



The EU AI Act, after the Omnibus

What applies, to whom, and by when — read out of the Official Journal text rather than summarised from it. Eight sector chapters, 38 classified use cases, and the 31-step sequence the work actually runs in.

THE DATES STILL AHEAD

- | | |
|------------|--|
| 2 Dec 2026 | Two new prohibitions and the legacy marking deadline

added by the Omnibus |
| 2 Dec 2027 | Annex III high-risk obligations apply

moved from 2 Aug 2026 |
| 2 Aug 2028 | Annex I product-safety high-risk obligations apply

moved from 2 Aug 2027 |
| 2 Aug 2030 | Legacy large-scale IT systems must comply |

11

changes made by Regulation (EU) 2026/1744

2

application dates deferred

38

use cases classified against the Act

8

sectors covered in full

What this is, and what it is not

SOURCE

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

LAST VERIFIED

Dates 11 Aug 2026
Omnibus 7 Sept 2026
Checklist 18 Aug 2026
Penalties 11 Aug 2026

NOT LEGAL ADVICE

Where this document and the
Official Journal disagree, the
Official Journal governs.

Most published material on the EU AI Act still carries the dates the Regulation had before July 2026. Regulation (EU) 2026/1744 deferred two of them, added one that did not exist, and changed eleven things in total. A plan built on the adopted text is now wrong in both directions: it will have you working to a deadline that has moved, and missing one that has not.

This playbook is built from the amended text. Every date, obligation and classification in it is generated from the same verified dataset that powers euaiact.com, and every legal claim is anchored to the provision it rests on, quoted rather than paraphrased. Where a quotation appears, it has been matched character for character against the Official Journal; a build that cannot match one fails rather than publishing it.

It follows the order the work happens in

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 — so the inventory comes first and the conformity route comes last, whatever order the articles are numbered in.

Three things it deliberately does not do

- **It does not tell you whether a given system complies.** That turns on facts about the system. What it does is name the question the Regulation obliges you to be able to answer, and cite the provision that asks it.
- **It does not treat over-classification as the safe error.** Calling something high-risk when the Act carves it out costs real money and never produces a complaint, so nobody catches it. Each sector chapter names the carve-outs as carefully as the duties.
- **It does not restate the Act.** The full text, article by article, is on euaiact.com with the amended wording marked. This is the layer above it: what the text means for a plan.

What is in here

FRONT

Executive summary	3
-------------------------	---

PART I · THE CALENDAR

Every application date, as amended	5
What is already in force	6
What is still ahead	7

PART II · REGULATION (EU) 2026/1744

The deadlines and the definitions	10
The new duties and powers	11
What it left alone	12

PART III · SCOPE AND RISK

Three questions, in this order	15
The five outcomes, and what each costs	17

PART IV · THE EIGHT SECTORS

HR and recruitment	19
Financial services and insurance	21
Healthcare and life sciences	24
Education and training	26
Public sector and law enforcement	28
Retail, marketing and media	30
Manufacturing and infrastructure	32
Technology and software vendors	34

PART V · THE WORK, IN ORDER

Establish what you hold, and screen against the prohibitions	37
Classify each system	38
Build the high-risk requirements	39
Meet the deployer duties, and clear the route to market	40
Operate, monitor and keep it current	41

PART VI · ENFORCEMENT

Who supervises what	43
How an amount is arrived at	44

BACK MATTER

Sources and method	45
--------------------------	----

WHERE TO GO NEXT

Who maintains this	46
Check your own exposure	47

Where the Act stands today

Four of the Act's eight application dates have passed. The obligations that bite hardest — the high-risk requirements in Chapter III — are the ones still ahead, and they moved: Annex III use cases to 2 Dec 2027, Annex I products to 2 Aug 2028.

Already enforceable

The Article 5 prohibitions, the Article 4 AI literacy duty, the general-purpose AI model obligations in Articles 53 and 55, and the Article 50 transparency duties. No transition remains on any of them.

Next to land

Two further prohibitions on 2 Dec 2026, together with the legacy deadline for marking AI-generated content under Article 50(2). Then the Annex III high-risk regime a year after that.


What the deferral is not

A reduction. Articles 9 to 15 keep every requirement they were adopted with. What moved is when they apply — the stated reason being that harmonised standards and notified bodies were not ready, not that the bar was too high.

READING THE BARS

Every one of the 38 use cases in this playbook is classified against a named provision, not a judgement. The distribution is what a real portfolio looks like: most things are not caught, and the few that are, are caught clearly.

How 38 common use cases classify



“Depends on the detail” is not a hedge. It marks the cases where the Act genuinely does not settle the answer from a description alone, and the sector chapters say what decides each one.

The judgement that matters most

Whether you are the provider or the deployer of a system is the single most consequential classification in the Act, and the one organisations most often get wrong. Putting your own name on a bought system, modifying it substantially, or changing its intended purpose makes you its provider — with the full Chapter III obligations that follow. Article 25 decides this, and it decides it on what you did, not on what the contract says.

PART I

The calendar

Four of the Act's eight application dates have passed. Of the four still ahead, three were moved or created by the Omnibus — and most published material still carries the dates they replaced.

01 Every application date, as amended

02 What is already in force

03 What is still ahead

Every application date, as amended

Solid nodes are in force. Hollow nodes are ahead. An amber ring marks a date Regulation (EU) 2026/1744 moved or created — those are the three dates that most published material still has wrong.



What is already in force

1 AUG 2024

No operator obligations yet: this is the reference date only.

PROVISIONS

[Article 113](#)

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](#)

2 AUG 2025

Providers of general-purpose AI models placed on the EU market. Models already on the market before this date have until 2 August 2027 to comply.

PROVISIONS

[Article 53](#)

[Article 55](#)

2 AUG 2026

Providers and deployers of AI systems that interact with people or generate synthetic content: chatbots, voice agents, generative image and video tools, and publishers of AI-written editorial content.

PROVISIONS

[Article 50](#)

The Regulation enters into force

Regulation (EU) 2024/1689 entered into force on 1 August 2024, starting the clock on every later deadline but imposing no immediate obligations on providers or deployers.

- The staggered application periods in Article 113 begin to run
- Member States begin designating national competent authorities
- The European Commission begins preparing guidance and delegated acts

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

General-purpose AI model obligations apply

Since 2 August 2025, providers of general-purpose AI models have had to publish training-data summaries, maintain technical documentation, and respect EU copyright law, with extra systemic-risk duties above the compute threshold.

- Article 53 documentation, copyright policy and public training-content summary for all GPAI models
- Article 55 systemic-risk duties (model evaluation, adversarial testing, incident reporting, cybersecurity) for models presenting systemic risk
- Governance provisions: the AI Office, the AI Board, and national competent authority designation
- The Member State penalty regime under Article 99, other than the GPAI-specific fines

Transparency obligations and the general application date

From 2 August 2026 the Article 50 transparency rules apply: people must be told when they are interacting with an AI system, and synthetic audio, image, video and text must be marked in a machine-readable format.

- Article 50(1): disclosure that a user is interacting with an AI system, unless it is obvious
- Article 50(2): machine-readable marking of synthetic audio, image, video and text
- Article 50(3): notification when an emotion recognition or biometric categorisation system is in use
- Article 50(4): deep fake disclosure, and disclosure for AI-generated text published to inform the public
- Member State obligation to have at least one operational AI regulatory sandbox
- The remainder of the Regulation, other than the high-risk provisions listed below

What is still ahead

2 DEC 2026

Providers and deployers of generative AI systems. The new prohibitions carry no grandfathering: systems already on the market are in scope.

PROVISIONS

[Article 5](#)

[Article 50](#)

Two new prohibitions and the legacy marking deadline

From 2 December 2026 two further practices are prohibited outright (AI that generates non-consensual intimate imagery, and AI that generates child sexual abuse material) and generative systems already on the market before 2 August 2026 must carry machine-readable marking.

- Article 5(1)(ba): AI systems generating or manipulating realistic depictions of the intimate parts or sexually explicit conduct of identifiable people without their explicit consent
- Article 5(1)(bb): AI systems generating or manipulating child sexual abuse material within the meaning of Directive 2011/93/EU
- New Article 5(1a) and (1b), governing provider liability for reasonably foreseeable misuse and the safeguards defence
- Article 50(2) machine-readable marking for generative systems placed on the market before 2 August 2026

~~2 Aug 2026~~

→

2 Dec 2026

This date did not exist in the adopted text. The Omnibus created it, both to add the two new prohibitions and to give pre-existing generative systems a transitional window for synthetic-content marking.

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.

2 AUG 2028

Manufacturers of products covered by the Annex I harmonisation legislation where an AI system is a safety component or is itself the product.

PROVISIONS

[Annex I](#)
[Article 6\(1\)](#)

Annex I product-safety high-risk obligations apply

AI systems that are safety components of products already regulated under EU harmonisation law (medical devices, machinery, vehicles, lifts, toys and the rest of Annex I) come into scope on 2 August 2028.

- Chapter III requirements for AI safety components within Annex I products
- Integration of AI conformity assessment into the existing sectoral conformity route
- Obligations for public authorities already operating high-risk systems before this date

~~2 Aug 2027~~

→

2 Aug 2028

Deferred by 12 months, keeping the Annex I product-safety route eight months behind Annex III so sectoral conformity assessment has time to absorb it.

2 AUG 2030

EU agencies and Member State authorities operating the Annex X large-scale IT systems.

PROVISIONS

[Article 111](#)

Legacy large-scale IT systems must comply

AI components inside the EU large-scale IT systems listed in Annex X, such as the Schengen Information System and Eurodac, that were placed on the market before 2 August 2027 have until 2 August 2030 to be brought into conformity.

- Full Chapter III conformity for in-scope Annex X system components

PART II

What the Omnibus changed

Regulation (EU) 2026/1744 entered into force on 27 July 2026 and made eleven changes that alter what an operator has to do. It reduced no substantive requirement.

01 The deadlines and the definitions

02 The new duties and powers

03 What it did not change

The deadlines and the definitions

High-risk deadline (Annex III use cases)

Article 113

~~2 August 2026~~

→

2 December 2027

A 16-month deferral for the eight Annex III use-case areas. The stated reason was that the infrastructure the deadline depends on, harmonised standards and designated notified bodies, was not ready.

High-risk deadline (Annex I products)

Article 113

~~2 August 2027~~

→

2 August 2028

A 12-month deferral, keeping the product-safety route eight months behind Annex III so sectoral conformity assessment has time to absorb it.

What counts as a safety component

Article 6(1a) to (1c)

~~No definition; any AI in a regulated product risk of being caught~~

→

AI used solely for non-safety purposes is not a safety component

This narrows the Annex I route rather than deferring it. AI used solely for user assistance, performance optimisation, service efficiency, automation, convenience or quality control does not qualify as a safety component, unless its failure would endanger health and safety. A product needing third-party assessment only for non-health-safety risks, such as radio spectrum or electromagnetic interference, no longer meets the Article 6(1)(b) condition.

AI literacy duty

Article 4

~~Providers and deployers shall ensure a sufficient level of AI literacy~~

→

Providers and deployers shall take measures to support AI literacy

Softened from a duty of result to a duty of effort. The obligation still binds every operator regardless of risk tier, and has applied since 2 February 2025, but it is now assessed on the measures taken rather than the level achieved.

Marking of existing synthetic content

Article 111(4), applying the duty in Article 50(2)

~~No transitional window~~

→

Legacy systems must comply by 2 December 2026

Generative systems already on the market before 2 August 2026 got a transitional window to implement machine-readable marking of synthetic output. The deadline is set by a new transitional paragraph, Article 111(4); the duty it defers is the existing Article 50(2). The only amendment to Article 50 itself is to paragraph 7, on codes of practice.

The new duties and powers

New prohibition: non-consensual intimate imagery

Article 5(1)(ba)

~~Not prohibited~~

→

Prohibited

A ninth ban, inserted rather than deferred. It sits in the €35 million / 7% penalty tier alongside the original eight prohibitions.

New prohibition: child sexual abuse material

Article 5(1)(bb)

~~Not prohibited~~

→

Prohibited

A tenth ban, on the same date and in the same penalty tier. Anyone still working from a list of eight prohibitions is working from the superseded text.

AI Office enforcement powers

Articles 75a-75d

~~No direct supervisory powers over these operators~~

→

Information requests, evaluations, remedial measures and fines, with periodic penalty payments

A new enforcement chapter. The AI Office can impose periodic penalty payments of up to 5% of average daily turnover for each day a breach continues, which on a long-running infringement can exceed the headline maximum fine.

Penalties for small mid-cap enterprises

Article 99(6a), with new definitions in Article 3(14a) and (14b)

~~Only SMEs, including start-ups, faced the lower figure~~

→

SMCs also face the lower of the fixed sum and the turnover percentage

A new category. SMEs already had the lower figure under Article 99(6); this extends the same treatment to small mid-cap enterprises, companies that have outgrown the SME thresholds. "SMC" is defined by new Article 3(14b) as a small mid-cap enterprise under point (2) of the Annex to Recommendation (EU) 2025/1099.

Processing special-category data for bias detection

Article 4a

~~Only the narrow Article 10(5) basis~~

→

A broadened legal basis for bias detection and correction

Addresses a practical deadlock: testing a high-risk system for discriminatory outcomes can require the very special-category data the GDPR restricts.

Scope of a notified body's designation

Annex XIV

~~Not present~~

→

New annex: a list of codes, categories and corresponding types of AI systems

Notification under Article 30 now uses a fixed taxonomy of AIA codes, so a notified body's designation states exactly which types of system it may assess. A body designated for one category cannot assess another.

42 of 119 articles changed. No requirement was reduced.

A deferral is widely read as a softening, and it is the most expensive misreading in circulation: it leads organisations to stand work down that they will have to restart with less time. Counted against the amended text, twelve of the 18 articles that create the high-risk duties were not touched at all, and none of the six that were lost a requirement.

The duties that stand exactly as adopted

Art 9	Risk Management System
Art 12	Record-Keeping
Art 13	Transparency and Provision of Information to Deployers
Art 14	Human Oversight
Art 15	Accuracy, Robustness and Cybersecurity
Art 16	Obligations of Providers of High-Risk AI Systems
Art 26	Obligations of Deployers of High-Risk AI Systems
Art 47	EU Declaration of Conformity
Art 48	CE Marking
Art 49	Registration
Art 71	EU Database for High-Risk AI Systems Listed in Annex III
Art 73	Reporting of Serious Incidents

And what the six amendments in this chapter actually did

Each either narrows who has to do the work in full, or removes a duplication. None removes an obligation.

Art 10 Data and Data Governance

Ties the data-governance criteria to the new Article 4a, and deletes the old Article 10(5) basis for processing special categories of data, which moves to Article 4a.

Art 11 Technical Documentation

Lets SMEs, start-ups and small mid-caps supply the Annex IV elements in a simplified form. The elements themselves are unchanged.

Art 17 Quality Management System

Makes the quality management system proportionate to the size of the provider, while keeping the required degree of rigour.

Art 27 Fundamental Rights Impact Assessment for High-risk AI Systems

Lets a deployer cross-reference a data protection impact assessment already done under the GDPR instead of repeating the work.

Art 43 Conformity Assessment

Folds the Article 17 quality management assessment into the Annex I Section A route, aligning the two conformity procedures.

Art 72 Post-market Monitoring by Providers and Post-market Monitoring Plan for High-risk AI Systems

Adds a Commission template and guidance for the post-market monitoring plan, due by 2 September 2027.

Two things outside this chapter that also stand

Annex III was not shortened: no high-risk use case was taken off the list, and the Commission keeps its Article 7 power to add to it. And data protection law applies untouched — the new Article 4a basis for processing special categories of data to detect and correct bias works inside the GDPR rather than beside it, and is a permission, not a duty.

PART III

Does the Act reach you?

Scope turns on three facts, established in order. Being outside the Union does not take you out of it, and buying rather than building does not either.

01 Three questions, in this order

02 The five outcomes, and what each costs

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.

The five outcomes, and what each costs

High-risk

19 of 38 use cases

Permitted, but carries the full Chapter III obligations: risk management, data governance, documentation, logging, human oversight, accuracy and cybersecurity.

Prohibited

8 of 38 use cases

Banned outright. There is no compliance route — the practice has to stop or change.

Depends on the detail

6 of 38 use cases

The Act does not settle this from the description alone. The note says what decides it.

Transparency duty

3 of 38 use cases

Not high-risk, but people must be told. Disclosure and content-marking duties under Article 50.

No specific duty

2 of 38 use cases

Not named in the Act's risk tiers. The AI literacy duty in Article 4 still applies, as does law outside the Act.

How the counts were produced. Each of the 38 use cases in this playbook was classified against a single named provision, recorded with that citation, and checked by the build: a classification that cites a provision the Act no longer contains fails the build rather than printing. The distribution is not a survey of the market — it is what happens when a representative set of ordinary business uses is put through the Act's own tests.

PART IV

The eight sectors

The Act does not name industries. It names uses. These chapters work backwards from a sector to the provisions that reach it, and quote them.

-
- 01 HR and recruitment
 - 02 Financial services and insurance
 - 03 Healthcare and life sciences
 - 04 Education and training
 - 05 Public sector and law enforcement
 - 06 Retail, marketing and media
 - 07 Manufacturing and infrastructure
 - 08 Technology and software vendors
-

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.

1 · Prohibited

5 · High-risk

What the Act names here

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.

Dates that bind this sector

2 Feb 2025 Prohibited practices and AI literacy apply

2 Dec 2027 Annex III high-risk obligations apply — [date changed by the Omnibus](#)

Common use cases, classified

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\(a\)](#)

What gets misread here

“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](#)

Commission guidance that reads on this sector. Guidelines on prohibited AI practices (February 2025), interpreting Article 5. Guidelines on transparency obligations for providers and deployers (2025), interpreting Article 50. These are not binding, but they are the Commission's own reading and the first place a supervisory authority will look.

Financial services and insurance

Two points of Annex III do most of the work here, and both contain carve-outs that decide whether a system is in scope at all. Reading them loosely is how banks end up treating fraud models as high-risk and insurers treating motor pricing as caught.

1 · Prohibited

2 · High-risk

1 · Transparency duty

2 · Depends on the detail

What the Act names here

AI systems intended to be used to evaluate the creditworthiness of natural persons or establish their credit score, with the exception of AI systems used for the purpose of detecting financial fraud

Annex III, 5(b)

AI systems intended to be used for risk assessment and pricing in relation to natural persons in the case of life and health insurance

Annex III, 5(c)

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.

2 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.

1 transparency duty — disclosure and marking

Article 50 applies whatever the risk tier: people must be told when they are interacting with an AI system, and synthetic audio, image, video and text must be marked in a machine-readable format and detectable as artificially generated. In force since 2 August 2026.

Dates that bind this sector

2 Feb 2025 Prohibited practices and AI literacy apply

2 Dec 2027 Annex III high-risk obligations apply — [date changed by the Omnibus](#)

Common use cases, classified

Credit scoring and creditworthiness

High-risk

Point 5(b) names evaluating the creditworthiness of natural persons or establishing their credit score.

It applies to natural persons. Scoring a company is not covered by this point.

[Annex III, 5\(b\)](#)

Financial fraud detection

Depends on the detail

Point 5(b) expressly excludes AI systems used for the purpose of detecting financial fraud, so fraud detection is not high-risk on that basis.

The carve-out is narrow. If the same model also decides creditworthiness or account access, that use is assessed on its own terms.

[Annex III, 5\(b\)](#)

Life and health insurance pricing

High-risk

Point 5(c) names risk assessment and pricing in relation to natural persons for life and health insurance.

Only life and health. Motor, property and commercial lines are not named in point 5(c).

[Annex III, 5\(c\)](#)

Motor, property or commercial insurance pricing

Depends on the detail

Point 5(c) is limited to life and health insurance, so other lines are not high-risk under it.

Still assess it against the prohibitions and against national insurance and anti-discrimination law, which the Act does not displace.

[Annex III, 5\(c\)](#)

Customer service chatbot

Transparency duty

A system intended to interact directly with people carries a disclosure duty, but is not high-risk merely for being conversational.

It becomes high-risk if it decides something in Annex III — creditworthiness, for instance, rather than explaining a product.

[Article 50\(1\)](#)

General-purpose customer scoring on social behaviour

Prohibited

Social scoring leading to detrimental treatment in unrelated contexts, or treatment that is unjustified or disproportionate, is prohibited.

[Article 5\(1\)\(c\)](#)

What gets misread here

“Our fraud detection scores customers, so it is high-risk like credit scoring.”

Point 5(b) expressly excludes systems used for detecting financial fraud. The carve-out is in the same sentence that creates the obligation.

[Annex III, 5\(b\)](#)

“All our insurance pricing models are high-risk.”

Point 5(c) reaches life and health insurance only. Motor, property and commercial lines are not named in it.

[Annex III, 5\(c\)](#)

“Scoring businesses for lending is covered too.”

Point 5(b) is about natural persons. Commercial credit decisions about a company are not caught by it.

[Annex III, 5\(b\)](#)

Commission guidance that reads on this sector. Guidelines on prohibited AI practices (February 2025), interpreting Article 5. Guidelines on transparency obligations for providers and deployers (2025), interpreting Article 50. These are not binding, but they are the Commission’s own reading and the first place a supervisory authority will look.

Healthcare and life sciences

Most clinical AI is high-risk through the product-safety route in Article 6(1), not through Annex III — which means the classification turns on device law rather than on how clinical the use feels. Annex III adds two further hooks for public services and emergency response.

3 · High-risk

1 · Depends on the detail

What the Act names here

the AI system is intended to be used as a safety component of a product, or the AI system is itself a product, covered by the Union harmonisation legislation listed in Annex I

[Article 6\(1\)](#)

AI systems intended to evaluate and classify emergency calls by natural persons or to be used to dispatch, or to establish priority in the dispatching of, emergency first response services

[Annex III, 5\(d\)](#)

to evaluate the eligibility of natural persons for essential public assistance benefits and services, including healthcare services

[Annex III, 5\(a\)](#)

What follows from this

3 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.

Dates that bind this sector

2 Dec 2027 Annex III high-risk obligations apply — [date changed by the Omnibus](#)

2 Aug 2028 Annex I product-safety high-risk obligations apply — [date changed by the Omnibus](#)

Common use cases, classified

Clinical decision support or diagnostic AI

High-risk

Where the AI is a safety component of a medical device, or is itself such a device requiring third-party conformity assessment, Article 6(1) makes it high-risk.

This route depends on the device classification under the medical devices regulation, not on how clinical the use feels.

[Article 6\(1\) with Annex I](#)

Emergency call triage and dispatch

High-risk

Point 5(d) names evaluating and classifying emergency calls and establishing priority in dispatching emergency first response services.

[Annex III, 5\(d\)](#)

Eligibility for public healthcare or benefits

High-risk

Point 5(a) covers systems used by or for public authorities to evaluate eligibility for essential public assistance benefits and services, including healthcare.

[Annex III, 5\(a\)](#)

Clinical documentation and administration

Depends on the detail

Administrative support is not named in Annex III and is usually not a safety component, so it commonly falls outside the high-risk tiers.

If output feeds a diagnostic or triage decision rather than a record, assess it as clinical support instead.

[Article 6](#)

What gets misread here

“It touches patients, so it is high-risk.”

The Article 6(1) route depends on whether the product requires third-party conformity assessment under the device legislation in Annex I. Clinical feel is not the test.

[Annex I](#)

“Ambient scribes and coding tools are clinical, so they are caught.”

Documentation and administration are not named in Annex III and are usually not safety components. The answer changes if the output drives a diagnostic or triage decision rather than a record.

[Article 6](#)

Commission guidance that reads on this sector. Guidelines on the definition of an AI system (February 2025), interpreting Article 3, Article 6. These are not binding, but they are the Commission’s own reading and the first place a supervisory authority will look.

Education and training

Annex III point 3 names four distinct education uses, covering the whole path from admission to assessment to invigilation. As in the workplace, inferring emotion is prohibited rather than regulated.

1 · Prohibited

3 · High-risk

What the Act names here

AI systems intended to be used to determine access or admission or to assign natural persons to educational and vocational training institutions at all levels

[Annex III, 3\(a\)](#)

AI systems intended to be used to evaluate learning outcomes, including when those outcomes are used to steer the learning process

[Annex III, 3\(b\)](#)

AI systems intended to be used for monitoring and detecting prohibited behaviour of students during tests

[Annex III, 3\(d\)](#)

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.

3 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.

Dates that bind this sector

2 Feb 2025 Prohibited practices and AI literacy apply

2 Dec 2027 Annex III high-risk obligations apply — [date changed by the Omnibus](#)

Common use cases, classified

Admissions and placement

High-risk

Point 3(a) names determining access or admission, or assigning people to educational and vocational training institutions at all levels.

[Annex III, 3\(a\)](#)

Automated grading and learning assessment

High-risk

Point 3(b) covers evaluating learning outcomes, including where those outcomes steer the learning process.

[Annex III, 3\(b\)](#)

Exam proctoring

High-risk

Point 3(d) names monitoring and detecting prohibited student behaviour during tests.

[Annex III, 3\(d\)](#)

Emotion recognition on students

Prohibited

Inferring emotions in education institutions is prohibited, alongside the workplace.

A narrow medical or safety exception exists in the same point.

[Article 5\(1\)\(f\)](#)

What gets misread here

“Engagement detection helps students, so it is a legitimate wellbeing feature.”

Inferring emotions in education institutions is prohibited. The provision's only exception is use put in place for medical or safety reasons.

[Article 5\(1\)\(f\)](#)

“Adaptive learning just personalises content, so nothing attaches.”

Point 3(b) reaches evaluating learning outcomes including where those outcomes steer the learning process — which is what adaptive systems do.

[Annex III, 3\(b\)](#)

Commission guidance that reads on this sector. Guidelines on prohibited AI practices (February 2025), interpreting Article 5. These are not binding, but they are the Commission's own reading and the first place a supervisory authority will look.

Public sector and law enforcement

The densest part of the Act. Four Annex III areas and most of the Article 5 prohibitions are aimed at public authorities, and several of the bans apply to law enforcement specifically rather than to everyone.

2 · Prohibited

4 · High-risk

What the Act names here

AI systems intended to be used by public authorities or on behalf of public authorities to evaluate the eligibility of natural persons for essential public assistance benefits and services

[Annex III, 5\(a\)](#)

the use of 'real-time' remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement

[Article 5\(1\)\(h\)](#)

AI systems intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law

[Annex III, 8\(a\)](#)

What follows from this

2 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.

4 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.

Dates that bind this sector

2 Feb 2025 Prohibited practices and AI literacy apply

2 Dec 2027 Annex III high-risk obligations apply — [date changed by the Omnibus](#)

Common use cases, classified

Predicting who will commit a crime

Prohibited

Risk assessment of a natural person committing a criminal offence, based solely on profiling or on personality traits, is prohibited.

The prohibition does not reach systems supporting a human assessment already grounded in objective, verifiable facts directly linked to criminal activity.

[Article 5\(1\)\(d\)](#)

Real-time face recognition in public spaces

Prohibited

Real-time remote biometric identification in publicly accessible spaces for law enforcement is prohibited as a rule.

Article 5(1)(h) lists exhaustive exceptions — targeted search for victims, an imminent threat to life, and locating a suspect for certain serious offences — each subject to authorisation.

[Article 5\(1\)\(h\)](#)

Retrospective biometric identification

High-risk

Remote biometric identification systems are high-risk under point 1(a); only the real-time law-enforcement use is prohibited.

[Annex III, 1\(a\)](#)

Public benefits eligibility and fraud

High-risk

Point 5(a) covers evaluating eligibility for essential public assistance benefits and services, and granting, reducing, revoking or reclaiming them.

[Annex III, 5\(a\)](#)

Visa, asylum and border risk assessment

High-risk

Point 7 covers risk assessment, examination of asylum, visa and residence applications, and identification in the migration context.

[Annex III, 7](#)

Assisting judicial decisions

High-risk

Point 8(a) covers systems assisting a judicial authority in researching and interpreting facts and law and applying the law to a concrete set of facts.

Ancillary administrative activity is outside the point; the target is the decision-making itself.

[Annex III, 8\(a\)](#)

What gets misread here

“Facial recognition is banned, full stop.”

The prohibition is specific: real-time, remote, in publicly accessible spaces, for law enforcement — subject to listed exceptions. Retrospective identification is high-risk under Annex III point 1(a), not prohibited.

[Article 5\(1\)\(h\)](#)

“Any crime-risk tool is prohibited.”

The ban covers assessing the risk of a person committing an offence based solely on profiling or personality traits. It does not reach systems supporting a human assessment already grounded in objective, verifiable facts linked to criminal activity.

[Article 5\(1\)\(d\)](#)

Commission guidance that reads on this sector. Guidelines on prohibited AI practices (February 2025), interpreting Article 5. These are not binding, but they are the Commission’s own reading and the first place a supervisory authority will look.

Retail, marketing and media

Little in retail is high-risk — recommendation and personalisation are not named in Annex III at all. What does bite here are the manipulation prohibitions and, from the 2026 amendment onward, the duty to mark synthetic content.

3 · Prohibited

1 · Transparency duty

1 · No specific duty

What the Act names here

an AI system that deploys subliminal techniques beyond a person's consciousness or purposefully manipulative or deceptive techniques

Article 5(1)(a)

exploits any of the vulnerabilities of a natural person or a specific group of persons due to their age, disability or a specific social or economic situation

Article 5(1)(b)

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

Article 50(2)

What follows from this

3 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.

1 transparency duty — disclosure and marking

Article 50 applies whatever the risk tier: people must be told when they are interacting with an AI system, and synthetic audio, image, video and text must be marked in a machine-readable format and detectable as artificially generated. In force since 2 August 2026.

Dates that bind this sector

2 Feb 2025 Prohibited practices and AI literacy apply

2 Aug 2026 Transparency obligations and the general application date

Common use cases, classified

Product recommendation and personalisation

No specific duty

Recommendation is not named in Annex III and is not a safety component, so no high-risk obligation attaches under the Act.

Consumer, data protection and platform law still apply; this answers the AI Act question only.

[Article 6](#)

Persuasion below conscious awareness

Prohibited

Subliminal, purposefully manipulative or deceptive techniques that materially distort behaviour and cause significant harm are prohibited.

[Article 5\(1\)\(a\)](#)

Targeting based on age, disability or hardship

Prohibited

Exploiting vulnerabilities due to age, disability or a specific social or economic situation, so as to distort behaviour and cause significant harm, is prohibited.

[Article 5\(1\)\(b\)](#)

AI-generated marketing content

Transparency duty

Providers of systems generating synthetic content must mark outputs in a machine-readable format detectable as artificially generated.

[Article 50\(2\)](#)

Inferring sensitive traits from biometrics

Prohibited

Biometric categorisation that infers race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation is prohibited.

Categorisation on other attributes is high-risk under Annex III point 1(b) rather than prohibited.

[Article 5\(1\)\(g\)](#)

What gets misread here

“Personalisation is profiling, so it must be high-risk.”

Recommendation and personalisation are not named in Annex III. The AI Act imposes no high-risk obligation on them — though consumer, data protection and platform law still apply.

[Annex III](#)

“Persuasive design is just marketing.”

Where techniques are subliminal, purposefully manipulative or deceptive, materially distort behaviour and cause significant harm, they are prohibited outright.

[Article 5\(1\)\(a\)](#)

Commission guidance that reads on this sector. Guidelines on prohibited AI practices (February 2025), interpreting Article 5. Guidelines on the definition of an AI system (February 2025), interpreting Article 3, Article 6. Guidelines on transparency obligations for providers and deployers (2025), interpreting Article 50. These are not binding, but they are the Commission’s own reading and the first place a supervisory authority will look.

Manufacturing and infrastructure

Industrial AI is reached through the product-safety route rather than Annex III — and this is the sector the 2026 amendment changed most, by narrowing what counts as a safety component and moving machinery to a sectoral approach.

2 · High-risk

1 · Depends on the detail

What the Act names here

Critical infrastructure: AI systems intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, or in the supply of water, gas, heating or electricity

[Annex III, 2](#)

AI systems that are solely used for non-safety related aspects of user assistance, performance optimisation, service efficiency, automation or convenience or quality control shall not qualify as safety components

[Article 6\(1a\)](#)

What follows from this

2 high-risk — the full Chapter III set, from 2 Aug 2028

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.

Dates that bind this sector

2 Aug 2028 Annex I product-safety high-risk obligations apply — [date changed by the Omnibus](#)

Common use cases, classified

Safety function in machinery or a product

High-risk

A safety component of a product covered by the Annex I legislation, where that product needs third-party conformity assessment, is high-risk.

The Digital Omnibus moved the machinery regulation from Section A to Section B of Annex I and inserted Article 6(1a) to (1c), which narrow what counts as a safety component.

[Article 6\(1\) with Annex I](#)

Managing critical infrastructure

High-risk

Point 2 names safety components in the management and operation of critical digital infrastructure, road traffic and the supply of water, gas, heating or electricity.

[Annex III, 2](#)

Predictive maintenance and quality control

Depends on the detail

Article 6(1a), inserted by the 2026 amendment, states that systems used solely for non-safety aspects such as performance optimisation, automation or quality control do not qualify as safety components.

Article 6(1b) pulls it back in where failure or malfunction would endanger health and safety.

[Article 6\(1a\)](#)

What gets misread here

“Anything on the factory floor is a safety component.”

Article 6(1a), inserted by the Digital Omnibus, excludes systems used solely for performance optimisation, automation or quality control. Article 6(1b) pulls back in anything whose failure would endanger health and safety.

[Article 6\(1a\)](#)

“Machinery AI follows the same route it did in 2024.”

The Omnibus moved the machinery regulation from Section A to Section B of Annex I, limiting the AI Act's application to those machines and relocating the requirements into the machinery regulation itself.

[Digital Omnibus](#)

Commission guidance that reads on this sector. Guidelines on the definition of an AI system (February 2025), interpreting Article 3, Article 6. These are not binding, but they are the Commission's own reading and the first place a supervisory authority will look.

Technology and software vendors

If you build the model or ship the system, obligations attach to you as provider — and general-purpose models carry their own regime in Articles 53 and 55, separate from the high-risk tiers entirely.

- 1 · Transparency duty
- 1 · No specific duty
- 2 · Depends on the detail

What the Act names here

draw up and keep up-to-date the technical documentation of the model, including its training and testing process and the results of its evaluation

Article 53

This shall not include AI systems intended to be used for biometric verification the sole purpose of which is to confirm that a specific natural person is the person he or she claims to be

Annex III, 1(a)

What follows from this

1 transparency duty — disclosure and marking

Article 50 applies whatever the risk tier: people must be told when they are interacting with an AI system, and synthetic audio, image, video and text must be marked in a machine-readable format and detectable as artificially generated. In force since 2 August 2026.

Dates that bind this sector

2 Aug 2025	General-purpose AI model obligations apply
2 Dec 2027	Annex III high-risk obligations apply — date changed by the Omnibus

Common use cases, classified

Providing a general-purpose AI model

Depends on the detail

Providers of general-purpose AI models carry documentation, copyright-policy and training-data-summary duties under Article 53, which are separate from the high-risk regime.

A model meeting the Article 51 systemic-risk thresholds carries the further obligations in Article 55.

[Article 53](#)

Biometric verification (1:1 unlock)

Depends on the detail

Point 1(a) expressly excludes biometric verification whose sole purpose is confirming that a person is who they claim to be, so face or fingerprint unlock is not high-risk on that basis.

One-to-many identification against a gallery is a different system and is caught by point 1(a).

[Annex III, 1\(a\)](#)

Building a synthetic media tool

Transparency duty

As provider you must mark outputs machine-readably as artificially generated or manipulated.

Article 5(1)(ba) and (bb), inserted in 2026, prohibit systems generating intimate imagery without consent or child sexual abuse material — those are bans, not disclosure duties.

[Article 50\(2\)](#)

Developer and productivity assistants

No specific duty

Internal productivity tooling is not named in Annex III and is not a safety component.

Article 4 still requires measures supporting AI literacy for staff using it, whatever the tier.

[Article 6](#)

What gets misread here

“Face unlock is biometric identification, so it is high-risk.”

Point 1(a) expressly excludes verification whose sole purpose is confirming someone is who they claim to be.

One-to-many identification against a gallery is a different system and is caught.

[Annex III, 1\(a\)](#)

“We fine-tuned someone else's model, so they remain the provider.”

Substantial modification, putting your name on the system, or changing its intended purpose makes you the provider for the purposes of the Act.

[Article 25](#)

Commission guidance that reads on this sector. Guidelines on prohibited AI practices (February 2025), interpreting Article 5. Guidelines on the definition of an AI system (February 2025), interpreting Article 3, Article 6. Guidelines on transparency obligations for providers and deployers (2025), interpreting Article 50. These are not binding, but they are the Commission's own reading and the first place a supervisory authority will look.

PART V

The work, in order

31 steps, sequenced as the work actually runs rather than as the Regulation is numbered. Each cites the provision that requires it and names who owns it.

-
- 01 Establish what you hold
 - 02 Screen against the prohibitions
 - 03 Classify each system
 - 04 Build the high-risk requirements
 - 05 Meet the deployer duties
 - 06 Clear the route to market
 - 07 Operate, monitor and keep it current
-

Establish what you hold, and screen against the prohibitions

01 Establish what you hold

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.

- ☐ **Build a register of every AI system you develop, buy, embed or reach through an API**
Include the intended purpose, the business process it sits in, whether output reaches people outside the organisation, and the vendor. Features added to tools already in use are the most commonly missed category, because no procurement decision marked their arrival.
[Article 3\(1\)](#) · Whoever owns technology risk, with procurement and engineering

- ☐ **Fix your role for each system: provider, deployer, importer or distributor**
Duties attach to roles, not to company types, and one organisation is routinely several at once. A deployer becomes a provider if it puts its own name on a system, substantially modifies it, or changes its intended purpose, which is the single most common scoping mistake.
[Article 3\(3\)-\(7\)](#) · [Article 25](#) · Legal, on engineering's description of what was changed

- ☐ **Confirm territorial scope for each system**
You are in scope if you place a system on the EU market wherever you are established, or if the output is used in the Union. Being outside the EU is not itself an exclusion. Military, defence, national security and purely personal non-professional use are excluded.
[Article 2](#) · Legal

02 Screen against the prohibitions

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.

- ☐ **Test every system against all ten Article 5 prohibitions**
Eight have been unlawful since 2 February 2025. Two more, non-consensual intimate imagery and child sexual abuse material, apply from 2 December 2026. Any material listing eight prohibitions describes the superseded text.
[Article 5](#) · [Prohibited practices guide](#) · Legal, with the system owner · due 2 Dec 2026

- ☐ **Where a carve-out is relied on, record which one and why it applies**
Several prohibitions have narrow statutory qualifications: notably real-time remote biometric identification for law enforcement. Relying on a carve-out is a documented legal position, not an assumption.
[Article 5\(2\)-\(7\)](#) · Legal
Applies if you rely on any Article 5 carve-out

Classify each system

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.

☐ Check the Annex I product-safety route

A system is high-risk if it is a safety component of a product covered by the Union harmonisation legislation in Annex I, or is itself such a product, and requires third-party conformity assessment under that legislation.

[Article 6\(1\)](#) · [Annex I](#) · Product compliance · due 2 Aug 2028

☐ Check the Annex III use-case route

Eight areas: biometrics, critical infrastructure, education, employment, essential services, law enforcement, migration and border control, and the administration of justice. This route carries the earlier deadline.

[Article 6\(2\)](#) · [Annex III](#) · Legal, with the business owner · due 2 Dec 2027

☐ If you rely on the Annex III derogation, document and register the assessment

A system in an Annex III area is not high-risk if it does not pose a significant risk of harm, on the grounds listed in Article 6(3). Relying on this requires a documented assessment before placing on the market, and registration in the EU database.

[Article 6\(3\)-\(4\)](#) · [Article 49\(2\)](#) · Provider

Applies if you conclude an Annex III system is not high-risk

☐ Check Article 50 separately, in every tier

Transparency duties attach regardless of risk tier wherever a system interacts with people, generates synthetic audio, image, video or text, produces deep fakes, or infers emotion or biometric categories. A minimal-risk system can still owe them.

[Article 50](#) · [Article 50 guide](#) · Product, with legal · due 2 Aug 2026

☐ If you place a general-purpose AI model on the market, assess systemic risk

General-purpose AI models are regulated in their own right under Chapter V, separately from systems built on them. Assess whether the model meets the systemic-risk threshold, and notify the Commission if it does.

[Article 51](#) · [Article 52](#) · Model provider

Applies if you place a general-purpose AI model on the EU market

Build the high-risk requirements

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.

- ☐ **Establish a documented, continuous risk management system**
 Not a one-off assessment: a process that runs across the lifecycle, identifying and evaluating known and reasonably foreseeable risks, and adopting mitigation measures. It has to be documented and maintained.
[Article 9](#) · Provider · due 2 Dec 2027

- ☐ **Apply data governance to training, validation and testing data**
 Covers design choices, collection and origin, preparation, assumptions, availability and suitability, and examination for bias. Article 10(5) permits processing special categories strictly for bias detection and correction, subject to safeguards, and Article 4a broadened that basis.
[Article 10](#) · [GDPR interplay](#) · Data and ML engineering, with privacy · due 2 Dec 2027

- ☐ **Produce technical documentation to the Annex IV specification**
 It must exist before the system is placed on the market and be kept current. Annex IV sets the minimum contents; this is the document a market surveillance authority asks for first.
[Article 11](#) · [Annex IV](#) · Provider · due 2 Dec 2027

- ☐ **Enable automatic logging over the system's lifetime**
 Logs must allow traceability appropriate to the intended purpose. Retention is at least six months unless other Union or national law requires longer.
[Article 12](#) · Engineering · due 2 Dec 2027

- ☐ **Write instructions for use that let a deployer meet its own duties**
 Including the intended purpose, accuracy and robustness levels, known limitations, foreseeable misuse, the human oversight measures required, and expected lifetime and maintenance.
[Article 13](#) · Provider · due 2 Dec 2027

- ☐ **Design human oversight into the system, not around it**
 The person overseeing must be able to understand capacities and limits, monitor for anomalies, resist automation bias, correctly interpret output, decide not to use it, and intervene or stop. A review step that cannot actually override the system does not satisfy this.
[Article 14](#) · [Human oversight guide](#) · Product and engineering · due 2 Dec 2027

- ☐ **Set and declare accuracy, robustness and cybersecurity levels**
 Declared metrics go in the instructions for use. Robustness covers resilience to errors and inconsistencies; cybersecurity covers attempts to alter use, behaviour or performance, including data poisoning and adversarial examples.
[Article 15](#) · Engineering and security · due 2 Dec 2027

- ☐ **Put a quality management system in place**
 A documented system covering regulatory compliance strategy, design and verification procedures, data management, the risk management system, post-market monitoring, incident reporting and accountability.
[Article 17](#) · Provider · due 2 Dec 2027

Meet the deployer duties, and clear the route to market

05 Meet the deployer duties

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.

- ☐ **Use the system in accordance with its instructions, and assign competent human oversight**
Oversight must be assigned to named people with the competence, training and authority to exercise it, and the authority matters as much as the training.
[Article 26\(1\)-\(2\)](#) · Deployer · due 2 Dec 2027

- ☐ **Monitor operation, keep logs, and report serious incidents**
Suspend use and inform the provider where you have reason to consider that use in accordance with the instructions may present a risk. Keep logs for at least six months.
[Article 26\(5\)-\(6\)](#) · [Article 73](#) · Deployer · due 2 Dec 2027

- ☐ **Inform affected people**
Workers and their representatives must be informed before a high-risk system is put into use in the workplace. People subject to a decision informed by an Annex III system must be told the system was used.
[Article 26\(7\), 26\(11\)](#) · HR and the business owner · due 2 Dec 2027

- ☐ **Complete a fundamental rights impact assessment where required**
Required for deployers that are public bodies, or private entities providing public services, and for certain creditworthiness and life or health insurance pricing uses.
[Article 27](#) · Deployer, with legal · due 2 Dec 2027
Applies if you are a public body, provide public services, or score credit or insurance risk

06 Clear the route to market

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.

- ☐ **Determine the conformity assessment route**
Annex VI internal control, or Annex VII involving a notified body. For most Annex III systems internal control is available where harmonised standards were applied in full; biometrics is the notable exception.
[Article 43](#) · [Conformity assessment guide](#) · Provider · due 2 Dec 2027

- ☐ **Draw up the EU declaration of conformity and affix CE marking**
The declaration is kept for ten years after placing on the market and made available to authorities on request.
[Article 47](#) · [Article 48](#) · [Annex V](#) · Provider · due 2 Dec 2027

- ☐ **Register in the EU database before placing on the market**
Applies to providers of Annex III high-risk systems, to those relying on the Article 6(3) derogation, and to certain public-authority deployers.
[Article 49](#) · [Article 71](#) · Provider · due 2 Dec 2027

- ☐ **Appoint an authorised representative in the Union**
Required for providers established outside the Union, by written mandate, before placing a high-risk system on the market.
[Article 22](#) · Provider · due 2 Dec 2027
Applies if you are established outside the EU

Operate, monitor and keep it current

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

☐ Run post-market monitoring against a documented plan

Proportionate to the nature and risks of the system: collect, document and analyse performance data across the lifetime, and feed it back into the risk management system.

[Article 72](#) · Provider · due 2 Dec 2027

☐ Establish a serious incident reporting path that meets the deadlines

Report to the market surveillance authority of the Member State where the incident occurred. The deadlines are short and vary with the nature of the incident, so the path has to exist before it is needed.

[Article 73](#) · Provider, with incident response · due 2 Dec 2027

☐ Mark synthetic content in a machine-readable format

Providers of generative systems must mark output as artificially generated or manipulated. Systems already on the market when the transparency rules began have until 2 December 2026 to comply.

[Article 50\(2\)](#) · Provider · due 2 Dec 2026

Applies if your system generates synthetic audio, image, video or text

☐ Take measures to support AI literacy, and record what you did

Applies to every provider and deployer regardless of risk tier, and has since 2 February 2025. The Digital Omnibus softened this from ensuring a level of literacy to taking measures to support it, which changes what you should be recording, not whether you owe it.

[Article 4 · AI Literacy Program](#) · HR and the business owner · due 2 Feb 2025

☐ Re-check guidance published before 27 July 2026

Regulation (EU) 2026/1744 changed deadlines, added two prohibitions, softened Article 4 and created new AI Office enforcement powers. Deadlines and prohibition counts in older material are likely to be wrong.

[Digital Omnibus changes](#) · Whoever maintains your compliance documentation

PART VI

Enforcement

Three supervisors, three ceilings and nine factors. The Omnibus moved part of the market from national authorities to the AI Office.

01 Who supervises what

02 How an amount is arrived at

Who supervises what, after the Omnibus

Regulation (EU) 2026/1744 added Articles 75a to 75d, giving the AI Office direct investigative and fining powers over part of the market that previously answered only to national authorities.

ARTICLES 74, 99

Ceiling

€35m / 7%, €15m / 3% or €7.5m / 1% depending on the breach

National market surveillance authorities

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.

- Full market surveillance powers under Regulation (EU) 2019/1020
- Require corrective action, withdrawal or recall
- Impose the administrative fines Member States set within the Article 99 ceilings

ARTICLES 75, 75A–75D, 101

Ceiling

Article 99(3)–(7) amounts apply correspondingly, plus periodic penalties up to 5% of average daily turnover per day

The AI Office

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.

- Open investigations on its own initiative or on a complaint
- Require information by simple request or by binding decision
- Carry out remote and on-site inspections, and seal premises, books and records for the duration of an inspection
- Accept binding commitments and close proceedings, with reopening possible
- Impose periodic penalty payments until compliance is restored

ARTICLE 100

Ceiling

€1.5m for a prohibited practice, €750,000 for other breaches

European Data Protection Supervisor

Fines against Union institutions, bodies, offices and agencies that fall within the scope of the Regulation.

- Impose administrative fines on Union bodies
- Take account of the same Article 99(7)-style factors

What stays with national authorities

The AI Office's remit under Article 75(1) is exclusive but not general. These are carved back out of it, and answer to the Member State authority as before — so one provider can face two supervisors for two products.

- Systems covered by the Annex I, Section A harmonisation legislation
- Certain Annex III high-risk systems
- Law enforcement, border management and financial institution systems falling under Article 74(6)
- Systems used in the administration of justice

How an amount is arrived at

ARTICLE 99(7)

The ceilings are the headline, but they are ceilings. These are the factors an authority is directed to weigh in deciding where beneath one an amount lands.

ARTICLE 99(6A)

Added by the Omnibus: small mid-cap companies now join SMEs and start-ups in facing the **lower** of the fixed sum and the turnover percentage, rather than the higher.

The nine factors

- The nature, gravity and duration of the infringement and of its consequences, taking into account the purpose of the system and, where appropriate, the number of people affected and the level of damage
- Whether the same operator has already been fined by other market surveillance authorities for the same infringement
- The size, annual turnover and market share of the operator
- Any other aggravating or mitigating factor applicable to the circumstances, such as financial benefits gained or losses avoided
- The degree of cooperation with the national competent authorities
- The degree of responsibility of the operator, taking into account the technical and organisational measures it implemented
- The manner in which the infringement became known to the authority: in particular whether the operator notified it, and if so to what extent
- Whether the infringement was intentional or negligent
- Any action taken to mitigate the harm suffered by the affected persons

And the one that is not a ceiling

Article 75c lets the AI Office levy a periodic penalty of up to 5% of average daily worldwide turnover for every day non-compliance continues. It is not capped by the Article 99 figures, and it compounds. The arithmetic at the ceiling:

ANNUAL TURNOVER	PER DAY	30 DAYS	90 DAYS	ONE YEAR
€100m	€14k	€411k	€1.2m	€5.0m
€1bn	€137k	€4.1m	€12m	€50m
€10bn	€1.4m	€41m	€123m	€500m

Arithmetic on the Article 75c formula at its maximum, not a prediction of what any authority would impose.

Read factors seven and eight together

How the infringement came to the authority's attention, and whether it was intentional or negligent, are separate factors that point the same way: an organisation that finds and reports its own breach is being assessed on a different basis from one that was caught. That is the practical case for the post-market monitoring work in Phase 7 of the checklist, independent of whether anyone ever asks for it.

How this document was verified

PRIMARY SOURCES

Regulation (EU) 2024/1689 (OJ L, 12.7.2024)

Regulation (EU) 2026/1744 (OJ L 2026/1744, 24.7.2026)

Consolidated text as at 27 July 2026

COMMISSION GUIDANCE

Guidelines on prohibited AI practices (February 2025)

Guidelines on the definition of an AI system (February 2025)

Guidelines on transparency obligations for providers and deployers (2025)

The legal text behind this playbook is held as 119 articles, 180 recitals and 14 annexes, each retrieved from the Publications Office and compared against the Official Journal character for character. The comparison is part of the build, not a one-time exercise: a quotation that no longer matches, or a citation to a provision that no longer says what is claimed, fails the build rather than printing.

That matters most for the negative claims. Several pages here tell you something is **not** caught — that Annex III does not name recommendation engines, that motor insurance sits outside point 5(c), that clinical transcription is not listed. A reader acts on a negative by doing nothing, so each is tested against the current text of the annex on every build rather than taken on trust.

What is ours, and what is the Act's

Everything in quotation marks is the Regulation's own wording. Everything else — the ordering, the classifications, the commentary on what gets misread — is Holistic AI's reading of it, and is offered as information rather than legal advice. Where the two disagree, the Official Journal governs.

What the build checks

CHECKED ON EVERY BUILD	COUNT
Provisions held and compared against the Official Journal	313
Quotations matched character for character to the provision they cite	21
Corrections to common misreadings, each tied to a provision	18
Use-case classifications, each citing the point that decides it	38
Checklist steps, each citing the article that requires it	31
Application dates, traced to Article 113 as amended	8

The full text, article by article with the amended wording marked, is at [euaiact.com](#), together with the scope finder, the risk classifier and the 31-step checklist as working tools.



Who maintains this, and why

euaiact.com is run by Holistic AI as an independent reference. The full legal text with the Omnibus amendments marked, the scope finder, the risk classifier, the checklist and all eight sector guides are free and ungated there — this playbook is generated from the same verified dataset that powers them.

What the company does

Holistic AI is an enterprise AI governance platform — a Leader in Avasant's 2025 Responsible AI Platforms RadarView — and one place to discover, assess, mitigate, monitor and govern risk across every AI agent, model and application an organisation runs. The problem it addresses is the one this playbook keeps running into — you cannot classify what you have not found, and most enterprises learn about a new AI system when something breaks or a regulator asks.

See every system

Automated discovery across code, cloud and documentation, including shadow AI, through 15+ integrations — GitHub, GitLab, AWS, Azure, GCP, Databricks, ServiceNow — into one inventory. Phase 1 of the checklist in Part V, done continuously rather than once.

Evidence as a by-product

Coverage for the EU AI Act, NIST AI RMF and ISO 42001, with full audit trails and version history, so the gap register and the classification record this playbook asks you to keep are produced by the work rather than reconstructed for an audit.

Independently recognised

- Named a Cool Vendor for AI Security by Gartner®
- Vendor in the 2024 Gartner® Innovation Guide for Generative AI in Trust, Risk and Security Management
- Vendor in the Gartner® Hype Cycle™ for Digital Government Services, 2024
- IDC ProductScape for Worldwide Generative AI Governance Platforms, 2025
- Leader in Avasant's 2025 Responsible AI Platforms RadarView
- Recognized among the AI Breakthroughs of 2025

Used by

AON · Allegis · eBay · GE Healthcare · GSK · Johnson Controls · ManpowerGroup · MAPFRE · MindBridge · PMI · Siemens · SLB · Starling Bank · UKG · Unilever · Publicis

Platform assurances: SOC 2 Type 2, ISO/IEC 27001 certified.

Check your own exposure

This playbook tells you what the Act says. It cannot tell you what you hold, which is the only question that decides what you owe. Two ways to answer it — one you can do yourself in a few minutes, one we run with you.

FREE, NO SIGN-UP

See which obligations already bind you

Work through your systems against the prohibitions in Article 5, the two high-risk routes in Article 6, the Article 50 transparency duties and the general-purpose model thresholds — and get back which of your obligations are live today and which are still ahead, each against the amended date. The 38 classified use cases in Part IV are the same catalogue it matches against.

[Check what applies to you →](#)

euaiact.com/eu-ai-act-readiness

WITH HOLISTIC AI

The EU AI Act readiness assessment

Where the self-service check tells you which tier a system falls in, the assessment produces the evidence you would have to show for it. It runs in four steps:

- | | |
|---|--|
| 01 Scope — what you actually operate | 02 Classify — which tier each system falls in |
| 03 Gap — what the Act requires that you do not yet have | 04 Sequence — what has to be true, and by when |

What you end up holding

- A written classification for every system assessed, with the deciding provision cited
- A gap register mapped to the specific articles that create each obligation
- A sequenced remediation plan against the amended application dates
- Evidence you can put in front of a market surveillance authority or a customer's due-diligence team

[Start the assessment →](#)

[Get a demo →](#)

holisticai.com/eu-ai-act-readiness

holisticai.com/demo



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, 8 sector guides and the 31-step compliance checklist in this playbook.

119

articles, verified against the
Official Journal

180

recitals

14

annexes

8

sector guides

YOUR NEXT STEP

Find out where you actually stand

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 to inventory, classify and evidence AI against the EU AI Act.

Start the EU AI Act assessment →

holisticai.com/eu-ai-act-readiness

Get a demo →

holisticai.com/demo

Or check which obligations already bind you, free and without signing up: euaiact.com/eu-ai-act-readiness