Eden Lord
CONSULTING

PUBLISHED MAY 2, 2026

Five questions to ask before signing an AI vendor contract

Most AI vendor evaluations fail before they begin. Not because the buyers aren't smart. Because the framework being used was built for traditional software and doesn't survive contact with AI in regulated work.

Here's what that looks like. Demos drive the agenda. Feature matrices win the day. Security review shows up late. Total cost of ownership gets reduced to license cost. Six months into deployment, the tool doesn't fit how people actually work.

Five questions, asked in this order, would have caught it.

1. Does this fit how our people actually work?

Map the workflow before you talk to any vendor. Write down who does what, what tools they use, where the friction is, what depends on what. If you skip this, the vendor's demo silently becomes your reference point — and the vendor's conception of your workflow is rarely accurate.

The right pilot tests real work your team does this week. Not the cases the vendor curated. Not the scenarios the vendor's solutions engineer practiced on. The cases on your team's actual desk on a Tuesday in October.

Watch for one specific failure mode. Demos showcase the tool performing well on tasks the vendor selected. Pilots show the tool performing on tasks you selected. The gap between the two is the gap between buying a demo and buying a tool.

2. What's the vendor's regulatory posture?

Either they have it or they don't. SOC 2 Type II at minimum. HIPAA and BAA willingness for healthcare. FERPA and COPPA for education work involving minors. Professional responsibility frameworks for legal. Data and model governance. Incident response history. Mature posture means documentation produced on request. Not verbal assurances. When regulators ask, they want documentation. Not promises.

On PHI specifically. Protected Health Information is the dominant regulatory exposure in healthcare AI. Where does the data sit. Who has access to it. What gets logged. What happens to it in the model. What gets retained. What gets deleted on request. Whether the vendor trains on customer data and under what terms. Whether prompts and outputs are stored, and where. Whether de-identification is real or theatrical. These questions deserve specific, documented answers — not vendor reassurances.

Be wary of vendors who guarantee proper handling and compliance. Compliance is not something a vendor can guarantee on your behalf. They can document their controls. They can sign a BAA. They can produce audit reports. They cannot guarantee that the way your organization deploys their tool, integrates it with your workflows, trains your staff on it, and monitors its use will be compliant — because that part isn't theirs to guarantee. Vendors who promise compliance as a feature are either telling you what they think you want to hear or don't understand the regulatory framework themselves. Neither is the vendor you want.

The same pattern exists across other regulated domains. In legal, vendors who guarantee privilege preservation. In education, vendors who guarantee FERPA compliance through a checkbox in their settings. In financial services, vendors who guarantee fair-lending alignment because the model is "audited." All of these are tells. Compliance is something organizations achieve through deployment discipline. Vendors enable it. They don't deliver it.

A note on IP. Who owns what — your data, your prompts, the model outputs, and any models trained or fine-tuned on your inputs — is the next regulatory question vendor evaluations will need to address systematically. It is genuinely unsettled, the contract language is evolving fast, and most evaluations don't ask the right questions yet. That topic deserves its own treatment. There will be one from this practice in the coming months.

3. What's the real total cost — including exit?

Most models stop at licenses plus implementation. The categories that drive long-term cost get skipped: integration, ongoing operations, and exit if the tool needs to be replaced.

A complete cost model has four categories. Direct costs — licenses, seats, usage. Integration costs — connecting to existing systems, identity management, data pipelines, custom development. Operational costs — training, change management, ongoing administration over the life of the engagement. Exit costs — what happens to your data, workflows, and dependent processes if you switch vendors in eighteen months.

Exit costs are the most consistently undercounted category, and often the largest. Tools that ingest your data, train models on it, and produce outputs that downstream processes depend on become genuinely hard to replace. The exit question has to be asked at procurement time, when you still have leverage to negotiate exit terms. Not at the moment exit is needed.

4. How viable is this vendor over our planning horizon?

Vendor viability in AI requires reading several signals at once. Financial position — funding stage, runway, revenue trajectory — is the starting point. It is not sufficient. The AI vendor market is full of well-funded companies whose strategic position is weaker than their balance sheet suggests.

Strategic position is the substantive question. What's this vendor's differentiation in a market where incumbents are rapidly adding AI features? Vendors whose differentiation depends on capabilities incumbents will replicate in twelve months are fragile, no matter the balance sheet. Ask: who's likely to acquire this vendor in the next twenty-four months, and what would that mean for our relationship with the tool?

Roadmap alignment is the last viability question. Vendors pivoting toward enterprise may deprioritize your use case. Vendors pivoting toward consumer may abandon enterprise altogether. The roadmap conversation should be specific, recent, and ideally with the product leadership rather than sales.

5. What's the exit path if this doesn't work?

Asking exit questions before signing is one of the strongest signals of evaluation maturity. The organizations that do this get better contract terms, better implementation discipline from the vendor, and meaningful protection if the relationship needs to end.

Exit planning has four dimensions. Contractual — notice periods, data return obligations, transition support written into the engagement. Technical — data portability, workflow migration paths, the actual work of unwinding integrations. Operational — retraining, change management, the human cost of switching. Strategic — whether viable alternatives exist and what switching to them would require.

Vendors who answer exit questions clearly should be preferred over vendors who deflect. Deflection is a tell. It usually means the vendor hasn't built the operational infrastructure to support customer exits, which means when an exit becomes necessary, it'll be unilateral on your side and unsupported by them.

For healthcare, exit carries the PHI question raised in Question Two. For legal, exit carries questions about privilege and work product preserved in vendor systems. For education, exit carries questions about student data retention and parental consent obligations. The exit conversation isn't generic. It's specific to what's regulated about your work.

The order matters

Demos belong in Question Four — vendor viability — but most evaluations start there. Workflow first. Regulatory second. Cost third. Viability fourth. Exit fifth. Vendors who fail Question One don't get evaluated on Question Two. Vendors who fail Question Two don't consume time on Question Three. By the time you reach demos, the candidates left have already cleared the hurdles that most evaluations encounter only after a vendor is signed.

This sequence produces a different conclusion than the standard one. Vendors with impressive demos but poor regulatory posture or poor workflow fit get filtered before evaluation resources are spent on them. Vendors with quieter demos but mature posture and strong workflow fit advance further than they typically would.

The right vendor is sometimes the one the demo cycle would have eliminated first.

Eden Lord, Principal
eden@edenlord.com