

Alert | Data Privacy & Cybersecurity/ Innovation & Artificial Intelligence



June 2026

Deepfakes, Chatbots, AI-Generated Text: European Commission Details Transparency Obligations Under the AI Act

On 8 May 2026, the European Commission published [draft guidelines on the implementation of the transparency obligations](#) under Article 50 of Regulation (EU) 2024/1689 (the AI Act). The draft guidelines describe how the AI Act's four transparency obligations are intended to apply to (i) interactive AI systems; (ii) providers of AI systems that generate or manipulate synthetic content; (iii) deployers of emotion recognition and biometric categorization systems; and (iv) deployers of deepfakes and AI-generated text on matters of public interest. While non-binding, European Commission guidelines carry considerable practical importance in the application of EU law.

Key Takeaways

- **Deepfake rules apply even without intent to deceive.** Content that looks or sounds like a real person must be labeled – even if no deception was intended and even if no real individual is depicted. (*Art. 50(4) AI Act*)
- **AI systems must disclose themselves clearly and upfront.** Users need to know they are interacting with AI at the moment of contact – not through a reference buried in fine print. This includes agentic AI that acts autonomously on behalf of users. (*Art. 50(1) AI Act*)
- **The exception for editorially reviewed AI text is narrow.** Simply having a human “check” AI-generated content is not sufficient. Only genuine, substantive editorial oversight with clear accountability qualifies for an exemption from labeling obligations. (*Art. 52(1) AI Act*)

The stakeholder consultation period for the draft guidelines closed on 3 June 2026. The underlying obligations of the AI Act as currently in force will become applicable on 2 August 2026, subject to targeted transitional relief for the Article 50(2) marking and detection obligations contemplated by the [draft AI Omnibus](#).

Background and Scope

Article 50 of the AI Act establishes four transparency obligations addressing different actors along the AI value chain:

1. **Providers of interactive AI systems** must design those systems so that users are informed of the artificial, non-human nature of the interaction (Article 50(1)).
2. **Providers of AI systems**, including general-purpose AI systems, that generate or manipulate synthetic audio, image, video, or text content must implement machine-readable marking and ensure outputs are detectable as AI-generated or manipulated (Article 50(2)).
3. **Deployers of emotion recognition and biometric categorization systems** must inform exposed individuals of the system's operation (Article 50(3)).
4. **Deployers of deepfakes** and of AI-generated text published to inform the public on matters of public interest must disclose the artificial origin of the content (Article 50(4)).

The draft guidelines confirm that these obligations may apply cumulatively to a single AI system, with authoritative interpretation reserved for the Court of Justice of the European Union.

Interactive AI Systems and Disclosure Standards

The European Commission interprets the disclosure obligation broadly. Whenever an AI system interacts directly with people, it must identify itself as AI. This explicitly includes agentic AI – systems that carry out tasks autonomously. It does not matter whether the AI communicates with the person who gave the instruction or with third parties it encounters along the way. If a provider cannot reliably rule out that its AI agent will come into contact with people, the operator must ensure the agent discloses its artificial nature in every situation where such contact is likely (*Art. 50(1) AI Act*),

The draft guidelines are equally clear about what does not constitute sufficient disclosure. A reference in the terms and conditions or product documentation is not sufficient. Technical labels such as metadata or watermarks alone do not meet the requirement either – users typically do not notice them at the point of interaction. Vague terms such as “assistant” or technical descriptions like “this system uses LLMs” also fall short.

Instead, the draft guidelines recommend a combination of formats: clearly visible plain-language notices, audio cues, and persistent visual indicators. The requirements are particularly strict when children or other vulnerable groups are part of the intended audience.

An Expanded Reading of Deepfakes

An important set of clarifications in the draft guidelines concern the deepfake definition under Article 3(60) of the AI Act and the corresponding labeling obligation in Article 50(4). The draft guidelines confirm that the assessment of whether content falsely appears authentic or truthful does not depend on

the deployer’s intention to deceive or mislead. Accordingly, the absence of fraudulent intent would not defeat the labeling requirement.

The reference to existing persons, objects, places, entities, or events is similarly read broadly. The draft guidelines note that it would be sufficient for the simulated subject to resemble someone or something that could exist or could once have existed in reality. A realistic synthetic depiction of a fictitious but natural-looking person would therefore constitute a deepfake within the meaning of Article 3(60) of the AI Act, even where no identifiable rights-holder was implicated. By contrast, clearly unrealistic content — for example fantasy scenes or depictions that defy the laws of nature or biology, such as flying humans, dragons, or elephants driving cars — would generally fall outside the scope of that definition.

Finally, the audience test described in the draft guidelines diverges from the established EU consumer-law standard. Whereas the “obviousness” exception under Article 50(1) is calibrated to a reasonably well-informed, observant, and circumspect average member of the target audience, the Article 50(4) deepfake assessment must account for the potentially diverse composition of the actual exposed audience, including foreseeable downstream exposure to children, elderly persons, or audiences with lower digital and AI literacy. Deployers should therefore not rely solely on the media literacy of their primary audience where the content is likely to reach vulnerable or less AI-literate groups as well.

AI-Generated Text on Matters of Public Interest

Anyone who publishes AI-generated or AI-edited text with the purpose of informing the public on matters of public interest must disclose that the text has been artificially generated or manipulated (Article 50(4) AI Act). The draft guidelines define “public interest” broadly: it encompasses, among other things, public administration, fundamental rights, health, the environment, and consumer protection, as well as economic, political, scientific, or cultural developments of societal relevance. At the same time, the draft guidelines interpret the exception for content subject to human review and editorial responsibility narrowly.

Text is considered “published” if it is accessible to an indefinite and relatively large number of readers. Private messages and internal documents are not covered.

What Does This Mean for Newsrooms, Businesses, and Creators?

Any organization seeking to rely on the exemption from the labeling obligation must satisfy two conditions simultaneously:

1. **Genuine substantive review by a human.** A person with relevant expertise must have reviewed the text specifically for its content. Spell-checking or grammar correction alone is not sufficient, nor is a cursory editorial sign-off.
2. **Clearly attributable editorial responsibility.** An identifiable person or organization must bear responsibility for the content. Their name and contact details must be publicly accessible, and the responsible party must have the authority to approve, amend, or reject the content.

Significance for Deployers

The draft guidelines may require organizations to reassess their existing AI governance frameworks. The broad definition of deepfakes, combined with the intent-independent and audience-oriented assessment, has far-reaching consequences: organizations using synthetic media for creative or commercial purposes

that depict realistic people, places, objects, or events should generally assume that the labeling obligation applies.

For artistic, creative, satirical, or fictional content, the draft guidelines do not provide a full exemption. A reduced disclosure obligation applies – one that need only go far enough to avoid impairing the work or its impact.

Note that compliance with the AI Act alone does not satisfy all applicable labeling obligations. Organizations may simultaneously be subject to requirements under competition law, platform rules, consumer protection laws, contractual disclosure obligations, and personality and intellectual property rights. Each of these legal frameworks may impose its own labeling requirements that go beyond those of the AI Act.

Practical Implications

- Marketing teams running campaigns with AI-generated visuals, customer testimonials, or product depictions should review their labeling practices.
- Corporate communications departments using AI-drafted investor materials or sustainability reports should assess whether their editorial review meets the substantive requirements set out in the draft guidelines.
- Organizations deploying agentic AI in customer-facing contexts must ensure the AI identifies itself directly at the point of interaction – not through a reference buried in the terms and conditions.
- Platform operators and their users should note that automated AI labels applied by platforms do not relieve the operator of its own disclosure obligation. Platform-side labels may complement, but cannot replace, the operator's own disclosure.

(Art. 50(4) AI Act)

Outlook

The stakeholder consultation period closed on 3 June 2026, and the European Commission may release final guidelines in the near term. In parallel, the AI Office is finalizing a Code of Practice on the marking and labeling of AI-generated content. Adherence to a Code of Practice deemed adequate by the AI Office may offer a route to demonstrating compliance with Articles 50(2), (4), and (5) of the AI Act. Non-signatories may demonstrate compliance through alternative means but may face heavier evidentiary burdens and more frequent information requests from market surveillance authorities.

Penalties for infringement may reach up to EUR 15 million or 3% of total worldwide annual turnover, whichever is higher.

Under the AI Act as currently in force, the transparency obligations apply from 2 August 2026. The AI Omnibus proposal, on which the European Parliament and the Council have reached political agreement, contemplates targeted transitional relief for the Article 50(2) marking and detection obligations for artificially generated content. The Council has announced a revised deadline of 2 December 2026 for those transparency obligations, though formal adoption remains pending. Organizations within scope may wish to audit their AI systems, disclosure practices, and editorial workflows in advance of the application date.

Authors

This GT Alert was prepared by:

- **Philip Radlanski** | +49 30 70 017 1337 | Philip.Radlanski@gtlaw.com
- **Carsten Kociok** | +49 30 70 017 1119 | Carsten.Kociok@gtlaw.com

Abu Dhabi[†]. Albany. Amsterdam. Aspen. Atlanta. Austin. Berlin[†]. Boston. Charlotte. Chicago. Dallas. Delaware. Denver. Dubai[†]. Fort Lauderdale. Houston. Las Vegas. London^{*}. Long Island. Los Angeles. Mexico City⁺. Miami. Milan[°]. Minneapolis. Munich[†]. New Jersey. New York. Northern Virginia. Orange County. Orlando. Philadelphia. Phoenix. Portland. Riyadh^{*}. Sacramento. Salt Lake City. San Diego. San Francisco. São Paulo[›]. Seoul[∞]. Shanghai. Silicon Valley. Singapore[°]. Tallahassee. Tampa. Tel Aviv[^]. Tokyo[∞]. Warsaw[~]. Washington, D.C. West Palm Beach. Westchester County.

This Greenberg Traurig Alert is issued for informational purposes only and is not intended to be construed or used as general legal advice nor as a solicitation of any type. Please contact the author(s) or your Greenberg Traurig contact if you have questions regarding the currency of this information. The hiring of a lawyer is an important decision. Before you decide, ask for written information about the lawyer's legal qualifications and experience. Greenberg Traurig is a service mark and trade name of Greenberg Traurig, LLP and Greenberg Traurig, P.A. [†]Greenberg Traurig's Abu Dhabi office is a branch of Greenberg Traurig, P.A., which is registered with the Abu Dhabi Global Market Registration Authority (Registration No. 29906) and licensed to carry out legal services and regulated as a DNFBP by the ADGM Financial Services Regulatory Authority. [†]Greenberg Traurig's Berlin and Munich offices are operated by Greenberg Traurig Germany, LLP, an affiliate of Greenberg Traurig, P.A. and Greenberg Traurig, LLP. [†]Greenberg Traurig's Dubai office is operated by Greenberg Traurig Limited, a company registered in the Dubai International Financial Centre (Registration No. CL7238), regulated as a DNFBP by the Dubai Financial Services Authority and licensed by The Government of Dubai Legal Affairs Department. ^{}Operates as a separate UK registered legal entity. ⁺Greenberg Traurig's Mexico City office is operated by Greenberg Traurig, S.C., an affiliate of Greenberg Traurig, P.A. and Greenberg Traurig, LLP. [»]Greenberg Traurig's Milan office is operated by Greenberg Traurig Studio Legal Associato, an affiliate of Greenberg Traurig, P.A. and Greenberg Traurig, LLP. [«]Greenberg Traurig operates in the Kingdom of Saudi Arabia through Greenberg Traurig Khalid Al-Thebity Law Firm, a professional limited liability company, licensed to practice law by the Ministry of Justice. [›]Greenberg Traurig's São Paulo office is operated by Greenberg Traurig Brazil Consultores em Direito Estrangeiro – Direito Estadunidense, incorporated in Brazil as a foreign legal consulting firm. Attorneys in the São Paulo office do not practice Brazilian law. [∞]Operates as Greenberg Traurig LLP Foreign Legal Consultant Office. [°]Greenberg Traurig's Singapore office is operated by Greenberg Traurig Singapore LLP which is licensed as a foreign law practice in Singapore. [^]Greenberg Traurig's Tel Aviv office is a branch of Greenberg Traurig, P.A., Florida, USA. [∞]Greenberg Traurig's Tokyo Office is operated by GT Tokyo Horitsu Jimusho and Greenberg Traurig Gaikokuhojimubengoshi Jimusho, affiliates of Greenberg Traurig, P.A. and Greenberg Traurig, LLP. [~]Greenberg Traurig's Warsaw office is operated by GREENBERG TRAUIG Nowakowska-Zimoch Wysockiński sp.k., an affiliate of Greenberg Traurig, P.A. and Greenberg Traurig, LLP. Certain partners in GREENBERG TRAUIG Nowakowska-Zimoch Wysockiński sp.k. are also shareholders in Greenberg Traurig, P.A. Images in this advertisement do not depict Greenberg Traurig attorneys, clients, staff or facilities. No aspect of this advertisement has been approved by the Supreme Court of New Jersey. ©2026 Greenberg Traurig, LLP. All rights reserved.*