

RUNTIME AI COMPLIANCE ENFORCEMENT

Compliance, enforced in production.

A compliance policy does not stop a €35M fine. Enforcement does. ComplyEdge runs on every AI request, blocks violations on the sub-100ms deterministic path, and writes a legally citable record on every decision.

€35M

or 7% of global annual revenue
(EU AI Act Art 99)

The regulatory clock is now binding.

The EU AI Act does not ask whether your AI is safe. It penalises evidence you cannot produce, at audit time, months after the decision.

Feb 2, 2025	Article 5 prohibited practices in force
Aug 2, 2026	GPAI provider fines begin · Article 50 transparency
Dec 2, 2027	Annex III high-risk enforcement · Reg (EU) 2026/1744

Your AI governance policy says one thing. Your deployed models do another.

Regulators do not audit your documentation. They audit what ran, when, and what your system decided.

"Can you prove every AI decision your product made, right now, for any request in the last six months?"

Three approaches. Three audit failures.

GRC platforms & policy docs

Vanta, OneTrust, internal wikis.

Document what your AI should do.

Static, updated quarterly, disconnected from what runs in production.

✗ No runtime evidence. The auditor sees policy, not enforcement.

LLM-as-judge guardrails

A second model scores the first. A probability, 2 to 5 seconds of latency, and a different answer across runs.

✗ "The model said 0.87" is not a legal defence.

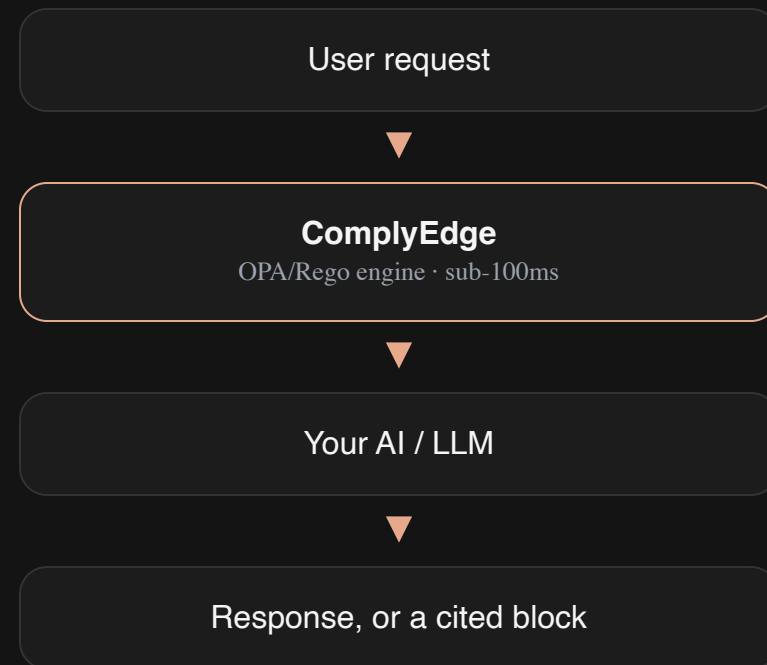
Homegrown regex & allow-lists

Hand-coded filters maintained by an overstretched engineer. Breaks on every model update.

✗ No reproducible record. Cannot cite an article.

Runtime enforcement. Every request. Legal citation on every decision.

ComplyEdge sits between your AI and your users. Every prompt and response is checked on the sub-100ms deterministic path against the EU AI Act and US regulations. Violations are blocked at the enforcement layer, with the exact article, rule ID, timestamp, and input hash on every decision.



Three steps. Three lines of code.

STEP 1 · INSTALL

```
$ pip install complyedge
```

STEP 2 · WRAP

```
from complyedge import \
    compliance_check

@compliance_check(
    jurisdiction="EU")
def agent(prompt):
    return llm.generate(prompt)
```

STEP 3 · ENFORCED

- Every input and output checked
- Violations blocked on the fast path
- Full audit record written
- Article and rule ID on every decision

Not a probability score. A legal record.

Every blocked request generates a Compliance Decision Record, the exact artifact a regulator demands during an investigation.

Article citation. Rule ID. Timestamp. Input hash. Reproducible, and hash-chained so any change is detectable.

COMPLIANCE DECISION RECORD

BLOCKED

OUTCOME Request blocked, violation detected

LEGAL BASIS EU AI Act, Article 5(1)(c)

VIOLATION Social scoring of natural persons

RULE ID rego-art5-1c-001

TIMESTAMP 2026-05-21 · 18:34:01 UTC

AUDIT LOG ✓ Recorded · hash-chained · verifiable

64 rules across 4 jurisdictions.

Open source and article-cited. Every rule is verifiable in the public repository, with the statute it enforces.

EUROPEAN UNION

64 YAML rules and 69 OPA/Rego policies covering Article 5 (prohibited), Article 50 (transparency), GPAI (Art 51, 53, 55), Articles 9 to 16 (risk, logging, transparency), and GDPR.

UNITED STATES & GLOBAL

16 US compliance rules plus universal protections: HIPAA, SOX, COPPA, TCPA, BIPA, CCPA, Colorado AI Act, NYC LL144, ECPA, PCI DSS, prompt injection.

64

YAML rules

69

Rego policies

16

US rules

4

jurisdictions

Benchmark-validated. Open source. Reproducible.

sub-100ms

deterministic hot path
(layer-1 microbenchmark p99)

100%

detection on blocked categories
60-prompt runtime benchmark

0

false positives
on the safe-harbour set

3

lines of code
no infra, no compliance team

Every number reproduces from the live API. Full corpus, prompts, and results in [scripts/benchmark/](#), Apache 2.0.

The deadline was the starting gun, not the finish.

Article 5 has been enforceable since Feb 2025. GPAI fines and Article 50 transparency began Aug 2, 2026. Since Aug 2, anyone can lodge a complaint with a market surveillance authority. The high-risk regime follows in Dec 2027. Readiness is the record you can produce today.

Feb 2, 2025

Article 5 prohibited practices in force.

Aug 2, 2026

GPAI fines · Article 50 transparency ·
complaint channels live.

Dec 2, 2027

Annex III high-risk enforcement, per
Reg (EU) 2026/1744.

Start protecting your AI today.

```
$ pip install complyedge
```

complyedge.io

Try free · no credit card

github.com/ComplyEdge

Star, fork, contribute

pypi.org/project/complyedge

Latest release · Python 3.9+