

Does your chatbot need an AI disclosure? What Article 50 actually requires by August 2026.

The EU AI Act does not just regulate high-risk systems. From 2 August 2026, Article 50 makes ordinary chatbots, AI writers, and synthetic media tell people they are AI. Here is which of the four duties applies to you and what the disclosure has to say.

Published 2026-06-30 · A Mue guide

Most of the EU AI Act coverage is about high-risk systems: hiring, credit, biometrics. That makes it easy to assume the Act does not touch you if you only run a support chatbot or use an AI writer. It does. Article 50 is the transparency layer, and it applies to ordinary tools that talk to people or generate content.

It does not ban anything and it does not impose the heavy high-risk conformity regime. It asks one thing: that people are not misled about whether they are dealing with, or looking at the output of, an AI system. The duty splits into four clauses, and which ones bind you depends on what your system does and whether you are its provider or its deployer.

The four duties, in plain terms

- Article 50(1), direct interaction: if a system talks to people, the provider must make sure they know it is an AI system, unless that is already obvious. This is the chatbot, voice-assistant, and AI-phone-agent clause.
- Article 50(2), synthetic-content marking: providers of systems that generate synthetic audio, image, video, or text must mark the output, in a machine-readable form, as artificially generated or manipulated. A visible label on its own is not enough.
- Article 50(3), emotion and biometric systems: deployers of an emotion-recognition or biometric-categorisation system must tell the people exposed to it that it is running.
- Article 50(4), deepfakes and public-interest text: deployers must disclose deepfake image, audio, or video as artificially generated, and must disclose AI-generated text published to inform the public on matters of public interest.

Provider or deployer changes who is on the hook

The marking duty in 50(2) belongs to the provider, the party that builds or puts the generative system on the market. The disclosure duties in 50(3) and 50(4) belong to the deployer, the party that uses the system in practice. Plenty of companies are both, which is why you work from the system and what it does, not from a single label for your business. A team that fine-tunes and ships its own model is a provider; a team that runs someone else's model to publish deepfakes is a deployer; a team that does both wears both hats.

What the disclosure has to say, and when

Under Article 50(5) the information must be clear and distinguishable, and given at the latest at the first interaction or exposure. In practice that means the notice goes in front of the person before they send their first chatbot message or at the point they first see the synthetic content, not buried in a privacy policy they will never open. The wording does not have to be legalistic; it has to be honest and visible.

The deadline the omnibus did not move

The transparency obligations apply from 2 August 2026. The May 2026 digital omnibus pushed several high-risk deadlines later, to December 2027 and August 2028, but it left Article 50 in place. So the transparency layer is now the nearer deadline for most companies, which is the opposite of how the timeline read a year ago. If you assumed AI Act compliance was a 2027 problem, the chatbot disclosure is the part that arrives first.

The cheap way to be ready

You do not need a law firm to write a one-line notice. Classify each AI surface you run against the four clauses, write the disclosure for the ones that apply, place it where the person meets the system, and keep a dated record of where you put it. The generator on this page does the first three in a few minutes and gives you the copy in three languages plus a machine-readable marking note for the 50(2) case. Doing it now, calmly, is far cheaper than doing it under a complaint.

One honest caveat

This is an organisational aid, not legal advice. Whether a specific duty binds your exact system, and how it interacts with sector rules, is a question for a qualified adviser. The value of acting early is that the simple version, written without a deadline overhead, beats the expensive version written after someone asks why the notice was missing.

Frequently asked questions

Does a chatbot need to say it is AI under the EU AI Act?

Yes. Under Article 50(1) of Regulation (EU) 2024/1689, the provider of a system that interacts directly with people, such as a chatbot, voice assistant, or AI phone agent, must make sure those people are informed they are dealing with an AI system. The only exception is where that is already obvious to a reasonably well-informed person in the circumstances. The information has to be given at the latest at the first interaction, and the obligation starts applying on 2 August 2026.

When do the Article 50 transparency rules start applying?

The Article 50 transparency obligations apply from 2 August 2026. That date was not changed by the May 2026 digital omnibus, which moved several high-risk deadlines later but left the transparency layer in place. So while parts of the high-risk regime now bind in December 2027 and August 2028, the duty to disclose chatbots, synthetic content, and deepfakes is the one

with the nearer deadline.

Who has to make the disclosure, the builder or the user of the system?

It depends on the duty. The synthetic-content marking duty in Article 50(2) sits with the provider, the party that builds or puts the generative system on the market. The disclosure duties in 50(3) for emotion recognition and biometric categorisation and in 50(4) for deepfakes and AI-generated public-interest text sit with the deployer, the party that uses the system in practice. Many companies are both at once, which is why you classify the system by what it does rather than by a single label for your business.

Is a visible "made with AI" label enough?

For interaction and deepfake disclosure aimed at people, a clear visible notice is the point. But for the synthetic-content marking duty in Article 50(2), a visible label on its own is not enough: providers must also mark output in a machine-readable format so that systems downstream can detect it was artificially generated or manipulated, using techniques such as the C2PA content-credentials standard. The disclosure to people and the machine-readable marking are two separate requirements.

Generate your Article 50 transparency notice

Free, no signup: agent.mue.app/tools/eu-ai-act-transparency-notice-generator

agent.mue.app/articles/does-your-chatbot-need-an-ai-disclosure

