



**Nemko**  
Digital

# **AI Transparency: The First Real Test of AI Governance**

When the European Union adopted the [Digital Omnibus on AI](#) in July 2026, much of the attention focused on the revised compliance calendar for [high-risk AI systems](#). The changes responded partly to delays in harmonised standards, national governance structures, and conformity assessment capacity, giving businesses additional time to prepare for some of the AI Act's most demanding requirements. However, they did not amount to a general postponement of AI Act compliance.

The transparency obligations in [Article 50](#) became **applicable on 2 August 2026**. The Digital Omnibus did not change the substance of the principal duties in Article 50(1) to (6), although it introduced a limited transition period for one specific requirement. Providers of generative AI systems placed on the market before 2 August 2026 have until **2 December 2026** to implement the Article 50(2) requirement to mark synthetic outputs in a machine readable and detectable format.


That distinction matters. While some high-risk AI requirements have moved further into the future, most Article 50 obligations already apply. The December deadline extension is narrow: it applies only to the machine-readable marking duty, only to providers, and only to generative AI systems placed on the market before 2 August 2026. It is not an extension for deployer labelling obligations or for systems entering the market from 2 August 2026 onwards.

There is also no general requirement to label content retroactively. According to the European Commission, deepfakes generated before 2 August 2026 do not have to be retroactively labelled, although voluntary labelling is encouraged where possible.

For businesses, the immediate task is therefore not simply to ask whether they use AI. They must determine which role they hold for each system, **which Article 50 duty follows from that role, and whether the December transition applies.**

## Four obligations, two different addressees

Article 50 contains four principal transparency duties, divided between providers and deployers. Treating them as one general obligation risks making companies responsible for measures that belong to their supplier, while overlooking responsibilities arising from their own use of AI.



This means that a company using a third-party chatbot under the supplier's name does not automatically inherit the provider's Article 50(1) obligation to design the system in such a way that users are informed that they are interacting with AI. That duty sits with the provider.

However, the answer changes when a company develops an AI system, has one developed for it, or places it on the market or puts it into service under its own name or trademark. Under the AI Act's definition, that company is the provider, even if another party supplied the underlying technology. A white-labelled chatbot presented as "Ask [YourBrand]" may therefore place the customer-facing business in the provider role. The practical question is not only who built the system, but whose name is on it when it reaches users.

The same organization may hold different roles across different use cases. It might be a deployer when employees use an external generative AI service, but a provider when it launches a branded AI assistant or AI-enabled product. Roles must therefore be assessed for each system rather than assigned once at company level.

## The requirements are important, but so are the exemptions

Article 50 is not a blanket requirement to label every AI interaction, edit, or piece of AI assisted content. Its exemptions and limitations are central to determining what organizations actually need to do.

For systems that interact directly with individuals, the provider does not have to provide a separate notice where the AI involvement is already obvious to a reasonably well informed, observant, and circumspect person in the relevant circumstances. Whether an interaction is sufficiently obvious remains a contextual assessment.

For synthetic-content marking, Article 50(2) does not apply where the AI system performs an assistive function for standard editing or does not substantially alter the input supplied by the deployer or its meaning. Standard grammar correction, spelling assistance, or limited editing functionality may therefore fall outside the marking duty. Generating a new image, substantial passage of text, audio recording, or video is more likely to require machine-readable marking.

The deployer duty for public-interest text is also narrower than a general rule for AI generated communications. It applies to AI-generated or manipulated text published for the purpose of informing the public on matters of public interest. It does not apply where the content has undergone human review or editorial control and a natural or legal person holds editorial responsibility for its publication. Ordinary marketing copy will not necessarily fall within this category, and meaningful human editorial review may exclude otherwise relevant public-interest text from the disclosure requirement.

Deepfakes forming part of evidently artistic, creative, satirical, fictional, or analogous works are not fully exempt, but the disclosure requirement is limited. Disclosure must be provided in an appropriate manner that does not hamper the display or enjoyment of the work.

These limitations show why organizations need a structured assessment rather than a universal "AI-generated" label. Over-labelling ordinary editing assistance or appropriately reviewed business communications can create unnecessary processes without improving compliance.

## Transparency must be timely, clear, and accessible

Article 50(5) adds practical requirements to the transparency obligations. Information must be provided in a clear and distinguishable manner, no later than the first interaction or exposure, and must meet applicable accessibility requirements.

For organizations, transparency is therefore not only a legal question but also a product, content, and user experience question.

A chatbot notice needs to appear when the person begins interacting with the system, while a label for a deepfake or relevant public-interest content should accompany the person's exposure to it. Relevant teams may need to consider the wording, placement, accessibility, and continued visibility of disclosures as content moves across channels.

Non-compliance can attract fines of up to €15 million or 3% of total worldwide annual turnover. The AI Act provides that, for undertakings, the percentage-based or fixed amount may apply depending on the circumstances, while SMEs are subject to the lower rather than the higher of the applicable amounts.

## Emotion recognition requires a separate warning

The Article 50(3) obligation to inform individuals about emotion-recognition systems should not be mistaken for permission to use them. Article 5(1)(f) prohibits AI systems used to infer a person's emotions in workplaces and educational institutions, except where the use is intended for medical or safety reasons. Before considering transparency measures, organizations should therefore determine whether the use case is permitted at all.

This is particularly relevant for HR, education, call-centre monitoring, and employee wellbeing applications. A product may be promoted using terms such as "engagement," "sentiment," "stress," or "wellbeing," but its functionality may still involve inferring emotions. The assessment should focus on what the system does, not only how the supplier describes it.

## How the market is responding

The [EU Code of Practice on Transparency of AI-Generated Content](#) has become an important part of the implementation landscape. It covers provider marking under Article 50(2), deployer labelling under Article 50(4), and the presentation requirements in Article 50(5). It contains separate sections for providers and deployers, including measures concerning machine-readable marking, deepfake disclosure, public-interest text, and label presentation.


Signing the Code is voluntary, but the underlying Article 50 obligations are binding. The Commission and the AI Board have assessed the Code as an adequate voluntary tool for demonstrating compliance. Signatories can rely on its measures to demonstrate compliance across EU Member States, while organizations following another route must be able to show that their alternative measures are adequate. This makes signing the Code a practical governance decision rather than a symbolic commitment.

Approximately [190 organizations had signed](#) the Code by the end of July 2026. The Commission's published examples of provider-side signatories include Anthropic, Google, Meta, Microsoft, OpenAI, Mistral AI, Cohere, and Synthesia. Deployer-side examples include Getty Images, Lenovo, Lufthansa, and Iberdrola.

What is particularly interesting is that, despite broad support for the Code, the market has not converged on a single technical solution. Instead, providers are implementing transparency through a combination of watermarking, provenance information, and user facing disclosures.

- One approach is watermarking, where a signal is embedded directly into AI generated content. [Anthropic has announced](#) watermarking for Claude-generated text, while [Google's SynthID](#) technology embeds imperceptible watermarks into supported AI-generated text, images, audio, and video. In both cases, the objective is to make AI-generated content detectable without altering the user experience.
- A second approach focuses on content [provenance](#) and metadata. OpenAI and Microsoft use Content Credentials based on the C2PA standard to attach information about how content was created, modified, or generated. Unlike watermarking, which embeds a signal into the content itself, provenance mechanisms provide structured information about the content's origin and history.
- A third approach relies on user-facing labels and disclosures. [Meta's "AI info" labels](#) are intended to help users understand when content has been generated or modified using AI. Similar principles are also reflected in the EU Code of Practice, which includes guidance on the labelling of deepfakes and certain AI-generated content.





Taken together, these examples suggest that the industry is converging on the **objective of transparency, but not on a single implementation method**. Watermarks, provenance metadata, and visible labels each address different aspects of the transparency challenge, and many providers are combining multiple approaches rather than relying on just one. For organizations adopting AI, this means transparency cannot be treated as a simple vendor feature. Different systems provide different levels of traceability, detectability, and disclosure, which makes governance and due diligence increasingly important.

# From inventory to decisions: organising for transparency

Transparency compliance starts with understanding both the organization's role and the systems it uses. An inventory should record suppliers, branding, deployment dates, relevant functionality, generated content types, affected individuals, and publication channels. The purpose is not simply to build a list of AI tools, but to identify where particular Article 50 obligations, exemptions, prohibitions, or deadlines apply.

## The assessment should address several concrete questions:

**Role:** Is the organization a provider, a deployer, or both?

**Branding:** Is a third-party system being offered or put into service under the organization's name or trademark?

**Timing:** Was the generative AI system placed on the market before 2 August 2026, bringing it within the limited transition to 2 December 2026?

**Functionality:** Does the system interact directly with individuals, generate synthetic content, infer emotions, categorise biometric data, or create deepfakes?

**Publication:** Is AI-generated text being published to inform the public on matters of public interest?

**Editorial control:** Has the text undergone meaningful human review, and does a person or organization hold editorial responsibility?

**Exceptions:** Is the functionality limited to standard editing, or is the content artistic, satirical, fictional, or analogous?

**Implementation:** Are disclosures clear, distinguishable, accessible, and delivered by the first interaction or exposure?

**Code decision:** Should the organization sign the EU Code of Practice, or can it demonstrate that an alternative approach is adequate?

These questions cut across legal analysis, technical verification, product design, content governance, and supplier management. Procurement can require suppliers to explain their marking and detection mechanisms. Product teams can verify when chatbot notices appear. Communications teams can distinguish ordinary marketing content from unreviewed public-interest publications. Compliance teams can document role determinations, exemptions, and the evidence supporting their conclusions.

*The objective is not to label every use of AI. It is to apply the right control to the right system, by the right actor, at the right time.*

## How We Help

At Nemko Digital, we help organizations determine where Article 50 creates an obligation and where an exemption, role distinction, or prohibition changes the answer. This includes identifying whether the organization is acting as a provider or deployer, assessing whether white-labelling or own-brand deployment changes that role, and determining which systems fall within the **2 December 2026 transition**.

Our support can include:

- Mapping AI systems, roles, brands, suppliers, users, outputs, and deployment dates
- Assessing Article 50 obligations and documenting applicable exemptions
- Reviewing chatbot notices, content labels, accessibility, and timing against Article 50(5)
- Testing whether machine-readable marks and provenance information are implemented and retained through relevant workflows
- Assessing emotion-recognition and biometric-categorisation use cases against both Article 5 prohibitions and Article 50 transparency duties
- Evaluating the advantages, commitments, and implementation implications of signing the EU Code of Practice
- Reviewing supplier documentation, contractual allocation of responsibilities, and evidence of compliance
- Conducting targeted Article 50 gap assessments and independent readiness reviews

The Digital Omnibus gave parts of the high-risk AI framework more time, but most transparency obligations are already applicable, and the limited transition for machine readable marking ends on 2 December 2026. The immediate challenge is therefore not simply understanding Article 50. It is making defensible decisions about roles, systems, exemptions, technical measures, and ownership before those decisions are tested by customers, regulators, or the market.

## About the author

---



### **Alicja Halbryt**

Before joining Nemko Digital, Alicja Halbryt worked for the Dutch Ministry of Economic Affairs as an AI standardisation expert at CEN/CENELEC JTC 21. She holds an MSc in Philosophy of Technology and an MA in Human-Centred Design. She is dedicated to shaping an ethical and human-aligned AI ecosystem.