

AI GOVERNANCE SERIES

07

Questions to answer before external users touch **AI-generated insight.**

Most teams can already put an AI answer in front of a customer. Very few can prove what that answer said, who approved it, and why.

AI GOVERNANCE

EU AI ACT

EXTERNAL ANALYTICS

COMPLIANCE

THE REPORTING HUB

thereportinghub.com · Oct 5, 2026

WHY THIS MATTERS NOW

The gap is not the model. It is the paperwork.

Most teams can already put an AI answer in front of a customer. Very few can prove what that answer said, who approved it, and why.

That gap is now dated. Under Article 50 of the EU AI Act, AI systems that interact with people or generate content must disclose it. The obligation applied from 2 August 2026. Systems placed on the market before 2 December 2026 get a short grace period, and it ends on that date. High-risk obligations moved out to December 2027 and August 2028, so the deadline moved. The exposure did not.

The readiness data says the same thing. In a 2026 survey of 525 US enterprises with 500 or more staff:

74% believe they would pass an AI compliance audit today.	27% report a fully mature AI governance program.	46% already run AI agents in production.	29% have taken action on EU AI Act requirements.	36% of boards discuss third-party AI risk.
---	--	--	--	--

Read those together. Half of these companies have agents live. One quarter can govern them. **Three quarters think they are fine.**

Internal AI forgives this. An employee who gets a wrong number asks a colleague. A customer who gets a wrong number forwards it to their own board, their auditor, or their lawyer.

The seven questions below are the ones to answer before that happens. Each has a plain pass test. If you cannot answer one of them in a sentence, you do not yet have an answer.

01

QUESTION · PROVENANCE

Can you show where the answer came from?

An external user reads a sentence. Behind it sit a dataset, a filter, a date range, and a calculation. If the system cannot name those four things, the sentence is an assertion, not an insight.

A

Dataset

B

Filter

C

Date range

D

Calculation

WHAT GOES WRONG

A customer quotes your AI answer back to you in a renewal meeting. Your analyst cannot reproduce the number, because nobody recorded which measure the model used.

PASS TEST

Click any AI-generated statement and see the source dataset, the filter applied, and the formula behind the figure. No manual lookup in a second tool.

Most internal assistants do not clear this bar today. Microsoft's own guidance for governed environments is to open the source report and confirm a consequential figure before you forward it. That is a reasonable rule for an employee. It does not survive contact with a customer portal, because nobody is standing between the model and the reader.

02

QUESTION · APPROVAL

Who approved it before it left the building?

There is a difference between AI that answers a question and AI that publishes. The moment the reader is outside your company, publishing is what is happening.

WHAT GOES WRONG

An agent generates a commentary paragraph on a client dashboard overnight. It is accurate, fluent, and says something about churn that your account team was not ready to say.

PASS TEST

Name the person or role who signs off AI-generated content before an external user sees it, and show where that sign-off is recorded. If the answer is "the model decides", you have no gate.

TWO WORKABLE DESIGNS**DESIGN A****Review before publish.**

A human clears each new AI output class once, then it flows.

DESIGN B**Scope so tightly that review is unnecessary.**

The agent can only answer from content a human already approved.

Both are defensible. **No gate at all is not.**

03

QUESTION · SCOPE

Can each audience see only its own data?

Row-level security on a report is a solved problem. An AI agent is a new surface over the same data, and it does not automatically inherit the discipline.

WHAT GOES WRONG

A customer asks the agent how they compare to other customers. The model has the benchmark data, because the semantic model has it, and it answers.

PASS TEST

One agent per audience, each scoped to its own content, configured by an administrator rather than by a prompt. Ask your vendor to show you the configuration screen, not the prompt.

PROMPT-LEVEL RESTRICTIONS	ARCHITECTURE-LEVEL RESTRICTIONS
Instructions	Constraints

Prompt-level restrictions are instructions. Architecture-level restrictions are constraints. Only the second one is defensible when a customer asks how you know their data stayed separate.

04

QUESTION · DISCLOSURE

Does the customer know it is AI?

This one is now a legal question, not a design preference.

Article 50 of the EU AI Act requires that people are told when they are interacting with an AI system, and that AI-generated content is marked as such in a machine-readable form. The obligation applied from 2 August 2026. Systems already on the market have until 2 December 2026.

2 AUG 2026**Article 50 obligation applies****2 DEC 2026****Grace period ends for systems
already on the market****DEC 2027 · AUG 2028****High-risk obligations****WHAT GOES WRONG**

Your portal labels nothing. A European customer's compliance team runs a review and finds AI-generated narrative presented as analyst commentary.

PASS TEST

AI-generated content carries a visible label and a machine-readable marker, and the disclosure survives export to PDF, email, and embed.

The second half is where most implementations fail. The label sits in the web interface and disappears the moment the content leaves it.

05

QUESTION · REPRODUCIBILITY

Can you reproduce the answer in 18 months?

Disputes arrive late. A customer challenges a figure long after the model version changed, the data refreshed, and the person who built the agent left.

WHAT GOES WRONG

You can see that an answer was given. You cannot show what it said, or what the data looked like when it said it.

PASS TEST

For any AI answer delivered to an external user, retrieve the question, the answer, the user, the timestamp, the data state, and the model configuration. Retention set deliberately, not by platform default.

THE ANSWER RECORD

Question	Answer	User
Timestamp	Data state	Model configuration

30 days

Check the default. Telemetry on common analytics platforms is often kept for 30 days. Contract disputes do not run on a 30-day clock.

06

QUESTION · RESIDENCY

Where does the inference actually run?

Your data governance story ends where your tenant ends. For most AI features, that is earlier than people assume.

WHAT GOES WRONG

A security questionnaire asks which subprocessors see customer data during AI inference. Nobody on your side knows, because the feature was enabled, not architected.

PASS TEST

Name the environment where inference runs, name every third party that sees the data in transit, and show it in writing. If the answer is your own cloud tenant with no egress, say so plainly, because it is the strongest answer available.

01

Name the environment where inference runs

02

Name every third party that sees the data in transit

03

Show it in writing

This is the question regulated buyers ask first and the one that **decides whether external AI gets approved at all**.

07

QUESTION · INCIDENT PLAN

What is the plan when it is wrong?

It will be wrong sometimes. Governance is not the claim that it will not be. It is the ability to act fast when it is.

44%

Only 44% of surveyed enterprises have AI-specific incident response procedures. The other 56% will improvise in front of a customer.

WHAT GOES WRONG

A wrong figure reaches 40 client portals on a Friday. On Monday you still do not know which 40.

PASS TEST

You can answer three questions within one business day.

- | | |
|---|---|
| 1 | Who received the incorrect output? |
| 2 | How do we correct it at source, not just hide it? |
| 3 | Who tells the customer, and by when? |

That is the same set an auditor will ask. Writing it down before an incident costs an afternoon. Writing it during one costs the account.

SELF-ASSESSMENT

Score yourself in ten minutes

Take each question to the person who owns the system. Mark it only if they can answer in one sentence, without checking.

#	QUESTION	PASS MEANS	✓
1	Provenance	Source, filter and formula visible on every answer	☐
2	Approval	A named gate, recorded, before anything external	☐
3	Scope	One agent per audience, set by an administrator	☐
4	Disclosure	Visible and machine-readable, and it survives export	☐
5	Reproducibility	Full answer record retrievable 18 months later	☐
6	Residency	Inference environment and subprocessors named in writing	☐
7	Incident plan	Who got it, how to fix it, who tells them	☐

0-2 NOT YET

Do not put AI in front of external users yet. Fix provenance and approval first. They unlock the rest.

3-5 COMMON

The common position. You have a working system and a reporting gap. The risk is not the AI. It is that you cannot evidence it under pressure.

6-7 READY

Governance is now a thing you sell, not a thing you survive. Say so in your sales process, because most of your competitors cannot.

WHERE REPORTING HUB FITS

These seven questions describe a category, not a product.

Any team can answer them with enough engineering time.

Reporting Hub exists because most teams do not have that time, and because the questions are easier to answer when governance is built into the delivery layer rather than added after it.

Provenance and disclosure	BI Genius shows its sources, its reasoning, and the DAX behind every answer.
Scope	One agent per audience, each scoped to its own content by an administrator.
Approval	Nothing reaches an external user without your team's review.
Residency	Deployed in your own Azure environment. Analytics and AI inference both run inside your tenant, with no data egress.
Delivery	Each customer gets their own branded portal, from one deployment, with a record of who saw what and when.

Built on Microsoft Foundry. Microsoft Solution Accelerator Partner. 4.5 out of 5 on G2, with 300,000+ active users across 30+ countries. BI Genius Starter is included on every plan, with no separate AI contract and no Fabric capacity to buy.

See it on your own data.

30-day free trial, no credit card required.

Start a Free Trial →

SOURCES REFERENCED

EU AI Act enforcement timeline Article 50 transparency rules Schellman 2026 AI governance research