



Transparency obligations

Four paragraphs, different duties, different operators. Which role owes each disclosure, what triggers it, and what discharging it actually looks like.

4

distinct duties

2

owed by providers

2

owed by deployers

Who owes the disclosure, and when

The four paragraphs of Article 50 put different duties on different operators, and the provider/deployer split is the part most often got wrong. Each duty below records which role owes it.

Who owes which disclosure

	PROVIDER	DEPLOYER
50(1) Interaction notice	☒	☐
50(2) Machine-readable mark	☒	☐
50(3) Emotion / biometric notice	☐	☒
50(4) Deep fake disclosure	☐	☒

The split is written into the paragraphs themselves. A deployer cannot discharge a provider's duty by disclosing twice, and a provider cannot push its own duty down the chain by documenting it.

50(1) — Tell people they are dealing with an AI system

Provider. Providers must design AI systems that interact directly with people so that those people are informed they are interacting with an AI system, unless that fact is already obvious to a reasonably well-informed observer.

Triggers. Any AI system intended to interact directly with natural persons: customer support chatbots, voice agents, AI receptionists, conversational assistants embedded in products.

- A disclosure shown at or before the first interaction, not buried in terms of service
- Wording a non-specialist understands: “You are chatting with an AI assistant”, not “powered by machine learning”
- Persistent labelling in voice channels, where there is no screen to carry a badge
- Meeting the accessibility requirements that Article 50(5) cross-applies

Exemptions: Where it is obvious from the point of view of a reasonably well-informed, observant and circumspect person, taking the circumstances and context of use into account; AI systems authorised by law to detect, prevent, investigate or prosecute criminal offences, subject to safeguards for third-party rights.

Article 50(1)

50(2) — Mark synthetic output in a machine-readable format

Provider. Providers of AI systems that generate synthetic audio, image, video or text must mark the output in a machine-readable format and make it detectable as artificially generated or manipulated.

Triggers. Generative systems producing audio, images, video or text, including general-purpose AI systems. The duty sits with the provider of the generating system, not the person who later publishes the output.

- A marking technique that is effective, interoperable, robust and reliable as far as is technically feasible
- Accounting for the specificities and limitations of each content type: text is far harder to mark durably than images
- Proportionate implementation cost, and the generally acknowledged state of the art as reflected in relevant technical standards

Exemptions: Where the AI system performs an assistive function for standard editing; Where the system does not substantially alter the input data supplied by the deployer, or its semantics; Where use is authorised by law for criminal-offence detection, prevention, investigation or prosecution.

Article 50(2)

50(3) — Notify people exposed to emotion or biometric categorisation

Deployer. Deployers of an emotion recognition system or a biometric categorisation system must inform the people exposed to it that it is operating, and process the personal data in line with the GDPR and the Law Enforcement Directive.

Triggers. Operating an emotion recognition or biometric categorisation system on identifiable people. Note that emotion inference in the workplace and in education is separately prohibited outright by Article 5.

- Notice to the exposed person, not only to the contracting customer
- A lawful basis and the full GDPR information duties, since biometric and emotion data are in play
- Records showing how notice was given, for market-surveillance requests

Exemptions: Systems permitted by law to detect, prevent or investigate criminal offences, subject to safeguards for third-party rights and in accordance with Union law.

Article 50(3)

50(4) — Disclose deep fakes and AI-written public-interest text

Deployer. Deployers who generate or manipulate image, audio or video content constituting a deep fake must disclose that it is artificially generated, and deployers publishing AI-generated text to inform the public on matters of public interest must disclose that too.

Triggers. Publishing a deep fake, or publishing AI-written text intended to inform the public on a matter of public interest. This duty is on the publisher, and it is separate from the provider's marking duty under 50(2).

- A disclosure a reader or viewer actually sees, alongside the content
- For artistic, creative, satirical or fictional work: disclosure limited to noting that such content exists, in a way that does not spoil the work
- An editorial policy recording who decides when the duty bites

Exemptions: Use authorised by law to detect, prevent, investigate or prosecute criminal offences; For text: where the content has been through human review or editorial control and a natural or legal person holds editorial responsibility for the publication.

Article 50(4)

Disclosure is cheap, and it is checked

HOLISTIC AI

The enterprise AI governance platform

Holistic AI governs AI across the lifecycle: an inventory of every system, the classification behind each one, the evidence a supervisory authority will ask for, and continuous monitoring once it is live.

EU AI Act readiness assessment →

holisticai.com/eu-ai-act-readiness

Get a demo →

holisticai.com/demo

Classify one system

The risk calculator walks the same sequence the Act does and names the provision behind each answer.

euaiact.com/eu-ai-act-risk-calculator

Check your exposure

The readiness check scores the whole estate against the obligations that already bind and the ones that are coming.

euaiact.com/eu-ai-act-readiness



The EU AI Act, as a working tool

euaiact.com is maintained by Holistic AI as an independent, source-verified reference: the complete legal text with the Omnibus amendments marked, a scope finder, a risk classifier, eight sector guides and a 31-step compliance checklist — all free.

YOUR NEXT STEP

Find out where you actually stand

Classify one system in a few minutes, or have Holistic AI assess the whole portfolio and produce the evidence behind it.

Classify a system →

euaiact.com/eu-ai-act-risk-calculator

Start the assessment →

holisticai.com/eu-ai-act-readiness

Or book a walkthrough of the platform: holisticai.com/demo