13 questions to ask an AI governance consultant before you hire one.

“AI governance consultant” is a title anyone can print this year. Underneath it are people doing genuinely scarce work — and people reselling panic with a framework’s name on it. These questions are built to tell them apart. One rule before the first call: the deliverable you need is almost always a document your lawyer, broker, or assessor can use. Recorded 18 Jul 2026 · current version at certian.com/guides/ai-governance-consultant

Q.00

The one to ask yourselves first: “Do we need a consultant, or do we need a decision?”

Why it matters: Much first-engagement work is inventory + policy + owner-naming — startable by an able internal lead plus your IT provider (the register’s guides are free). Hire for speed, an outside eye, or a deliverable someone else will rely on. Consultants describe their worst engagements as the ones where the client wouldn’t name an owner or decide a risk appetite. Do the internal wrestling first; it’s what makes the engagement finishable.

Q.01

“What will exist when you’re done that doesn’t exist now?”

Why it matters: Accept: a named artifact — an AI-use inventory, a policy set, a documentation package for renewal or procurement. Be wary of “alignment” or maturity language without an operational artifact. Expect “every engagement is custom.” Counter: “then describe the exact artifacts you delivered for your last client.”

Q.02

“Where does your operational expertise end — and when will you tell us to take the work to a lawyer, a broker, or a third-party assessor?”

Why it matters: The border is a relay, not a wall: a good consultant drafts the inventories, data maps, and technical policy language that counsel’s review, the broker’s placement, and an assessor’s review are built on — and knows where their leg of the race ends. Legal interpretations, coverage promises, or implied binding assessments are chairs they don’t hold. The strong answer names the handoffs unprompted.

Q.03

“What does your credential actually test — and what doesn’t it?”

Why it matters: Credentials here are new; some substantial, some a webinar with a certificate. Vet what the holder had to demonstrate, and notice whether they volunteer the limits. There is no chair this field requires you to hold — which is exactly why the vetting burden sits with you.

Q.04

“Which frameworks do you work from — and which specific controls would you exclude for a business our size?”

Why it matters: NIST AI RMF and ISO/IEC 42001 are the common scaffolding; “we tailor to your risk appetite” is rehearsed patter. Make them name controls that don’t apply to you, and why. A framework without specific, named exclusions is a template with a markup.

Q.05

“How does the governance documentation you build integrate with what we already run — privacy workflows, data-subject requests, incident response, existing controls?”

Why it matters: AI governance doesn’t live on an island: an AI failure lands in your existing incident process; AI data handling lands in your existing privacy obligations. A standalone committee and policy that map to nothing you already operate is a redundant silo. The strong answer starts with what you have.

Q.06

“Will the person pitching this engagement be the one writing the deliverables and running the interviews?”

Why it matters: The principal wows the sales call; the drafting lands with a junior — or a generative-AI prompt. Neither is disqualifying if disclosed and priced accordingly: lock in whose work you’re buying, who reviews it, and what happens to the rate when the answer changes.

Q.07

“Do you resell, refer, or take fees from any tool, platform, insurer, or law firm you might recommend?”

Why it matters: Referral economics aren’t disqualifying — undisclosed ones are. The strong answer is a complete list, volunteered in writing. (Standard we hold ourselves to: every commercial relationship on Certian’s pages is disclosed on the page where it appears, including the founder’s own practice.)

Q.08

“Who owns the work product — and can our lawyer, broker, or assessor use it without you in the room?”

Why it matters: Deliverables that only function with the consultant attached are a subscription wearing a deliverable’s clothes. Inventories, data maps, and documentation should be yours, editable, and legible to the third parties they exist to feed.

Q.09

“What does this cost, structured how — and what’s the milestone that proves value?”

Why it matters: Fixed scope with named deliverables beats open retainers for a first hire — but scope honestly: a sprawling environment with no data mapping may genuinely need a phased baseline, and saying so with a defined phase-one milestone is scoping, not padding. No milestone at all is a forecast about their revenue, not your risk.

Q.10

“The rules moved three times this year. How do you stay current — and what did you change your advice about?”

Why it matters: Anyone tracking this field honestly has reversed something in 2026 — deadlines moved; a state act was repealed and rewritten. “Nothing’s changed our approach” is the wrong answer even when delivered confidently — especially then. Dated record changes: certian.com/changelog.

Q.11

“Walk us through an engagement that didn’t go well. What happened?”

Why it matters: Standard vetting, extra force in a field this young: nobody has ten years of AI-governance references. Honest war stories from adjacent work are a fine answer; a flawless record in a two-year-old discipline is not.

Q.12

“Do you carry professional liability insurance — and does your policy cover the AI advisory work you’re proposing?”

Why it matters: Carriers have been adding AI language at renewal since January 2026, and coverage that predates the practice may not map to it. A consultant who has checked their own policy against the work they’re selling has taken the subject matter personally.