

A due diligence checklist for the EU AI Act.

\$4.1 Billion

Maximum potential exposure. Enterprise tier.

The 12 questions that decide which side of the fine you land on.

Twelve questions every compliance buyer should ask their vendor before **August 2, 2026**. Each one carries a specific financial exposure, scaled by company size. Below is the math.

WHY THIS MATTERS NOW

The architecture is being decided now. The fines arrive later.

On **August 2, 2026**, high-risk AI system obligations under Articles 9 through 17 and Article 26 become enforceable. The compliance platform a company chooses today is the platform a regulator audits in 2027.

ENFORCEMENT BEGINS Aug 2, 2026 Articles 9–17, Article 26	PENALTY CEILING \$41M €35M statutory, or 7% of turnover
--	---

Every vendor in the market claims readiness. Few can prove it under audit. The twelve questions that follow separate the claim from the proof, and each one carries a number.

THE PREMISE

A signed policy is not a verifiable result. Below is the math, by company size.

THIS IS NOT THE FIRST TIME EUROPE HAS MEANT IT

GDPR was the warm-up. The fines were already real.

Before the AI Act, there was GDPR. Regulators who were called toothless issued some of the largest penalties in corporate history. **The AI Act penalties are written to be higher.**

Meta	\$1.4B €1.2B · 2023	Largest GDPR fine to date. Unlawful EU–US data transfers.
TikTok	\$620M €530M · 2025	Unlawful transfer of European user data to servers in China.
Cumulative	\$8B+ €7B+ · since 2018	Total GDPR penalties across thousands of enforcement actions.

The AI Act ceiling is \$41M (€35M) or 7% of global turnover. GDPR maxed at 4%.

PART ONE

6

The Six Regulatory Invariants.

Six questions that describe what good architecture looks like. Each one is a property the EU AI Act requires and most vendors cannot deliver. **Ask these before you sign.**

- 01 Deterministic rules
- 02 Character-level sourceSpans
- 03 Unified raw JSON
- 04 Cryptographic receipt
- 05 Public immutable ledger
- 06 Byte-for-byte reproducibility

Q1 DETERMINISM

Do you use pure deterministic rules, with no LLM in the evaluation path?

If the evaluation layer is a language model, the verdict is probabilistic. Run the same document twice, get two different answers. **A regulator cannot audit a coin flip.**

Article 15 · Accuracy, robustness, and consistency of high-risk systems

POTENTIAL EXPOSURE IF THE ANSWER IS NO

SMALL / MID < \$585M turnover \$12–18M	LARGE \$5.8–11.7B turnover \$175–350M	ENTERPRISE > \$23B turnover \$700M–1.6B+
--	--	--

A probabilistic verdict is not a verdict. It is an opinion the model may revise tomorrow.

Q2 SOURCESPANS

Does every finding carry character-level sourceSpans?

A finding that points to "page 4, paragraph 2" is a reference. A finding that points to an exact byte range is evidence. **The difference is whether the audit can be defended.**

```
KLARITEE FINDING char_offsets: [18442, 18506]
```

Article 12 · Record-keeping and traceability of system events

POTENTIAL EXPOSURE IF THE ANSWER IS NO

SMALL / MID < \$585M turnover \$12–18M	LARGE \$5.8–11.7B turnover \$175–290M	ENTERPRISE > \$23B turnover \$525M–1.05B
--	--	--

Page-level attachments fail Article 12. Byte coordinates are the only defensible anchor.

Q3 RAW JSON

Can you export the entire evidence chain as one unified JSON document?

A dashboard shows evidence. A flattened, machine-readable JSON document produces it. **The regulator's team needs to parse the chain in their own environment, not screenshot yours.**

Article 11 · Technical documentation in machine-processable form

POTENTIAL EXPOSURE IF THE ANSWER IS NO

SMALL / MID < \$585M turnover \$9–18M	LARGE \$5.8–11.7B turnover \$140–235M	ENTERPRISE > \$23B turnover \$470–935M
---	--	--

Evidence shown is not evidence produced. One file. Every clause, every offset, every score.

Q4 CRYPTOGRAPHIC RECEIPT

Does every evaluation produce a cryptographic receipt?

A SHA-256 receipt binds the input, the engine version, and the output into one permanent record. **Without it, the audit trail is a database row an administrator can quietly edit.**

RECEIPT BINDS
input + engine_version + timestamp + output → SHA-256

Article 12 · Tamper-evident logging and record integrity

POTENTIAL EXPOSURE IF THE ANSWER IS NO

SMALL / MID < \$585M turnover \$12–18M	LARGE \$5.8–11.7B turnover \$175–350M	ENTERPRISE > \$23B turnover \$700M–1.4B
--	--	---

A database row is not a receipt. A hash is. One binds. The other can be rewritten.

Q5 PUBLIC IMMUTABLE LEDGER

Is your audit trail anchored to a public, immutable ledger?

If the record lives only in the vendor's private database, an administrator can rewrite history. **If the vendor is acquired, changes pricing, or goes dark, the audit trail goes with them.**

Article 12 · Independent, vendor-proof verifiability

POTENTIAL EXPOSURE IF THE ANSWER IS NO

SMALL / MID < \$585M turnover \$8–14M	LARGE \$5.8–11.7B turnover \$115–235M	ENTERPRISE > \$23B turnover \$350–820M
---	--	--

A private database is a record you are trusted to keep. A public ledger is a record you cannot alter.

Q6 REPRODUCIBILITY

Can you guarantee byte-for-byte reproducibility two years from now?

Article 12 is a reconstruction requirement. A regulator can reopen a 2026 evaluation in 2028 and demand the same result. **If the model has been updated, the endpoints changed, the prompts rewritten, the hash will not match.**

Article 12 · Lifecycle reconstruction of system functioning

POTENTIAL EXPOSURE IF THE ANSWER IS NO

SMALL / MID < \$585M turnover \$18M	LARGE \$5.8–11.7B turnover \$350M+	ENTERPRISE > \$23B turnover \$1.2B+
---	---	---

Reproducibility is not a marketing line. It is the requirement. Same input, same engine, same hash. Forever.

PART TWO

6

The Six `Record<string, any>` Questions.

Part One asked what good architecture looks like. Part Two is the stress test. One TypeScript pattern disables type checking at the exact boundary where your compliance score is produced. **These six questions find it.**

- 07 In production?
- 08 In the scoring engine?
- 09 Insurance coverage?
- 10 Disclosed in the TOS?
- 11 Line-by-line audit?
- 12 Will the regulator ask?

Q7 PRODUCTION PRESENCE

Is **Record<string, any>** present
anywhere in your production codebase?

Not staging. Not test. Production. The code that produces the compliance
output your regulators will eventually read.

THE PATTERN

Record<string, any> → type checking disabled

Article 15 · Technical robustness of the production system

POTENTIAL EXPOSURE IF THE ANSWER IS YES

SMALL / MID
< \$585M turnover

\$9–18M

LARGE
\$5.8–11.7B turnover

\$175–350M

ENTERPRISE
> \$23B turnover

\$585M–1.17B

"It is only in one place" is not an answer. **Ask where. Ask for the file path.**

Q8 THE SCORING ENGINE

HIGHEST EXPOSURE

Is that type in the engine that calculates the score, or only the interface that displays it?

These are not the same risk. Where the unguarded type sits determines whether the problem is cosmetic or structural.

DISPLAY LAYER

A formatting concern. The number was sound before it was shown.

SCORING ENGINE

The number itself is unguarded. The verdict was never safe.

Article 15 · Integrity of the scoring logic itself

POTENTIAL EXPOSURE IF IT IS IN THE SCORING ENGINE

SMALL / MID
< \$585M turnover

\$18–41M

LARGE
\$5.8–11.7B turnover

\$350–820M

ENTERPRISE
> \$23B turnover

\$1.6–4.1B+

A cosmetic flaw is forgivable. An unguarded score is the compliance verdict itself.

Q9 INSURANCE COVERAGE

If this pattern produced a faulty verdict and you were fined, would cyber insurance cover it?

Cyber insurers are now reviewing AI system liability clauses. **A fine caused by a vendor's type architecture may fall outside the policy.** Ask your insurer the same question. Get the answer in writing.

Coverage exclusion risk · All company sizes

POTENTIAL EXPOSURE IF COVERAGE IS DENIED

EXPOSURE, EVERY TIER

The full fine, plus defense costs

Unlike the other questions, this one does not scale by company size. **If the policy excludes it, the entire penalty and the cost of fighting it land on the balance sheet.**

"We assume it is covered" is not coverage. **Get the exclusion list. Read clause by clause.**

Q10 TERMS OF SERVICE

Does the terms of service disclose this pattern, or limit liability for what it affects?

If the disclaimer is buried in the indemnification clause, the risk was transferred to the buyer at signature. **Most compliance buyers never read clause 14.**

Article 13 · Transparency obligations toward the deployer

POTENTIAL EXPOSURE IF LIABILITY WAS SILENTLY SHIFTED

SMALL / MID < \$585M turnover \$12–23M + private lawsuits	LARGE \$5.8–11.7B turnover \$235–585M	ENTERPRISE > \$23B turnover \$820M–2.3B+
--	--	--

A buried disclaimer is still a disclaimer. **If you signed it, you accepted the risk.**

Q11 LINE-BY-LINE AUDIT

Can you provide a line-by-line audit report of where this type does and does not appear?

A claim is not an audit. **Ask for the file paths, the line numbers, and the code context.** A vendor who has done the work can produce it. A vendor who has not will offer reassurance instead.

WHAT A REAL AUDIT RETURNS

221 files · 15 chunks · 0 instances · file paths confirmed

Article 11 · Documented, inspectable technical evidence

POTENTIAL EXPOSURE IF NO AUDIT CAN BE PRODUCED

SMALL / MID
< \$585M turnover

\$9–18M

LARGE
\$5.8–11.7B turnover

\$175–290M

ENTERPRISE
> \$23B turnover

\$470–935M

"Everyone uses that pattern" is not an audit. A file-path report is.

Q12 THE REGULATOR'S QUESTION

Will an EU AI Act regulator ask about the type architecture of your compliance system?

Article 9 requires technical documentation. Article 12 requires logging integrity. Article 17 requires a quality management system. **Each one reaches the architecture that produces the verdict.**

The answer is yes.

Articles 9, 12, 17 · Documentation, logging, and quality management

POTENTIAL EXPOSURE WHEN THE QUESTION IS ASKED

SMALL / MID < \$585M turnover \$12–23M	LARGE \$5.8–11.7B turnover \$235–470M	ENTERPRISE > \$23B turnover \$600M–1.4B
--	--	---

The question is not whether the regulator will ask. **It is whether you will have an answer.**

THE EUROPEAN AI OFFICE

The regulator who checks this has the tools to check it.

The European AI Office is not a complaint box. It holds direct technical and financial enforcement powers, and it can use them without waiting for a court.

Source code access

Can request training data, technical documentation, and the source code itself.

Independent audits

Can conduct its own model evaluations rather than trusting the vendor's report.

Direct fines

Can impose corrective measures and penalties directly, up to 7% of global turnover.

Binding decisions

Issues binding rulings and coordinates national market surveillance authorities.

Mandatory cooperation. Direct model access. A 7% ceiling. The teeth are real.

TWELVE QUESTIONS, SUMMED

The full exposure, by company size.

Cumulative potential exposure across all twelve questions. Not a forecast. The ceiling the architecture leaves open if the answers are wrong.

Small / Mid-Size Under \$585M global turnover	\$117–410M
Large \$5.8–11.7B global turnover	\$1.17–3.5B
Enterprise Over \$23B global turnover	\$3.5B–7% of turnover

At the largest firms, 7% of turnover can exceed \$6 to \$12 billion.

CONSIDER WHAT YOU ALREADY KNOW

You have had years of warning.
You have had the questions.
You have had time to ask.

So is it unlikely, or highly likely,
that the EU follows through?

GDPR was called toothless. Then came the \$1.4 billion fine. **The AI Act was written with sharper teeth and a higher ceiling.** The only variable left is whether your vendor can answer twelve questions.

WHAT A VENDOR SHOULD BE ABLE TO HAND YOU

Every vendor claims a moat. Ask to see the receipt.

Klaritee answers all twelve questions the same way: with a downloadable artifact a regulator can re-verify in their own environment. **The architecture is the proof.**

THE ONLY MOAT THAT SURVIVES AN AUDIT

M.O.A.T. Mathematically Observable Audit Trail

AVAILABLE NOW

The Prove It proof package · a full production run, hash-verifiable

FORTHCOMING

resonancelabs.ai exposure calculator **COMING SOON**

Run the twelve questions against the platform you already own. The ones it cannot answer are the ones a regulator will press.

A CLOSING OBSERVATION

The rest of the market is helping companies write summaries about compliance.

Klaritee is compiling the receipt the regulator will eventually ask for.
One produces a document. The other produces proof.

Klaritee compiles *mathematically verifiable truth*.

Run the twelve questions against the platform you already own.