14 questions to ask an AI vendor before you sign.

The AI product you are buying carries legal obligations that do not stay on the vendor's side of the contract. Disclosure duties, documentation duties, and insurance exclusions attach to your deployment of their system. Asked before signature, these are diligence questions. Asked after, they are a claim file. The liability questions here are exactly that — questions for the vendor and for your own counsel. Recorded 17 Jul 2026 · current version at certian.com/guides/vendor-ai-questions

Q.01

"Which models are actually under the hood — yours, or someone else's — and will you name them in the contract?"

Why it matters: Most AI products are an interface over a third-party model, and every downstream obligation depends on knowing whose. Expect "that's proprietary architecture." The counter-move: you don't need the architecture — you need the legal identity of each model provider, as a named subprocessor, and notice when it changes.

Q.02

"Will you state in writing who builds, who modifies, and who deploys — the facts our counsel needs for the role analysis?"

Why it matters: The EU splits duties between providers and deployers; Texas and Colorado write obligations for developers and deployers by name. A vendor will rarely bind itself to a statutory role label — so ask for the facts instead: who develops and substantially modifies the system, who places it on which markets, who controls deployment. Your counsel makes the role call from that record.

Q.03

"Do you train on our data — and where is that stated in a document that can't change without notice?"

Why it matters: An answer buried in a settings toggle or an updateable policy page is a variable, not an answer. Get the commitment into the agreement itself: what is used for training, what is retained, for how long, and what notice you receive if any of it changes. Negotiable before signature; nearly immovable after.

Q.04

"What training-data disclosures have you — or your model provider — actually published?"

Why it matters: EU-market general-purpose model providers must publish a training-content summary on the Commission's mandatory template — in force since August 2025, with fining powers to €15M or 3% of worldwide turnover activating 2 August 2026 (Art. 53(1)(d), Art. 101). California's AB 2013 has applied since January 2026. These are public documents; a vendor who can't point to theirs is telling you what documentation you'll get when you need it.

Q.05

"If your product talks to our customers, whose job is the AI disclosure — yours in the interface, or ours in deployment?"

Why it matters: The EU requires that people interacting with an AI system be informed (Art. 50(1), from 2 Aug 2026); Korea requires advance notice of AI use (AI Basic Act, Art. 31). Disclosure can live in the vendor's interface, your deployment layer, or fall between the two — and "between the two" is the default unless the contract says otherwise.

© 2026 Certian, Inc. · Educational record — not legal advice · Statements subject to verification · certian.com/guides/vendor-ai-questions

Q.06

“Do your outputs carry machine-readable provenance marks — which standard, and will your engineers answer that in writing?”

Why it matters: The EU’s marking duty and California’s SB 942 both become operative 2 August 2026. “We support content provenance” is a roadmap sentence — and the account executive on the demo can’t inspect metadata live anyway. Put it in the technical questionnaire: which standard (C2PA, IPTC, watermarking), on which output types, and a sample output with provenance metadata intact, in writing from their engineering team.

Q.07

“Which certifications do you hold — scope, certificate number, and which public registry confirms it?”

Why it matters: An ISO/IEC 42001 certificate covers a defined management-system scope, which may or may not include the product you’re buying. A SOC 2 report covers the controls it was scoped to cover — not an AI-governance credential by default. Scope statement and registry entry are checkable facts; a badge on a website is neither.

Q.08

“What documentation will you hand us — the package the law already obliges developers to provide?”

Why it matters: Colorado’s rewritten act (SB 26-189, eff. 1 Jan 2027, dates may move) requires developers to give deployers understandable technical documentation before use — intended uses, training-data categories, known limitations, human-review instructions. EU general-purpose model providers owe downstream integrators an information package (Art. 53(1)(b)). If a statute already obliges your vendor to produce this bundle, the question is whether you get it too — as a contract deliverable with a delivery date.

Q.09

“What can change without telling us — models, weights, subprocessors — and what counts as a ‘material update’ you must notify?”

Why it matters: Colorado’s rewritten act requires notice of material updates and substantial modifications within a reasonable time. Beyond any statute: a silent model swap changes your risk profile, your documentation’s accuracy, and possibly the role analysis from Q.02. Define “material” in the contract; the vendor’s default definition is “whatever we didn’t tell you about.”

Q.10

“When the compliance requirements change, who pays for the platform changes?”

Why it matters: The rules under this product are still moving — Colorado takes effect in 2027, Article 50’s final Guidelines landed only in late July 2026. When the next rule requires a new disclosure surface, is that a compliance update included in your subscription, or a “premium enterprise feature” priced at your next renewal? A commitment that regulatory-compliance updates are included in the base subscription is cheap to ask for now.

Q.11

“What counts as an incident, when do we hear about it, and what do you owe us then?”

Why it matters: Your insurer, your customers, and your own record depend on what the vendor’s notification commitment was — and standard paper defines the trigger narrowly (“a confirmed, material breach”). The scope of what counts as an incident, and how fast notice arrives, is precisely the clause your counsel redlines. The definition, the timeline, and the remediation owed must exist in writing before signature.

Q.12

“Where does your paper put liability for all of the above — caps, carve-outs, IP and indemnity?”

Why it matters: The question the previous eleven exist to inform — and a question for your counsel, not an answer this guide provides. A first draft’s liability allocation is a starting position; read where the caps sit, what the indemnities cover, and what’s carved out, then let counsel redline. Raised at diligence, this is leverage; raised at renewal, it’s an admission.

Q.13

“Which of these risks does your own insurance actually cover — and does your E&O policy now carry an AI exclusion?”

Why it matters: An indemnity is worth the balance sheet — or the insurance — behind it. Carriers have attached AI exclusions to commercial policies at renewal since January 2026, including professional lines. A vendor whose own E&O excludes AI risk is offering an indemnity their insurer may not stand behind. Their answer belongs in your counsel’s file and your broker’s.

Q.14

“What do we get at exit — our data, our fine-tunes, our logs, and the records we’d need to show a regulator two years from now?”

Why it matters: Disclosure and documentation duties outlive the subscription. If a regulator asks in 2028 about a decision the system touched in 2026, the records either left with the vendor or they didn’t. Exit terms are the cheapest clause to negotiate at signature and the most expensive at termination.