13 questions to ask your IT provider about AI controls.

AI arrived in your business whether anyone approved it or not. Your IT provider manages your networks, devices, and business tenant — your essential partner in finding it and putting boundaries around it. These questions establish what they can see, what they can control, and what they can hand you when your insurer, your lawyer, or your biggest customer asks. Recorded 18 Jul 2026 · current version at certian.com/guides/it-provider-ai-controls

Fair scope, before the first question: your provider sees what you pay them to manage. Personal phones, home computers, and tools nobody told them about sit outside that boundary — a technical fact, not a service tier; claims of total visibility outrun what any toolset can do. These questions are not a test of omniscience — they are the agenda for deciding, together, where the boundary should sit and what it costs to move it. A question outside your current agreement calls for a scope conversation, not disappointment.

Q.01

“What AI is actually running on our managed devices and in our business tenant right now?”

Why it matters: The inventory question: approved tools, AI features switched on inside existing software, and the unofficial layer visible on managed equipment. Expect: “we’d have to run a discovery scan for that.” Fair answer once — it’s the follow-through that’s the data.

Q.02

“Which of our existing subscriptions have added AI features since we bought them — and who decided whether they’re on?”

Why it matters: Features arrive by product update, defaulted on, processing your data. Someone is deciding — the question is whether it’s your provider, you, or the vendor’s default. The strong answer comes with settings screenshots.

Q.03

“For the AI tools we approve, what do the vendors’ own terms say about our data — and what do our admin settings actually control?”

Why it matters: Scoped honestly: your provider can read what a vendor publishes and configure what the admin console exposes — retention, training opt-outs, tenant boundaries. They cannot see a vendor’s backend. For the contract-side version, see certian.com/guides/vendor-ai-questions.

Q.04

“Do we have an AI acceptable-use policy — and if we implement one, would our staff recognize the tools named in it?”

Why it matters: A policy naming tools nobody uses is shelf-ware. Division of labor: the provider supplies the technical reality; management writes and owns the policy from it.

Q.05

“Can you restrict or approve AI tools by technical policy — and what can’t be controlled?”

Why it matters: DNS and browser controls, endpoint policy, tenant settings, app-approval flows — core provider craft; the strong answer comes with a dashboard. The answer that includes its own limits is the one you can build on.

Q.06

“If an employee uses a personal phone or home computer to put our data into an AI tool, what can anyone actually do?”

Why it matters: The honest answer for most small businesses: technically, very little — the network edge isn’t a wall anymore. The real controls are upstream: which data staff can reach, what the policy says, whether anyone’s been trained.

Q.07

“If we put AI in front of customers, who — among everyone we pay — owns the disclosure duties?”

Why it matters: Disclosure obligations for AI that interacts with people are in force or arriving (EU from 2 Aug 2026; state laws on the register’s map). The chatbot may be the web agency’s build, a SaaS vendor’s host — while the duty may sit with your business. Find the owner on purpose.

Q.08

“When our business insurance renews with new AI language in it, what documentation can your systems generate for us?”

Why it matters: Carriers have attached AI exclusions to commercial policies at renewal since January 2026. The useful things are exports a provider’s systems can actually produce: inventory, control settings, policy deployment records. What it means for coverage is your broker’s call.

Q.09

“If staff try to use unapproved AI tools on our managed network, do we have the visibility to know?”

Why it matters: Scoped to reality: what logging exists today, what would trigger an alert, and what additional tooling you’d need to buy. “We’d need to add a tool for that” is a legitimate answer with a price tag.

Q.10

“When you evaluate software for us, is AI part of the vetting — and what do you ask vendors?”

Why it matters: A provider using a real vetting list is doing invisible work worth paying for. The register keeps a 14-question version at certian.com/guides/vendor-ai-questions.

Q.11

“What AI questions are your other clients asking — and what should we be asking that we haven’t?”

Why it matters: A provider seeing thirty businesses has pattern knowledge no single client has. The answer’s specificity is the tell.

Q.12

“What of this is included in our current agreement, and what requires a separate project scope?”

Why it matters: Inventory work, policy support, and advanced controls are often outside standard support — that’s legitimate; assessment work has real cost. Define the boundaries and the costs before an insurance deadline forces a rushed deployment.

Q.13

“What do you need from us to do any of this well?”

Why it matters: Most of the above fails without leadership sponsorship: a named internal owner, decisions about risk appetite, backing for the policy, a budget that matches the ask. A specific list back is the sign of a provider thinking in shared-responsibility terms.