Article 50 Labelling Implementation Guide

Prepared from your self-assessment | Regulation (EU) 2024/1689 (AI Act), Article 50

This document is general information prepared from a self-assessment — it is not legal advice.

1. Overview

By **2 August 2026**, everything you publish or operate that touches EU audiences needs to make its AI involvement obvious to the person on the other end. Concretely: your chatbots and voice assistants must tell people up front that they are AI (today this is only "partial or inconsistent"); your generative AI product sold under your own brand must mark its outputs in a machine-readable way (today you embed no markers at all); your realistic AI images, audio and video must carry a visible label that survives download (today you apply no visible labels); and your public-interest AI text must either be labelled or pass a documented human editorial review (today review happens only "sometimes"). All of these disclosures must be clear, distinguishable, present at first exposure, and accessible. Your pre-August-2026 archive does not need retrospective labelling — but the moment you republish or reuse it, it does. Underpinning all of it, you need a written map of where AI touches your content and who is responsible. Enforcement starts the same day, with penalties up to **EUR 15,000,000 or 3% of worldwide annual turnover**, whichever is higher. One system you brand yourself gets a short extension: see the transition note under Art. 50(2).

Your triggered requirements at a glance

ID	Article	Your role
A50-ORG-CONTENT-MAP	Derived operational req.	Deployer
A50-1-CHATBOT-DISCLOSURE	Art. 50(1)	Provider
A50-2-MACHINE-READABLE-MARKING	Art. 50(2)	Provider
A50-4-DEEPFAKE-LABEL	Art. 50(4), 1st subpara.	Deployer
A50-4-PUBLIC-INTEREST-TEXT	Art. 50(4), 2nd subpara.	Deployer
A50-5-DISCLOSURE-QUALITY	Art. 50(5)	Provider & deployer
A50-SCOPE-NON-RETROACTIVE	Commission guidance on Art. 50(4)	Deployer

2. A50-ORG-CONTENT-MAP — Derived operational requirement (Commission code of practice & guidelines)

What to implement

SAMPLE

You answered that you use third-party AI tools, brand a generative AI product as your own, publish AI text and AI-edited media, and run customer-facing chat. That combination means you are a **provider for some systems and a deployer for others simultaneously** — and each channel carries a different obligation. Before you can label anything correctly, you need an AI content map.

Build a single inventory covering every channel where AI generates or alters something you publish: website imagery, blog and news copy, social posts, product marketing, reports, chatbot and voice assistant transcripts, and your own-brand generative product. For each entry record: the tool used, whether output is text / image / audio / video, whether it can depict real people, places or events, whether it touches public-interest topics, who publishes it, and whether you are provider, deployer or both.

Decide the role on the basis of **what you actually do**, not what a vendor contract says. Where you take a third-party model and offer it to others under your own brand, you are acting as a provider and Art. 50(2) attaches to you — not only to the original developer.

Because you use third-party tools and are likely to use agencies or freelancers, write the labelling responsibilities into those relationships: who applies the label, who reviews, who has final publication sign-off.

Acceptance criteria

- ☐ AI content inventory completed across all channels (web, social, marketing, reports, chatbots, own-brand product)
- ☐ Role determination recorded per system — provider, deployer, or both — based on actual task distribution, not only contract wording
- ☐ Responsibilities for review, approval and publication defined contractually with agencies and freelancers
- ☐ Each inventory row mapped to the specific Article 50 obligation(s) that attach to it
- ☐ Named owner for keeping the inventory current

Notes

This is the audit backbone. Every other section below depends on it being accurate — you cannot label what you have not found.

---

3. A50-1-CHATBOT-DISCLOSURE — Art. 50(1)

What to implement

You operate customer-facing chatbots and voice assistants, and users are told they are talking to AI only "partially or inconsistently". Art. 50(1) requires that people interacting directly with an AI system are informed of that fact — unless it is obvious to a reasonably well-informed, observant

SAMPLE

and circumspect person given the context. "Partially or inconsistently" is the gap to close: the disclosure has to be reliable on every entry point, not just the main one.

For text chat: an on-screen notice at the top of the conversation window, e.g. *"You are chatting with an AI assistant."* Persist it in the header so it is visible after scrolling, and repeat it if a session resumes after a break.

For voice: a spoken statement in the opening turn, before the assistant takes any user input. Do not rely on a distinctive synthetic voice to make it "obvious" — if a caller could reasonably mistake it for a person, say it.

Audit every surface where the assistant appears: website widget, in-app, WhatsApp or messaging integrations, phone line, embedded flows on partner sites. Each needs its own notice. If a human agent can take over mid-conversation, mark the handover in both directions so users know which they are speaking to.

Do not rely on the terms of service, a privacy policy or an FAQ page to carry this.

Acceptance criteria

- ☐ On-screen or spoken notice at the start of every interaction (e.g. "You are chatting with an AI assistant")
- ☐ Notice appears at first interaction — not buried in terms, FAQs or a help page
- ☐ Every deployment surface covered: web widget, app, messaging channels, voice line, partner embeds
- ☐ Human-handover transitions signalled in both directions, where applicable
- ☐ Written record of where and how each notice is displayed (screenshots / call script excerpts, dated)

Exemptions

- AI systems authorised by law for criminal offence detection or prevention, with safeguards. This is unlikely to apply to customer service chat.

---

4. A50-2-MACHINE-READABLE-MARKING — Art. 50(2)

What to implement

You offer a generative AI product under your own brand and currently embed **no machine-readable markers**. That is the single largest technical gap in this assessment. Art. 50(2) requires providers of systems generating synthetic audio, image, video or text to ensure the outputs are marked in a machine-readable format and detectable as artificially generated or manipulated — so

SAMPLE

far as technically feasible. Offering a third-party model under your own brand puts this obligation on you, not only on the original developer.

Work through it output type by output type:

Output type	Typical marking approach
Image	C2PA-style provenance manifest; invisible watermark
Video	C2PA manifest; frame-level watermark
Audio	Inaudible watermark; provenance metadata
Text	Provenance metadata where the delivery format supports it

Start by asking your upstream model vendors what marking they already emit — many now support watermarking or C2PA signing. If the vendor marks at generation time, your job is to **stop stripping it**: check that your own processing pipeline (resizing, transcoding, re-encoding, CMS upload, CDN optimisation) preserves the marker end to end. Test this, because most default image pipelines destroy metadata.

Where the vendor provides nothing, add your own marking layer at the point your product returns output to the user.

This obligation is machine-readable marking — it is separate from, and does not replace, the visible label required in the next section.

Acceptance criteria

- ☐ Technical marking solution implemented for each output type your branded product generates (watermark / C2PA-style provenance / metadata)
- ☐ Marking method documented, including its robustness and known limitations
- ☐ Pipeline test confirming markers survive your processing, storage and delivery chain
- ☐ Where marking is not technically feasible for a given output type, that assessment written down with reasons
- ☐ Confirmation recorded that the obligation is being handled by you as own-brand provider, not assumed to sit with the original developer

Transition deadline

Systems **placed on the market before 2 August 2026** must comply with the marking and detection obligations by **2 December 2026**. If your branded generative product is already live, that four-month extension applies to it. New systems launched on or after 2 August 2026 must comply from launch — do not plan a launch into the gap assuming the later date covers it.

SAMPLE

5. A50-4-DEEPFAKE-LABEL — Art. 50(4), first subparagraph

What to implement

You use AI to generate or substantially edit images, audio and video, that content can depict real people, places or events in a way that appears authentic, and you apply **no visible labels today**. Art. 50(4), first subparagraph requires you to disclose that such content has been artificially generated or manipulated.

The label must be **visible and recognisable at first glance** — burned into or overlaid on the asset itself, not sitting in a caption field that disappears when the file is shared. Practical approach:

- **Images:** a persistent overlay or burned-in strip, e.g. *"AI-generated"* or *"Digitally altered with AI"*, positioned in a consistent corner, with sufficient contrast to remain legible.
- **Video:** an on-screen label in the opening seconds and, for longer or clippable content, a persistent corner bug so the disclosure survives excerpting.
- **Audio:** a spoken or clearly audible statement at the start of the clip.

The critical test your current setup would fail: **the label must persist on the asset after download**. Metadata-only marking does not satisfy this requirement — even after you complete the Art. 50(2) work above, the two obligations are cumulative. Equally, footnotes, asterisks and general site-wide notices ("some content on this site is AI-assisted") are not sufficient.

The EU has proposed label icons. Using them is **voluntary** — any clear, visible, audience-appropriate label satisfies the duty. Pick one wording and one placement rule and apply it consistently across channels; inconsistency is what gets noticed.

Build the label into the production step, not the publishing step, so the asset leaves your pipeline already labelled and stays labelled wherever it is reposted.

Acceptance criteria

- ☐ Visible label applied directly on or with the content, recognisable at first glance
- ☐ Label verified to persist on the asset after download, re-upload and social sharing
- ☐ Confirmed that metadata-only marking is not being relied on as the disclosure
- ☐ Confirmed that no footnote, asterisk or site-wide notice is being used as the sole disclosure
- ☐ House standard documented: wording, placement, contrast, minimum size, per format
- ☐ Standard issued to agencies and freelancers producing media on your behalf

Exemptions

SAMPLE

- **Evidently artistic, creative, satirical or fictional works:** disclosure is still required, but may be made in a manner that does not hamper the display or enjoyment of the work — for example in an opening or closing credit rather than an overlay across the frame. This narrows *how* you disclose, not *whether* you do.

---

6. A50-4-PUBLIC-INTEREST-TEXT — Art. 50(4), second subparagraph

What to implement

You publish AI-generated text intended to inform the public on matters of public interest, and human editorial review happens only **"sometimes"**. Art. 50(4), second subparagraph gives you a choice per piece: disclose that the text is artificially generated, **or** put it through human review or editorial control with a named person or entity holding editorial responsibility.

"Sometimes" means you currently satisfy neither route reliably. Fix it by making the choice explicit in the workflow rather than incidental:

Route A — disclose. A visible AI-generation notice on the published text, at the top or immediately adjacent to the headline/byline area, in the audience's language. This is the appropriate default for anything published without genuine editorial handling.

Route B — editorial responsibility. Only usable where the review is real. Per the Commission's draft guidelines, the exemption covers AI-assisted drafts with **genuine editorial revision**, and routine AI reports such as weather, brief financial or sports items where a plausibility check is performed. It does **not** cover near-fully AI-written content that receives only linguistic polish, unchecked automated translations, or unverified aggregation feeds. A formal skim-through does not qualify — the reviewer must genuinely check accuracy, plausibility and sources.

Operationally: add a mandatory field in your CMS forcing the publisher to select "AI-disclosed" or "editorially reviewed by [name]". No third option, no default. Where Route B is selected, the record must name the responsible natural or legal person.

Acceptance criteria

- ☐ Either a visible AI-generation disclosure on the published text, **or** the documented editorial route, chosen and recorded per item
- ☐ Documented human editorial review process defined: genuine checking of accuracy, plausibility and sources, with clearly assigned responsibility
- ☐ Editorial workflow records naming the responsible person or entity for each item published under the exemption
- ☐ Guidance issued distinguishing qualifying review (genuine revision, plausibility-checked routine reports) from non-qualifying handling (linguistic polish only, unchecked translations, unverified aggregation)

SAMPLE

- [ ] CMS or workflow gate preventing publication until one route is selected

Exemptions

- Use authorised by law to detect, prevent, investigate or prosecute criminal offences.
- Content that has undergone human review or editorial control with a natural or legal person holding editorial responsibility (Route B above).

---

7. A50-5-DISCLOSURE-QUALITY — Art. 50(5)

What to implement

Art. 50(5) sets the quality bar for every disclosure produced under the sections above: clear and distinguishable, delivered **at the latest at the time of first interaction or exposure**, and conforming to applicable accessibility requirements. It applies to your chatbot notice, your visible media labels and your public-interest text disclosures alike.

Three checks:

Clarity. Plain wording in the language of the audience — if you publish in several EU languages, the disclosure is translated too. Avoid internal jargon ("synthetic asset", "GenAI-assisted"); say AI-generated.

Timing. The disclosure is seen before or at the moment of exposure. That means: the chatbot notice before the user's first message, not after the bot's first answer; the video label in the opening seconds, not the end card; the text disclosure above the fold, not at the foot of the article.

Accessibility. Labels must be perceivable, not just present. Text overlays need adequate contrast and size; image labels need equivalent alt text or an accessible caption so screen-reader users get the disclosure; audio-only disclosures need a text equivalent; chatbot notices need to be exposed to assistive technology, not rendered as a decorative graphic.

Run this as a single review pass across all disclosure types once they exist, rather than as a per-channel afterthought.

Acceptance criteria

- [ ] Every disclosure reviewed for clarity in the audience's language, including translated variants
- [ ] Timing verified for each disclosure type: present at first interaction or exposure, not after
- [ ] Accessibility check completed — readable, perceivable, and exposed to assistive technology for users with disabilities
- [ ] Findings and fixes recorded, with a re-check scheduled whenever a channel or template changes

SAMPLE

8. A50-SCOPE-NON-RETROACTIVE — Commission guidance on Art. 50(4)

What to implement

You hold an archive of AI-generated content published before 2 August 2026. Per Commission guidance, that content does **not** need retrospective labelling — the Act applies from 2 August 2026. You do not have to go back and re-label the archive.

The risk is reuse. Republishing, substantially modifying, or reusing archived AI content in a new campaign or commercial context after that date triggers the labelling obligation afresh. A 2025 AI-generated hero image dropped into a 2027 campaign is, at that point, subject to the visible-label duty in Section 5.

So the work here is tagging, not labelling. Flag AI-generated assets in your DAM, CMS or media library with a persistent marker — a controlled tag such as `ai-generated: yes/no/unknown` — so that anyone pulling an asset for reuse sees immediately that a labelling check is required. Where provenance is genuinely unknown for older assets, tag them `unknown` and treat that as "check before reuse" rather than leaving the field blank.

Then add the check to the workflows that trigger it: campaign briefs, republication, social resharing, template and landing-page refreshes.

Acceptance criteria

- [ ] Archive policy in place flagging AI-generated assets so reuse decisions trigger a labelling check
- [ ] Labelling check added as a gate in campaign and republication workflows
- [ ] Assets of uncertain provenance tagged as unknown and routed to the same check
- [ ] Teams briefed that no retrospective labelling of unpublished-again archive content is required — the trigger is reuse

9. Suggested order of work

#	Requirement	Why here	Depends on
1	A50-1-CHATBOT-DISCLOSURE	Quick win. You already disclose partially — extending a notice to every surface is a copy-and-config change, largely doable this quarter.	—
2	A50-ORG-CONTENT-MAP	The backbone. Nothing below can be scoped until you know every channel and your role in each.	—

SAMPLE

#	Requirement	Why here	Depends on
3	A50-SCOPE-NON-RETROACTIVE	Cheap and time-sensitive: tag the archive while the inventory work is fresh, before provenance knowledge is lost.	#2
4	A50-4-PUBLIC-INTEREST-TEXT	Process change, not engineering. Converting "sometimes reviewed" into a two-route CMS gate is achievable without vendor dependencies.	#2
5	A50-4-DEEPFAKE-LABEL	Needs a house label standard plus production-pipeline changes and agency briefing. Start early; hard deadline 2 Aug 2026.	#2
6	A50-2-MACHINE-READABLE-MARKING	Longest lead time — vendor discussions, watermarking/C2PA integration, pipeline preservation testing. Begin in parallel with #5 even though it may finish last.	#2, vendor input
7	A50-5-DISCLOSURE-QUALITY	Cross-cutting review pass — run once the disclosures from #1, #4 and #5 exist, then re-run on template change.	#1, #4, #5

Dependency and deadline notes

- Items 1–5 and 7 must be complete by **2 August 2026**. Item 6 has a later deadline of **2 December 2026** for your branded generative system if it was placed on the market before 2 August 2026 — but only for that system, and only for the marking obligation. Do not let the extension delay kick-off; vendor and pipeline work is the slowest item on this list.
- Items 5 and 6 are **cumulative, not alternatives**. Completing machine-readable marking does not discharge the visible-label duty, and vice versa.
- Item 7 is not a one-off. Schedule a re-check whenever you add a channel, change a template, or add a language.

Prepared with AI assistance; reviewed by the A50Check team.

SAMPLE