

# The European Commission issues draft guidelines on the transparency requirements under the AI Act

21 May 2026

On 8 May 2026, the European Commission issued draft guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of the AI Act (the “guidelines”). These are intended to provide practical guidance for organisations that are providers or deployers of AI systems, to ensure compliance with Article 50 AI Act. A [public consultation on the guidelines](#) is open until 3 June 2026. The guidelines complement the Commission's already published draft [Code of Practice on Transparency of AI-Generated Content](#). The Code of Practice focuses on the implementation of transparency measures from a technical perspective, while the guidelines expand on the scope and application of the requirements.

Following the agreement reached between the Council of the EU and the European Parliament on 7 May 2026 to delay certain provisions of the AI Act ([see our recent blog post](#) for more detail), transparency obligations will be the next to take effect on 2 August 2026, except for the obligations

relating to the watermarking of synthetic AI content by providers, which have been postponed until 2 December 2026.

## The transparency obligations under Article 50 AI Act

Article 50 sets out four transparency requirements with very broad application, covering providers or deployers of all AI systems (not just high-risk or general-purpose AI) which:

1. interact directly with natural persons;
2. generate or manipulate synthetic image, video, audio or text content;
3. involve emotion recognition or biometric categorisation functionality; or
4. generate or manipulate deep fakes or text published to inform the public on matters of public interest.

Providers and deployers must provide transparency information in a clear and distinguishable manner, by the time the relevant individual first interacts with the content or AI system, so the user can easily identify and understand it.

## Key takeaways from the guidelines

The guidelines aim to provide practical clarification on the four Article 50 obligations.

### 1. Informing people they are interacting with AI (Art 50(1) AI Act)

When an AI system is intended to interact directly with individuals, providers must ensure users know that they are interacting with AI. Users can be notified in whichever format is appropriate, including via text, audio or visual information. This allows users to make informed decisions

about the outputs of the AI systems they interact with. This requirement emphasises the importance of clear user-facing notices, which can be provided in text, audio, visual or multi-modal formats, particularly where AI systems are used by broad, diverse audiences.

The only exceptions to this disclosure requirement are if (i) the use of AI is obvious; or (ii) the AI system is authorised to help prevent crime. When relying on the “obviousness exception”, providers must be able to show that the interaction is obviously artificial to a well-informed, observant member of their target audience.

## 2. Marking and detecting AI-generated or manipulated content (Art 50(2) AI Act)

If an AI system generates synthetic content or alters pre-existing content, providers must implement technical solutions enabling people to verify the origin of the content and clearly distinguish what has been affected by AI.

This obligation is two-fold. First, providers must ensure the outputs of their AI system are marked in a machine-readable format. Second, they must provide a means for detecting this mark, so that natural persons are able to identify and distinguish AI-affected content.

This obligation is expected to have a significant impact on generative AI systems, including general purpose AI which is used broadly to produce text, images, audio and other types of content. It also highlights the importance of traceability throughout AI lifecycles.

## 3. Emotion recognition and biometric categorisation (Art 50(3) AI Act)

Deployers must inform individuals when they are exposed to AI systems used for emotion recognition or biometric categorisation. Emotion recognition is the use of AI to infer a person’s emotions, based on biometric data, and biometric categorisation involves aligning people to specific categories on the basis of biometric data.

The purpose of this obligation is to ensure individuals are aware of systems which are particularly intrusive from a privacy perspective. There is no designated format in which this

information must be provided, but it must be clear and distinguishable.

As emotion recognition systems are classified as high-risk under the AI Act, this transparency obligation operates alongside other safeguards and requirements applicable to high-risk AI systems. Furthermore, any processing of personal data by these systems will be subject to the applicable data protection law. In particular, processing of biometric data inevitably falls within the scope of the GDPR, including variable interpretations of this concept.

#### 4. Labelling deep fakes and certain text publications (Art 50(4) AI Act)

Deployers must clearly label (i) deep fakes, and (ii) text generated or manipulated by AI and published to inform the public about matters in the public interest. This additional disclosure requirement is in place to combat the high risk of manipulation that comes with this type of content.

However, if a deep fake is "*evidently artistic, creative, satirical, fictional or analogous*", the disclosure obligation is limited to labels applied in an "*appropriate manner that does not hamper the display or enjoyment of the work*".

## Next steps

Organisations should be taking steps now to prepare for the imminent introduction of these requirements. The key next steps to take before these obligations enter into force on 2 August 2026 (and 2 December 2026 for Article 50(2)), include:

- Carrying out an applicability assessment to identify the relevant AI systems that are in-scope;
- preparing transparency disclosures in line with the requirements under the guidelines and, where relevant, the draft Code of Practice; and
- updating user-facing documentation as well as internal governance procedures.

We can likely expect a final version of the guidelines to be adopted in advance of the August application date, with the ongoing public consultation closing on 3 June 2026.

## Conclusion

The guidelines signal a clear expectation by the EC of a “transparency by design” approach, where transparency must be built-into the full lifecycle of AI systems, from the moment a user first interacts with the system through to marking relevant outputs.

Considering the operational challenges that compliance with these wide-ranging obligations may pose and the substantial penalties for non-compliance (with fines up to £15m or 3% of an undertaking's worldwide annual turnover), providers and deployers should use the guidelines to begin preparing now, without waiting for any further designation from a regulator or issuance of guidelines.


Authored by Dan Whitehead, Etienne Drouard, Sara Marinoni, and Michaela Glass.

## Contacts



Dan Whitehead

Partner

 London  
 [Email me](#)



Etienne Drouard

Partner

 Paris  
 [Email me](#)



Sara Marinoni

Senior Associate

 London

 [Email me](#)



Michaela Glass

Associate

 London

 [Email me](#)