AI Act Transparency Checklist

Does Article 50 of the EU AI Act apply to your product — and does your product meet its transparency obligations?

FREE SELF-CHECK · 10 MINUTES · VERSION 1.2 · 29 JULY 2026

For SaaS and AI companies with users in the EU. As a rule you are in scope wherever you are registered, if you place an AI system on the EU market or its output is used in the EU. Article 50 — the tier that catches almost every AI-powered product — applies from 2 August 2026.

> **IF YOU HEARD THE AI ACT WAS DELAYED**
>
> Part of it was. In 2026 the EU deferred the heavy high-risk rules to late 2027 and 2028. **The transparency duties in this checklist were not deferred. They apply from 2 August 2026.**
>
> One exception worth knowing: if your product was already on the EU market before 2 August 2026, you have until **2 December 2026** to implement the technical marking in Section 4. Every other duty below still starts on 2 August.

Scope note: Article 50 only. High-risk classification (Annex III) is a separate, heavier regime — see Section 6. Items marked as good practice are not express requirements.

1. Provider or deployer?

Article 50 splits: **50(1) and 50(2) fall on providers, 50(3) and 50(4) on deployers.** Settle this first — over-compliance means building marking infrastructure you never owed.

- ☐ **Provider** — you place an AI system on the EU market under your own name or trademark. Building your feature on someone else's model via API usually still makes you the provider of your system.
- ☐ **Deployer** — you use an AI system under your own authority, professionally. Running a third-party chatbot or writing copy with someone else's tool is deployment.
- ☐ **Contractors and agencies** — yours where they act under your control; where the agency decides for itself whether to use AI, it becomes the deployer. Say which in the contract.

Out of scope: purely personal, non-professional use, and AI for the sole purpose of scientific R&D. The trap: open-source AI systems are not exempt.

2. Do the transparency duties apply to you?

One tick is enough to bring you in scope. The role in brackets carries the duty.

1. **Direct interaction** (provider) — chatbot, voice assistant, AI hotline or agent in your product, site or support channel. (Art. 50(1))
2. **Generating synthetic content** (provider) — your product generates or substantially edits audio, images, video or text under your own name, including via a third-party model. If you merely use someone else's tool, the duty is theirs — see items 4–5. (Art. 50(2))
3. **Emotion recognition / biometric categorisation** (deployer) — you analyse emotions, tone or facial expressions, or categorise people from biometric data. (Art. 50(3))
4. **Deepfakes** (deployer) — AI content depicting real people, places or events that could pass as authentic. (Art. 50(4))
5. **AI text on matters of public interest** (deployer) — published to inform the public (news, industry updates, blog) without human editorial responsibility. (Art. 50(4))

No ticks? Article 50 most likely does not apply — but Section 3 does, and that changes the day you ship your first AI feature. Skip Sections 4–5. The baseline below is what enterprise vendor questionnaires ask about long before a regulator does; the free call at the end is the right length for exactly that.

3. Applies to every company using AI

REGARDLESS OF ARTICLE 50

- ☐ **AI literacy** (Art. 4, since 2 February 2025): staff operating your AI systems understand how they work, their risks and their limits, and you can show how you ensure it. No dedicated fine attaches, but training records are cheap evidence of good faith.
- ☐ **AI inventory** — a register of AI systems and features in use, with the role noted against each. Good practice, but it is what makes every other assessment possible.

4. Self-check by scenario

ONLY IF ARTICLE 50 APPLIES — SEE SECTION 2

Users interact with your AI — provider, Art. 50(1)

- ☐ Users are told they are talking to AI at first contact — in the widget, in voice flows, in email agents, not just in the terms. The bot is not named or designed to pass as human.
- ☐ You are not over-relying on the "obvious from the context" exception: the test is a reasonably well-informed, observant and circumspect person in the actual context of use, not your team.

Your product generates content — provider, Art. 50(2)

- ☐ Outputs are marked as artificially generated in a machine-readable format (metadata, watermarking, provenance or equivalent) — effective, interoperable and robust as far as technically feasible, surviving copy-paste and format conversion.
- ☐ You have checked the assistive-editing exception: no marking duty where the AI performs an assistive function for standard editing or does not substantially alter the input data. Autocorrect is out, a generative rewrite is in.
- ☐ You have decided, and recorded, whether you sign the provider section of the Code of Practice, and if not, what you do instead. If the system was placed on the market before 2 August 2026, the marking implementation lands before 2 December 2026.
- ☐ API customers know the marking mechanics so they can meet their own duties downstream.

You publish AI-generated content — deployer, Art. 50(4)

- ☐ Deepfake-type content carries a clear, visible disclosure at the moment the user sees it.
- ☐ You have assessed the creative carve-out honestly: for evidently artistic, satirical or fictional work, disclosure must not hamper enjoyment of the work. A realistic ad is not satire.
- ☐ You know which published visuals and copy are AI-generated or AI-edited — including assets from agencies, freelancers and stock libraries — and have checked provenance through warranties, metadata or a direct question. Most risk sits in content you did not know was AI-made.

Emotion recognition or biometric categorisation — deployer, Art. 50(3)

- ☐ The people analysed are expressly informed the system is in use — the notice goes to them, not the audience — and the processing has a valid GDPR basis, biometric data being special-category data.

AI text on matters of public interest — deployer, Art. 50(4)

- ☐ Either a named person holds editorial responsibility after review (which lifts the duty), or the AI origin is disclosed. "Someone skimmed it" is not editorial responsibility — record it.

Documentation and suppliers

- ☐ You have documented what you check, who is responsible and since when — good practice, but "nothing" is not a workable answer to a regulator or an enterprise customer. Vendor and agency contracts require disclosure of AI use in delivered content and allocate the provider/deployer role expressly.

5. How the disclosure has to be made

ART. 50(5) — THE PART MOST TEAMS SKIP

- ☐ Provided **at the latest at the first interaction or exposure** — not on second use, not after sign-up.

- ☐ **Clear and distinguishable.** Per the Guidelines these do not qualify: a snippet in a footer, a faint label on an image, a label flashing for an instant on video, or disclosure buried in the terms.
- ☐ **Accessible to persons with disabilities** where relevant — survives a screen reader, does not rely on colour alone. An express element of Art. 50(5), almost universally missed.

Already-published content: material both generated and published before 2 August 2026 needs no retroactive labelling. For deepfakes the cut-off is the date of generation; for public-interest text, the date of publication — so text written in July and published in September is caught.

6. Red flags: you may be beyond Article 50

Transparency is the lightest tier. If either of these describes you, a heavier regime is likely in play, and assessing it is outside the scope of this checklist.

- ☐ AI influencing employment decisions (CV screening, evaluation, promotion, termination); credit, insurance or essential services; education (admissions, assessment, proctoring); or anything touching law enforcement, migration or critical infrastructure.
- ☐ Content depicting identifiable real people. Labelling it as AI satisfies Article 50 — not image rights, personality rights, defamation or GDPR, which apply independently and are often the greater exposure.

These obligations start later — 2 December 2027 for stand-alone systems, 2 August 2028 for AI in regulated products — but later is not lighter: the conformity work takes quarters, and such a system is still caught by Article 50 from 2 August 2026.

7. Dates and penalties

2 Feb 2025	Prohibited practices banned; AI literacy duty (Art. 4) applies.
2 Aug 2025	GPAI obligations, governance and the penalties chapter apply.
2 Aug 2026	**Article 50 transparency duties apply**, together with enforcement powers.
2 Dec 2026	Art. 50(2) marking for systems placed on the market before 2 Aug 2026; NCII/CSAM prohibition.
2 Dec 2027	High-risk obligations, stand-alone Annex III systems.
2 Aug 2028	High-risk obligations, AI embedded in Annex I products.

Article 50 sits in the middle tier of Article 99: up to €15 million or 3% of worldwide annual turnover, **whichever is higher**. For SMEs and start-ups, Article 99(6) inverts this to **whichever is lower** — on €2 million turnover, €60,000 rather than €15 million. A proportionality cap, not an exemption.

Next step

If you ticked any trigger in Section 2 and any gap in Sections 4–5, the fixed-price **AI Transparency Compliance Check ($700, one week)** turns this into a plan: a provider/deployer role map, a memo mapping each obligation to your product, a gap table, and ready-to-use disclosure texts your developers can implement straight away.

Book a free 30-minute intro call →

EMAIL oleh@levytskyilegal.com

WEB levytskyilegal.com

Oleh Levytskyi — technology and privacy lawyer (CIPP/E) advising SaaS, AI and digital businesses on GDPR, AI governance and commercial contracts.

Sources: Regulation (EU) 2024/1689 (AI Act) as amended by Regulation (EU) 2026/1744 (Digital Omnibus on AI), published in the Official Journal on 24 July 2026 — the 2 December 2026 marking deadline sits in the new Art. 111(4), and the high-risk regime moves to 2 December 2027 (Annex III) and 2 August 2028 (Annex I). European Commission Guidelines on the implementation of Article 50, 20 July 2026. Code of Practice on Transparency of AI-Generated Content, assessed as adequate 9 July 2026 — non-binding, but the recognised route to demonstrating compliance with Art. 50(2), (4) and (5). General information, not legal advice; current as of 29 July 2026.