Holistic AI

EU AI ACT · SECTOR GUIDE

HR and recruitment

Recruitment is the most explicitly covered commercial use in the Act. Annex III point 4 names it directly, which means most hiring tools are high-risk by description rather than by argument — and one common HR use is prohibited outright.

THE DATE THAT BINDS THIS SECTOR

2 Dec 2027 Annex III high-risk obligations apply

moved from 2 Aug 2026

1	5	3
prohibited use case	high-risk use cases	provisions name this sector

Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744

WHAT THE ACT SAYS ABOUT YOU

The provisions that name this sector

The Act does not name industries, it names uses. These are the passages that reach HR and recruitment, quoted from the Official Journal. Everything else in this guide follows from them.

> AI systems intended to be used for the recruitment or selection of natural persons, in particular to place targeted job advertisements, to analyse and filter job applications, and to evaluate candidates
>
> Annex III, 4(a)

> AI systems intended to be used to make decisions affecting terms of work-related relationships, the promotion or termination of work-related contractual relationships, to allocate tasks based on individual behaviour or personal traits or characteristics
>
> Annex III, 4(b)

> the use of AI systems to infer emotions of a natural person in the areas of workplace and education institutions
>
> Article 5(1)(f)

What follows from this

1 prohibited — no compliance route

A prohibition is not a high bar to clear, it is a closed door. There is no conformity assessment, no documentation package and no derogation that makes the practice lawful: it stops, or it changes until it is no longer the practice the Article describes. Article 99(3) puts these at the top fine tier, €35m or 7% of worldwide turnover.

5 high-risk — the full Chapter III set, from 2 Dec 2027

Each carries risk management across the lifecycle (Article 9), data and data governance (Article 10), technical documentation to Annex IV (Article 11), automatic logging (Article 12), transparency and instructions for use (Article 13), human oversight (Article 14), and accuracy, robustness and cybersecurity (Article 15) — then conformity assessment, the EU declaration, CE marking and registration in the EU database. Deployers owe their own duties under Articles 26 and 27 on top.

YOUR USE CASES

Common use cases, classified

Each verdict below cites the single provision that decides it. Where the Act does not settle the answer from a description alone, the entry says so and names what does decide it — that is a finding, not a hedge.

CV screening and candidate ranking

High-risk

Annex III point 4(a) names systems used to recruit or select people, and specifically to analyse and filter job applications and evaluate candidates.

Annex III, 4(a)

Targeted job advertising

High-risk

Point 4(a) names placing targeted job advertisements explicitly — this is high-risk even though no application has been made yet.

Annex III, 4(a)

Promotion, termination and task allocation

High-risk

Point 4(b) covers decisions affecting terms of work-related relationships, promotion or termination, and allocating tasks based on behaviour or personal traits.

Annex III, 4(b)

Performance and behaviour monitoring

High-risk

Point 4(b) reaches monitoring and evaluating the performance and behaviour of people in work-related relationships.

Annex III, 4(b)

Emotion recognition on staff

Prohibited

Inferring emotions of a natural person in the workplace is a prohibited practice, not a high-risk one — there is no compliance route for it.

Article 5(1)(f) allows a narrow exception where the system is put in place for medical or safety reasons.

Article 5(1)(f)

Interview or screening chatbot

High-risk

If it evaluates or filters candidates it falls under point 4(a) like any other screening tool; the conversational interface changes nothing.

Annex III, 4(a)

THE EXPENSIVE MISTAKES

What gets misread here

Over-classification is the quieter failure. Treating something as high-risk when the Act carves it out costs real money and never produces a complaint, so nobody ever catches it.

> “Emotion or sentiment analysis on staff is high-risk, so we can do it with a conformity assessment.”
>
> It is prohibited, not high-risk. There is no conformity route — the practice stops. The only exception in the provision is use put in place for medical or safety reasons.
>
> Article 5(1)(f)

> “Advertising is marketing, so job adverts sit outside the hiring obligations.”
>
> Point 4(a) names placing targeted job advertisements explicitly. The obligation attaches before anyone has applied.
>
> Annex III, 4(a)

> “We bought the tool, so the obligations are the vendor's.”
>
> Deployers carry their own obligations, and putting your name on a system, substantially modifying it, or changing its intended purpose makes you the provider of it.
>
> Article 25

THE CALENDAR

The dates that bind you

Most published material still carries the dates the Regulation had before July 2026. These are the amended ones.

2 FEB 2025

All providers and deployers placing AI on the EU market or whose output is used in the EU.

PROVISIONS

Article 5
Article 4

Prohibited practices and AI literacy apply

Since 2 February 2025 the first eight prohibited AI practices in Article 5 have been unlawful across the EU, and every provider and deployer has owed an AI literacy duty under Article 4. Two further prohibitions were added later and apply from 2 December 2026.

- The Article 5 bans on social scoring, untargeted facial scraping, emotion inference at work and in education, subliminal manipulation, exploitation of vulnerability, predictive policing on profiling alone, biometric categorisation for sensitive traits, and, subject to narrow carve-outs, real-time remote biometric identification in public for law enforcement
- The Article 4 AI literacy duty: since amended by the Omnibus from a duty to **ensure** a sufficient level of AI literacy to a duty to take measures to support it
- Chapter I definitions and scope provisions

2 DEC 2027

Providers, importers, distributors and deployers of high-risk AI systems listed in Annex III.

PROVISIONS

Annex III
Article 6(2)

Annex III high-risk obligations apply

Providers of Annex III high-risk systems (including AI used in employment, education, credit scoring, life and health insurance pricing, essential services, law enforcement and migration) must meet the full Chapter III requirements from 2 December 2027.

- Article 9 risk management system across the lifecycle
- Article 10 data governance and training-data quality
- Article 11 and Annex IV technical documentation
- Article 12 automatic logging and record-keeping
- Articles 13–15 transparency to deployers, human oversight, accuracy, robustness and cybersecurity
- Article 17 quality management system
- Conformity assessment, EU declaration of conformity, CE marking, and registration in the EU database

~~2 Aug 2026~~ → 2 Dec 2027

The Omnibus entered into force on 27 July 2026 and deferred these obligations by 16 months. The conformity-assessment infrastructure the deadline depends on, harmonised standards and designated notified bodies, was the stated reason.

Already in force for everyone, whatever your risk tier

1 Aug 2024	The Regulation enters into force
2 Feb 2025	Prohibited practices and AI literacy apply
2 Aug 2025	General-purpose AI model obligations apply
2 Aug 2026	Transparency obligations and the general application date

DOES IT REACH YOU AT ALL

Three questions, in this order

Scope is settled by three facts, and they have to be established in sequence: being outside the Union does not take you out, and being a buyer rather than a builder does not either.

01 · TERRITORY

Article 2

Is the system, or its output, used in the Union?

You are in scope if you place a system on the EU market wherever you are established, and also if you are established outside the Union but the output produced by the system is used inside it. Military, defence and national-security purposes are excluded, as is purely personal non-professional use, and so is free and open-source software except where it is high-risk or caught by Article 5 or 50.

02 · ROLE

Article 3(3)–(7), Article 25

Are you the provider, deployer, importer or distributor?

Duties attach to roles, not to company types, and one organisation is routinely several at once. The trap is Article 25: a deployer becomes the provider of a system — taking on the whole provider obligation set — if it puts its own name or trademark on it, modifies it substantially, or changes its intended purpose. Contracts do not reallocate this.

03 · TIER

Articles 5, 6, 50; Annexes I and III

Which tier does each system fall in?

Prohibited under Article 5; high-risk because it is a safety component of a product covered by Annex I, or because it is named in Annex III; subject to transparency duties under Article 50; or none of these. The tiers are decided per system and per intended purpose, so the same model can sit in two tiers in two products.

OUT OF SCOPE

Article 2(3) to (12)

What Article 2 puts outside the Regulation entirely

- Systems placed on the market, put into service or used exclusively for military, defence or national-security purposes, whatever the type of entity carrying out those activities.
- The obligations of deployers who are natural persons using an AI system in the course of a purely personal non-professional activity. Note what this does and does not reach: it lifts the deployer's obligations, not the provider's.
- Research, testing and development before a system is placed on the market or put into service — but testing in real-world conditions is expressly not covered by that exclusion.
- Systems and models developed and put into service for the sole purpose of scientific research and development, including their output.
- Systems released under free and open-source licences, except where they are placed on the market or put into service as high-risk, or as a system falling under Article 5 or Article 50.
- Third-country public authorities and international organisations using AI systems under international law-enforcement or judicial cooperation with the Union — but only where that country or organisation provides adequate safeguards for fundamental rights.

The Article 6 filter, added by the Omnibus

Not every AI component in a regulated product is a safety component. Article 6(1a) now excludes AI used solely for performance optimisation, service efficiency, automation, convenience or quality control. Article 6(1b) pulls back in anything whose failure or malfunctioning would endanger health and safety. Between them they decide whether a factory or medical-device deployment carries the Chapter III weight, and they are the first thing to check before assuming it does.

IF YOU GET IT WRONG

Who enforces this, and for how much

Regulation (EU) 2026/1744 added Articles 75a to 75d, moving part of the market from national authorities to the AI Office. Which one you answer to is decided by the system, not by you.

SUPERVISOR	REMIT	CEILING
National market surveillance authorities Articles 74, 99	The default supervisor for everything not reallocated: high-risk systems, the Article 5 prohibitions, and the Article 50 transparency duties, in each Member State where the system is on the market.	€35m / 7%, €15m / 3% or €7.5m / 1% depending on the breach
The AI Office Articles 75, 75a–75d, 101	Exclusive supervision of AI systems built on a general-purpose AI model where the model and the system come from the same provider, and of AI systems integrated into very large online platforms and search engines under the DSA. Also the sole enforcer against general-purpose AI model providers.	Article 99(3)–(7) amounts apply correspondingly, plus periodic penalties up to 5% of average daily turnover per day
European Data Protection Supervisor Article 100	Fines against Union institutions, bodies, offices and agencies that fall within the scope of the Regulation.	€1.5m for a prohibited practice, €750,000 for other breaches

WHAT TO DO

The work, in the order it happens

Not the Regulation's chapter order. You cannot classify a system before you know your role in it, and you cannot scope obligations before you know what you hold. The full 31-step version, with the provision behind each step, is on euaiact.com.

01 Establish what you hold

3 steps

Every obligation in the Regulation depends on facts about a specific system, so nothing can be assessed until the inventory exists. This phase has no deadline of its own and blocks every phase after it.

02 Screen against the prohibitions

2 steps

Article 5 is absolute: a prohibited practice cannot be made lawful by documentation, consent or safeguards. Screen for it before spending effort anywhere else, because a hit means the system cannot be placed on the market or used at all.

03 Classify each system

5 steps

Classification decides which regime applies and therefore how much work follows. The two high-risk routes carry different deadlines, so the route matters as much as the conclusion.

04 Build the high-risk requirements

8 steps

This is the heavy phase and the reason the deferral to December 2027 is a schedule change rather than a reprieve. Each requirement produces evidence that a conformity assessment will look for.

05 Meet the deployer duties

4 steps

Most organisations deploy far more systems than they provide, and deployer duties are frequently missed because the provider's documentation is mistaken for the whole obligation.

06 Clear the route to market

4 steps

The gate between a compliant build and lawful placing on the market. Which route applies depends on the classification route and on whether harmonised standards were applied.

07 Operate, monitor and keep it current

5 steps

Compliance is a standing condition, not a launch milestone. These obligations run for as long as the system is on the market.

SOURCES

How this guide was verified

PRIMARY SOURCES

Regulation (EU) 2024/1689, consolidated at 27 July 2026, as amended by Regulation (EU) 2026/1744.

LAST VERIFIED

11 Aug 2026

Every quotation in this guide was matched character for character against the Official Journal text, and every classification cites the provision that decides it. The comparison runs on each build rather than once: a quotation that no longer matches, or a citation to a provision that no longer says what is claimed, stops the file being written.

That matters most for what this guide says is **not** caught. A reader acts on a negative by doing nothing, so each one is tested against the current text of the annex every time.

Commission guidance that reads on this sector

- Guidelines on prohibited AI practices (February 2025) — interprets Article 5
- Guidelines on transparency obligations for providers and deployers (2025) — interprets Article 50

Not binding, but the Commission's own reading and the first place a supervisory authority will look.

Information, not legal advice. Where this guide and the Official Journal disagree, the Official Journal governs.

WHERE TO GO NEXT

Check your own systems

This guide tells you what the Act says about HR and recruitment. It cannot tell you what you run, which is the only thing that decides what you owe.

FREE, FIVE MINUTES, NO SIGN-UP

Classify one system

Answer a short series of questions about a single system and get its tier, the provision that decides it, and the obligations that follow — the same classification logic behind the use cases in this guide.

Open the risk classifier →

euaiact.com/eu-ai-act-risk-calculator

WITH HOLISTIC AI

The EU AI Act readiness assessment

Where the classifier answers one system, the assessment covers the portfolio and produces the evidence you would have to show for it: a written classification per system with its deciding provision, a gap register mapped to the articles that create each obligation, and a remediation plan sequenced against the amended dates.

Holistic AI is an enterprise AI governance platform — a Leader in Avasant's 2025 Responsible AI Platforms RadarView, and a Gartner® Cool Vendor for AI Security — used by AON, eBay, GE Healthcare, Siemens, Starling Bank and Unilever.

Start the assessment →

holisticai.com/eu-ai-act-readiness

Get a demo →

holisticai.com/demo

The full legal text, the other seven sector guides and the 31-step checklist are at euaiact.com, free and ungated.

Holistic AI

The EU AI Act, as a working tool

euaiact.com is maintained by Holistic AI as an independent, source-verified reference: the complete legal text with the Omnibus amendments marked, a scope finder, a risk classifier, eight sector guides and a 31-step compliance checklist — all free.

YOUR NEXT STEP

Find out where you actually stand

Classify one system in a few minutes, or have Holistic AI assess the whole portfolio and produce the evidence behind it.

Classify a system →

euaiact.com/eu-ai-act-risk-calculator

Start the assessment →

holisticai.com/eu-ai-act-readiness

Or book a walkthrough of the platform: holisticai.com/demo

Corrections, sources and feedback: hello@holisticai.com
Partnerships: rahul.patel@holisticai.com