

THE AI OFF SWITCH

7 questions to ask any vendor before you let AI touch your operations — and the answers that should make you walk away

Every software vendor you work with is adding AI to their product this year. Most of them will present it as an upgrade you should be excited about. Before you say yes to any of it, ask these seven questions — out loud, in the meeting, with someone taking notes.

You don't need to be technical to ask them. You just need to be willing to sit through the silence that follows.

1. "Where does our data go, exactly — and does it train your models?"

The answer you want: A specific, written answer. Your data stays in your system, the AI processes it under enterprise terms that prohibit training, and they can show you the clause.

Walk away if: They say "it's all secure" without specifics, or the answer requires a follow-up call with someone who never materializes.

2. "What can the AI do without a human signing off? Show me the list."

The answer you want: A short, explicit list of actions the AI can take alone, and a longer list of actions that stop and wait for a named person to approve. The dividing line should be yours to set, not theirs.

Walk away if: The answer is "it's fully autonomous" delivered as a selling point, or "it can't really *do* anything" (then what are you paying for?).

3. "Show me the log of everything the AI did last week."

The answer you want: They open a screen and show you: every action, timestamped, with what changed and who — or what — did it. If they can produce this in the sales meeting, it exists. If they can't, it doesn't.

Walk away if: "We're working on audit features." Logging is not a roadmap item. It's the foundation.

4. "How do I turn it off — right now, mid-task?"

The answer you want: One switch, immediate effect, including work already in progress. Bonus points if they demonstrate it live.

Walk away if: Turning it off requires a support ticket.

5. "Does the AI follow *our* permission rules, or its own?"

The answer you want: The AI operates under the same access controls as your employees. If a junior clerk can't approve a payment, neither can the AI acting in that role.

Walk away if: The AI has its own special access "to work properly." That's a master key with a friendly name.

6. "When it makes a mistake, who is accountable — and how would we even know?"

The answer you want: Mistakes are visible in the log, consequential actions required human approval anyway, and the vendor names who owns remediation.

Walk away if: They tell you their AI doesn't make mistakes. Every AI makes mistakes. Vendors who say otherwise are telling you they haven't looked.

7. "If we leave you next year, what do we keep?"

The answer you want: Your data in a usable format, at minimum. If it's custom-built software: the source code, the documentation, all of it — yours.

Walk away if: Leaving means starting over. That's not a software vendor. That's a landlord.

THE PATTERN BEHIND ALL SEVEN

Notice what these questions have in common: none of them ask how *smart* the AI is. Smart is the part every vendor will happily demo. The questions that matter are about **control** — who holds it, whether it's real, and whether it survives the relationship ending.

AI you can't see, can't limit, and can't shut off isn't an employee. It's a liability with good marketing.

Slingr builds custom business software you own outright — including the AI agents we build into it. They work like employees: scoped permissions, every action logged, human sign-off on what matters, and an off switch that works.
slingr.io