

# The 2 August 2026 AI Act deadline: does the GPAI rule land on you?

The 2 August 2026 date the omnibus left in place is the one for general-purpose AI model providers, the foundation-model makers. If your company builds AI into a product rather than training the model, the GPAI obligations are not yours. Here is who the date binds, and the lighter duty that does reach you.

Published 2026-07-05 · A Mue guide

A lot of AI Act timelines got a year longer this spring, so the 2 August 2026 date reads to many companies like a countdown they are behind on. For most of them it is not. The obligations that this specific date starts enforcing are aimed at a narrow group: the providers of general-purpose AI models, the handful of firms that train foundation models and place them on the market. If your company uses one of those models inside a product rather than building the model itself, the general-purpose AI (GPAI) rules are not the ones that bind you, and it is worth knowing that before you budget for a compliance scramble you do not owe.

## What the 2 August 2026 date actually is

The GPAI model obligations have technically applied since 2 August 2025, but for the first year there was no one enforcing them. That grace year ends on 2 August 2026, when the European Commission, acting through its AI Office, gains the supervision and enforcement powers over GPAI providers: from that date it can request documentation, evaluate a model, require compliance and risk mitigation, and issue fines. The digital omnibus agreed in May 2026 pushed the high-risk deadlines back, but it left this date where it was, which is why 2 August 2026 still matters even though so much else moved.

## Who the GPAI obligations land on

The duties in Articles 53 to 55 sit with the provider of the general-purpose model: the party that develops it and puts it on the EU market. In practice that is the foundation-model makers, the OpenAI, Anthropic, Google, Meta, and Mistral tier of company. Their obligations are to keep technical documentation on the model, give downstream developers the information they need to build on it safely, put a copyright policy in place, and publish a summary of the training data, with a heavier set of systemic-risk duties on the largest models. None of that is a checklist a company integrating a model through an API can meaningfully perform, because it is about how the model was built, not how it is used.

If you call a model through an API, fine-tune it lightly on your own data, or ship a product that has a vendor model inside it, you are a deployer or a downstream provider of an AI system, not a provider of a GPAI model. The GPAI provider obligations are simply not addressed to you.

This is the distinction the coverage tends to blur when it treats 2 August 2026 as a single deadline that everyone using AI has to meet.

## **The line that can turn a user into a provider**

There is one way to inherit the GPAI provider duties without setting out to train a model, and it turns on scale. The Commission guidelines treat a substantial modification of a general model, a fine-tune on the order of the compute that went into the original training, as making you the provider of that modified model, with the obligations that follow. The everyday cases stay well below that line: prompting, retrieval, a small fine-tune on your own examples, or wiring a model into an app all keep you on the deployer side. It is only when you are retraining a general model at something like its original scale that the provider duties can attach. If that is what you are doing, it is worth taking advice on exactly where the threshold falls, because the guidance is where the detail lives.

## **The duty that does reach you: Article 50**

Ruling out the GPAI obligations does not mean 2 August 2026 is a non-event for you, because a different duty starts the same day and is aimed squarely at companies using AI in a product. The Article 50 transparency rules apply from 2 August 2026, a date the omnibus did not move. If your system interacts with people, you have to tell them they are dealing with an AI. If it generates or manipulates audio, image, video, or text, you have to mark that output as artificially generated in a machine-readable way. That is a notice and a label, not a conformity program, and it is the part of the deadline that genuinely lands on most businesses. Our transparency-notice generator drafts the wording tied to the categories Article 50 names.

## **And if your own system is high-risk**

The other thing 2 August 2026 used to mean was the start of the full high-risk regime, and that is the part that moved. The omnibus pushed stand-alone Annex III high-risk systems to 2 December 2027 and AI built into products already regulated under EU product law to 2 August 2028. So a high-risk system now has a later deadline than a limited-risk chatbot, which is the reverse of how the timeline read a year ago. Most systems are not high-risk in the first place, so the honest first step is to place yours before you assume any of the heavy dates apply. The classifier walks the Annex III categories and tells you which deadline, if any, is actually yours.

## **A starting point, not legal advice**

The short version is that 2 August 2026 is two deadlines wearing one date: an enforcement start for GPAI model providers, and a transparency start for everyone running AI that talks to people or makes content. Sorting which one is yours is most of the work, and for the large majority of companies it is the lighter one. The scope of the GPAI rules, the modification threshold, and how Article 50 applies to a specific product are matters for qualified counsel, but knowing which side of the provider line you are on tells you whether you are facing an afternoon of transparency work or something larger, and you would rather know that before the

date than after a question you cannot answer.

## Frequently asked questions

### **Does the 2 August 2026 GPAI deadline apply to my company if we only use ChatGPT or Claude?**

Almost certainly not the GPAI part. The general-purpose AI (GPAI) model obligations fall on the provider that trains and places the model on the market, the foundation-model makers. If you build a product on top of a model through an API, or run a vendor tool, you are a deployer or a downstream provider of an AI system, not a GPAI model provider. The date that touches you on 2 August 2026 is the Article 50 transparency duty, not the GPAI obligations in Articles 53 to 55.

### **What actually happens on 2 August 2026 for GPAI models?**

The GPAI model rules have technically applied since 2 August 2025, but for the first year there was no one enforcing them. On 2 August 2026 the European Commission, acting through the AI Office, gains its supervision and enforcement powers over GPAI providers: it can request documentation, evaluate a model, require risk mitigations, and fine. The May 2026 digital omnibus moved several high-risk deadlines back a year but left this date in place.

### **Could fine-tuning a model make my company a GPAI provider?**

It can, but only at real scale. Using or lightly adapting a model keeps you a deployer or a downstream provider of an AI system. The Commission guidelines treat a substantial modification, a fine-tune on the order of the original training effort, as making you the provider of that modified model, with the GPAI duties that follow. A typical prompt-and-integrate setup, or a small fine-tune on your own data, does not cross that line. If you are training or heavily retraining a general model, take advice on where the line falls.

### **If the GPAI rules are not mine, what do I owe on 2 August 2026?**

If your product interacts with people or generates content, the Article 50 transparency duties: tell people when they are dealing with an AI system, and mark AI-generated or manipulated audio, image, video, or text as synthetic. That is a notice and a label, not a conformity program. Any high-risk obligations, if your system is high-risk at all, were pushed to 2 December 2027, or 2 August 2028 for AI inside regulated products, by the omnibus.

### **What is the fine for a GPAI provider that does not comply?**

GPAI model providers sit on a separate penalty track. Article 101 sets their ceiling at up to 15 million euro or 3 percent of worldwide annual turnover, whichever is higher, distinct from the general Article 99 tiers that can reach 35 million euro or 7 percent. It is aimed at the model makers, which is another sign the GPAI regime is built around the providers, not the businesses using their models.

## **Classify your system in five questions**

Free, no signup: [agent.mue.app/tools/eu-ai-act-risk-classifier](https://agent.mue.app/tools/eu-ai-act-risk-classifier)

[agent.mue.app/articles/does-the-2-august-2026-gpai-deadline-apply-to-you](https://agent.mue.app/articles/does-the-2-august-2026-gpai-deadline-apply-to-you)

