




How to Ship Article 50 Marking Before December 2, 2026: A CTO's Guide to the EU Transparency Code

If your AI system was on the EU market before 2 August 2026, the machine-readable marking grace period ends on 2 December 2026. The deadline most teams carry in their heads is a year too far out.

Article 50 of the EU AI Act became applicable on 2 August 2026, and the Commission's [Article 50 FAQs](<https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act>) give systems already placed on the market before that date until 2 December 2026 to meet the Article 50(2) duty to mark synthetic audio, image, video and text in a machine-readable way. That is the last date. Nothing in the text moves it to 2027.

This gets misfiled because two clocks look similar. 2 August 2026 was the general application date for Article 50. 2 December 2026 is the narrower deadline for the marking



and detection duty on systems that were already on the market. If you shipped a generative feature before August and have not touched provenance since, the second date is yours, and in practice it is the tighter one: four months of runway, not two years.

Four adjectives in Article 50(2) that function as design constraints

The [Service Desk text of Article 50](<https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-50>) sets the standard as machine-readable marking that is “effective, interoperable, robust and reliable” as far as technically feasible. Effective means a downstream system can actually recover the signal. Interoperable means you cannot invent a private format and call it done. Robust means it survives normal handling. Reliable means it does not fire on content you did not generate.

The Commission’s FAQ names the acceptable techniques: watermarks, metadata identifications, cryptographic methods for proving origin, logging methods and fingerprints. That list is permissive. You are not obliged to pick C2PA, and you are not obliged to pick invisible watermarking. You have to pick something from that family, document why it suits your output type, and make it detectable.

Four content types are in scope: audio, image, video, text. Text is the awkward one, because marking short text outputs is the least mature of the four, and “as far as technically feasible” is doing visible work in that sentence. My read: for text, the defensible position in December 2026 is metadata plus server-side logging that lets you answer “did our system produce this string” on request, with a written record of why stronger in-content marking was not technically feasible at your output lengths. That is an argument rather than a settled interpretation, and I would want counsel to look at it before anyone relies on it.

Transparency breaches carry fines up to €15 million or 3% of total worldwide annual turnover, whichever is higher, per the Commission’s Article 50 FAQs. For a small company that is still an existential number.

Signing the Code of Practice is optional, the obligation is not

The Commission published the final [Code of Practice on Transparency of AI-Generated Content](<https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-co>



ment) on 10 June 2026, and its own [FAQ on transparent AI systems](<https://digital-strategy.ec.europa.eu/en/faqs/guidelines-and-code-practice-transparent-ai-systems>) is explicit that the Code is voluntary. Signing it is not the source of the legal obligation. The duty comes from Article 50 and binds you either way.

What signing buys you is a recognised route to demonstrating compliance. Non-signatories must demonstrate compliance through alternative documentation and controls, which means you construct and defend your own evidence package. For a small company with no regulatory affairs function, that is a real cost difference. If your output types are unusual enough that the Code did not anticipate them, a documented approach of your own may be cleaner than fitting into a framework that does not quite describe what you do.

The Code took roughly seven months to draft: kick-off plenary on 5 November 2025, first draft 17 December 2025, second draft 3 March 2026, draft guidelines 8 May 2026, final text 10 June 2026. Worth knowing, because it tells you how much settled practice sits underneath the document. Not much.

Five steps, starting with the role mapping nobody can do for you

1 **Decide whether you are provider, deployer, or both**

Marking under Article 50(2) is principally a provider obligation. Deployers carry separate duties: disclosing deepfakes under Article 50(4) and disclosing AI-generated text published to inform the public on matters of public interest under Article 50(5). Many smaller companies will land on both sides: provider for the system they ship, deployer for the content they publish with someone else's tool. Write the mapping down per product surface. Everything downstream depends on getting this line right, and nobody else can draw it for you.

2 **Check the exemptions before you build anything**

Article 50(2) does not apply where the AI system performs an assistive function for standard editing, or does not substantially alter the deployer's input data or its semantics. If your feature is spell-check-adjacent, or reformats what the user supplied, you may be out of scope entirely. Article 50(6) excludes systems intended for military, defence or national security purposes from paragraphs 1 to 4. Article 50(5) exempts public-interest text that underwent human review or editorial control where a natural or legal person holds editorial responsibility. Each exemption you can document is work you do not have to do in November.



How to Ship Article 50 Marking Before December 2, 2026: A CTO's Guide to the EU Transparency Code

3 Pick one technique per output type and write the justification

Audio, image, video and text do not get the same answer. The FAQ list (watermarks, metadata identifications, cryptographic methods for proving origin, logging methods, fingerprints) lets you mix. The deliverable is the implementation plus a short memo per output type explaining why that technique is effective, interoperable, robust and reliable for your case, and what its technical limits are. The memo is the thing a regulator reads.

4 Build the detection side as well as the marking side

Article 50(2) couples marking with detectability. A watermark nobody can read is not compliance. Decide now who needs to verify your marks (platforms, customers, auditors, your own support team) and expose something they can actually use. This is the step teams skip, because marking feels like the deliverable and detection feels like someone else's problem.

5 Adopt the Commission icons where disclosure is user-facing

The Commission published a set of icons that creators, publishers and deployers may use to disclose artificial images, audio including deepfakes, and text. Using the official set costs nothing and removes an argument about whether your disclosure was clear. It is the cheapest line item here.

Note what is absent from that list: a procurement cycle. All of it can be done with engineering time and a lawyer's afternoon. I have no cost figures for provenance tooling and I am not going to invent any. The binding constraint for most teams is the role mapping in step one.

Shipping under a permissive licence does not exempt you, either. The Commission's Article 50 FAQs state that Article 50(2) expressly covers providers of general-purpose AI systems, and open-source status is not a carve-out from the marking duty. If you maintain an open-weights model or an open generative tool with EU users, you are a provider.

Metadata is easy to attach and easy to lose

The failure mode that worries me most happens after the code is merged. I have written before about [why most platforms strip Content Credentials on upload](<https://www.arturmarkus.com/c2pa-explained-how-content-credentials-actually-work-and-why-5-of-6-platforms-strip-them/>), and that behaviour does not change because a deadline arrived. If your whole marking strategy is a metadata field, the mark survives exactly as long as nothing touches the file.



"As far as technically feasible" is a genuine defence when the stripping happens downstream of you. It is a weaker one if you chose the most fragile available technique while a more durable option existed for your output type. My read: regulators will care less about whether a mark survived one specific platform and more about whether you chose deliberately and wrote down the reasoning, known limits included.

The second failure is role confusion inside a single product. A team maps itself as "deployer" because it uses a third-party model, then publishes AI-generated articles on public-interest topics without disclosure, assuming the upstream provider's marking covers it. It does not. Article 50(5) is a separate duty with a separate exemption, and that exemption turns on human review and named editorial responsibility, not on whether the file carries a watermark.

WHERE I LAND

*Between now and 2 December, I would spend more hours on documentation than on implementation. The techniques the Commission lists are all shippable by a competent team in a few sprints. What takes longer, and what is harder to reconstruct afterwards, is the written record of which role you occupy per surface, which exemptions you claim and why, and why you selected each technique for each of the four content types. **Compliance here is mostly an evidence problem wearing an engineering costume.***

When building marking is the wrong call

If your feature genuinely sits inside the assistive-editing exemption, do not build marking into it just to be safe. You would create a signal saying "this content is AI-generated" on output that mostly is not, which is a reliability problem under the same article that asks for reliable marking. Document the exemption instead.

If your system is intended for military, defence or national security purposes, Article 50(6) excludes it from paragraphs 1 to 4. That is a narrow carve-out and worth a legal read rather than a product manager's judgment.


And if you are a pure deployer with no provider role anywhere in your stack, Article 50(2) is not your clock. Your obligations are 50(4) and 50(5), disclosure duties rather than marking duties, and the engineering work is closer to a UI change than a pipeline change. I went



How to Ship Article 50 Marking Before December 2, 2026: A CTO's Guide to the EU Transparency Code

through the date confusion around this in a [note on what everyone gets wrong about the omnibus timeline](https://www.arturmarkus.com/5-things-everyone-gets-wrong-about-the-eu-ai-act-omnibus-watermarking-is-due-december-2-2026/).

So: ****2 December 2026 is a hard date for marking anything you shipped before August.**** You can skip the Code, you cannot skip the obligation, and the exposure is €15 million or 3% of worldwide turnover, whichever is higher. Eight weeks is enough if you start with the role mapping and treat the memo as the deliverable.



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 →