ii) electronic-storage requirements for records you generate or store on my behalf? Specifically: is the storage write-once-read-many, with an enforceable retention horizon that survives operator action?

7. Can your operators (or your operators' operators) delete records before the retention window expires? The right answer is no, and the mechanism — bucket-level retention lock, cloud-provider object lock, third-party WORM — should be named.

8. How do you reconcile failed writes to immutable storage? A vendor that says "we update the failed row" does not understand immutability. The correct pattern is to emit *new* reconciliation rows that reference the failed row's ID without modifying it.

Section 3 · PII handling

9. What taxonomy do you apply to personal information at ingestion? Vendors with no taxonomy ship the same payload to every downstream service. Vendors with a taxonomy can describe what `high-PII` means and how it is treated differently from low-sensitivity context.

10. Does any field tagged `high-PII` (SSN, full account number, wallet seed phrase, biometric, authentication artifact) ever reach an LLM-bound payload? The answer should be no, by structural enforcement (middleware) rather than developer discipline.

11. Do you run an independent egress canary on LLM-bound payloads? Defense-in-depth pattern: a second guard that scans every outbound payload for residual PII. The pattern should be implemented independently from the first layer.

12. If a PII egress event occurs, what is your detection, alert, and notification path? Look for: fail-closed posture, automatic alerting, structured incident logging, named CCO loop-in.

Section 4 · AI governance + attribution

13. Which named foundation model handles which class of request, and how is the binding enforced? "We use AI" is not an answer. "We use Claude Sonnet 4.6 for X, Claude Haiku 4.5 for Y, and the binding is enforced by an ESLint rule that blocks hardcoded model strings" is.

14. Is the AI ever in a fiduciary position — recommending, deciding, or communicating directly with my client? The right posture for an SEC-registered RIA is HITL Tier 2: AI assists; humans validate and sign off; AI never makes a fiduciary decision and never communicates with the client directly.

15. What is recorded about each AI call, where is it stored, and for how long? At minimum: actor identity, model identifier, prompt hash or redacted prompt, response artifact, timestamp, retention horizon. If any field is missing, you have an attribution gap.

16. Is there a "principal chain" capturing which human authorized which AI session that invoked which tool? The single piece of information that answers "who is accountable if this goes wrong."

Section 5 · Open source + exit costs

17. What portion of your stack is open source, and what license governs it? Open source matters for two reasons: independent inspection (you can audit the patterns without an NDA), and exit cost (you can take a fork with you).

18. What is my exit cost from your platform if I decide to leave in 6 months? 24 months? 60 months? Concrete: format of my data on export, schema documentation, ongoing access to my historical records, and any contractual hold periods.

19. Are your Architecture Decision Records, dispatch protocols, and substrate patterns available for

review by qualified reviewers? Under NDA is acceptable. Wholesale refusal is a signal.

Section 6 · Subprocessor discipline

20. Where is your subprocessor list published, how often is it updated, and what is your client-notification window for material changes? Every Reg S-P-amended-rule-ready vendor has answered this. Vendors who treat their subprocessor stack as proprietary are imposing your compliance posture without telling you.

How to use this checklist

- **Score honestly.** A vendor that answers 18 of 20 well and acknowledges the two gaps is a better partner than one that ships marketing language on all 20.
- **Document the answers.** Save the vendor's responses verbatim; they become part of your Rule 204-2 records-retention substrate.
- **Re-score annually.** Vendor posture drifts; annual re-scoring catches the drift before an examiner does.
- **Surface gaps to your CCO.** Anything that scores poorly is an item for your Rule 206(4)-7 annual review.

Contributed by Protocol Wealth, LLC (CRD #335298). This document is intended for industry discussion and is not investment advice, legal advice, or a recommendation to adopt any specific vendor. Engineering substrate transparency · aggregate substrate material · not advisory performance. Comments and corrections welcome to contact@protocolwealthllc.com.