Sample verdict. The company is fictional, so no client's details are shown. The legal text is real and quoted word for word. A real verdict is written on the answers you give to seven questions.

Which Article 50 duties apply to Fernwick Media?

Illustrative subject: a digital news publisher with readers across the EU. It runs a reader assistant it built in-house, publishes short AI-written news briefs, and illustrates some stories with AI-generated images. All three went live before 2 August 2026.

// THE ANSWER, BEFORE ANYTHING ELSE

VERDICT · PAGE ONE, AS PROMISED

Four duties apply — *and each one is narrower than it first looks.*

You carry the chatbot notice and the machine-readable marking for the assistant you built, because you are its provider. You carry the labels on photorealistic scenes and on briefs that go out unreviewed, because you publish them. The marking of the images and briefs themselves sits with the companies that make those tools, and the briefs your duty editor approves need no label at all.

// THE FACTS IT RESTS ON, AS YOU STATED THEM

QUESTION PUT TO YOU	ANSWER ON RECORD
Reader assistant	"Ask Fernwick Media", built in-house on a third-party model via API, live since April 2026 under your name. Answers reader questions and summarises articles for the reader who asks.
Is it labelled as AI?	Only in the Terms of Use. The chat window says nothing.
News briefs	Short summaries of your own articles, including local council decisions, written by a subscribed third-party tool. Overnight briefs publish automatically; daytime briefs are approved by the duty editor.
Images	From a commercial image generator: photorealistic scenes of real local places, and clearly drawn editorial cartoons.
Editorial responsibility	The publishing company; the editor-in-chief is named on the About page.

If any line above is wrong, say so and the verdict is rewritten at no charge.

Sample verdict.

The map: every duty, who carries it, from when

DUTY	WHO CARRIES IT	ON THESE FACTS	FROM
50(1) · tell readers they are talking to an AI	You, as provider of the assistant	applies · gap	2 August 2026
50(2) · machine-readable marking of the assistant's output	You, as provider of the assistant	applies	2 December 2026
50(2) · marking of the images and the briefs	The providers of those tools	not yours	—
50(4) · label on photorealistic scenes	You, as deployer	applies	2 August 2026
50(4) · label on editorial cartoons	—	not a deep fake	—
50(4) · label on overnight briefs	You, as deployer	applies	2 August 2026
50(4) · label on editor-approved briefs	You, as deployer	exception met	—

Article 50(1): the reader assistant

> Providers shall ensure that AI systems intended to interact directly with natural persons are designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system, unless this is obvious from the point of view of a natural person who is reasonably well-informed, observant and circumspect, taking into account the circumstances and the context of use.
>
> Regulation (EU) 2024/1689, Article 50(1), first sentence

Why you are the provider. The Commission's guidelines give this exact case as an example of a provider: an organisation "that has developed an interactive AI system (e.g. chatbot) in-house and puts it into service in the Union for its own use and under its name or trademark." Building on a third-party model does not move the duty to the model's vendor.

Why the Terms of Use are not enough. The same guidelines list "Disclosures contained only in terms and conditions, URLs, or documentation" among the techniques that are insufficient on their own: they "may complement, though not replace, in-context disclosure." Article 50(5) sets the timing: "at the latest at the time of the first interaction or exposure."

What closes the gap. One clear notice in the chat window, before the first answer. The guidelines say a single, prominent notification "is likely to suffice in most instances".

Sources: Regulation (EU) 2024/1689, Art. 50(1), 50(5) · Commission Guidelines on Article 50, C(2026) 5054 final, 20.7.2026, paras. (38), (40) and the provider examples in section 2.3.

Sample verdict.

Article 50(2): marking what the assistant writes

> Providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video or text content, shall ensure that the outputs of the AI system are marked in a machine-readable format and detectable as artificially generated or manipulated.
>
> Regulation (EU) 2024/1689, Article 50(2), first sentence

It reaches the assistant. The assistant generates text, and the guidelines confirm the paragraph covers text as well as audio, image and video. Its summaries do not escape through the editing exception: the guidelines list "AI-generated summaries of text" among the "semantic changes that require marking under Article 50(2) AI Act".

The date is still ahead of you. The assistant was on the market before 2 August 2026, so the transitional rule applies:

> Providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video or text content, that have been placed on the market before 2 August 2026 shall take the necessary steps in order to comply with Article 50(2) by 2 December 2026.
>
> Regulation (EU) 2026/1744, amending Article 111 of Regulation (EU) 2024/1689

Where this paragraph stops

The images and the briefs come from tools other companies built and sell. You use them as they are; you did not put them on the market under your name. Article 50(2) speaks to *providers*, so the machine-readable marking of those outputs is the job of the companies that make the tools. It is worth asking each vendor, in writing, how they mark what they generate. It is not a duty this paragraph places on you.

Article 50(4): what you publish

> Deployers of an AI system that generates or manipulates image, audio or video content constituting a deep fake, shall disclose that the content has been artificially generated or manipulated.
>
> Regulation (EU) 2024/1689, Article 50(4), first subparagraph, first sentence

> 'deep fake' means AI-generated or manipulated image, audio or video content that resembles existing persons, objects, places, entities or events and would falsely appear to a person to be authentic or truthful;
>
> Regulation (EU) 2024/1689, Article 3(60)

The photorealistic scenes are deep fakes. They show real local places, rendered to look like photographs, next to news stories. That is the definition, line by line. Each one needs a visible disclosure.

The editorial cartoons are not. Among its examples of content that does not constitute a deep fake, the guidelines list an "AI-generated cartoon of a pre-existing image depicting a historical event." A drawing that reads as a drawing does not falsely appear authentic.

Sample verdict.

Article 50(4): the news briefs

> Deployers of an AI system that generates or manipulates text which is published with the purpose of informing the public on matters of public interest shall disclose that the text has been artificially generated or manipulated. This obligation shall not apply where the use is authorised by law to detect, prevent, investigate or prosecute criminal offences or where the AI-generated content has undergone a process of human review or editorial control and where a natural or legal person holds editorial responsibility for the publication of the content.
>
> Regulation (EU) 2024/1689, Article 50(4), second subparagraph

Your briefs are squarely in scope. The guidelines give, as an example of text that must be labelled unless reviewed, an "AI-generated summary of a human-authored article on a newspaper's website discussing a recent decision by a town council." That is your brief on a council decision, almost word for word.

The exception has two conditions, and both must hold

CONDITION	OVERNIGHT	DAYTIME	WHY
i · Human review or editorial control	not met	met	Overnight briefs publish with no one reading them. Daytime briefs are approved by the duty editor, who can change or reject them.
ii · Someone holds editorial responsibility	met	met	The publishing company holds it, and the editor-in-chief is named on the About page.

The guidelines set the bar for the first condition: "Fact-checking the accuracy of the content is a minimum requirement that should be part of that review." A glance before publishing is not the same thing. If the duty editor checks the facts, the daytime briefs need no label; the overnight ones do, until someone reviews them.

The assistant's summaries sit outside this paragraph. The guidelines place a "News summary generated by a chatbot that is only available to the user that prompted the chatbot" outside Article 50(4). They are still covered by your marking duty under Article 50(2).

What follows, in order

- **Now** (these duties have applied since 2 August 2026): a notice in the chat window before the first answer; a visible label on every photorealistic scene; a label on overnight briefs, or route them through the duty editor before they publish.
- **By 2 December 2026**: machine-readable marking on everything the assistant generates.
- **Alongside**: keep the editor-in-chief named where readers can find it. The guidelines say the identity of whoever holds editorial responsibility "should be made publicly available on an easily findable location".

Sources: Regulation (EU) 2024/1689, Art. 50(4), 113 · Commission Guidelines C(2026) 5054 final, sections 6.2.1 and 6.2.3, paras. (133), (134), (138).

Sample verdict.

What you are not required to do

This is the section nobody writes, and the reason a verdict is worth paying for.

- You are not required to add machine-readable marking to the images or the briefs yourself. Article 50(2) addresses the providers of the tools that generate them.
- You are not required to label the editorial cartoons as deep fakes. A drawing that reads as a drawing does not meet the definition in Article 3(60).
- You are not required to label briefs your duty editor has fact-checked and approved, as long as editorial responsibility stays with a named person or company.
- You are not required to label the assistant's summaries under Article 50(4). They reach only the reader who asked.
- You are not required to repeat the chatbot notice on every message. A single prominent notice before the first interaction is likely to suffice in most cases; riskier contexts, such as vulnerable users, call for more.
- You are not required to act on Article 50(3). Nothing in your products recognises emotions or sorts people by biometric traits.

One paragraph to forward

> Our reader assistant falls under Article 50(1) and 50(2) of the EU AI Act: it needs an in-chat AI notice before the first answer now, and machine-readable marking of its output by 2 December 2026. Under Article 50(4), our photorealistic AI images need a visible label, and our AI-written briefs need a label unless an editor fact-checks and approves them before publication. Our editorial cartoons and editor-approved briefs need no label, and the marking of the images and briefs themselves is the duty of the tools' providers.
>
> paste as it stands · no rewriting needed

Your own verdict

One product, seven questions by email, no calls. Three working days after your answers, you get a written verdict on which Article 50 duties apply and what follows if they do. $200, through Upwork escrow. Questions first: andrei@logium.co.

General information on how the Regulation is written and how its conditions map onto stated facts. Not legal advice on your position, and not a substitute for a licensed attorney in your jurisdiction. Legal sources are reconfirmed against the primary text on the day each verdict is written.