

Which EU AI Act tier is your AI system in? The five questions that decide it.

Most of the EU AI Act only bites if your system is high-risk, and most systems are not. Before you budget for the full conformity regime, work out which of the four tiers you are actually in. Here are the five questions that settle it and the deadline each tier carries.

Published 2026-06-30 · A Mue guide

Most EU AI Act coverage is written about the high-risk regime: the risk management, the technical file, the conformity assessment, the post-market monitoring. It reads as if every AI system has to do all of it. Most do not. The Act is risk-based, and the heavy regime only attaches to systems in a specific list. Before you budget for any of it, the first job is to find out which of the four tiers your system is actually in.

The four tiers, in plain terms

- Prohibited (Article 5): a short list of banned uses such as social scoring, untargeted facial-image scraping, and emotion recognition at work. No safeguards make them lawful.
- High-risk (Article 6 and Annex III): systems used in areas like biometrics, critical infrastructure, education, employment, access to essential services and credit, law enforcement, migration, and justice. These carry the full regime.
- Limited-risk (Article 50): systems that interact with people or generate content. The duty is transparency: tell people they are dealing with AI, and mark synthetic media.
- Minimal-risk: everything else. No system-specific obligations, though the AI literacy duty in Article 4 applies to every operator regardless of tier.

The five questions that settle it

You do not need to read all 113 articles to place a system. Five questions get you there in order, because the tiers are a cascade: you stop at the first one that fits.

- One: does it do anything on the Article 5 banned list? If yes, it is prohibited and no amount of paperwork makes it lawful. This is rare but worth ruling out first.
- Two: is it a safety component of a product already regulated under EU product law (machinery, medical devices, toys, and the rest of Annex I)? If yes, it is high-risk on that basis.
- Three: is its intended use in one of the Annex III areas (biometrics, critical infrastructure, education, employment and worker management, essential private and public services including credit and insurance, law enforcement, migration, or administration of justice)? If yes, it is high-risk unless a narrow Article 6(3) exemption applies.
- Four: does it interact directly with people, or generate or manipulate audio, image, video, or text? If yes, and it did not land higher, it is limited-risk and owes Article 50 transparency.
- Five: do you train or substantially modify a general-purpose model? If yes, you carry the

GPAI duties in Articles 53 to 55 as a separate track, on top of whatever tier your applications fall into.

If a system clears questions one through four without a match, it is minimal-risk: no system-specific duties, just the literacy baseline. Most internal tools and a good share of customer-facing ones land here or at limited-risk, not high-risk. That is the part the coverage tends to bury.

Why the tier is the whole budget question

The gap between tiers is not gradual. A limited-risk system owes a transparency notice you can write in an afternoon. A high-risk system owes a risk-management system, data-governance controls, a technical file, logging, human oversight, a conformity assessment before it goes live, and post-market monitoring after. Mistaking a limited-risk system for high-risk wastes a year of effort; mistaking a high-risk one for limited-risk leaves you exposed when someone asks for the technical file you never built. Getting the tier right is what makes the rest of the planning honest.

The deadline rides on the tier too

- 2 February 2025: the prohibited-practice ban and the AI literacy duty already apply.
- 2 August 2025: the general-purpose AI model duties, governance, and penalties apply.
- 2 August 2026: the Article 50 transparency duties apply. The May 2026 digital omnibus did not move this date.
 - 2 December 2027: the Annex III high-risk duties apply, pushed back by the omnibus from the original 2 August 2026.
 - 2 August 2028: the extended deadline for high-risk systems embedded in products under existing EU product law (Annex I), moved by the omnibus from 2 August 2027.

So a limited-risk chatbot has a nearer deadline (August 2026) than a high-risk hiring tool (December 2027), which is the opposite of how the timeline read a year ago. The high-risk dates above reflect the May 2026 omnibus, a political agreement that still needs formal adoption and can move again through Commission guidance. For the detail, see our note on the postponed high-risk deadline.

A starting point, not legal advice

The classifier mirrors the structure of the Act so you can see where you likely sit and what work it implies, but the precise scope of Annex III, the Article 6(3) exemptions, and how sector rules interact are matters for qualified counsel. The value of placing the system early is that it tells you whether you are facing an afternoon of transparency work or a year of conformity work, and you would rather know that before the deadline than after a complaint.

Frequently asked questions

What are the four EU AI Act risk tiers?

Prohibited, high-risk, limited-risk, and minimal. Prohibited practices under Article 5 are banned outright. High-risk systems under Article 6 and Annex III carry the full regime: risk management, data governance, technical documentation, human oversight, conformity

assessment, and post-market monitoring. Limited-risk systems under Article 50 mainly owe transparency. Minimal-risk systems carry no system-specific obligations, though the AI literacy duty in Article 4 still applies to everyone.

Is my AI system high-risk under the EU AI Act?

Probably not, unless it is used in one of the Annex III areas (biometrics, critical infrastructure, education, employment and worker management, access to essential services and credit, law enforcement, migration, or administration of justice) or it is a safety component of a product already regulated under EU product law. A support chatbot, an internal summariser, or a marketing copy generator is usually limited-risk or minimal-risk. The classifier walks five questions to place yours and lists the exact articles that apply.

When do the EU AI Act obligations start applying?

It depends on the tier. The prohibited-practice ban and the AI literacy duty have applied since 2 February 2025. General-purpose AI model duties apply from 2 August 2025. Article 50 transparency duties apply from 2 August 2026, a date the May 2026 digital omnibus did not move. Most Annex III high-risk duties apply from 2 December 2027, and product-embedded high-risk systems from 2 August 2028, both pushed back by that omnibus from their original 2026 and 2027 dates.

Does the EU AI Act treat general-purpose AI separately?

Yes. If you train or substantially modify a general-purpose (foundation) model, you carry the GPAI provider duties in Articles 53 to 55 on top of whatever tier your downstream applications fall into. That is a separate track, which is why the classifier asks about it directly rather than folding it into the four tiers.

Classify your system in five questions

Free, no signup: agent.mue.app/tools/eu-ai-act-risk-classifier

agent.mue.app/articles/which-eu-ai-act-tier-is-your-ai-system-in

