# The AI Act's fine book just opened

*On 2 August 2026, the European Commission's power to fine general-purpose AI providers went live — up to 3% of global annual turnover, or €15 million, whichever is higher. Nine of the ten leading frontier labs have signed the accompanying Code of Practice; Meta refused. The AI Office has about 165 staff to police a landscape whose four largest US builders alone plan to spend $725 billion on AI infrastructure this year. The next twelve months will show whether Europe wrote a rulebook it can actually enforce, or whether the enforcement calendar quietly slips past a Commission that lacks the hands to read the paperwork.*

AUTHOR
AT **Alicante Tech Vibes *Research Labs***
~9 MIN READ · ALL FIGURES SOURCED & DATED

| | | |
|---|---|---|
| **2 AUG 2026**<br>DATE ENFORCEMENT POWERS ACTIVATED | **3%**<br>MAX FINE, SHARE OF GLOBAL TURNOVER | **€15M**<br>FLOOR FOR MAX FINE WHEN 3% IS LOWER |
| **$10^{25}$ FLOP**<br>COMPUTE THRESHOLD, SYSTEMIC-RISK GPAI | **~165**<br>AI OFFICE STAFF, INCL. 40 NEW HIRES | **9 of 10**<br>FRONTIER LABS THAT SIGNED THE CODE |

## I. The switch from “in force” to “enforceable”

The rules did not change on 2 August 2026. The consequences did. GPAI obligations under Chapter V of the EU AI Act have applied since 2 August 2025; from this month, the Commission may open investigations, demand documentation and model access, and impose fines — up to 3% of a provider's annual worldwide turnover or €15 million, whichever is higher, under Article 101.

The Commission's own press notice on 31 July framed the transition tersely: as of 2 August, “the European Commission's AI Office, together with national authorities, will begin enforcing” the Act. In parallel, transparency duties on the deployer side kicked in: interactive systems must disclose they are AI, deepfakes must be labelled, and synthetic content must carry machine-readable marks. The GPAI layer — the one that binds the labs that build models rather than the businesses that use them — is the one with money behind it.

Two grandfathering gates matter. Models placed on the EU market before 2 August 2025 have until 2 August 2027 to reach full compliance. Everything trained and released after that first date already had to meet the documentation, copyright and safety obligations by design; the Commission simply could not fine them for missing the mark until now.

## II. Nine signed. Meta didn't.

Alongside the binding text, the AI Office negotiated a voluntary General-Purpose AI Code of Practice, finalised on 10 July 2025. Signatories get a “presumption of conformity”: regulators assume the signatory is meeting the underlying obligations unless something concrete suggests otherwise. Refusers get the same law and heavier scrutiny.

By August 2026 the signatory list read like a roster of Western frontier labs: Anthropic, Google, Microsoft, OpenAI, IBM, Amazon, Mistral AI, Cohere and Aleph Alpha have signed all three chapters (Transparency, Copyright, Safety and Security). Elon Musk's xAI signed only the Safety chapter. Meta declined outright. Anthropic said the Code “advances the principles of transparency, safety and accountability”; Meta's Joel Kaplan, its Chief Global Affairs Officer, called the Code an “over-reach” that would “throttle the development and deployment of frontier AI models in Europe.”

The split matters because it defines two regulatory postures inside the same market. Signatories accept a safe harbour: their documentation templates and safety-framework mechanics are pre-baked into a text the Commission helped write. Meta accepts the same underlying obligations without the safe harbour, and bets that its lawyers can defend a bespoke compliance regime under Commission scrutiny.

## III. The office that has to police it has 165 people

The most sourceable skeptic case is not about the law — it is about the desk. The AI Office, the Commission unit inside DG CNECT that runs GPAI supervision, sits at roughly 165 staff after a recruitment of about 40 contractual agents announced this year. An independent review by Pour Demain, cited by Lawfare, judged the office “significantly underresourced relative to its mandate” and recommended scaling to at least 160 staff by 2030 — a target the office arguably reaches on the day it was told it needed to grow past it.

Compare that with the object of supervision. The four largest US hyperscalers alone plan roughly $725 billion of AI infrastructure spending in 2026. A single Big Tech legal-and-policy team is larger than the entire unit that must issue requests for information, run model evaluations, hire external evaluators, and defend enforcement in court.

That resource asymmetry does not mean the office will fail; the Digital Services Act moved from silent first year to symbolic fines within eighteen months. But it does explain why the honest technical read on 2 August 2026 is **latent authority** rather than an imminent fine wave. The Commission has the tools. It does not yet have the hands.

**FIG. 1: AI OFFICE HEADCOUNT VS. WHAT IT IS MEANT TO SUPERVISE (2026, INDICATIVE)**


SCALING: BAR WIDTHS REFLECT RATIO, NOT DIRECTLY COMPARABLE UNITS (HEADCOUNT VS. CAPEX). SOURCES: FORKAST (2026); VALUE ADD VC, “$725B AI CAPEX 2026” (2026).

## IV. Where the euros actually move

The direct fiscal impact of a first-day-of-enforcement moment is small; the second-order economic channel is where the money really flows. Three transmission mechanisms deserve tracking.

**Compliance-cost split by size.** For a five-person European startup building a limited-risk GPAI-derived product, first-year AI Act compliance falls in the €15,000–€31,000 range, rising to €46,500–€94,000 for a high-risk application (annual ongoing costs of €5–35k). These are estimates from a Vienna consultancy that specialises in the field, and readers should treat them as order-of-magnitude, not audited. The gradient matters more than the number: for OpenAI or Google the same fixed costs are trivial; for a two-person shop building an EU-facing agent product they are a meaningful share of first-year runway. The Act is size-neutral in text and regressive in practice.

**Legal certainty as a moat.** Signatories bought a written safe harbour. Non-signatories — principally Meta, and any smaller European lab that finds the Code's Copyright chapter incompatible with its data pipeline — must build bespoke compliance and defend it under adversarial scrutiny. This is a durable structural cost that scales with the number of member-state regulators who choose to test it.

**Market-access risk on the open-weight tail.** The Commission has flagged DeepSeek as a "highly likely" candidate for systemic-risk designation, whether or not it clears the $10^{25}$ FLOP threshold on measured training compute. That signals a policy read in which capability-in-use, not just compute-at-training, drives classification — a shift that would tighten the perimeter for every downloadable open-weight model above frontier capability, from every jurisdiction.

> ***"Over-reach will throttle the development and deployment of frontier AI models in Europe."***
>
> JOEL KAPLAN, META CHIEF GLOBAL AFFAIRS OFFICER, EXPLAINING META'S REFUSAL TO SIGN THE GPAI CODE OF PRACTICE. REPORTED BY EURONEWS, 23 JULY 2025.

## V. The second-order effects nobody is pricing

Three consequences travel further than the fine schedule suggests. The first is **the Brussels-effect question in AI**. GDPR became a global privacy floor because compliance was cheaper to unify than to fork. If the AI Act's documentation templates — the summary of training data, the copyright-policy statement, the systemic-risk framework — become the default for products that touch the EU market, US and Chinese labs will effectively be writing to a European rubric even outside Europe. Signatories already are.

The second is **the labour re-mix in the compliance stack**. The Act does not directly displace developers, but it creates measurable demand for a new occupation: the GPAI compliance officer, the copyright-policy lawyer, the systemic-risk evaluator. On the industry side the seats already exist inside frontier labs; on the regulator side, only 165 people are on the payroll to match them. Expect a two-year hiring imbalance that will bid up compliance-side wages in Brussels and Frankfurt before it corrects.

The third is **the export-control adjacency**. The $10^{25}$ FLOP threshold in Article 51 is a policy number that now interacts with US commerce-department compute caps, Chinese domestic chip supply, and the DeepSeek precedent. A single number chosen by Brussels is quietly aligning with a single number chosen by Washington. That is not an accident; it is a de-facto trans-Atlantic definition of "frontier."

## VI. How it plays out: three scenarios

Twelve-to-twenty-four-month horizon. Probabilities are analytical judgment, informed by the Commission's DSA enforcement precedent and the resource picture in Section III.

### 1 · Two-tier consolidation ~55% 12–24 MONTHS · SLOW-COOK ENFORCEMENT

The AI Office opens Requests for Information against several GPAI providers, publishes soft guidance, and settles most disputes without headline fines. Signatories quietly absorb costs, Meta faces disproportionate inspection frequency and eventually litigates one RFI. Watch for the office's first published RFI docket, any Meta EU market-share erosion in downstream products, and Code-signatory churn.

### 2 · Symbolic first fine ~25%

BY AUG 2027 · DSA-STYLE ESCALATION

The Commission lands a mid-size fine on a non-signatory GPAI provider — most probably a Chinese open-weight model designated under Article 51 — to establish precedent. Effect resembles the DSA's first-year enforcement trajectory: a modest number in euros, an outsized effect on legal-team behaviour worldwide. Watch for any Article 51 systemic-risk designation letter, an AI Office model-evaluation contract disclosure, and litigation filings in the Court of Justice of the EU.

### 3 · Enforcement stall ~20%

18–24 MONTHS · POLITICAL ROLLBACK

The Digital Omnibus simplification package waters down GPAI obligations; the AI Office's staffing plan slips; Draghi-style competitiveness rhetoric wins the framing battle in the European Council. No substantive enforcement action opens against a frontier provider before 2027. Watch for revisions to the Digital Omnibus text, hiring-freeze reporting in the AI Office, and any withdrawal or partial-withdrawal from the Code of Practice by a current signatory.

**WHAT WOULD FALSIFY THIS THESIS**

The core thesis is that 2 August 2026 marks a genuine regime shift, not a bureaucratic milestone. That thesis is wrong if: (a) no formal GPAI enforcement action — RFI, model-access request, or fine — is publicly opened before 2 August 2027; (b) at least one current signatory withdraws from the Code and faces no measurable regulatory consequence; and (c) EU market share of Meta-derived GPAI products holds or grows against signatory alternatives over the same period. The strongest counter-case, argued by Draghi-inflected critics, is that Europe has again written a rulebook that produces compliance paperwork without changing behaviour — and that the AI Office may generate “paperwork that satisfies a legal requirement without generating actionable safety information.” The falsification test is designed to fire if that critique proves right.

### WHAT TO WATCH, IN ORDER

1. **Q4 2026 —** First publicly disclosed Request for Information from the AI Office to a GPAI provider (check: AI Office news page and provider 8-K/annual-report disclosures).
2. **Q1 2027 —** Any Article 51 systemic-risk designation letter, especially covering a non-EU open-weight model such as DeepSeek or Qwen (check: EC Digital Strategy news portal).
3. **Q2 2027 —** Meta EU market share in GPAI-derived consumer products (Llama-based assistants, WhatsApp AI features) versus signatory alternatives (check: quarterly enterprise-AI trackers, Meta 10-Q).
4. **By Aug 2027 —** Grandfathering deadline: pre-Aug-2025 models must reach full compliance. Watch for last-minute retirements or documentation-only releases (check: model cards, provider changelogs).
5. **Any month —** Signatory list movement on the AI Office's Code page — withdrawals, additions, or a shift by xAI to full-Code signature (check: the Commission's Code page).

## SOURCES


1. European Commission, “Commission starts enforcing AI Act rules and new transparency requirements on 2 August,” 31 July 2026 · digital-strategy.ec.europa.eu
2. European Commission Press Corner, “Commission starts enforcing AI Act rules and new transparency requirements on 2 August,” IP/26/1714, 31 July 2026 · ec.europa.eu
3. Future of Life Institute AI Act tracker, “Enforcement of Chapter V under the EU AI Act” · artificialintelligenceact.eu
4. AI Office, “General-Purpose AI Code of Practice, Final Version,” 10 July 2025 · code-of-practice.ai
5. Latham & Watkins, “EU AI Act: GPAI Model Obligations in Force and Final GPAI Code of Practice in Place,” 2026 · lw.com
6. Wikipedia, “General-Purpose AI Code of Practice,” updated January 2026 · en.wikipedia.org
7. Euronews, “Meta won't sign EU's AI Code, but who will?,” 23 July 2025 · euronews.com
8. Forkast, “The Enforcement Desk: How 40 New Hires Will Define EU AI Oversight,” 2026 · forkast.news
9. Lawfare, “How Much Power Does the EU AI Office Actually Have?,” 2026 · lawfaremedia.org
10. Wavect, “Real EU AI Act Compliance Cost for a 5-Person Startup,” 2026 · wavect.io

11. Pinsent Masons Out-Law, “EU AI Act rules on GPAI models under DeepSeek review,” 2026 · pinsentmasons.com
12. Value Add VC, “$725B on AI — Where Big Tech Is Spending in 2026” · valueaddvc.com
13. Chatham House, “The EU's new AI code of practice has its critics but will be valuable for global governance,” August 2025 · chathamhouse.org
14. Beam AI, “EU AI Act 2026: GPAI Enforcement & 3% Fines Begin,” 2026 · beam.ai
15. Sisal Hubs / Medium, “Draghi's Competitiveness Warning Reveals Europe's Regulatory Burden Crisis,” 2026 · medium.com

METHOD NOTE: FIGURES ARE ATTRIBUTED AND DATED INLINE; THE €725B CAPEX FIGURE AND THE COMPLIANCE-COST RANGES ARE THIRD-PARTY ESTIMATES AND ARE FLAGGED AS SUCH. WHERE INDEPENDENT SOURCES DISAGREE, THE DISAGREEMENT IS REPORTED RATHER THAN AVERAGED.

---

Built by the Alicante Tech Vibes community, a tope. 2026-08-24