5 Things Everyone Gets Wrong About the EU AI Act Omnibus: Watermarking Is Due December 2, 2026

The Digital Omnibus moved two high-risk deadlines to 2027 and 2028. It left the 2 December 2026 watermarking deadline exactly where it was, and much of the commentary since 27 July has not noticed.

Regulation (EU) 2026/1744 entered into force on 27 July 2026 and pushed stand-alone Annex III high-risk systems to 2 December 2027 and product-embedded Annex I systems to 2 August 2028. Nothing else that matters to generative AI moved. If your generative system was on the market before 2 August 2026, the Article 50(2) marking duty applies to it on 2 December 2026. That is a boring story, which is probably why it lost the news cycle to “the AI Act is delayed.”

The cost of getting this wrong is not theoretical. The fine tiers in [Regulation (EU) 2024/1689](https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng) were untouched by the

Omnibus: €35M or 7% of worldwide turnover for Article 5 breaches, €15M or 3% for most other infringements, €7.5M or 1% for supplying misleading information. Article 101 sets a separate ceiling for general-purpose AI model providers breaching Chapter V: €15,000,000 or 3% of total worldwide annual turnover, whichever is higher. A team that stood down its transparency work in August 2026 on the strength of a LinkedIn summary has roughly four months of runway left, and that runway is measured in output labelling infrastructure rather than policy documents.

THE CLAIM	THE REALITY
"The Digital Omnibus delayed the AI Act." Compliance can slow down across the board until 2027.	The Omnibus amended deadlines selectively. It did not move the 2 August 2026 application date for Article 50 transparency obligations. Only the two high-risk tracks shifted, to 2 December 2027 (Annex III stand-alone) and 2 August 2028 (Annex I embedded), per the European Parliament Legislative Train.

The second myth is that prohibited practices start later now, along with everything else. They do not. Article 5 prohibitions and the Article 4 AI literacy duty have been in application since 2 February 2025, and the Omnibus did not touch either of them. Eighteen months of exposure has already accrued, at the €35M / 7% tier.

The third is that GPAI obligations were suspended pending the standards work. Chapter V duties for general-purpose AI model providers became applicable on 2 August 2025 and were not rolled back: technical documentation, downstream model information, training-data summaries, and EU copyright and text-and-data-mining compliance. The penalty ceiling under Article 101 is €15M or 3% of global turnover.

THE CLAIM	THE REALITY
Watermarking is years away. Nobody is marking synthetic output yet.	Generative systems placed on the market on or after 2 August 2026 get no grace period at all: Article 50(2) marking applies from the date of placement. Systems already on the market before that date have until 2 December 2026, per Paul Weiss.

The fifth myth is the tidiest one: that the Commission's Code of Practice on Transparency of AI-generated Content is the operative rule, so signing it is the compliance step. The [Code of Practice is voluntary](https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content). Article 50(2), (4) and (5) bind you whether you sign or not. The Commission confirmed the Code as adequate and published final transparency guidelines in July 2026, which makes a signature useful evidence of good faith and nothing more.

Three of the four generative AI dates have already passed

Strip out the noise and the AI Act's generative-AI surface has four dates on it. Two are historical: 2 February 2025 for Article 5 and Article 4, 2 August 2025 for GPAI model obligations. One is very recent: 2 August 2026 for Article 50 transparency, which is when the obligation started applying to newly placed systems. The last is 2 December 2026, the end of the grace period for generative systems that were already on the market.

2 Dec 2026
Article 50(2) marking due for generative systems on the market before 2 Aug 2026
`Paul Weiss / Faegre Drinker client alerts, July 2026`

€35M / 7%
top Article 99 fine tier, unchanged by the Omnibus
`Regulation (EU) 2024/1689, Art. 99`

€15M / 3%
Article 101 ceiling for GPAI providers breaching Chapter V
`regulation-ai.eu Article 101 commentary, 2026`

2 Dec 2027
new date for stand-alone Annex III high-risk systems (was 2 Aug 2026)
`EP Legislative Train, 2026`

One detail in the transparency rules gets dropped from almost every summary I have read, and it is the one most likely to catch a company that does not think of itself as an AI company. Article 50 also requires disclosure of AI-generated text published to inform the public on matters of public interest, unless that text passed human editorial control with a responsible person. That is a rule about your content pipeline rather than your model. A marketing team publishing AI-drafted explainers on, say, energy pricing or health claims sits inside this obligation, and the exemption is procedural: someone has to actually own the review.

The other underreported detail runs in your favour. Deepfakes generated before 2 August 2026 carry no mandatory retroactive labelling duty. The Commission encourages voluntary labelling in its [quick facts on transparency rules](https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai), but you are not required to go back through an archive and stamp it. I read that as a deliberate concession to the practical impossibility of retroactive provenance, and it means your December work points forward: mark new output, leave the archive alone.

✓ **The one thing to do instead of re-reading the Omnibus**
Inventory your generative systems by placement date. Anything on the EU market before 2 August 2026 needs machine-readable marking working by 2 December 2026. Anything placed after that date needed it on day one, so if you shipped in August or September without it, you are already non-compliant and the remediation clock started then, not in December.

The high-risk delay was a standards problem, and transparency was not

The [EPRS briefing](https://www.europarl.europa.eu/RegData/etudes/BRIE/2026/782651/EPRS_BRI\(2026\)782651_EN.pdf) makes this explicit: the new high-risk dates are conditional, tied to Commission confirmation that harmonised standards are actually available, with 2027 and 2028 functioning as outer caps rather than fresh start dates. You cannot require conformity assessment against standards that do not exist yet. That is a coherent reason to move a deadline.

Article 50(2) has no equivalent dependency. Marking synthetic output in a machine-readable format is engineering that exists today, from providers who ship it in production. There was no standards gap to wait on, so there was nothing to concede.

My reading of the conditionality is that it cuts both ways, and the commentary has only picked up the friendly half. If the Commission confirms standards availability earlier than expected, high-risk obligations can bite before December 2027. Companies treating 2 December 2027 as a fixed planning horizon are betting on the outer cap. I would not build a roadmap on that assumption, though I will say plainly that I have seen nothing telling me how the Commission intends to exercise that confirmation, and anyone who claims to know the timing is guessing.

WHERE I LAND

The "AI Act delayed" narrative persists because a selective amendment is harder to summarise than a blanket one, and because the two dates that moved were the two dates most companies had written into their board slides. High-risk was the headline number for eighteen months, so when it slipped, the whole regulation felt like it slipped with it. My take: the Omnibus is better read as the Commission protecting the parts it can actually enforce. Transparency, prohibitions and GPAI documentation are things a regulator can check from outside the company. High-risk conformity assessment is not, and it depends on standards bodies the Commission does not control. ***The deadlines that survive are the enforceable ones.***

The AI Office has already issued €47 million in fines

This is where I think the risk calculus changed in 2026. The AI Office issued [€47 million across three companies](https://www.arturmarkus.com/eu-ai-office-issues-first-e47-million-in-fines-against-three-companies-for-high-risk-ai-violations-hiring-platform-gets-e18m-credit-scorer-e14m-retail-chain-e15m/) for high-risk violations, split €18M against a hiring platform, €14M against a credit scorer and €15M against a retail chain. Those cases tell you the AI Office will name companies and put numbers on paper rather than issue guidance and hope.

Transparency breaches under Article 50 are also cheaper to prove than high-risk breaches. A regulator does not need to audit your training data or reconstruct your risk management system. Someone downloads an image from your product, checks for embedded provenance metadata, and finds it or does not. Short evidentiary paths get used first. That is my inference rather than something the Commission has stated, but it is the pattern I would expect from an enforcement body with limited headcount and a mandate to show results.

The broader picture I sketched in [my piece on the G20 and

Brussels](https://www.arturmarkus.com/the-g20-just-endorsed-no-new-ai-regulators-brussels-ignored-it-the-same-week-and-was-right-to/) still holds. The EU is sequencing its AI regulation rather than retreating from it, and the Omnibus fits that: move what cannot be enforced yet, keep what can.

What a defensible position looks like on 3 December

Three things, in the order I would do them.

First, classify by placement date. The distinction between "before 2 August 2026" and "on or after" is the only thing that determines whether you have a grace period, and it is a legal question about market placement rather than about when you wrote the code. Get it written down with dates and system names, because it is the first document anyone will ask for.

Second, make the marking machine-readable and verify it survives your own pipeline. Marking that gets stripped by a re-encode, a thumbnail generator or a CDN transform is marking that does not exist from the regulator's side. I would treat provenance persistence as a test case in CI rather than a one-time integration check.

Third, name the human in the editorial loop for public-interest text. The Article 50 exemption for AI-generated text is conditional on human editorial control with a responsible person. "Our team reviews everything" does not meet that. A named owner per publication channel does.

What I would not do is spend December drafting a policy document about your transparency commitments. The Code of Practice is voluntary and the Commission already confirmed it as adequate, so a signature buys goodwill and no safe harbour. Article 50(2), (4) and (5) apply either way. The deliverable is working metadata in shipped output, and that is an engineering task with a fixed date on it.

What I find genuinely uncertain is how the Commission will treat partial compliance on 2 December: a company with marking on image output but not audio, or with provenance that works in the product but not through a partner integration. Nothing in the material I have read addresses proportionality for incomplete implementations, and I would rather say that than pretend there is a known answer. If you are somewhere in that grey zone and deciding what to prioritise before December, [that is the kind of trade-off worth talking through properly](https://www.arturmarkus.com/contacts/) rather than guessing at from a blog post.

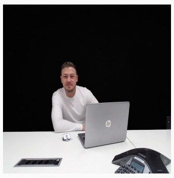
WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting →