

Article 50 disclosure cheat sheet

What the transparency article asks of you, by system type. Already enforceable: the June 2026 Omnibus did not move it.

WHAT ARTICLE 50 REQUIRES


Article	System type	What you must do	Who must act	Applies
50(1)	Chatbots and assistants	Tell people they are interacting with an AI system, before or at the first interaction. Unless that is obvious to a reasonably well-informed person. Read the exception narrowly.	Provider	Since 2 Aug 2026
50(4)	Deepfakes and other generated images, audio, video	Disclose that the content is artificially generated or manipulated, where it is seen.	Deployer	Since 2 Aug 2026
50(4)	AI-generated text published to inform the public	Disclose that the text is AI-generated. Not required if it was human-reviewed and a person holds editorial responsibility.	Deployer	Since 2 Aug 2026
50(2)	Providers of generative systems	Mark generated audio, image, video or text in a machine-readable format, so it can be detected as AI-generated. Content credentials, a watermark, or both. A visible badge alone is not machine-readable.	Provider	From 2 Dec 2026

WHO IS RESPONSIBLE

Provider	Ships the system under its own name, even when the model underneath is someone else's.
Deployer	Uses the system or publishes the content, whatever tool produced it.

Model output passed through unchanged? The 50(2) marking generally sits with the model provider. If you generate, assemble or re-encode media yourself, check rather than assume.

DATES



Breaches fall under Article 99(4): up to EUR 15,000,000 or 3% of worldwide annual turnover, whichever is higher. For SMEs and start-ups, whichever is lower.

THREE DISCLOSURES YOU CAN ADAPT

Chat greeting 50(1) You are chatting with an AI assistant from [Product], not a person. Its answers can be wrong, so check anything important. In the first message, before the user types.	Deepfake label 50(4) This video was generated with AI. [Name] did not record it. On the video itself, where it is seen.	Public-interest text 50(4) This article was generated with AI and published without editorial review. With the text, where readers see it.
---	---	--

CHECKLIST BEFORE YOU SHIP

☐ Visible before or at the first interaction Not in a tooltip, not only in the privacy policy.	☐ Says "AI" in plain words Is it a person or not? Vague wording does not answer that.
☐ Works on mobile Check the widget at 360 pixels wide.	☐ Accessible Real text, not an image of text, with enough contrast.
☐ In the user's language The language of the interface the person is using.	☐ Logged and documented A dated note per feature: wording, screenshot, exemptions.

AI vendor questionnaire template

The EU AI Act questions enterprise buyers typically send an AI vendor. Answer per feature, name the Article, date every answer.

Vendor	Product or feature	Completed by	Date
01	Role. Are you the provider or a deployer of this system, feature by feature? Say what would change it: the buyer's own name on it, a substantial modification, a new purpose. Art. 3(3), 3(4), 25		
02	Risk tier, with reasoning. Which risk tier does each feature fall into, and why? One line per feature in a dated classification memo, including the Article 5 check. Art. 5, 6, Annex III		
03	Article 6(3) position. A feature sits in an Annex III area but is not high-risk: where is the assessment? Article 6(4) requires it to be documented. Otherwise: not applicable, and why. Art. 6(3), 6(4)		
04	Article 50 transparency. Where are users told they are interacting with AI, and in what words? Who marks generated content? Quote the wording; attach a screenshot of the first interaction. Art. 50(1), 50(2)		
05	Human oversight. How do people review, correct or override the system's outputs? If not high-risk, say so first, then describe the controls you run anyway. Art. 14		
06	Data governance. Do you train models on our data? Does the model provider retain inputs? A plain yes or no, which data, under which contract term. Point to the DPA.		
07	Logging. Which logs does the system keep, and how can we access them? For high-risk systems, logging is a requirement in its own right. Art. 12		
08	Incident handling. How will you tell us about incidents and about changes to a feature's purpose or classification? For high-risk systems, serious incident reporting applies on top. Art. 73		
09	Model provider. Which general-purpose AI model does each feature use, and who provides it? Link to the model provider's published documentation. Art. 53		
10	Documentation available. Which of these can you share? ☐ Feature inventory and role statement ☐ Dated classification memo ☐ Article 6(3) assessment, if relied on ☐ Disclosure screenshot, marking note ☐ AI literacy record (Article 4) ☐ Controls description or instructions for use ☐ Sub-processor list and DPA ☐ Model provider documentation		

Answer honestly. There is no AI Act certificate for limited-risk or minimal-risk systems. A self-assessment is not a certification.