KILDE

AI Act Article 50 Compliance Kit — Edition 1.2

Verification & evidence edition · current to 21 August 2026

General information and document templates only. Not legal advice; no compliance outcome is guaranteed. Licensed for single-organisation use — see the Licence at the end of this document. © 2026 Kilde.

START HERE — Which Article 50 duties apply to your product?

AI Act Article 50 Compliance Kit · v1.0 · verified against sources dated 15 July 2026 · This kit is general information and implementation templates, not legal advice; no attorney-client relationship is created. For questions about your specific situation, consult counsel — this kit will make that conversation shorter.

Article 50 contains four separate duties. Most apps owe one or two, not all four. This triage takes about ten minutes and tells you which chapters of this kit to open. Answer each question for each distinct AI feature in your product (a chatbot and an image generator are assessed separately).

First, two words that decide everything

- You are a PROVIDER of an AI system if you develop it (or have it developed) and place it on the market under your name — this includes wrapping OpenAI, Gemini, or Claude in your own app. Your app is your AI system, even though the model underneath is someone else's.
- You are a DEPLOYER if you use an AI system in the course of business (rather than selling it).
- Small companies are usually both: provider of the app you sell, deployer of the AI content you publish on your own channels. Duties 1–2 below attach to providers; duties 3–4 attach to deployers.

The four questions

Q1 — Does your product talk to people? (chat interface, voice agent, AI email/support agent, character or companion app, sales bot, AI assistant embedded in your SaaS) → YES → Duty 50(1) applies to you as provider: the person must be informed they are interacting with AI, clearly, at the latest at first interaction. → Narrow exit: you may skip disclosure only where it is already obvious to a reasonably well-informed, observant person in the circumstances. The bar is higher than you think — "our users are tech-savvy" is not the test, and the safe pattern (a one-line disclosure) costs you nothing. Open Chapter 02 (copy library) and record your reasoning in Chapter 05 (evidence log). If you want to rely on the obviousness exemption, complete the self-assessment memo in Chapter 02 §6 first.

Q2 — Does your product generate or meaningfully edit content? (images, video, audio/voice, music, or text produced for the user) → YES → Duty 50(2) applies to you as provider: outputs must carry machine-readable marking identifying them as AI-generated — metadata/watermarking, not a visible caption. Open Chapter 03 (marking guide) — including the section on why your upstream model provider's marking does not automatically cover your app, and how to test whether marks survive your own pipeline. → Exits: purely assistive features that don't substantially alter the input or its meaning (grammar fixing, denoising, standard photo touch-up) are exempt — document the call in the evidence log. Timing: systems placed on the market on or after 2 Aug 2026 must mark immediately; systems already on the market before that date have until 2 Dec 2026 [VERIFY-ON-FINAL — confirm exact transition scope when the Omnibus OJ text is checked; the tracker covers this].

Q3 — Do you operate emotion recognition or biometric categorisation on people? (mood detection in calls, engagement scoring from video, categorising people by biometrics) → YES → Duty 50(3) applies to you as deployer: the exposed people must be informed the system is operating, and GDPR obligations run in parallel. This is the rarest duty for app makers — if it's yours, treat it as the highest-sensitivity one and involve counsel;

the kit's Chapter 04 §4 covers the notice pattern only.

Q4 — Do you publish AI-generated content to an audience? (marketing images/videos with real-looking people or events, AI voiceovers of real voices, AI-written articles on matters of public interest, deepfakes of any kind) → YES → Duty 50(4) applies to you as deployer: deepfake-type content must be visibly disclosed as artificially generated; AI text informing the public on matters of public interest must be disclosed unless a human exercised substantive editorial review and someone holds editorial responsibility. Open Chapter 04 (labeling patterns + the editorial-responsibility register).

Am I even covered? (the non-EU question)

If your system is placed on the EU market or its output is used in the EU, you are covered — regardless of where your company sits. A Hong Kong, US, or Singapore app with EU users is in scope. There is no establishment threshold and no small-company exemption from the duties (there IS a small-company cap on fines — see Chapter 01 §5). If you are certain you have no EU users and no EU-circulating output, and you geo-enforce that, document the decision in the evidence log and reassess quarterly.

Your 10-minute output

Fill this in and paste it at the top of your evidence log (Chapter 05):

Feature	Provider or deployer?	Duties triggered	Exemption claimed?	Kit chapters to complete
e.g. Support chatbot	Provider	50(1)	No	02, 05
e.g. Avatar generator	Provider	50(2)	No	03, 05
e.g. Blog written w/ AI	Deployer	50(4)-text	Yes — editorial review	04, 05

Dates that matter (all now settled, in-force law — kit current to 20 August 2026)

- 2 Aug 2026 — already in force: duties 50(1), (3), (4) and 50(2)-for-new-systems apply now. Enforcement powers are live (fines up to €15M/3%, SME lower-cap); no public enforcement actions yet as of this edition — the calm before routine supervision is the cheapest time to close your gap.
- 2 Dec 2026 — the live deadline: machine-readable marking for generative systems that were on the market before 2 Aug, under Regulation (EU) 2026/1744 (in force 27 July 2026). Chapter 03 is the runway.
- Code of Practice: finalized 10 June, deemed adequate by the Commission (8 July) and AI Board (9 July) — signatories "can rely on its measures to demonstrate compliance" for 50(2)/(4)/(5). ~190 signatories; accession remains open with no deadline. Chapter 06 covers whether to sign.
- The Commission's final Guidelines on Article 50 were adopted 20 July 2026 — this kit edition reflects them throughout (the earlier draft-guidance hedges are resolved).

Chapter 01 — The Obligations Map

Plain-language reference for the four Article 50 duties: who owes what, the exemptions as actually written, the dates, and your penalty exposure. Every claim here traces to the statute or the sources in FACTS (citations available on request); items the pending final Guidelines could still shift are flagged [VERIFY-ON-FINAL] and covered by the update tracker. Not legal advice.

§1 · The four duties at a glance

Duty	You owe it if you are…	What you must do	Deadline	Kit chapter
50(1) Interaction disclosure	Provider of any AI system that interacts directly with people	Ensure the person is informed they're dealing with AI — clearly, at latest at first interaction	2 Aug 2026	02
50(2) Machine-readable marking	Provider of any system generating synthetic audio/image/video/text (GPAI included; wrappers included)	Mark outputs in machine-readable form as artificially generated — effective, interoperable, robust, reliable as far as technically feasible	2 Aug 2026 (new systems) · 2 Dec 2026 (systems on market before 2 Aug) [VERIFY-ON-FINAL]	03
50(3) Emotion/biometric notice	Deployer of emotion-recognition or biometric-categorisation systems	Inform exposed persons the system operates; GDPR applies in parallel	2 Aug 2026	04 §4
50(4) Content labeling	Deployer publishing deepfakes, or AI text informing the public on matters of public interest	Disclose artificial generation visibly; text duty waived under substantive human editorial review + named editorial responsibility	2 Aug 2026	04

Also universal: 50(5) — whatever you disclose must be clear, distinguishable, delivered at the latest at first interaction or exposure, and meet accessibility requirements. "It's in our terms of service" fails this test on its face.

§2 · The exemptions, precisely (where teams get overconfident)

- Obviousness (50(1)): measured against a "reasonably well-informed, observant and circumspect" average person in context. The final Guidelines (20 July 2026) keep the two-step test and read the exemption restrictively: mere public awareness that AI exists is not enough, and the threshold rises where your