AI-GOVERNANCE-JOBS.COM / INSIGHTS

How to Ace the AI Governance Take-Home

A real interview scenario, the exact things that separate a strong candidate from a fluffy one, and the questions a sharp hiring manager will ask next. For candidates and recruiters both.

Most AI governance interviews fall apart in the same spot. The candidate knows the buzzwords. They can say "NIST" and "EU AI Act" in a sentence. But ask them to actually govern a live system with real customer data and real business pressure, and the answers go soft.

So here is a take-home that has no place to hide. Give it to candidates. Or study it, if you are the one sitting on the other side of the table. Either way, read the green flags and red flags at the bottom. That is where the interview is really won or lost.

THE SCENARIO — SEND THIS TO THE CANDIDATE

The context. Your company is launching an internal chatbot powered by a large language model. It handles tier-one technical issues. To do that, it ingests historical customer chat transcripts, some of which contain personal data, along with account information. And it does one more thing: it suggests account upgrades on the fly, based on who the customer is and what their account looks like.

Your task. Write a two-page executive briefing for the Chief AI Officer. Evaluate the risk. Propose a compliance framework. Hit these three pillars, and do not pad them:

1. **Name the risks.** Identify the three biggest risks baked into this deployment. Think across data privacy, model behavior, and how regulators would classify what you are building.
2. **Align to a framework.** Tell us which parts of the NIST AI Risk Management Framework or ISO 42001 matter most for *this* system. Not the whole framework. The parts that

earn their place here.

3. **Test it before it ships.** Propose two concrete guardrails or tests to run before launch. Concrete. Something an engineer could build on Monday.

The Format

Time to complete	48 to 72 hours. This is a take-home, not a pop quiz. Respect their time and their day job.
Deliverable	A two-page document, or a five-slide deck. Their choice. How they use the space tells you something.
The pitch	First 15 minutes of the next round, they present the briefing to you. You play the Chief AI Officer. Watch how they handle a room, not just a document.

What Actually Separates the Strong From the Fluffy

Three things reveal whether someone has done this work or just read about it. Here is what each one looks like when it is real, and when it is smoke.

1. Do They Understand Data Privacy in the Real World?

RED FLAG

"We will just delete all the personal data." That sentence tells you they have never built a training or fine-tuning pipeline. It almost never works that cleanly, and a candidate who says it does not know what they do not know.

GREEN FLAG

They talk about programmatic data masking. Data minimization as a principle, not a slogan. Or a retrieval-augmented setup that keeps sensitive data out of the model weights entirely. They know the difference between data the model *reads* and data the model *becomes*.

2. Do They Know AI Risk From Plain IT Risk?

RED FLAG

They dress up basic cybersecurity as AI governance. Firewalls and access controls matter, but if that is the whole answer, they are solving last decade's problem.

GREEN FLAG

They name the failure modes that are specific to these systems: hallucination, prompt injection, model drift. And here is the tell. They notice that the upsell feature quietly changes the risk picture. Automated upselling based on account status is consumer profiling, and under something like the EU AI Act, that can push the whole system into a higher-risk category. A candidate who catches that is thinking like a regulator, not a coder.

3. Can They Govern Without Grinding the Team to a Halt?

RED FLAG

They pile on compliance hurdles that would stop engineering cold. Governance that halts shipping is not governance. It is a veto with a nicer name, and the business will route around it.

GREEN FLAG

They propose automated guardrails that catch toxic inputs and outputs on the fly, so a lawyer does not have to read every chat transcript by hand. They design controls that let the team keep moving. That is the mark of someone who has actually shipped governance inside a real company, not just written policy about it.

Three Questions to Ask When They Pitch

The document shows you how they think alone. These questions show you how they think under pressure. Ask them during the 15-minute pitch and watch closely.

1. "The engineering lead tells you the RAG setup you recommended adds real latency, and it is hurting the customer experience. The CEO wants this shipped Friday. What do you actually do?" *You are testing whether they can hold a line on risk while still being someone the business wants in the room.*

2. "Six months after launch, the bot starts recommending upgrades to customers who clearly cannot afford them, and one complains publicly. Your framework was supposed to prevent this. Where did it break, and what did you miss?" *You are testing for intellectual honesty and whether they build for drift and monitoring, not just launch day.*

3. "Explain to me, as if I run the sales team and I am skeptical, why this upsell feature counts as profiling and why that should change anything about how we build it." *You are testing whether they can translate regulation into plain language for a hostile audience. This is the whole job.*

FOR RECRUITERS

A candidate can memorize a framework. They cannot fake having governed a live system with real stakes. This exercise surfaces the difference in about an hour of your time. Use it to screen for judgment, not vocabulary.

Hiring for AI governance, risk, or compliance? Post the role where the people who can answer these questions are already looking.

AI-Governance-Jobs.com — The board built for AI governance and AI compliance roles.
Insights for the people building the guardrails.

Read the latest version online: **https://www.ai-governance-jobs.com/resources/ai-governance-take-home/**
Source: AI-Governance-Jobs.com · GRC Careers · free AI governance career resources.