

# EU AI Act: Second Draft of Code of Practice on Transparency and Watermarking Published

April 7, 2026

On 18 December 2025, we [shared an update](#) about the first draft of Code of Practice, which set critical and broadly applicable transparency and watermarking rules under the European Union Artificial Intelligence Act.

On 5 March 2026, the group of experts responsible for the code published the [second draft of the Code of Practice on Transparency of AI-Generated Content](#) (second draft code).

Where the first draft provided a high-level directional framework, this second draft represents a significant step forward. The document is moving towards an operationally concrete instrument that businesses can realistically map against their existing transparency and content provenance frameworks. While the final code will be a voluntary instrument, it is expected to become a key benchmark against which regulators assess compliance with Article 50.

## Updated requirements for providers

### Digitally signed metadata is now mandatory

This is the most operationally demanding new requirement. The second draft code mandates digitally signed, timestamped metadata containing:

- Whether content is AI-generated or manipulated.
- Interoperable identifier cross-referenceable by other layers.
- Information on how to access the provider's detection tool, all secured with certificate management.

Operationalising this requirement will require dedicated PKI infrastructure, integration with content provenance standards (likely C2PA or equivalent) and ongoing certificate life cycle management.

Fingerprinting or logging mechanisms for AI-generated or manipulated content which allow for checking whether an output has been generated or manipulated by a generative AI system are now framed explicitly as an **optional additional measure** which can be implemented at the discretion of the signatory in addition to digitally signed metadata and imperceptible watermarking.

### Reclassification of GPAI model provider obligations from mandatory to voluntary

General-purpose AI (GPAI) model providers are now “encouraged” rather than required to implement relevant marking techniques at the model level. This increases the burden on system providers integrating such models, who cannot assume upstream models will arrive with built-in watermarking capabilities. For providers making models directly available to system providers, it reduces direct regulatory exposure, but it is possible that industry practice will evolve, or commercial contracts may include provisions requiring watermarking regardless of what the second draft code requires.

### Mandatory, free of charge, EU-localised detection

# mechanisms

The second draft code requires signatories to provide a publicly available detection tool for the content generated or manipulated by their AI systems, which is free of charge, EU localised (implementable locally or hosted within the EU) and General Data Protection Regulation-compliant in its processing. Providers will also be required to make their detection tools available to authorities if the provider exits the market.

## Mandatory cooperation on provider-agnostic detection interface and interoperability infrastructure

The second draft code introduces a cooperative obligation to develop an interoperable, provider-agnostic detection interface. This is combined with the development of a shared repository of public watermarks, metadata repository addresses and detector addresses. The detection interface is required to be executable locally on a computer and must provide a common entry point to all detection mechanisms employed by providers of generative AI systems.

These cooperation obligations are framed as mandatory for signatories in the second draft code. For providers, particularly the largest AI system providers, it means current proprietary detection approaches will need to be designed (or redesigned) with interoperability in mind or stand-alone detection frameworks negotiated across industry from scratch.

## Updated requirements for deployers

### Update from common taxonomy to labelling requirements

The first draft required deployers to classify content as “fully AI-generated” or “AI-assisted” and apply different disclosure accordingly. This would have been a significant ongoing classification burden. Its removal simplifies compliance, and the second draft code replaces this with modality-specific placement rules for the “AI” acronym design standard. This provides more of a checklist approach to compliance with the marking requirement.

Deployers across media, advertising, platforms and publishing will need to redesign content workflows to ensure they are compliant with the requirements detailed in the second draft code. The detailed modality-specific rules (persistent display for short video, repeated disclaimers for long audio, etc.) offer greater compliance certainty in most cases, though deployers should note that implementation of labelling is contextual. Where the disclosure options mentioned in the second draft code are not available or would affect the display or enjoyment of the work, a deployer may conceive of alternative solutions.

### Reduced compliance framework burden

Splits training, compliance and monitoring requirements from the first draft into proportionate compliance, awareness and review, with:

- Internal compliance documentation.
- Awareness and training (which softened from mandatory training to reasonable and proportionate efforts).
- Review, feedback, and cooperation, introducing channels for flagging missing or incorrect disclosures.

### Prohibition on label removal for onward content dissemination

When making reference to disclosures/markings of generated or manipulated content, the second draft code states that disclosure should always “travel with the content”. The second draft code, like the first, does not make any suggestion as to how deployers could prevent end-users from removing content markings before sending generated or manipulated content onward.

The logic of this requirement is also not clear. Given that the provisions of the EU AI Act only apply to deployers when generating deep fakes and/or AI-generated and -manipulated text using an AI model, an end-user could, and very likely will, remove the AI labelling. This could be done for a number of reasons, including the end-user’s own enjoyment of the content or the incorporation of the whole, or a piece, of that generated content into another piece of work. Content that should be marked as synthetic in one context does not obviously need to retain that marking in subsequent contexts.

## Compliance documentation changes

Deployers are now no longer required to record date of review, approval reference and file identifiers – a departure from the previous draft, which required specific logs for reliance on the editorial exemption. These requirements are replaced with the need to publish:

- Identification of the person with editorial responsibility.
- An overview of organisational measures ensuring adequate review.
- A new requirement to publish contact details, where not already publicly available.

## What’s next?

Now that feedback from stakeholders has been incorporated into the draft, the European Commission will consider any further comments with a view to finalising the code by the beginning of June.

This final code will eventually serve as a key mechanism for demonstrating compliance with the EU AI Act’s transparency obligations (which will come into effect on 2 August 2026), though adherence to the code will not by itself constitute conclusive evidence of compliance. Companies operating in the EU should review the changes above, along with our [December 2025 Cooley alert](#), to assess the readiness of their current watermarking and labelling infrastructure and consider making comments if they feel strongly about any aspect of the second draft code.

This content is provided for general informational purposes only, and your access or use of the content does not create an attorney-client relationship between you or your organization and Cooley LLP, Cooley (UK) LLP, or any other affiliated practice or entity (collectively referred to as "Cooley"). By accessing this content, you agree that the information provided does not constitute legal or other professional advice. This content is not a substitute for obtaining legal advice from a qualified attorney licensed in your jurisdiction, and you should not act or refrain from acting based on this content. This content may be changed without notice. It is not guaranteed to be complete, correct or up to date, and it may not reflect the most current legal developments. Prior results do not guarantee a similar outcome. Do not send any confidential information to Cooley, as we do not have any duty to keep any information you provide to us confidential. When advising companies, our attorney-client relationship is with the company, not with any individual. This content may have been generated with the assistance of artificial intelligence (AI) in accordance with our AI Principles, may be considered Attorney Advertising and is subject to our [legal notices](#).

## Key Contacts

Tristan Lockwood London	tlockwood@cooley.com +44 20 7556 4115
----------------------------	--

Leo Spicer-Phelps London	lspicerphelps@cooley.com +44 (0) 20 7556 4334
Janet H. Kim Washington, DC	janetkim@cooley.com +1 202 728 7060
Patrick Van Eecke Brussels	pvaneecke@cooley.com +32 2 486 7501

This information is a general description of the law; it is not intended to provide specific legal advice nor is it intended to create an attorney-client relationship with Cooley LLP. Before taking any action on this information you should seek professional counsel.

Copyright © 2023 Cooley LLP, 3175 Hanover Street, Palo Alto, CA 94304; Cooley (UK) LLP, 22 Bishopsgate, London, UK EC2N 4BQ. Permission is granted to make and redistribute, without charge, copies of this entire document provided that such copies are complete and unaltered and identify Cooley LLP as the author. All other rights reserved.