AI Transparency (Article 50) Addendum

EU AI Act Article 50 puts marking and disclosure duties on the providers and deployers of generative AI systems — which may include you. This addendum states exactly what Wakemark provides toward those obligations, and — with equal weight — what it does not. Wakemark supplies the proof layer; the legal duty stays where the law put it.

Last updated: 2026-07-13

1. What this document is — and is not

- It is a precise functional description of Wakemark's provenance, custody, logging, verification, and audit-export surface, mapped to the vocabulary regulators use.
- It is not legal advice, a compliance certification, or an assumption of your regulatory obligations. Nothing in it — and nothing in the product — transfers an Article 50 duty from you to Wakemark or discharges one on your behalf.
- Wakemark's own phrasing rule applies throughout: Wakemark is designed to support Article 50 workflows. It does not ensure or guarantee compliance, and any Wakemark output that appeared to say otherwise would be wrong.

2. The Article 50 obligations, in brief

Regulation (EU) 2024/1689 (the EU AI Act), Article 50, creates two duties relevant here. Article 50(2) requires providers of AI systems that generate synthetic audio, image, video, or text to mark outputs in a machine-readable format, detectable as artificially generated or manipulated. A company that trains no model but integrates third-party models via API into its own user-facing system can still be a system provider carrying this duty. Article 50(4) requires deployers to disclose deepfakes and certain AI-generated text, and to keep records of their labelling practice. Which role you occupy — provider, deployer, or both — is a legal determination Wakemark never makes for you: the audit export requires you to declare your role and never infers it.

3. What Wakemark provides

The Commission's Code of Practice on Transparency of AI-Generated Content (final, 10 June 2026) expects at least two active machine-readable marking layers and treats logging as a supplementary third. Mapped honestly onto that structure:

Layer	What the Code expects	Where Wakemark honestly sits

Signed metadata (Sub-measure 1.1.1)	Digitally signed, tamper-evident metadata embedded in the output — C2PA Content Credentials is the industry reference.	Wakemark does NOT create this layer. The generating provider embeds it (or doesn't). Wakemark PRESERVES it — outputs are stored byte-exact, never transcoded, so whatever the provider embedded survives — then DETECTS it (a C2PA inspection at persist time, reported in four honest states: verified / present-unverified / none / not-inspected) and LOGS the result.
Invisible watermarking (Sub-measure 1.1.2)	Imperceptible watermarks that survive compression and cropping.	Wakemark neither embeds nor detects invisible watermarks, and every surface that could imply otherwise says so instead. If your outputs need this layer, it must come from the generating provider or a specialised tool.
Fingerprinting / logging (Sub-measure 1.1.3)	Logging or registry lookup — optional, supplementary, and explicitly not sufficient on its own.	This IS Wakemark's ledger: an append-only event record with a sha256 digest per asset and a public verifier. The audit export presents it in the Code's own posture — supplementary, and explicitly not sufficient on its own — never as satisfying the Article 50(2) marking duty.
Preservation of marks (Measure 1.2)	Existing marks must survive processing, including where marked content becomes an input to a new output.	Wakemark's strongest contribution: byte-exact custody plus a cross-provider lineage graph recording how a marked input became a new output — evidence only a layer inside the pipeline can produce.

The Code of Practice is voluntary, and by its own text adherence does not constitute conclusive evidence of compliance.

Concretely, the surface you can put in front of a reviewer:

- A receipt on every stored output — digest, provider, model, capability, C2PA state, lineage — rebuildable from the ledger.
- The public verifier (/verify) — anyone, without an account, can re-hash a file against the recorded digest and validate an embedded C2PA manifest in-browser; a public, redacted lookup API backs it.
- The audit export — a per-project, date-ranged report (JSON + PDF) with per-asset marking coverage, an explicit gap report (what is NOT marked, attributed per provider and model), provenance chains for derived outputs, and a first-class “what Wakemark cannot attest” section. You declare your role (provider / deployer / both); the report never assumes it.
- Customer-controlled retention that deepens rather than erases evidence — a purged asset leaves a tombstone: digest, provenance state, and lineage survive as evidence that the content existed and was purged per your policy.

4. What Wakemark does not do

- It does not mark content. Wakemark is not a C2PA issuer today and adds no metadata to your media (signing is a designed-for future capability, not a shipped one).
- It does not embed or detect invisible watermarks.
- It does not decide whether content is a deepfake — the Article 3(60) classification, and the disclosure that follows, are the deployer's.

- It holds no record of publication-time labelling or editorial review (Article 50(4)) — where, whether, and how you labelled content at publication, and who reviewed it, are records you must keep.
- It does not certify compliance, is not itself certified or audited, and produces no document that establishes compliance — each audit export says this about itself before it says anything else.

5. Who remains responsible

You do. If you are a provider under Article 50(2), the machine-readable marking duty is yours, and Wakemark's ledger — by the Code's own terms — does not satisfy it alone. If you are a deployer under Article 50(4), disclosure and its documentation are yours. Wakemark's role is to make your evidence durable, inspectable, and independently verifiable: the proof layer under your compliance work, not a substitute for it. Wakemark accepts no liability for your regulatory obligations, and this addendum creates none.

6. Dates and sources

Article 50 applies from 2 August 2026; systems already on the market before that date have until 2 December 2026 for the machine-readable marking duty under the Digital Omnibus grandfathering. Penalties reach €15M or 3% of worldwide turnover. Sources: Regulation (EU) 2024/1689 Articles 50, 3(60), 99, 113; Commission Code of Practice on Transparency of AI-Generated Content (final, 10 June 2026). As with any evolving regulation, confirm the current text for your own obligations.

7. Beyond the EU

Provenance, custody, and lineage answer “what is this and where did it come from” — a question no single regulation invented. The audit export is a jurisdiction profile over the same ledger (the current profile is the EU AI Act's Article 50); other jurisdictions' regimes can be additional profiles, not rewrites. Nothing in the data model is EU-specific.

THE OPERATOR

The service is operated by O ÷W¦† â ±1r, an individual (a natural person) based in Türkiye. There is currently no registered company behind the service — “Wakemark” is the product and service name, not a company name. For privacy, data-protection, or legal inquiries, contact us at hello@wakemark.ai.