

The AI Act in your company

What to check, in what order, and which obligations fall on your side. For companies operating in the EU.

How to use this guide. Three rules that everything else follows from:

1. **Assess the use case, not the product name.** The same product, say Microsoft's Copilot, used to edit text and used to support recruitment, gives you two completely different cases.
2. **Your role matters.** A provider has different obligations from a company that deploys an off-the-shelf tool. The same company can be a provider for one system and a deployer for another.
3. **The categories can overlap.** Prohibited practices, high-risk classification and transparency are separate tests, not rungs on a single ladder.

In brief: six things you need to know

- 1 **The AI Act may apply to your company whenever it uses an AI tool**, not only when it develops one. Your obligations depend on the **specific use case** and **your role**. Providers and deployers have different duties.
- 2 **The transparency obligations in Article 50 have applied since 2 August 2026.** They do not all apply to everyone. Article 50 has four paragraphs, each addressing a different role. Section 4 explains how to identify the obligations relevant to you.
- 3 **Prohibited AI practices have been subject to the AI Act since February 2025.** One prohibition concerns emotion recognition in the workplace and educational institutions, including systems used in relation to employees, candidates and students. Emotion analysis concerning customers is not generally covered by that prohibition, but it may trigger a duty to inform them and must comply with the GDPR. The maximum fine for a prohibited practice is EUR 35 million or 7% of worldwide annual turnover.
- 4 **The application dates for high-risk AI systems have been postponed.** The new date for stand-alone systems, including many recruitment, employee assessment, credit scoring and education use cases, is 2 December 2027.
- 5 **Since February 2025, providers and deployers have had to take measures to ensure an appropriate level of AI literacy** among people who work with AI systems on their behalf (Article 4). The measures must be proportionate. There are several ways to meet the requirement, and the simplest is training that ends with a record of attendance. An electronic one is perfectly sufficient.
- 6 **The AI Act does not require your company to obtain a certificate.** What matters is that you know which AI use cases you have and can document them properly. A register of AI use cases is the foundation of a quiet life here.

1. Does the AI Act really apply to my company?

Let us check this the simple way. Ask yourself the five questions below:

- Does your website or helpline use a chatbot or voicebot?
- Do you publish AI-generated text, images, video or audio, such as product descriptions, posts, adverts, voice-overs, dubbing or avatars?
- Does anyone in the company use ChatGPT, Copilot, Gemini, Claude or a similar tool for work?
- Do you use tools that support recruitment, employee assessment, customer scoring or fraud detection?
- Do you sell a product or service that contains an AI component?

To fall within the scope of the AI Act, **a single "yes" is enough**. Bear in mind that merely owning a tool bought years ago that has since gained AI features, Microsoft 365 for example, does not by itself mean you are in one of the cases above.

This guide is intended primarily for companies established in the EU or operating on the EU market. The AI Act can also apply to a provider or deployer based outside the EU when it places an AI system on the EU market or the output produced by the system is used in the Union.

2. Timeline

Date	What happens	Who is affected
2 February 2025	Prohibited AI practices and the duty to support AI literacy	Providers and deployers, depending on the obligation
2 August 2025	Obligations for providers of general-purpose AI models (GPAI)	Model providers
2 August 2026	Article 50 transparency obligations: disclosure of AI interaction and rules for AI-generated content	Depends on the role and use case
2 December 2026	End of the transition period for machine-readable marking by providers of older systems; new prohibited practices concerning non-consensual intimate content and child sexual abuse material	Providers and organisations using affected generative AI systems
2 December 2027	Stand-alone high-risk AI systems, including recruitment, HR, credit scoring, insurance and education use cases	Selected sectors and use cases
2 August 2028	High-risk AI systems embedded in regulated products, including machinery, medical devices and toys	Product manufacturers and other relevant operators

The application dates for high-risk AI systems were postponed by the Digital Omnibus, adopted by the Council of the EU in June 2026. The transparency obligations took effect on 2 August 2026 as planned, and they affect a much broader group of organisations.

The most sensible working assumption is to treat the AI Act as though it already applied in full. Adapting on that basis is what produces results that last.

3. Four regulatory categories to check

These are not four levels on a single ladder. Test separately for a prohibited practice, high-risk classification and a transparency obligation. The categories can overlap, and the same tool may be classified differently when used for different purposes.

Prohibited AI practices

This is where it stops being a grey area, because the maximum fine is EUR 35 million or 7% of worldwide annual turnover. Use cases that require particular care include:

- emotion recognition involving an employee, candidate or student; the prohibition concerns the workplace and educational institutions, subject to limited medical or safety exceptions;
- social scoring that assigns a broad "trust score" and uses it to impose unrelated or disproportionate consequences;
- systems that manipulate a person's decision by exploiting vulnerabilities and cause, or are reasonably likely to cause, significant harm.

For emotion recognition, the classification depends on whose emotions are being inferred, the data used for the inference and the context. Sentiment analysis of a transcript is not necessarily the same as an emotion recognition system using biometric data.

Emotion analysis concerning customers, for example in a call centre, does not generally fall within the workplace prohibition. It may still trigger the duty to inform the affected people under Article 50(3) and must comply with the GDPR. Industry materials often blur this distinction.

High-risk AI systems

These are, perhaps surprisingly, permitted, but subject to conditions. There has to be documentation, human oversight, demonstrable care about data quality, and logs on top of that.

Examples include automated CV screening, employee assessment and promotion systems, credit scoring, risk assessment in life and health insurance, and systems used to evaluate learning outcomes.

Transparency obligations

By the time you are reading this guide, the transparency obligation already applies to you. It is time to make sure you are marking content properly, and that does not mean only the tags visible to bots.

Examples include chatbots, voicebots, generated product descriptions and articles, synthetic voice-overs, AI dubbing, avatars and generated images used in advertising.

Minimal-risk uses

No additional AI Act obligations arise here, beyond the two cross-cutting duties: supporting AI literacy and complying with data protection rules.

Examples include spam filters, autocomplete, product recommendations and machine translation used internally.

4. What you need to have in place

Identify your role first. Article 50 has four paragraphs. Paragraphs 1 and 2 primarily address the **provider** of the system; paragraphs 3 and 4 address the **deployer**, meaning the organisation using it. If you buy an off-the-shelf tool from an external provider, some of the duties below rest with that provider rather than with you.

4.1. A chatbot must disclose that it is AI

Article 50(1) - obligation of the provider

An AI system intended to interact directly with people must clearly inform them that it is AI, and it must do so at the first interaction. The disclosure has to be clear and accessible, so not in the small print of your terms and conditions. There are exceptions where this is obvious anyway, but it is always safer to stick to the rule above.

Good: "Hello. I am Localize.pl's AI assistant and I reply automatically."

Poor: "Hi, I'm Anna, how can I help?"

If you offer a chatbot under your own name, including one built on another company's API, then at that point you are most likely its provider and this obligation belongs to you.

4.2. Machine-readable marking of generated content

Article 50(2) - obligation of the provider of the generative AI system

The duty to mark synthetic content in a machine-readable format, for example through provenance metadata, watermarking or compatible technical solutions such as C2PA, rests with the provider of the generative AI system, subject to the exceptions in Article 50(2).

What follows from this for you as a user of the tool? As a matter of due diligence you need to check whether the tool has any marking features at all, what its licence terms say, and whether your editing workflow strips the metadata. Look above all at the last step of the process and at what actually reaches the audience.

Treat the above as a **sensible practice and an element of risk management**, because it does not follow directly from Article 50(2). It is worth doing regardless.

For certainty, check the AI Act timeline for systems that were on the market before the regulation applied, because in some cases you may still have time.

4.3. Deep fakes

Article 50(4) - obligation of the deployer

A disclosure is required when material meets the **AI Act definition of a deep fake**. The devil is in the detail here, because the definition only says that a deep fake falsely appears to be an **authentic representation of an existing person, object, place, entity or event**. To put it plainly: a pink flying tomato the size of a city could never exist, but a dog playing the piano with its paws could, even if it is inhumanly difficult for a dog.

Synthetic voice-overs, avatars and AI dubbing are assessed against that same definition. Not every synthetic asset is a deep fake. A generated voice reading your own product copy does not usually pretend to be an authentic recording of a particular person, and there is then no basis for treating the material as a deep fake.

Where the test comes out negative, **voluntary disclosure is good communication practice**. It builds trust and removes the risk of being accused of misleading an audience. It is not an obligation, of course, but it is worth adopting.

4.4. Public-interest text

If you publish text that informs the public on matters of public interest and it is AI-generated, the position is clear: you have to disclose the fact. And straight away an **exception**: where the material has undergone human editorial review and a person takes responsibility for it, the obligation falls away. You do, however, need that process clearly described and documented.

4.5. AI literacy

Article 4 — applicable since 2 February 2025

Providers and deployers must take measures to ensure, to their best extent, an appropriate level of AI literacy among staff and other people dealing with the operation and use of AI systems on their behalf.

Here too, the key is being able to show evidence that those measures were taken. It can be training with a record of attendance, role-specific instructions, learning materials, onboarding, exercises, consultations and so on. At a minimum, use a test that leaves proof that competence was checked. **Always match the scope to roles, existing knowledge, context of use and risk.**

Does the regulation say what the training has to look like? No. There is no prescribed training format, no annual cycle and no compulsory attendance sheet. It is still worth defining **your own standard** internally as a matter of good practice, even though the absence of one is not a breach of the AI Act.

5. A register of AI systems

The regulation does not expressly order every company to "keep a register", but **without one it is far harder to demonstrate that the remaining obligations have been met**. If you intend to do only the bare minimum about the AI Act, this is the document to concentrate on.

Minimum set of columns:

Field	Why it matters
System or tool name	Identification
Provider and version	Accountability and change tracking
Purpose and use case	Basis for classification
Department and business owner	Responsibility
Our role: provider or deployer	Applicable obligations
Classification and rationale	Evidence of a reasoned decision
Personal data involved	Link to GDPR obligations

Field	Why it matters
Interaction with people or generation of content	Article 50 screening
Human oversight: owner and process	Articles 14 and 26 where applicable
Assessment date and next review date	Currency of the assessment

You also get a ready-made template as part of a Localize.pl audit, but you are perfectly free to build your own in a spreadsheet. It only has to be current and have an owner.

6. Seven risk areas worth checking

This is a list of **areas to check**, derived from the structure of the rules, and the order is arbitrary.

- 1 **"We do not use AI."** Look more closely and it turns out marketing generates product descriptions, HR filters CVs in a recruitment system, and support gets AI suggestions in the help desk. The absence of a register is what lets "Shadow AI" take hold.
- 2 **A chatbot that never says it is a bot.** Usually the simplest thing to fix, one message in the interface is enough, and the easiest for anyone outside the company to spot. Before you start looking for a solution, check whether you really are the provider of that chatbot.
- 3 **"Conversation quality analytics."** Check whether the system **infers emotions from biometric data** and **whose emotions** these are. Analysis involving employees may be prohibited, whereas analysis involving customers is only subject to the disclosure duty under Article 50(3) and the GDPR. These are two different situations with entirely different consequences.
- 4 **Customer data pasted into public models.** The main exposure here is a breach of data protection rules and of trade secrets. The answer is an internal company policy.
- 5 **"The supplier is responsible."** True, but if you develop a system, including on another company's API, and offer it **under your own name**, from that point you may be its **provider** under Article 3(3). Article 25 is a separate issue: it concerns **the transfer of provider obligations for high-risk AI systems**, among other things after rebranding, a substantial modification or a change of intended purpose. So these are two different bases. Establish first what system you offer and whether it really is high-risk.
- 6 **No record whatsoever of AI literacy measures.** Check which documents or registers you can produce to confirm **what was done and for whom**.
- 7 **Employee assessment without informing the workforce.** Independently of the GDPR, for high-risk AI systems Article 26(7) creates a separate information duty towards workers and their representatives.

7. A practical 30-day starter plan

Week	Action	Outcome
1	List every AI tool used in the company. Ask each department and review SaaS invoices.	Register v1
2	For each use case, identify your role and run three separate tests: prohibited practice, high-risk classification and Article 50.	Risk map
3	Address the Article 50 obligations that apply to your role: chatbot disclosure, provider marking checks and disclosures for material that meets the definition of a deep fake.	Role-based transparency measures
4	Appoint an owner, adopt a simple AI policy, plan proportionate AI literacy measures and record what was completed.	Article 4 evidence and basic governance

This is a genuinely realistic scope for a company of up to a few hundred people. To a large degree it **takes the risk out of the areas that are most visible from the outside**.

8. When to seek external support

If you use only a few tools, none of them concerns HR, finance or health, and you do not build your own AI products, you will almost certainly manage this in-house.

Otherwise an audit is worth considering. It makes particular sense when:

- you use AI in recruitment, employee assessment, customer scoring or insurance;
 - you offer an AI tool to customers under your own brand, which makes you a provider rather than a user;
 - you have corporate customers who ask their business partners about AI Act compliance;
 - you work with personal data at scale;
 - nobody in the company has the time, the appetite or the confidence to run the process.
-

9. Glossary

- **Provider** — a person or organisation that develops an AI system or a general-purpose AI model, or has one developed, and places it on the market or puts the system into service under its own name or trade mark.
 - **Deployer** — a person or organisation using an AI system under its authority in a professional or organisational context. Most companies using off-the-shelf AI tools act as deployers.
 - **GPAI** — a general-purpose AI model, such as GPT, Claude, Gemini or Llama.
 - **Deep fake** — AI-generated or manipulated image, audio or video content that resembles an existing person, object, place, entity or event and falsely appears authentic or truthful.
 - **FRIA** — a fundamental rights impact assessment required of certain deployers of high-risk AI systems.
 - **C2PA** — a technical standard for recording and verifying the provenance of digital content.
-

Prepared by Localize.pl. Legal position as of **29 July 2026**, following the adoption of the Digital Omnibus. This document is for information only and does not constitute legal advice.

The AI Act is an EU regulation and applies across the Union. Supervision rests with the national authority designated by each Member State, and the obligations under the regulation apply regardless of when that authority begins operating.

Key official sources: Regulation (EU) 2024/1689 (AI Act); Council of the EU AI timeline.