WILSON SONSINI

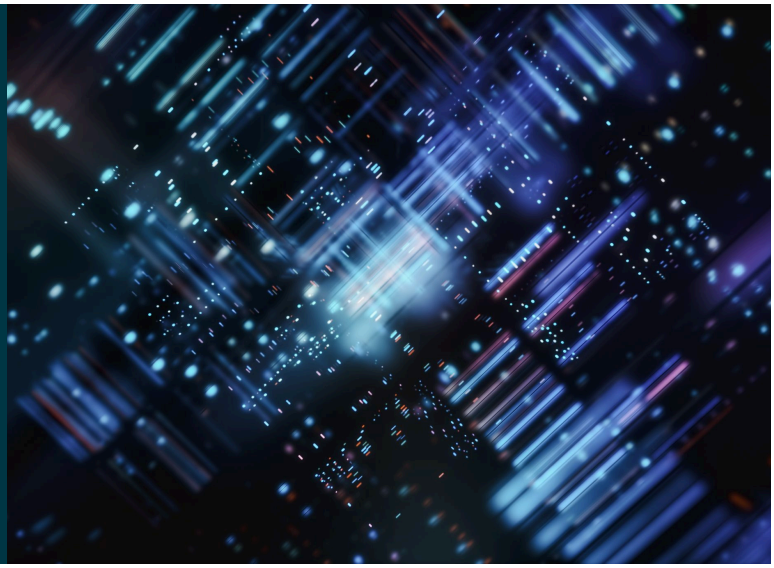

EU Commission Publishes AI Transparency Code of Practice

CONTRIBUTORS

Laura De Boel

Roberto Yunquera Sehwani

Karol Piwonski

ALERTS

July 14, 2026

On June 10, 2026, the European Commission published a Code of Practice on marking and labeling of AI-generated content (the Code) following a public consultation that took place last year. The Code is divided into two sections:

- Section 1 sets out provenance requirements applicable to providers offering generative AI systems.
- Section 2 sets out transparency rules for AI-generated content applicable to deployers of AI tools, such as requirements to visibly label deepfakes and AI-generated text on matters of public interest.

Background

The European Union's Artificial Intelligence Act (EU AI Act) sets out a series of transparency obligations for AI systems. The Code is a voluntary set of measures published by the EU Commission's AI Office, which allow providers and deployers who are signatories to the Code to demonstrate compliance with these obligations. Complying with the Code will create a presumption of compliance. Companies can adopt an alternative approach to the transparency obligations, but they will bear the burden of showing they achieve equivalent compliance.

On May 8, 2026, the EU Commission released draft guidelines on the implementation of the EU AI Act's transparency obligations. These draft guidelines discuss a wider set of obligations and focus on the scope of transparency obligations, whereas the Code details the methods companies can implement to comply with these obligations.

Provenance Requirements for Generative AI Systems

Mandatory Marking. The AI Act requires providers to implement a machine-readable marking on AI-generated audio, image, video, and text. The Code acknowledges that, under the current state of the art, there is no single solution that can provide sufficient reliability. The Code indicates that in most cases signatories will need to implement a multi-layered marking approach including at least two methods. The Code describes the following methods:

- Metadata. Where the content format supports metadata (e.g., images, audio, video, PDFs), providers must record AI-generation status in the metadata, digitally signed and time-stamped in a tamper-evident manner. This does not apply to free-form text, which cannot carry metadata.
- Watermarks. In addition, providers must embed an imperceptible watermark in AI-generated content, except for very short text (under 200 tokens), which is exempt.

The Code lists fingerprinting or logging as an optional third method that signatories could implement in addition to the above. Fingerprinting reduces content to a condensed digital descriptor (similar to a hash) that can later be checked against previously generated content, even if the content has since been altered. Logging, by contrast, requires recording each instance of content generation to maintain a record of the content that was generated by the AI system's output. The Code notes these measures must be implemented in a secure and privacy-preserving manner, although it does not explain in detail how these measures should operate in practice.

Detection Solutions. Providers must make available a detection solution enabling deployers, end-users, competent authorities, and other stakeholders to verify whether content has been generated or manipulated by their AI system. Key requirements include:

- The detection solution may be provided as a public specification (allowing anyone to build a detection tool), software, or a cloud-based API.
- Solutions must generally be free of charge, subject to a narrow exception for smaller providers facing high-volume requests.
- Detection results must be clear and comprehensible.
- Solutions must include safeguards to protect the privacy of submitted content, including a zero-retention policy.

Quality Requirements. Providers must ensure that both their marking and detection solutions satisfy the following key requirements:

- Effectiveness. Results must be understandable to natural persons.
- Reliability. Solutions must achieve low error rates across diverse content types, lengths, and contexts.
- Robustness. Solutions must maintain performance under common processing operations (e.g., recompression, cropping, paraphrasing) and adversarial attacks (e.g., watermark removal attempts).
- Interoperability. Providers must implement an interoperability solution for watermark detection such as a standardized API access method, a publicly readable signpost mechanism embedded in content, or participation in a consortium detection solution by February 2, 2027.

Providers may satisfy these requirements through third-party or upstream model-level marking solutions, while retaining ultimate responsibility for compliance. This means that in many cases, companies that market AI systems based on third-party AI models will need to seek assistance from their vendors to be technically able to comply with these rules.

AI systems placed on the market after August 2, 2026, must comply from the time they are placed on the market. Systems already on the market before that date have until December 2, 2026, to comply (a deadline extended from August 2026 by the recent amendments to the EU AI Act adopted via the AI Omnibus Regulation—see our client alert here).

Key Requirements for Deployers of AI Systems

Mandatory disclosure of deepfakes and AI-generated published text. Subject to limited exceptions, users of generative AI tools must disclose the artificial origin of:

- AI-generated content that constitutes a deepfake (i.e., image, audio, or video content resembling real persons, objects, places, entities, or events that would falsely appear authentic); and
- AI-generated or manipulated text published on matters of public interest (e.g., AI summary of a news article published on a newspaper's website, AI-edited investor reports on a listed company's website).

The Code introduces a publicly available EU icon as the primary disclosure mechanism. Deployers may alternatively use equivalent labels complying with the Code's design and placement specifications. For instance, these alternative labels must feature the capitalized acronym "AI" and be immediately perceivable without user interaction. Labels must appear at the beginning of and at regular intervals throughout videos, above or near the headline of published text, and alongside a plain-language disclaimer at the beginning of audio outputs.

Companies are also encouraged to include a second interactive layer providing additional information on the type and extent of AI involvement (e.g., text next to a deepfake picture disclosing that a face has been altered by AI).

Deployers must maintain documentation of how they implement disclosure obligations and put in place processes to verify correct labeling, proportionate to their size and resources. Deployers are encouraged to provide feedback channels for flagging mislabeled content and to cooperate with competent authorities. These obligations will apply from August 2, 2026.

Next Steps

The transparency requirements come into effect on a staggered basis. These deadlines set concrete milestones for a broad range of companies. Recommended actions include:

- *Audit current marking and disclosure practices.* Providers and deployers should map existing marking processes against the EU AI Act's mandatory requirements to identify gaps before the applicable deadlines.
- *Implement marking and detection infrastructure.* Providers of generative AI systems must implement marking and detection infrastructure and begin compliance testing ahead of the December 2026 deadline. AI system providers that rely on third-party models may want to check with their model providers on how to implement necessary safeguards at the model-level (such as watermarking or metadata).
- *Plan for the interoperability deadline.* Providers should assess which interoperability pathway best fits their technical architecture, ahead of the February 2027 deadline.
- *Document compliance processes.* Companies should maintain documentation in anticipation of requests from competent authorities.
- *Consider adherence to the Code.* Providers and deployers of generative AI systems may sign the Code by completing the process detailed by the AI Office here. While the Code is voluntary, the AI Office states that, for signatories, it will focus its enforcement on monitoring adherence to the Code, which offers greater predictability and reduced administrative burden. Companies that instead adopt their own approach to their transparency obligations will bear the burden of demonstrating that their measures are adequate. Providers and deployers may sign the Code at any time; however, the AI Office encourages signing up by July 22, 2026, to be included in the list of initial signatories.

Non-compliance with EU AI Act transparency requirements can result in enforcement action under the EU AI Act's supervisory framework, including requests for information, fines, and corrective orders.

For more information or if you have any questions regarding the EU AI Act, please contact Laura De Boel, Cédric Burton, Yann Padova, or Nikolaos Theodorakis from Wilson Sonsini's Data, Privacy, and Cybersecurity practice.

Wilson Sonsini's AI Working Group assists clients with AI-related matters. Please contact Laura De Boel, Maneesha Mithal, Manja Sachet, or Scott McKinney for more information.

Roberto Yunquera Sehwani, Karol Piwonski, Hugh Ó Laoide Kelly, and María Rodríguez Hernández contributed to the preparation of this alert.