PRIVARO.AI

EU AI Act: a practical compliance guide

The European AI Act is already in force and its obligations land in phases through 2027. This guide covers what applies to you depending on how you use AI, on which dates, what the fines are, and what you need documented before anyone asks for it.

Updated: September 2026

What the EU AI Act is and who it applies to

Regulation (EU) 2024/1689 is the world's first comprehensive law on artificial intelligence. It does not regulate the technology in the abstract: it regulates uses. Every AI system is classified by the risk it poses to people's rights, with obligations proportional to that risk.

It applies to providers placing AI systems on the EU market and to organisations deploying them in their operations, whether or not they are established in the EU, as long as the output is used within the Union. In practice: if your team uses an LLM to process European customer data, it affects you, even when the model belongs to someone else.

Key dates

Obligations do not arrive all at once. This is the official application timeline.

1 August 2024	Entry into force The Regulation enters into force and starts the clock on every later deadline.
2 February 2025	Prohibited practices and AI literacy Uses such as social scoring and untargeted facial scraping become banned. Organisations must also ensure a sufficient level of AI literacy among their staff.
2 August 2025	General-purpose AI models (GPAI) Transparency, technical documentation and copyright-policy duties for foundation model providers. Governance rules and penalties start to apply.
2 August 2026	General application The bulk of the Regulation applies, including Annex III high-risk systems and transparency duties for systems interacting with people.
2 August 2027	High-risk embedded in products Applies to AI systems that are safety components of products already covered by other Union legislation (Annex I).

The four risk tiers

Your first compliance task is classifying every AI system you use correctly.

	What it covers	What it requires
Unacceptable risk	Social scoring, subliminal manipulation, emotion recognition at work or school, biometric categorisation using sensitive traits.	Banned. There is no mitigation available: the use must stop.
High risk	Recruitment, credit scoring, education, essential services, justice, migration, and AI acting as a product safety component.	Risk management system, data governance, technical documentation, automatic event logging, human oversight, accuracy and cybersecurity, plus registration in the EU database.
Transparency risk	Chatbots, content generators, deepfakes, systems interacting directly with people.	Tell users they are interacting with an AI and label generated or manipulated content in a machine-readable way.
Minimal risk	Spam filters, video games, internal recommenders and most office use cases.	No specific obligations. Voluntary codes of conduct are encouraged.

Fines

- Up to EUR 35 million or 7% of global annual turnover for engaging in prohibited practices.
- Up to EUR 15 million or 3% for breaching high-risk, transparency or GPAI obligations.
- Up to EUR 7.5 million or 1% for supplying incorrect or misleading information to authorities.

Compliance checklist

Ten concrete steps you can start today, ordered by what regulators ask for first.

1. Inventory every AI system

Include the ones you did not buy as AI: assistants inside your CRM, code copilots, meeting transcribers. What is not inventoried cannot be governed.

2. Classify each system by risk tier

Document the reasoning, not just the conclusion. Justifying why something is not high risk matters as much as the classification itself.

3. Establish your role in each case

Providers and deployers carry very different duties. Fine-tuning a model on your own data can turn you into a provider.

4. Evidence your team's AI literacy

Mandatory since February 2025. Record who received what training and when.

5. Set up data governance

Which data enters each system, where it comes from, on what legal basis, and what happens to sensitive data before it reaches a model.

6. Log every AI interaction

The Regulation requires automatic event traceability for high-risk systems. A record of who asked what, to which model, with which data, is your baseline evidence.

7. Design human oversight

Define who can review, correct or reverse an AI-assisted decision, and show they genuinely can.

8. Meet user-facing transparency

Disclose when someone is talking to an AI and label generated content. These are the easiest duties to check from the outside.

9. Review contracts with model providers

Data retention, training use, processing location and sub-processors. Their failures end up being yours.

10. Build the evidence file

Technical documentation, risk assessments, logs and incidents in one exportable place. Conformity is demonstrated, not declared.

Where Privaro fits

Privaro does not deliver compliance on its own: it covers the technical part that is hardest to improvise, which is controlling and proving what data leaves for the models.

- Detects personal and sensitive data in every prompt before it reaches the model.
- Applies policies by data type, team and provider: tokenize, anonymize or block.
- Logs every interaction with user, model, detected entities and the decision applied.
- Scans model responses too, catching data leaks on the way out.
- Exports audit-ready reports for your data protection officer.
- Works with OpenAI, Anthropic, Gemini and others without changing your integration.

Frequently asked questions

Does the EU AI Act apply if my company is outside the EU?

Yes, if the system's output is used within the Union or you place the system on the EU market. The test is where the effect happens, not where you are headquartered.

Does using ChatGPT at work make me an AI provider?

Usually not: you would be a deployer, with lighter duties. But if you fine-tune the model, ship it under your own brand or substantially change its intended purpose, you can become a provider and take on the full set of obligations.

Does the EU AI Act replace GDPR?

No, they stack. GDPR still governs personal data processing while the AI Act adds duties about the AI system itself. One use case can breach both independently.

What if I use streaming responses and cannot mask the output?

With streaming, the output is audited in real time rather than masked. The detection is recorded for evidence and alerting even though the text already reached the user. Document that limitation in your risk assessment.

Official sources

- Regulation (EU) 2024/1689 — https://eur-lex.europa.eu/eli/reg/2024/1689/oj
- European Commission — AI Act overview — https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

This guide is informational and is not legal advice. For specific compliance decisions, consult your legal counsel.

https://privaro.ai/eu-ai-act-compliance