Quether

GUIDE

The EU AI Act, explained for leadership

What it requires today, what was postponed and what a company that only uses AI has to do

FORMAT

PDF guide, with the updated timeline and the compliance checklist

WHO IT'S FOR

Managing directors, finance directors and compliance leads at companies that use AI tools without developing them. If your company develops or markets AI systems, this is the starting point, not the destination.

WHAT IT TAKES

A fifteen-minute read. The checklist, half a morning's work.

LAST REVIEWED

26 August 2026

https://quether.com/en/resources/eu-ai-act

What you take away

A reading of Regulation (EU) 2024/1689 and the digital omnibus of July 2026 in the language of a leadership team: which category what you use falls into, which obligations already apply, which have moved to 2027 and 2028, and the minimum required of a mid-sized company.

+ The real timeline after the omnibus: what applies today and what moved to 2027 and 2028
+ The difference between provider and deployer, which decides almost everything else
+ How to place what you're already using in one of the four risk categories
+ The minimum required of a company that only uses AI, in eleven points
+ Where Spanish law stands today and what's missing for the penalty regime to be settled

> **BEFORE YOU GO ON**
>
> This is a management reading, not legal advice. Quether isn't a law firm. Before making a decision with legal consequences, check it with a professional and against the official text: the link to EUR-Lex is in the sources. The rules are on the move—the digital omnibus changed dates in July 2026—and this guide reflects their state on the review date shown above.

What it is and who it applies to

The EU Artificial Intelligence Act—Regulation (EU) 2024/1689, in force since 1 August 2024—doesn't regulate the technology: it regulates its uses. The same tool can be free of obligations in one case and high-risk in another, depending on what it's used for and who it makes decisions about. That's why the first question is never "which tool are we using?", but "what are we using it for, and about whom".

Because it's a regulation, it applies directly in every Member State with no need for transposition. What each country does have to do is designate authorities and set its penalty regime—and that's where Spain is still only halfway there, as explained further down.

WHO PLACES IT ON THE MARKET

Provider

Develops an AI system and markets it, or puts it into service under its own name or trademark. Bears the bulk of the obligations: risk management, data quality, technical documentation, record-keeping, human oversight by design and post-market monitoring.

WHO USES IT

Deployer

Uses an AI system under its own authority—the norm in a mid-sized company. Lighter but real obligations: using it in line with the instructions, ensuring human oversight, reporting serious incidents and, for high-risk systems, informing the workers affected.

THE TRAP

Watch out for the change of role

A deployer can become a provider without realising it: if it puts its brand on a high-risk system, substantially modifies it or changes its intended purpose. It then takes on the provider's obligations, in full.

The real timeline, after the omnibus

On 27 July 2026 Regulation (EU) 2026/1744, the Digital Omnibus on AI, came into force, simplifying the regulation and moving dates. The headline: what already applied still applies. What moved is the high-risk block.

1 August 2024 ALREADY APPLIES

Entry into force

The regulation exists, but its obligations are phased in. The clock starts ticking.

2 February 2025 ALREADY APPLIES

Prohibited practices and AI literacy

Unacceptable uses are banned. And Article 4 comes into force: every company that uses AI must take AI literacy measures for its staff. No size threshold: a twelve-person company is covered too.

2 August 2025 ALREADY APPLIES

Governance and general-purpose models

The governance rules and the obligations for providers of general-purpose AI models kick in—the large models that almost everything else is built on.

2 August 2026 ALREADY APPLIES

General application and transparency

The regulation becomes fully applicable and the transparency obligations of Article 50 kick in: telling people when they're interacting with an AI and labelling certain synthetic content. This was not postponed.

December 2026 NOT YET

New omnibus prohibition

The prohibition added by the omnibus comes into force: AI systems that generate non-consensual sexually explicit intimate content and child sexual abuse material.

2 December 2027 POSTPONED

High risk — Annex III

Stand-alone high-risk systems: biometrics, critical infrastructure, education, employment, essential services, migration and justice. It was due on 2 August 2026 and the omnibus moved it here.

2 August 2028 POSTPONED

High risk — Annex I

AI embedded in products already covered by sectoral safety legislation—machinery, lifts, medical devices. It was due in August 2027.

The postponement isn't a general moratorium

What moved were the high-risk obligations, which affect whoever develops or integrates that kind of system. The prohibitions, AI literacy and transparency—which is exactly what concerns a company that simply uses AI tools—are unchanged and already enforceable. Reading “the AI Act has been postponed” and filing the matter away is this year's most expensive misreading.

The four categories, and where yours falls

The regulation classifies by the risk of the use, not by the power of the tool. Go through all four in order: the first one that fits is yours.

CATEGORY	WHAT'S IN IT	WHAT IT MEANS
Unacceptable risk	Social scoring, manipulation that exploits vulnerabilities, emotion recognition in the workplace or in educational institutions, biometric categorisation based on sensitive characteristics, untargeted scraping of facial images.	Prohibited. There's no compliance regime: it can't be done.
High risk	Recruitment and staff management, access to education, credit scoring, life and health insurance, critical infrastructure, biometrics, and AI embedded in products with their own safety legislation.	The full regime. For whoever uses it: human oversight, use in line with the instructions, informing the workers affected and reporting incidents.
Transparency risk	Systems that interact with people, synthetic content generation, deepfakes, emotion recognition systems outside the prohibited cases.	Disclose. People have to know they're talking to an AI or that the content is generated. In force since August 2026.
Minimal risk	The vast majority of business use: drafting, summarising, translating, searching, classifying documents, assisting with programming.	No specific obligations under the regulation. Data protection, intellectual property and confidentiality still apply.

The practical conclusion for a mid-sized company tends to be reassuring and tends to surprise: almost everything it's already using falls under minimal risk or transparency. The point where that changes is always the same—when the system starts making decisions about people—and that point is almost always crossed unintentionally, starting with human resources.

Two uses that cross the line without warning

Automated CV screening and candidate ranking fall under Annex III: they're high-risk, with an application date of 2 December 2027. And analysing employees' emotional state from their voice, face or text has been outright prohibited in the workplace since February 2025, with very few medical and safety exceptions. Both are bought as just another feature of a human resources tool.

The obligation that already applies to you: AI literacy

Since 2 February 2025, Article 4 has required providers and deployers to take measures so that their staff have a sufficient level of AI literacy. There's no size or sector threshold: if your company uses AI, it applies.

The July 2026 omnibus softened the wording—it no longer requires ensuring a specific individual level, but taking, and being able to demonstrate, support and promotion measures—and strengthened the role of the Commission and the Member States in facilitating that training. The obligation to have something written down and demonstrable doesn't go away.

+ It doesn't require a certificate, a minimum number of hours or a particular provider.
+ It does require taking into account each person's technical knowledge, experience and context of use: an identical generic course for the whole workforce is precisely what doesn't fit.
+ And it requires being able to prove it: what was assessed, what was done, who took part and when. Without a record, the measure doesn't exist for the purposes of an inspection.

> **Training by role—who uses, who reviews, who decides—isn't a generous reading of the law. It's what the law asks for and, as it happens, the only thing that works.**

Where Spain stands

The regulation applies directly, but each state designates its authorities and sets its penalty regime. In Spain that's done by the draft Organic Law for the proper use and governance of artificial intelligence (Proyecto de Ley Orgánica para el buen uso y la gobernanza de la Inteligencia Artificial), approved by the Council of Ministers on 26 May 2026 and published in the Boletín Oficial de las Cortes Generales on 12 June 2026.

As of this guide's review date it's still going through parliament. In other words: the European framework already applies, but the Spanish penalty regime isn't settled yet. The penalty not being finalised doesn't suspend the obligation—it only delays who collects it and how.

CENTRAL AUTHORITY

AESIA

The Spanish Agency for the Supervision of Artificial Intelligence (Agencia Española de Supervisión de la Inteligencia Artificial), based in A Coruña, is the central supervisory and notifying authority.

DIVISION OF POWERS

It isn't the only one

It shares powers with the Spanish Data Protection Agency and with sectoral authorities, depending on the system's area. The same case can have more than one supervisor.

ORDER OF MAGNITUDE

Penalties

The European regulation sets penalties of up to €35 million or 7% of worldwide turnover for prohibited practices, with lower bands for other infringements. The Spanish bill adopts that range.

The minimum for a company that only uses AI

Eleven points. They don't cover the high-risk regime—if it applies to you, this isn't enough—but they cover what most companies turning over €3–50M are required to do today, and they can be wrapped up in a morning's work.

- [ ] Inventory: which AI tools are actually used in the company, including the ones nobody approved. Start by asking, not by auditing.
- [ ] For each one, what it's used for and who or what it makes decisions about.
- [ ] Classification of each use into one of the four categories, with the reason written next to it.
- [ ] A check that no use falls into prohibited practices—look especially at human resources and video surveillance.
- [ ] A transparency notice wherever a person interacts with an AI system or receives generated content.
- [ ] A human review point in every use whose output has consequences, with a record of who reviewed it.

- [ ] A policy on which data can and can't be uploaded to external tools, written down and communicated.
- [ ] Contracts reviewed: where it's stored, for how long, whether it's used for training and what happens on cancellation.
- [ ] An AI literacy plan by role—who uses, who reviews, who decides—with a record of who did it and when.
- [ ] A named person responsible for this. You don't need a committee; you need someone.
- [ ] A review date in the calendar. The regulatory calendar keeps moving: the high-risk block arrives in 2027.

Four misreadings doing the rounds

FALSE

"It's been postponed, we'll worry about it later"

The high-risk block was postponed. The prohibitions, AI literacy and transparency are in force and are exactly the ones that affect everyday business use.

FALSE

"It's for the big tech companies"

Article 4 has no size threshold, and deployer obligations apply to anyone using a system under their authority. There are simplifications for SMEs and mid-cap companies, but simplifying isn't exempting.

FALSE

"If we don't develop AI, it doesn't apply to us"

Using it already makes you a deployer. And if you put your brand on a high-risk system or change its purpose, you become a provider with all its obligations.

FALSE

"Complying with the GDPR is enough"

They're separate frameworks that coexist. Data protection looks at personal data; the AI framework looks at how the system is used and its consequences, whether or not personal data is involved. Complying with one doesn't cover the other.

Sources

Regulation (EU) 2024/1689 — Artificial Intelligence Act

The original text, in force since 1 August 2024.

https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Regulation (EU) 2026/1744 — Digital Omnibus on AI

The simplification reform. Published on 24 July 2026, in force since the 27th.

https://eur-lex.europa.eu/eli/reg/2026/1744/oj

European Commission — The AI Omnibus enters into force

The official note with the new high-risk dates.

https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force

European Commission — Regulatory framework for AI

The risk categories and the application timeline.

https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

Article 4 — AI literacy

The article that applies to every company using AI, with no size threshold.

https://artificialintelligenceact.eu/article/4/

Government of Spain — Draft law for the proper use and governance of AI

Approved by the Council of Ministers on 26 May 2026.

https://digital.gob.es/comunicacion/notas-prensa/mtdfp/2026/05/el-gobierno-aprueba-el-proyecto-de-ley-que-garantizara-una-super

Quether

Quether works alongside leaders of mid-sized companies to transform their organisation—and themselves—without losing their essence.

This material, always up to date, at https://quether.com/en/resources/eu-ai-act