Navigating the EU AI Act.



A PRACTICAL GUIDE FOR SMEs AND ORGANISATIONS

Foundations · Scope · Risk Classification · AI Literacy · Transparency · Governance · Accountability · Fundamental Rights · GPAI · Sectors · Compliance · Digital Omnibus

UNDER REGULATION (EU) 2024/1689

TARGET AUDIENCE

SMEs · Organisations · Training Providers (CVET)

Senior Management · Legal & Compliance · Industry · Stakeholders

APPLIES TO & SCOPE

Providers · Deployers · Importers · Distributors · Authorized Representatives · Product Manufacturer

EU market · AI systems by risk category

TIMELINE

In force: 1 August 2024

Phased application till 2028

EDITION

July 2026

Regulation (EU) 2024/1689

Co-funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2026, authors: PhD.Cand. Evita Milan (Steinbeis Transfer Hub Berlin) · Kaspars Mezerins (RTU Riga Business School) · Elmaara-Aiga Timma (Latvian IT Cluster / EDIH)

This version was published in August2026
For questions or comments, please contact: evita.milan@steinbeis.de

Steinbeis Transfer-Hub Berlin

Latvian Digital Innovation Hub
POWERED BY LATVIAN IT CLUSTER

RBS Riga Business School | RTU

This guide is provided for informational purposes only and does not constitute legal advice. References to Regulation (EU) 2024/1689 reflect the text in force as of publication date. Confirm specific obligations with qualified counsel.

Foreword

Organisations across Europe are facing a rapidly evolving digital landscape shaped by artificial intelligence and new regulatory requirements. While the EU AI Act creates important obligations for those developing or using AI systems, many organisations lack the resources and expertise needed to implement these requirements effectively. This guide helps organisations understand the practical implications of the EU AI Act, strengthen AI literacy, and adopt AI responsibly through accessible, practice-oriented guidance.

Objectives of this guide closely align with the EDIH mission of helping organisations build digital capabilities, adopt innovative technologies, and advance their digital transformation. Supporting the responsible and trustworthy adoption of AI, while strengthening the skills needed to comply with emerging European regulations, is an increasingly important part of that mission.

The EU AI Act is not only a regulatory framework, but also an opportunity to strengthen organisational capabilities, build trust, and create long-term value from AI adoption. Rather than being viewed as an administrative burden, the AI Act can serve as a practical mirror that helps organisations assess their readiness, governance practices, risk management processes, and alignment with European values. By investing early in AI literacy, transparency, accountability, and responsible AI practices, organisations can improve decision-making, increase stakeholder confidence, and gain a competitive advantage.

We hope this guide will help organisations navigate the EU AI Act with confidence, strengthen AI literacy, and make the most of the opportunities offered by digital transformation.



Aiga Irmeja
Chairperson of the Board
Latvian IT Cluster / European Digital Innovation Hub

ACKNOWLEDGEMENTS

Expressing our gratitude to those involved

We would like to thank everyone who helped us along the way in preparing, validating and testing this guide to ba as practical as possible for companies and organizations.

Our thanks go to Dr.Claudio Rivera, Dr.Edgars Voļskis and Jānis Paksis from Riga Business School for their comments and suggestions, which helped us improve the guide and direct it towards company perspective.

A special thank you goes to Dipl. Jur. Greta Sparzynski for her professional lawyer´s perspective and contribution to fostering a better understanding and better “translation” of this legislation.

We would like to extend our gratitude to the Latvian IT community, the Latvian IT Cluster, and its members as well as the National Artificial Intelligence Centre of Latvia for their support and engagement throughout the development of this guide.

We are also grateful to Dr.Jack Burgess for English and Frank Graage for German language proofreeding and improvements.

Finally, we would like to thank all those who supported us and participated in the workshops during the preparation of this guide. We appreciate your time, encouragement, and willingness to share your thoughts.

EXECUTIVE SUMMARY

The world's first comprehensive AI law is in force. *No business is too small.*

Most organisations will not need all of this guide. The difficulty is that you cannot know which parts apply to you until you have worked through the first three chapters - and until then, every obligation in the EU Artificial Intelligence Act (Regulation (EU) 2024/1689) looks equally urgent. Throughout this guide, the terminology used follows the AI Act.

Chapters 1 and 2 start with two decisive questions: whether your software meets the legal definition of an AI system, and whether you are a provider, deployer, importer, distributor, product manufacturer, authorised representative or affected person for each system you build or use. Your role is one of the key factors determining your obligations. The applicable requirements also depend on the type and intended purpose of the AI system or model, its classification under the Act, and, for certain procedural requirements and support measures, the size of the organisation. Chapter 3 then places each system in one of four risk tiers. Almost everything downstream follows from that risk classification - the cornerstone of the Act.

Chapters 4 to 7 set out the duties that attach once the tier is known: AI literacy and enforcement, transparency and record-keeping, general-purpose AI, and the surrounding EU digital rulebook. Chapters 8 to 10 turn those duties into working practice - a vendor due-diligence checklist for deployers, a six-step documentation roadmap for providers, and sector-specific guidance for HR, finance, education and healthcare.

Chapter 11 covers the changes introduced by the AI Digital Omnibus package, which entered into force in July 2026, which postponed the high-risk deadlines in law and extended organisations reliefs to companies under 750 employees.

IN FORCE

Since August 1

2024

Application phased through 2027 (or 2028 after Omnibus) across risk tiers and provider categories.

ARE SMES EXCLUDED FROM EU AI ACT OBLIGATIONS?

NO no exemption

SMEs receive lighter procedural treatment - not exemption from substantive requirements.

PENALTY BAND FOR prohibited AI practices

€35M or 7%turnover

€15M / 3% – other AI Act violations
€7.5M / 1% – misleading information to authorities

WHAT THIS MEANS IN PRACTICE

If your organisation develops, uses, imports, or distributes AI systems within the EU market, full obligations apply regardless of headcount or turnover/profit. What SMEs do receive is a lighter procedural framework - simplified documentation, priority access to regulatory sandboxes, and proportionate fee structures.



"The EU's AI Act is the first-ever comprehensive legal framework on Artificial Intelligence worldwide. So, this is a historic moment. The AI Act transposes European values to a new era." — Ursula von der Leyen, President of the European Commission

© European Union, 2026, CC BY 4.0, via Wikimedia Commons

This guide is provided for informational purposes only and does not constitute legal advice. References to Regulation (EU) 2024/1689 reflect the text in force as of publication date. Confirm specific obligations with qualified counsel.

CONTENTS

Inside this report

Foreword 03
Acknowledgements 04
Executive Summary 05

CHAPTER 01

01 Why organisations need to care about the EU AI Act **08**
1.1 Is your software "AI" - or not? 10
1.2 General-purpose AI Model vs AI System 11

CHAPTER 02

02 Does the Act apply to your business? **12**
2.1 Scope & Applicability 13
2.2 Obligation matrix 13
2.3 Roles, responsibility and liability 16
2.4 SME-specific provisions 17

CHAPTER 03

03 Understanding the risk-based framework **18**
3.1 The risk-based framework 19
3.2 Prohibited practices 20
3.3 High-risk AI systems 21
3.4 Provider & deployer duties 22
3.5 Limited risk: transparency 23
3.6 Minimal risk 24
3.7 Literacy 24

CHAPTER 04

04 AI literacy, audits & the regulatory ecosystem **26**
4.1 AI literacy & enforcement 27
4.2 Practical compliance 28
4.3 Audits and enforcement 29
4.4 Penalties & Fines 30

CHAPTER 05

05 Transparency and information requirements **31**
5.1 General Transparency Duties 32
5.2 Content labelling 33
5.3 Best practise 30
5.4 AI Value Chain Responsibilities 35

CONTENTS · CONTINUED

CHAPTER 06

06 General-purpose AI Models (GPAI) 36

6.1 What counts as GPAI 37
6.2 The GPAI layer structure 38
6.3 Use-case lens 39

CHAPTER 07

07 Other regulations influencing the EU AI Act 40

7.1 The wider ecosystem 41
7.2 GDPR 42
7.3 Product safety 43
7.4 Platforms & critical infrastructure 44
7.5 Regulatory sandboxes 45

CHAPTER 08

08 Practical compliance for SME deployers 46

8.1 Deployer Obligations Overview 47
8.2 Vendor Due Diligence 48

CHAPTER 09

09 Practical compliance for providers & high-risk deployers 49

9.1 Provider obligations overview 50
9.2 Technical documentation requirements 51
9.3 Risk management for small teams 52
9.4 Oversight 53
9.5 Data governance 53

CHAPTER 10

10 Sector-Specific Guidance 54

10.1 Sectors where the Act lands hardest 55
10.2 What to focus on, sector by sector 56

CHAPTER 11

11 The Digital Omnibus on AI 57

11.1 Introduction 58
11.2 New prohibitions and the safe harbour 59
11.3 Transparency, AI literacy and product rules 60
11.4 The code of practice on transparency of AI-generated content 61
11.5 Relief, governance and enforcement 62
11.6 Key takeaways, action plan, sources 63

Closing 67
Glossary 68

01

CHAPTER ONE

Why organisations need to care *about the EU AI Act.*

The EU AI Act entered into force on 1 August 2024, establishing the world's first comprehensive legal framework for artificial intelligence (AI). The Act applies to providers placing AI systems or GPAI models on the EU market or putting AI systems into service in the EU, to deployers established or located in the EU, and to certain providers and deployers established outside the EU where the output produced by the AI system is used in the EU. Importers, distributors and other operators may also be subject to specific obligations.



Compliance under the EU AI Act is an ongoing process of strategic alignment with evolving, risk-based legal obligations, not a single milestone. SMEs that treat it superficially risk penalties and reputational damage, while those that integrate compliance into product and governance decisions build trust, unlock smoother market access, and strengthen their long-term competitive position.

This Guide is developed to help organisations navigate the EU AI Act - translating its articles, annexes, and risk tiers into a practical path that small and mid-sized businesses can actually follow.

SECTIONS

1.1 The Definition Test

1.2 Model vs System

READ TIME

8 minutes

KEY ARTICLES

Art. 3(1) AI System

Art. 3 GPAI Model

01 1.1.INTRODUCTION

A single market with harmonised rules for trustworthy AI.

The EU AI Act creates a single market with harmonised rules for trustworthy AI in the EU. It establishes the world's first comprehensive legal framework for artificial intelligence - a risk-focused approach designed to protect fundamental rights and safeguard democracy, while fostering a competitive environment for innovation.

Compliance is not a static milestone but a strategic alignment of an organisation's AI portfolio with specific, risk-based legal obligations. Organisations that treat this as a tick-box exercise risk significant penalties and reputational damage; those that embrace it as a strategic driver will gain market trust and competitive advantage.

GLOBAL REACH

The reach of the EU AI Act is not confined to Europe: Pursuant to Article 2(1) of the AI Act, the Regulation applies to providers, deployers, importers, distributors, product manufacturers placing AI systems on the market under their name or trademark, authorised representatives of providers, and affected persons, as well as to certain providers and deployers established outside the EU where the output of an AI system is used in the Union.

Following the precedent of the GDPR, it has effectively become a new global benchmark for digital operations.

Equally important is recognising that compliance is not a one-off exercise. The regulatory framework is capable of evolving alongside technological developments, requiring organisations to monitor regulatory updates and reassess their AI portfolio where necessary, particularly in light of the Commission's power to amend the list of high-risk AI use cases set out in Annex III.

SMEs: not exempt, but not forgotten.

It is critical to understand that SMEs are **not exempt from legal obligations**. If you provide or deploy an AI system in the EU market, you are subject to the Act regardless of headcount - a two-person startup may be subject to the same categories of legal obligations as a large enterprise, depending on its role and the type of AI system involved.

What SMEs do receive is lighter treatment, not exemption: administrative fines are calculated as the lower of the fixed euro amount or the revenue percentage, while for large enterprises it is the higher of the two. SMEs also benefit from simplified documentation requirements, access to regulatory sandboxes, and proportionate obligations. These measures are intended to support innovation and reduce compliance burdens for smaller organisations, as reflected in Articles 62 and 63 on regulatory sandboxes and innovation support measures, as well as Article 99, which provides for awareness-raising and support activities for SMEs and start-ups.

Under the Digital Omnibus, these simplifications now extend to companies with up to 750 employees and €150 million in annual revenue.

2024	04	8+2	€35M
YEAR IN FORCE	RISK TIERS	PROHIBITED PRACTICES	MAX PENALTY BAND

SECTION 1.1 · ARTICLE 3(1)

Is your software "AI" - or not?

Before addressing compliance, every organisation must determine whether its software qualifies as an "AI System" under the Act. The Act does not apply to all automation - it draws a sharp threshold between standard software and regulated AI.

AI System (EU AI Act, Article 3 (1))
'AI system' means a **machine-based system** that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for **explicit or implicit objectives**, infers, from the input it receives, how to generate outputs such as **predictions, content, recommendations, or decisions** that can influence physical or virtual environments.

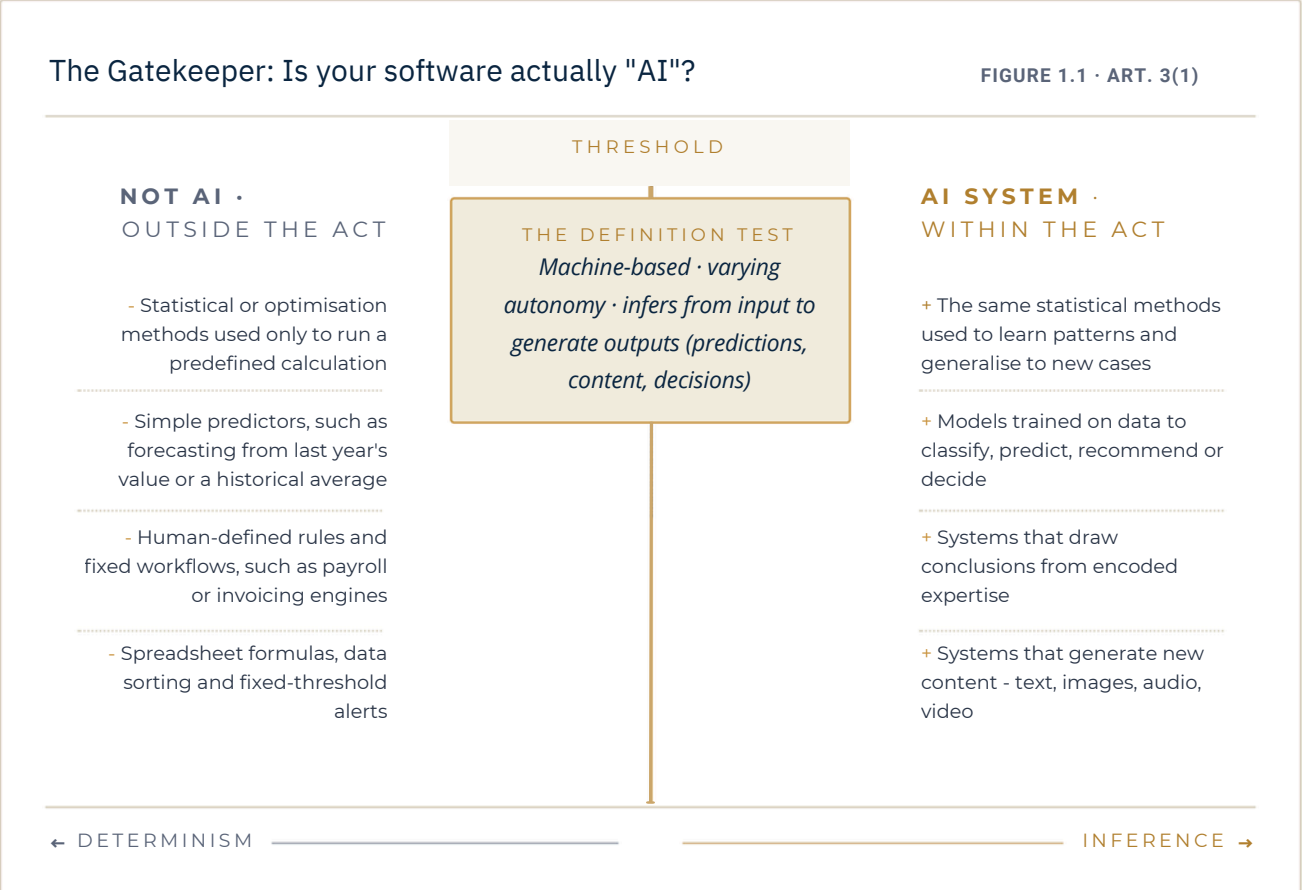


FIGURE 1.1 *The Gatekeeper Test under Article 3(1). The threshold sits at inference.*

This illustration provides a simplified overview. Under the AI Act, whether software qualifies as an AI system depends on the overall assessment of all elements of the legal definition, not solely on its inference capabilities.

According to the European Commission's Guidelines on the definition of an AI system, software intended to improve mathematical optimization, basic data processing systems, systems based on classical heuristics, and simple prediction systems **are generally not considered AI systems under the AI Act**, although each case should be assessed individually based on the specific architecture and functionality of the software.

SECTION 1.2 · ARTICLE 3

General-purpose AI Model vs AI System

Distinguishing a GPAI Model from an AI System determines which AI Act rules apply. Obligations then depend on your role, the intended use, the system's risk classification and, for GPAI models, whether they present systemic risk.

ARTICLE 3(1)

AI System

A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

Example · A customer-service chatbot built on GPT · An HR plug-in integrating an AI screening tool

ARTICLE 3 (63)

A general-purpose AI model

GPAI model, including where trained with large amounts of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks, and that can be integrated into a variety of downstream systems or applications.

Example · OpenAI's GPT-4 · Google's Gemini · Open-weight models such as Llama or particular Mistral models

The engine vs the car - a practical analogy.

THE RAW CAPABILITY
AI Model = the engine

→

THE DEPLOYED PRODUCT
AI System = the car

THE REALITY - MOST SMES BUILD AI SYSTEMS, NOT GPAI MODELS

Most SMEs do not place GPAI models on the market. Your business is almost certainly building, customising, or deploying an AI System. Fine-tuning does not by itself make you a GPAI model provider - but placing the resulting system on the market under your own name makes you the provider of that AI System.

When you integrate someone else's model, you become *responsible for the system.*

You are its provider. Your obligations depend on your role, the intended purpose and the risk classification.

ATTENTION - THE "DOWNSTREAM" TRAP

Picking up an off-the-shelf model and dropping it into a hiring workflow, a credit decision, or a medical-device function can quietly move you up the risk pyramid. System risk follows the use, not the underlying model. Obligations on the model itself sit with its provider - see Chapter 6.

Worked illustration - an SME chatbot

A — GPAI MODEL

OpenAI trains GPT from scratch. It is the raw capability. OpenAI carries GPAI provider obligations.

B — YOUR INTEGRATION

Your SME wraps GPT with a CRM connector, prompt scaffolding, guardrails, and a UI. You ship it as "ResolveBot".

C — YOUR AI SYSTEM

"ResolveBot" is now an AI System under Art. 3(1). You are the provider. Obligations follow your role and the use-case risk.

02

CHAPTER TWO

Does the Act apply to *your business?*

By the end of this chapter you will know:

- Whether the Act applies to your business at all
- Which of the six operator roles you hold, for each system
- What being an SME actually changes, and what it does not

SECTIONS
2.1 Scope · 2.2 Obligation · 2.3 Roles · 2.4 SME provisions

READ TIME
14 minutes

KEY ARTICLES
Art. 2 Scope
Art. 3(3)–3(8) Roles

02 2.1. SCOPE & APPLICABILITY

The Act catches more businesses than most organisations expect.

Obligations vary considerably depending on your role and what the AI system does. This two-step process - the in-scope checklist followed by the determination flowchart - will help you establish whether you are in scope and which role you hold.

Step 1 · The in-scope checklist

01 EU PRESENCE

Are you established or located in the EU and using an AI system under your authority?

02 MARKET ACCESS

Do you place AI systems on the EU market, or put them into service in the EU?

03 OUTPUT USED IN THE EU

Are you based outside the EU, but the output produced by your AI system is used, or intended to be used, in the EU?

04 SUPPLY CHAIN ROLE

Do you import or distribute AI systems built by non-EU providers to EU-based customers?

ANY "YES" = IN SCOPE

If you answered **"yes"** to any of the above, your business is likely in scope. Some systems and activities are excluded - the flowchart overleaf checks these. Your specific obligations depend on **your functional role in the AI value chain** rather than the size of your company.



STEP 2 · DECISION TREE

From software assessment to role identification.

AI System Compliance Assessment + AI Act Compliance Checker

FIGURE 2.1 · ART. 2

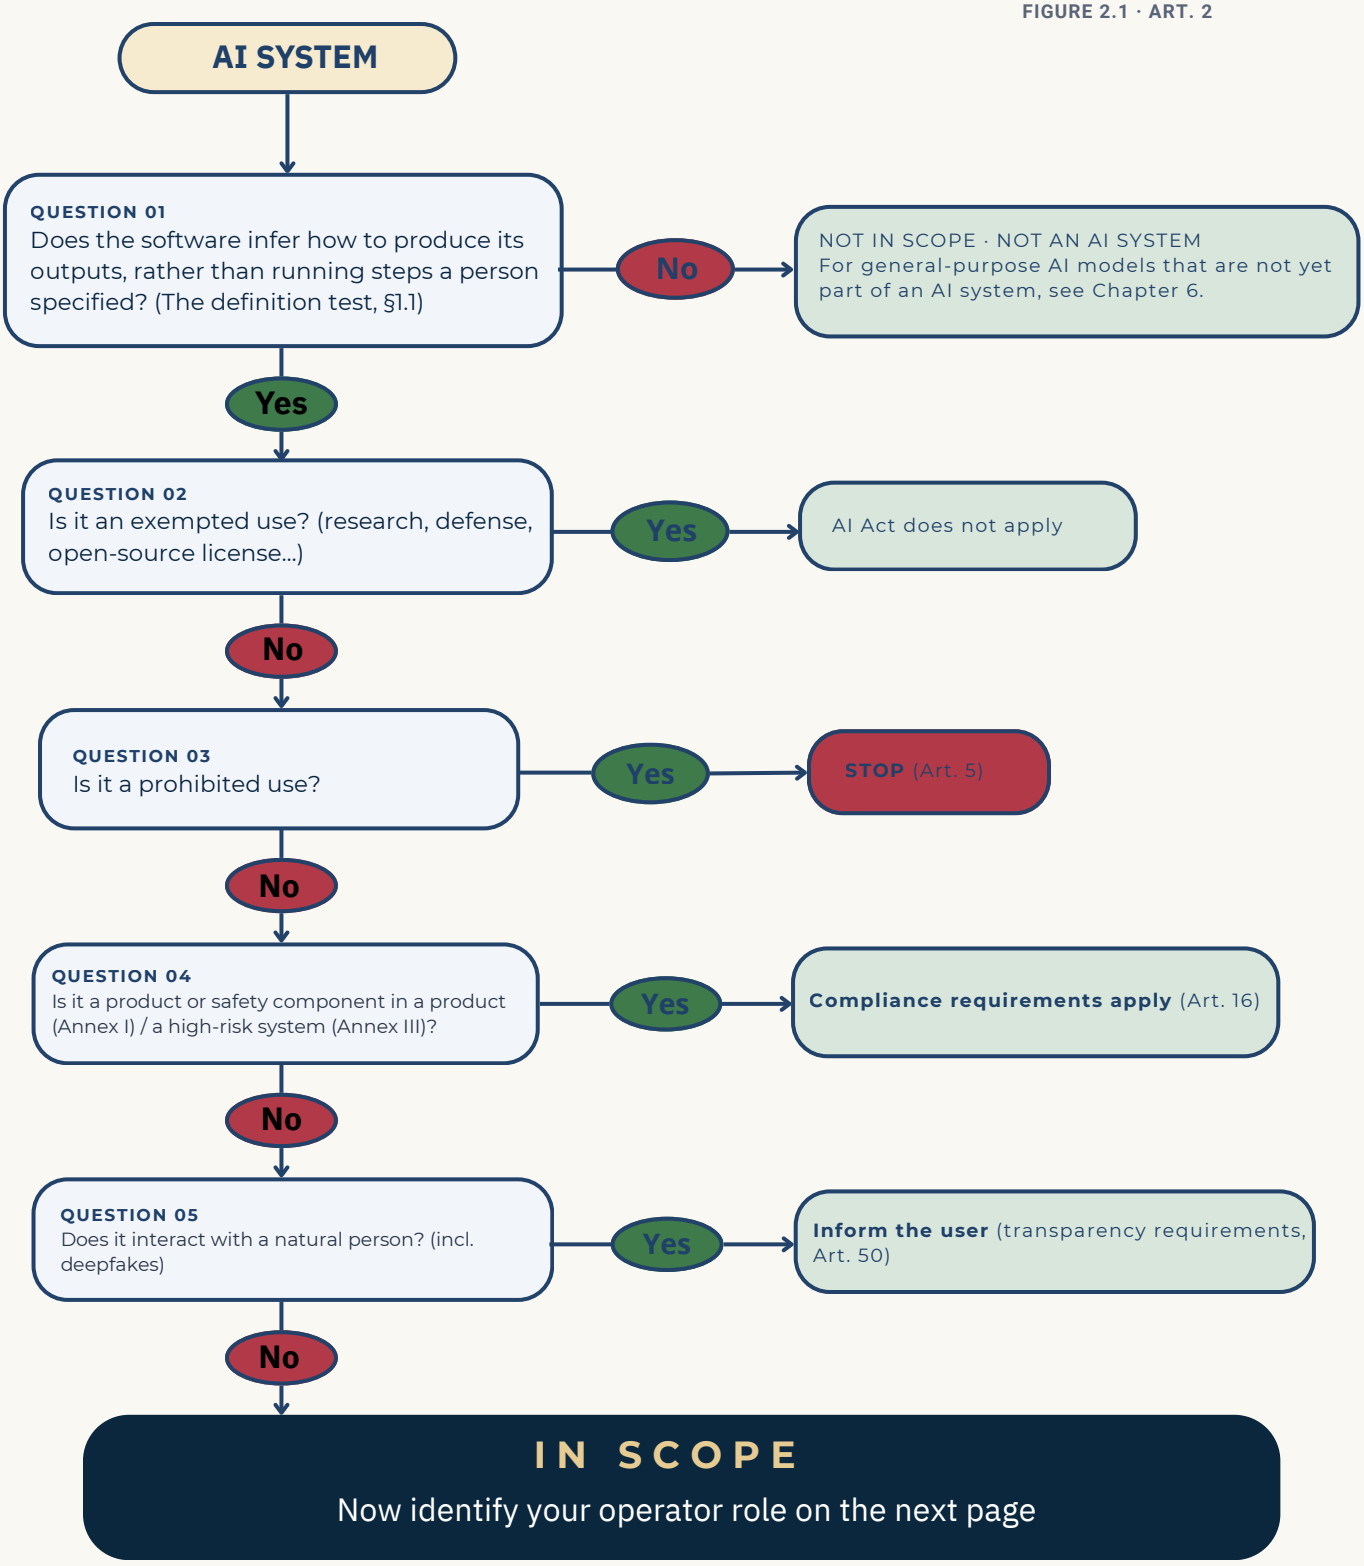


FIGURE 2.1 *Run every distinct AI system through this decision tree separately. The same SME can be a Provider for one product and a Deployer for another.*

SECTION 2.2 · OBLIGATIONS MATRIX

Roles, examples, and what each one owes.

Your duties follow your role, the system's intended purpose, and its risk tier, not your size. One SME can hold several roles, so assess each system separately. The obligations below apply mainly to high-risk systems - a limited-risk role owes far less.

ROLE	WHO YOU ARE	SME EXAMPLE	PRIMARY OBLIGATIONS
Provider ART. 3(3)	You develop, or have developed, an AI system and place it on the market or put it into service under your own name or trademark.	A software company building an AI-powered contract review tool and selling it to other businesses.	Heaviest obligations. Risk management, technical documentation, conformity assessment, CE marking, registration in the EU database.
Deployer ART. 3(4)	You use an AI system in your professional activity - under your own authority.	A recruitment agency using an AI tool to help screen job applications.	Use according to provider instructions. Human oversight, monitoring, logging, and incident reporting for high-risk use cases.
Importer ART. 3(6)	You are based in the EU and bring an AI system made by a non-EU provider onto the EU market.	An EU company licensing and reselling an AI tool built by a US vendor.	Verify conformity assessment, check documentation, identify yourself on the product. Maintain records for 10 years.
Distributor ART. 3(7)	You make an AI system available on the EU market without modifying it.	A reseller packaging and distributing a third-party AI analytics platform unchanged.	Verify CE marking and provider compliance. Cease distribution immediately if the system is found to be non-conforming.
Authorised Representative ART. 3(5) DUTIES ART. 22	You are EU-based and act on behalf of a non-EU provider to fulfil their EU compliance obligations.	An EU consultancy formally appointed by a non-EU AI vendor to manage their EU compliance.	Appointed in writing. Verify conformity documentation, retain copies 10 years, cooperate with authorities.
Product Manufacturer ART. 25(3)	You place a regulated product containing a high-risk AI safety component on the market under your own name.	A machinery maker embedding an AI safety component.	Treated as the provider of that high-risk AI system.

SECTION 2.3 - ROLES, RESPONSIBILITY AND LIABILITY

When a deployer quietly becomes *a provider.*

Article 25 sets out three situations in which a deployer legally becomes a provider - and inherits all of a provider's obligations: technical documentation, conformity assessment, registration and post-market surveillance.

SITUATION	WHAT TRIGGERS IT
1 Rebranding	Putting your own brand or trademark on a High-Risk AI system.
2 Substantial modification	Significant reworking of the system's technical architecture or code.
3 Purpose change	Changing a non-High-Risk system's purpose so that it becomes High-Risk.

Rebranding a limited-risk system makes you its provider too - but the obligations are Article 50 transparency, not the high-risk regime. Assess and document this before rebranding or modifying any third-party AI system. Full liability consequences: §11.6.

The liability ecosystem: shield and scales.

SHIELD - THE EU AI ACT
Focus: safety and market access. Mechanism: prevention and fines. Outcome: regulatory liability - the state versus your company.

SCALES - CIVIL LIABILITY
Focus: victim compensation. Basis: the revised Product Liability Directive (EU) 2024/2853 - software and AI are "products" under strict liability, transposition due 09.12.2026 - plus national tort law.

The separate AI Liability Directive (AILD) was officially withdrawn (OJ, 06.10.2025).

Who carries what.

FUNCTION	COMPANY / OWNER	INDIVIDUAL EMPLOYEE
Primary law	EU AI Act (regulatory)	Labour law / civil law
Maximum penalty	Up to EUR 35 million or 7% of global turnover	Disciplinary or professional liability
Trigger	System failure or lack of documentation	Gross negligence: ignoring AI warnings, failing to report errors

SECTION 2.4 · SME-SPECIFIC PROVISIONS

SME & small mid-cap (SMCs) provisions.

SMEs and SMCs are not exempt; size affects support and procedures, not applicability or risk classification. Most of these reliefs only apply once you have a high-risk system.

PROVISION	WHAT IT MEANS	HONEST LIMITATIONS
Simplified Technical Documentation HIGH-RISK ONLY	SMEs may use simplified templates instead of full technical documentation for high-risk systems.	Optional, not automatic. Commission templates are not yet finalised - do not wait for them before beginning compliance work.
Simplified QMS HIGH-RISK ONLY	SMEs and small mid-caps may fulfil the Article 17 quality management obligations in a simplified manner.	Does not remove other obligations. Only the QMS format is simplified - not the underlying compliance requirements.
Regulatory Sandbox Access ANY TIER	Priority access to national AI testing environments, generally free of charge for SMEs and startups.	Requires a formal application. Member States must establish sandboxes by 2 August 2027 - check your national authority to plan ahead.
Conformity Assessment Fee Reductions HIGH-RISK ONLY	Notified bodies must take SME size into account when setting conformity assessment fees.	A proportionate reduction, not a fixed discount. No standard rate has been set across Member States.

The following obligations apply to all organisations regardless of size:

- ✕ If your AI system is high-risk, all high-risk obligations apply in full, regardless of company size.
- ✕ If you are a Provider of a high-risk system, conformity assessment and CE marking are required regardless of size.

Implementing the EU AI Act requires strong **AI governance**, tools for documentation and monitoring, and the ability to accurately classify systems by risk.

Inventory & classify

Maintain a live register of every AI system you provide or deploy, with its risk tier and intended purpose.

Document & evidence

Hold technical documentation, data governance records, and conformity evidence at the maturity your risk tier requires.

Monitor & report

Run post-market monitoring, log incidents, and route serious incidents to the relevant national authority on time.

Train & assign

Assign clear accountability for AI compliance and train the humans who must oversee high-risk systems in use.

03 Understanding the risk-based *framework.*

Four risk tiers, from banned to unregulated in all but one respect. Classification to these four tiers depends on intended purpose, not technical design, and determines the scope of your obligations.

By the end of this chapter you will know:

- which of the four tiers each of your systems falls into
- whether anything you run is prohibited outright
- what the tier actually costs you in obligations

SECTIONS
3.1 The risk-based framework · 3.2 Prohibited practices · 3.3 High-risk · 3.4 Duties · 3.5 Limited risk· 3.6 Minimal risk · 3.7 AI literacy

READ TIME
16 minutes

KEY ARTICLES
Art. 5 Prohibited
Annex I / III
Art. 50 · Art. 6

03 3.1. THE RISK-BASED FRAMEWORK

Four tiers, one rule: the use case sets the burden.

The EU AI Act sorts every AI system into one of four risk tiers - from fully prohibited to minimal regulatory burden - using a risk-focused approach designed to protect fundamental rights while keeping innovation viable. Risk classification follows the **intended purpose:** the use for which the system is designed, as described in the provider's instructions, promotional materials, and technical documentation.

The EU AI Act risk pyramid

FIGURE 3.1 · ART. 5 · ANNEX I / III · ART. 50

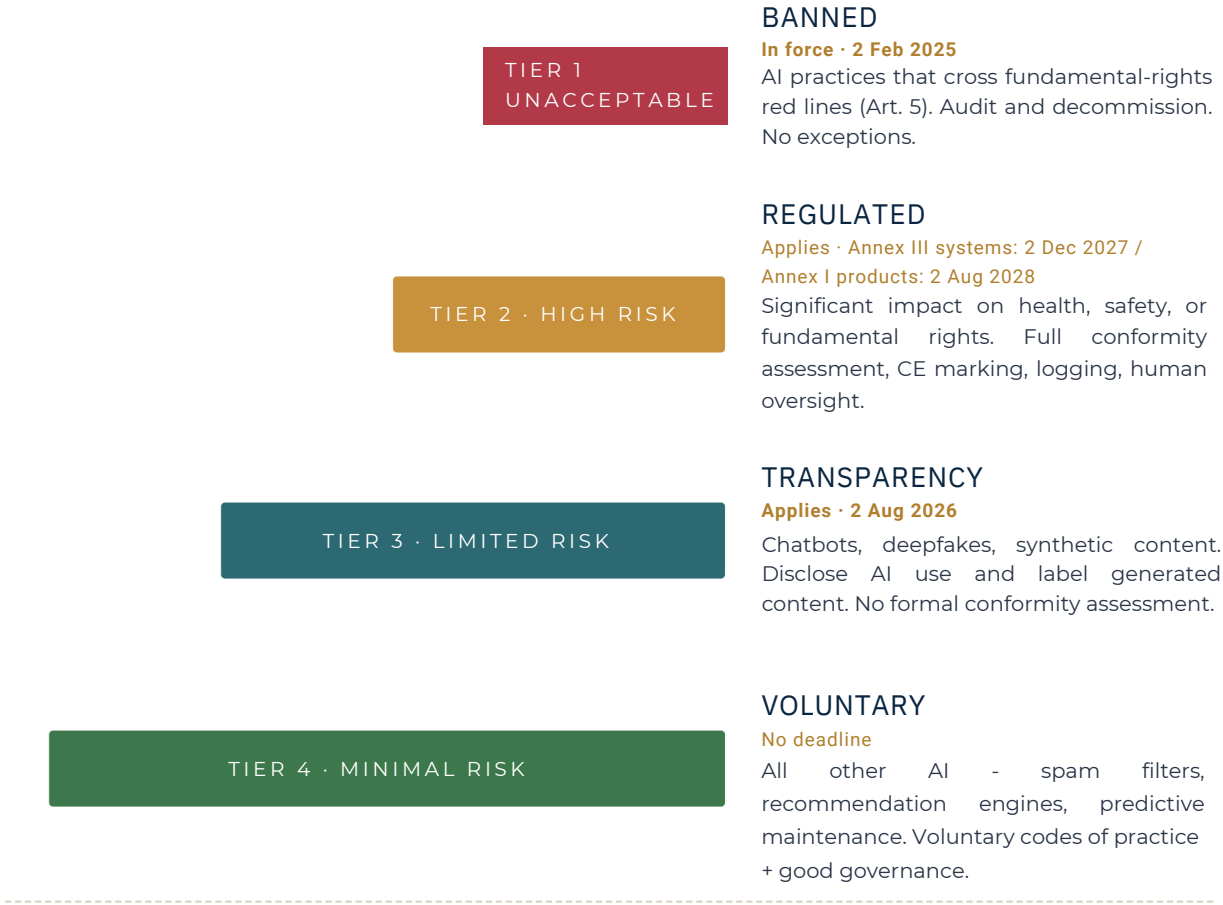


FIGURE 3.1 *Risk classification: tier assignment is driven by intended purpose.*

OBLIGATION BEYOND RISK TIERS

Article 4 is the only horizontal obligation that applies across all AI risk tiers. Whether your organisation uses minimal-risk, high-risk or General-Purpose AI, appropriate AI literacy is required.

SECTION 3.2 · ARTICLE 5

Tier 1 - eight (+2) practices that cross the line.

Article 5 identifies specific practices as inherently incompatible with EU values and the Charter of Fundamental Rights. These are not merely heavily regulated - they are **completely banned**. Every organisation must audit its AI portfolio against each of the eight categories below.

BANNED SINCE 2 FEBRUARY 2025 - ACT NOW

Operating one after the deadline exposes you to the highest penalty tier. For SMEs and SMCs, the fine is the lower of EUR 35 million or 7% of worldwide annual turnover; for larger companies, the higher of the two. The Digital Omnibus adds two more from 2 December 2026: AI generating non-consensual intimate imagery, and AI generating child sexual abuse material.

	PROHIBITED PRACTICE	DEFINITION	LIMITED EXCLUSION / SAFE HARBOUR
01	Subliminal, manipulative or deceptive techniques	Influencing behaviour beyond a person's consciousness in ways that cause significant harm.	-
02	Exploitation of vulnerabilities	Targeting individuals or groups by age, disability, or socio-economic situation to distort their behaviour.	-
03	Social scoring	Classifying persons by social behaviour or personality traits leading to unfavourable or disproportionate treatment.	-
04	Individual predictive policing	Assessing the risk a person will commit a crime based solely on profiling or personality traits.	-
05	Facial-recognition database scraping	Untargeted scraping of facial images from the internet or CCTV footage to build recognition databases.	-
06	Emotion recognition at work / in education	Inferring emotions of individuals in workplace or educational settings.	Excluded if strictly for medical or safety reasons (e.g. pilot fatigue, driver alertness).
07	Biometric categorisation by sensitive traits	Categorising people by race, political opinions, union membership, religious beliefs, or sexual orientation.	-
08	Real-time remote biometric ID (RBI)	Live facial recognition in publicly accessible spaces for law-enforcement purposes.	Narrow exceptions: preventing a terrorist attack, locating a missing child.

MARKETING TEAMS - PROCEED WITH CAUTION

Aggressive sales claims can inadvertently trigger high-risk or prohibited regulatory status. If your product's marketing suggests it classifies individuals by personality traits or assesses "social risk", legal and compliance must review the claims against Article 5 and Annex III before release.

SECTION 3.3 · ANNEX I & III

Tier 2 - Nine mandatory risk mitigation elements for high-risk systems

High-risk AI is permitted - but its legality is entirely contingent on meeting rigorous standards. **Risk classification** follows the system's i**ntended purpose**, as defined by the provider's instructions, promotional or sales materials, and technical documentation. Two categories exist: product-embedded (Annex I) and stand-alone software in sensitive sectors (Annex III).

Category A · Annex I

Product-embedded AI

AI systems intended to be used as a *safety component* of products covered by EU harmonisation legislation.

Examples · Medical devices · Machinery · Toys · Lifts · Pressure equipment · Marine equipment

Category B · Annex III

Stand-alone software in sensitive sectors

AI systems that impact fundamental rights in specific listed areas.

Examples · Biometric ID · Critical infrastructure · Education & vocational training · Recruitment & HR · Credit scoring & insurance · Justice & democratic processes

	REQUIREMENT	WHAT IT MEANS IN PRACTICE
01	Risk management system	A continuous process to identify, evaluate, and mitigate risks throughout the AI lifecycle.
02	Technical documentation	A comprehensive record proving the system meets all regulatory requirements - your "AI manual".
03	Data quality and governance	High-quality, representative datasets to reduce bias and discriminatory outputs.
04	Transparency towards users	Clear instructions explaining how the system works, its limitations, and its intended use.
05	Human oversight	Documented measures allowing humans to intervene, override, or shut the system down at any time.
06	Robustness and cybersecurity	Technical safeguards against vulnerabilities, adversarial attacks, and errors.
07	Quality management system	Documented internal processes that keep the system compliant during development and after deployment.
08	Activity logging	Automatic, tamper-evident recording of operations for traceability - 6 months minimum, longer if required.
09	EU database registration	Registration in the EU-wide AI database. Public for most Annex III areas; non-public or national for sensitive ones.

SECTION 3.4 ·PROVIDER & DEPLOYER DUTIES

Who does what - provider vs deployer.

High-risk obligations apply primarily to providers (developers and marketers), with deployers (users) sharing specific duties. The application dates were deferred by the 2026 Digital Omnibus to 2 December 2027 (Annex III systems) and 2 August 2028 (Annex I systems). If you rebrand or substantially modify a high-risk system, the provider regime attaches to you regardless of your original role.

Article 16 · Providers — HEAVIEST OBLIGATIONS

If you develop or place on market

- Establish a **continuous risk-management system** across the AI lifecycle.
- Ensure **high-quality training** datasets to minimise bias.
- Hold **detailed technical documentation** for regulatory authorities.
- Enable **automatic** logging (retain ≥ 6 months) for traceability.
- **Guarantee accuracy, robustness, cybersecurity** and meaningful **human oversight.**
- Run a **conformity assessment**, issue a **declaration** of conformity, affix the **CE marking** and register the system in the EU database.
- Supply clear instructions to deployers, report serious incidents, operate a QMS.

Article 26 · Deployers — USE-TIME OBLIGATIONS

If you use under your own authority

- **Monitor operations** in accordance with provider instructions.
- Assign competent **human-oversight** personnel .
- Suspend use **if risks emerge - notify providers** and authorities as required.
- Retain operational logs.
- **Inform** employees when workplace AI is deployed.
- For public bodies and entities in credit/insurance: conduct a **Fundamental Rights Impact Assessment** (FRIA) before deployment.

Implementation timeline

FIGURE 3.2 · APPLICATION MILESTONES



02 FEB 2025
Prohibitions (Art. 5) & AI Literacy (Art.4) live

02 DEC 2026
Most High-Risk (Annex III stand-alone) - HR, biometrics, credit

02 AUG 2028
Legacy embedded High-Risk in Annex I products (extended deadline)

WORKED CASE - HR & RECRUITING AI (ANNEX III)

An SME building a CV-screening tool for recruitment sits squarely in Annex III, Category B - High- risk regardless of company size. Provider obligations (Art. 16) and the nine mandatory requirements apply from 2 December 2027; the customer using it owes deployer duties under Article 26.

SECTION 3.5 · LIMITED RISK · TRANSPARENCY

Tier 3 - transparency, not bureaucracy.

Limited-risk AI- chatbots, deepfakes, emotion recognition outside the workplace and schools, and other AI generating synthetic content - is not subject to the full conformity assessment regime. But it carries mandatory transparency obligations from 2 August 2026. The core principle is simple: users must know they are interacting with a machine, and AI-generated content must be clearly labelled.

SCENARIO	OBLIGATION	PRACTICAL EXAMPLE
AI interacts directly with persons	INFORM Disclose that the user is interacting with an AI system, not a human.	A customer-service chatbot must identify itself as a machine at the start of the interaction.
AI generates synthetic content	MACHINE-READABLE Embed a machine-readable marker in generated text, images, audio, or video.	AI-generated marketing images must be watermarked or tagged so detection tools can identify them.
Biometric categorisation	INFORM Inform individuals that a biometric categorisation system is being applied to them.	An age-estimation or gender-based advertising tool must tell people it is categorising them.
Emotion recognition	INFORM Inform individuals that their emotional states are being inferred by an AI system.	A wellbeing app using facial analysis must disclose this clearly in its user interface.
Deepfakes & AI-fabricated content	DISCLOSURE Clearly label AI-generated deceptive content as artificial.	AI-generated "spokesperson" videos for advertising must carry a visible disclosure label.

WORKED CASES - INTERNAL CHATBOTS & GENAI MARKETING

An SME deploying a customer-service chatbot (Article 50 disclosure at start of conversation) or a GenAI tool for marketing assets (machine-readable watermarks, IP review) lands in Limited Risk. No conformity assessment - but skipping disclosure converts a low-friction obligation into a regulatory liability.

SECTION 3.6 MINIMAL RISK

Tier 4 - where most SME AI actually lives.

The vast majority of AI applications in use across SMEs fall into the minimal risk category: recommendation engines, spam filters, predictive maintenance, inventory optimisation, and most productivity or analytics applications. No mandatory obligations apply to these systems under the Act - except the universal AI literacy requirement covered in Chapter 4.

SME · 01



Predictive maintenance

Sensor-driven anomaly detection on production lines - the canonical Minimal-risk SME use case.

SME · 02



Inventory optimisation

Demand-forecasting AI for stock and replenishment - no Act obligations beyond literacy.

SME · 03



Productivity copilots

AI assistants for docs and spreadsheets – Minimal risk if not driving consequential decisions.

SME · 04



Spam & content filters

Email and content classifiers – the textbook Minimal-risk example across SME tech stacks



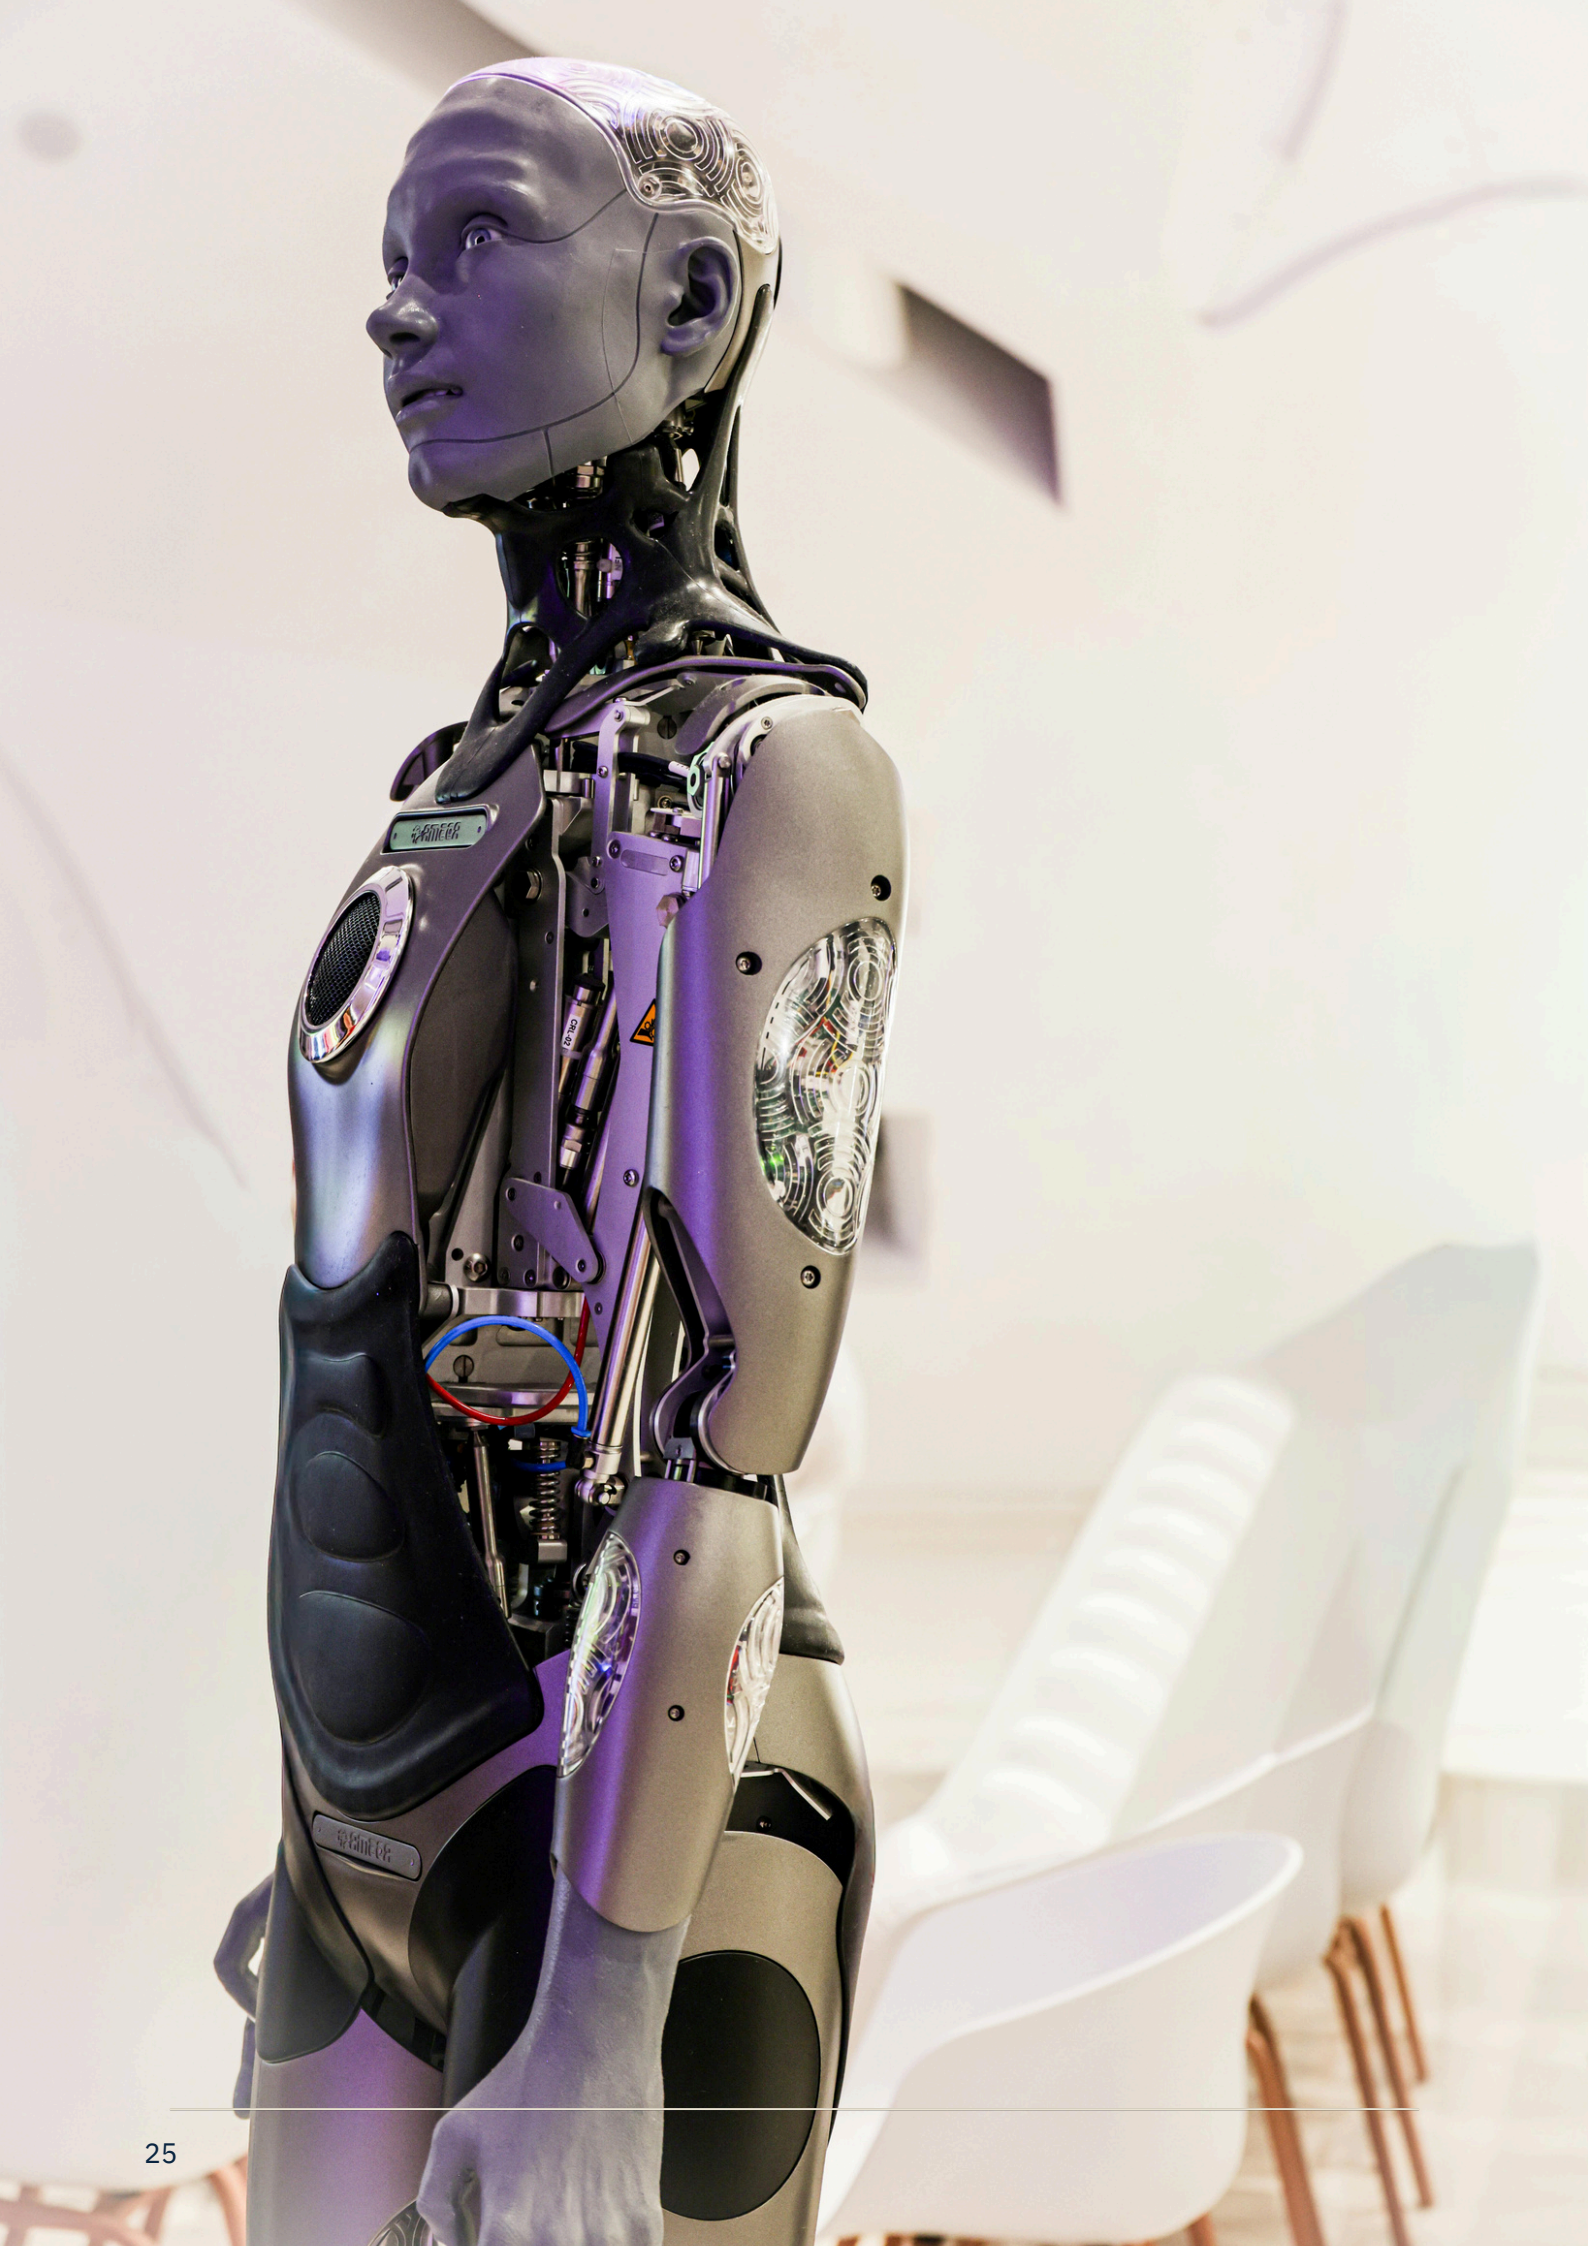
VOLUNTARY CODES OF CONDUCT

No mandatory rules apply - but organisations are encouraged to follow voluntary codes of conduct and document their AI governance processes. Good documentation protects you if a system's classification is ever questioned, and demonstrates a culture of responsible AI use to clients, partners, and regulators.

SECTION 3.7 · LITERACY

One requirement applies to every tier - AI literacy.

Regardless of which risk tier your systems fall into, Article 4 applies to every provider and deployer: you must take measures to support your staff's AI literacy. The 2026 Digital Omnibus made this an effort-based duty - you need not guarantee any particular level. The obligation has been active since 2 February 2025. Full detail in Chapter 4.



04

CHAPTER FOUR

AI literacy, audits & the *regulatory ecosystem.*



Article 4, which is already in force, focuses on the human element of the Act.

AI literacy is the foundation of responsible AI use. When using AI tools for work, employees must understand what AI can and cannot do, how to recognise limitations and risks, and when human judgement remains necessary.
Poor AI literacy can lead to:

- incorrect business decisions;
- data leakage and privacy breaches;
- regulatory non-compliance;
- overreliance on AI-generated outputs;
- biased outputs and statements;
- falsificated facts and lead to
- loss of cutomer and partner trust.

On the contrary, strong AI literacy improves risk management, trust, compliance and the overall effectiveness of AI adoption across the organisation.

SECTIONS
4.1 AI literacy & enforcement ·
4.2 Practical compliance ·
4.3 Audits and enforcement ·
4.4 Penalties & fines

READ TIME
14 minutes

KEY ARTICLES
Art. 4 AI Literacy

Art. 3(56) · Art. 99
Penalties

04 4.1 AI LITERACY & ENFORCEMENT

Article 4 - the part of the Act about *people*.

Article 4 makes AI literacy a legal duty (in force since February 2025), not a nice-to-have. It governs the human aspects of the legislation: every organisation that develops or uses AI must ensure the people working with it genuinely understand what they are doing and can exercise meaningful human oversight.

LEGAL DEFINITION · ART. 3(56)
AI literacy means the skills, knowledge and understanding that allow providers, deployers and affected persons to make an informed deployment of AI systems while recognising their opportunities, risks and potential harms.

Who does Article 4 actually affect?

Providers	Organisations that develop or commission AI systems and place them on the market or use them under their own name, including internally.	EXAMPLE An SME building an AI contract-review tool should ensure its development team has sufficient AI literacy regarding the model's capabilities, limitations and risks.
Deployers	Any organisation using AI systems under its authority (excluding purely personal use).	EXAMPLE A logistics firm using AI route optimisation, or an HR team using ChatGPT for drafting - every interacting staff member must be trained.
Workforce + contractors, freelancers, consultants	Art. 4 requires providers and deployers to ensure that staff and other persons dealing with AI on their behalf have sufficient AI literacy - not only employees.	EXAMPLE A marketing agency using a freelance copywriter who relies on AI tools must cover that person's AI literacy as part of the engagement.

AI literacy should take into account the technical knowledge, experience, education and training of the persons concerned, as well as the context in which the AI systems are used and the persons or groups on whom the AI systems are to be used.

When did it start?

02 Feb 2025
Article 4 live - AI literacy duty in force

02 Aug 2026
Enforcement of the AI Act starts at national and EU-level concerning general-purpose AI models, prohibitions, transparency rules and AI literacy

SECTION 4.2

Practical compliance - what it looks like for an organisation.

The AI Act takes a proportionate approach: organisations are expected to provide AI literacy appropriate to their role, the AI systems they use, and the associated risks. In practice, organisations using higher-risk AI systems should provide more comprehensive and role-specific AI literacy than organisations using AI only for low-risk tasks.
The four building blocks below provide a practical baseline for AI literacy across most roles, with additional role-specific training where required.

TRAINING TOPIC	WHAT STAFF MUST UNDERSTAND	MINIMUM STANDARD
Basic logics of machine learning	How does the specific AI tool work? What are its core capabilities and limitations?	e.g. "The tool predicts the next word - it does not 'think' or 'reason'."
Hallucinations & bias	Awareness that AI can produce plausible-sounding but factually wrong outputs, and that training-data bias affects results.	Staff must verify AI outputs before using them in decisions affecting people or processes.
Data privacy	Sensitive client, employee, or financial data must not be entered into public AI models (e.g. public ChatGPT prompts).	A documented policy on what data may and may not be used with AI tools.
Human oversight	When and how to apply human judgement before acting on AI output - especially for consequential decisions.	High-Risk: formal protocols. Minimal Risk: basic awareness + culture of verification.

AI literacy should be tailored to roles. Developers, HR staff, customer service teams and executives do not all need the same training.

HOW ORGANISATIONS CAN DEMONSTRATE AI LITERACY

The AI Act does not require external certifications. The following are examples of measures that can help demonstrate AI literacy, particularly when documented:

- Internal workshops led by a designated AI Coordinator or team lead.
- Completion of reputable free online courses (Digital Europe, Coursera, national digital agencies).
- Documented "lunch & learn" sessions covering AI risks relevant to your specific tools.
- A documented AI Use Policy supported by staff communication or acknowledgement.
- A brief onboarding module for new employees covering your organisation's AI tools and their risks.

The key is documentation: keep records of who was trained, when, and on what.

SECTION 4.3

Audits and enforcement - who actually shows up.

Enforcement is structured across two levels: European and national - with a clear division of responsibilities. Most organisations will primarily interact with their national market surveillance authority. Only providers of certain High-Risk AI systems will typically engage with a Notified Body, while GPAI providers may interact directly with the AI Office.

ENFORCEMENT BODY	SCOPE & REMIT	WHAT THEY WILL CHECK
EU AI Office ★ Brussels · Central oversight	The EU body responsible for overseeing General-Purpose AI models and promoting consistent application of the AI Act across Member States.	GPAI compliance Systemic risk Codes of practice Cross-border coordination
National Market Surveillance Authorities ⚑ Each Member State· The ones who knock	Each EU country appoints its own national AI regulator. These are the bodies organisations will actually encounter - reviewing documentation, investigating complaints from employees or customers, and conducting on-site inspections.	Technical docs QMS Activity logs Conformity cert Human oversight Literacy proof
Notified Bodies ✓ Independent third parties · Pre-market	Certain high-risk AI systems require assessment by a Notified Body before they can be placed on the market. In many cases, however, providers may use an internal conformity assessment where permitted by the AI Act.	Pre-market review Technical file Risk management CE marking

Under the microscope - the "Paper Trail of Trust"

01 · TECHNICAL DOCUMENTATION
Maintain technical documentation for each high-risk AI system, including its purpose, design, limitations and relevant data information.

02 · QUALITY MANAGEMENT SYSTEM
Document your QMS system. SMEs can use a simplified QMS, but the underlying obligation itself remains.

03 · AUTOMATIC LOGGING
Ensure high-risk AI systems automatically record required logs and retain them for an appropriate period and provide them upon request.

04 · HUMAN OVERSIGHT
Document how human oversight is implemented, including roles, responsibilities and intervention procedures where applicable.

05 · AI LITERACY RECORDS
Legally not required, but good practice : keep records of AI literacy activities, including dates, topics covered and participants.

SECTION 4.4 · ARTICLE 99

Three penalty tiers - designed to deter.

The AI Act sets three levels of administrative fines based on the severity of the infringement. Maximum fines are calculated using fixed amounts and turnover thresholds, with proportionate treatment for SMEs and eligible small mid-caps.

TIER 1 · PROHIBITED AI

€35M OR 7%

GLOBAL ANNUAL TURNOVER*

Use or placing on the market prohibited AI practices under Article 5 (social scoring, facial scraping, etc.).

*Subject to the proportionality rules for SMEs and eligible Small Mid-Caps.

TIER 2 · CORE RULES

€15M OR 3%

GLOBAL ANNUAL TURNOVER*

Non-compliance with high-risk AI requirements (e.g. risk management, technical documentation, logging, human oversight or conformity assessment requirements).

Proportionate caps for SMEs and eligible Small Mid-Caps

TIER 3 · INFORMATION

€7.5M OR 1%

GLOBAL ANNUAL TURNOVER*

Providing false, misleading or incomplete information to competent authorities or notified bodies.

Proportionate caps for SMEs and eligible Small Mid-Caps



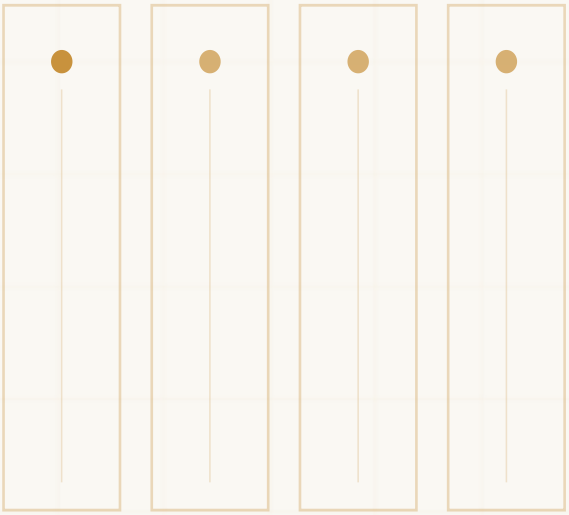
AI LITERACY + PENALTY RISK - THE AGGRAVATING-FACTOR TRAP

Article 4 does not have a separate penalty tier. However, organisations should not dismiss Article 4 simply because it lacks its own penalty category. Because an organisation's AI literacy measures may be considered when assessing the seriousness of an infringement and determining the level of an administrative fine.

05

CHAPTER FIVE

Transparency and information *requirements.*



Four interlocking duties - notify, label, log, communicate - with one shared goal: ensure users know when AI is involved, and that the evidence trail proves it. All four apply from **2 August 2026.**

SECTIONS

5.1 Transparency Duties ·
5.2 Content labelling ·
5.3 Best practise ·
5.4 AI Value Chain Responsibilities

READ TIME

18 minutes

KEY ARTICLES

Art. 50 Transparency

Art. 12 & 13 Logs & Information

05 5.1 TRANSPARENCY & INFORMATION

Four duties.
One audit-ready paper trail.

The EU AI Act sets clear, enforceable rules on transparency. Their purpose is to ensure that people are informed when the AI Act requires disclosure, such as when interacting with AI systems or encountering AI-generated or manipulated content. Organisations should be able to demonstrate compliance with these obligations. Most transparency requirements apply from 2 August 2026. The work should start now.

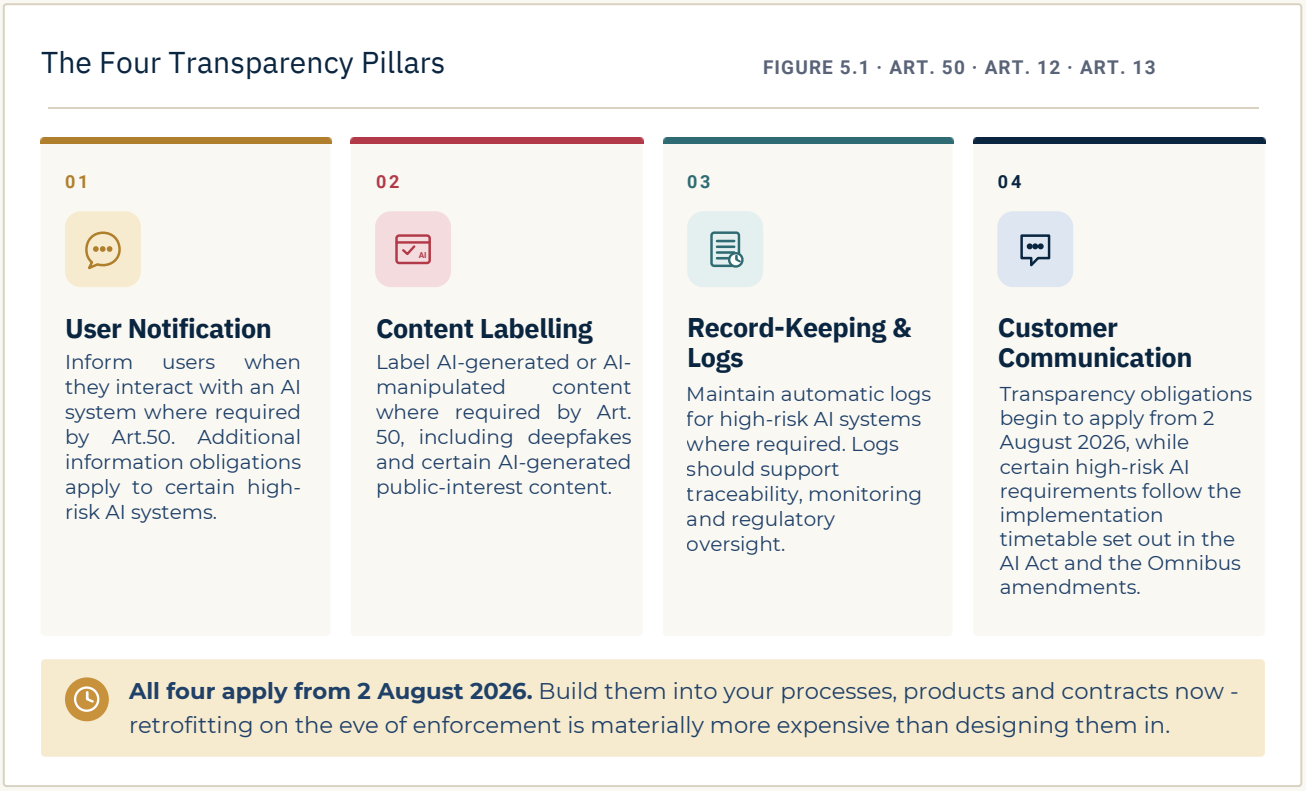


FIGURE 5.1 *The four interlocking transparency duties under Article 50 and adjacent provisions. They reinforce each other - together, they form the evidence trail that demonstrates compliance.*

SECTION 5.2 · ARTICLE 50(2) & (4)

Label every AI-generated output you publish.

Any AI-generated content- text, images, audio or video - presented to users must carry a clear and visible label. This applies regardless of whether the content was created with a first-party tool or a third-party GPAI model. The prohibition on unlabelled deepfakes is **absolute**.

Deepfake / Synthetic Media

"AI-created or altered."

Mandatory · no exceptions for art or entertainment unless artificial nature is otherwise obvious.

GPAI-generated Text

"Generated by AI" + machine-readable marker.

Training-data summary + identified biases disclosed by the GPAI provider; deployers label end-user output.

AI-generated Images

Visible watermark or embedded metadata.

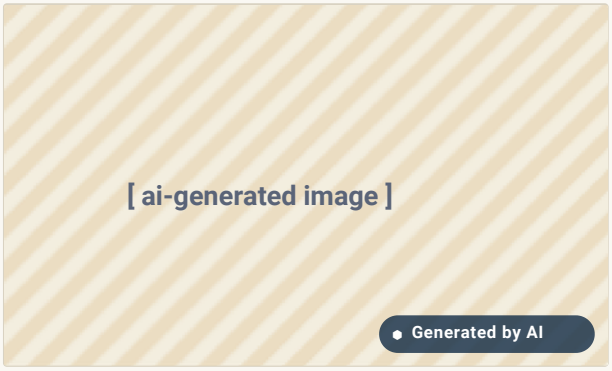
Required wherever images may be mistaken for authentic photographs.

AI-generated Audio / Video

Clearly visible or audible disclosure.

Proportionate to duration & reach - advertising, news and public-communication contexts in particular.

HOW A COMPLIANT LABEL LOOKS



Watermark in a corner - visible at glance, machine-readable in metadata.

AI · Generated by AI · verify before use

Dear customer,
Thank you for reaching out. Based on your order history we have prepared the following recommendation...

Persistent banner on any user-facing AI-generated message.

SME TIP Use watermarks, pop-up disclosures or platform-native labelling features in your AI tools. Test label withreal users - they must be easy to see and understand. Avoid tiny fonts or disclosures buried ins footers.

SECTION 5.3 · BEST PRACTICE

Go further than the law - trust scales.

The Act sets the legal minimum. Organisations that go further with proactive, transparent communication will be better positioned competitively and less exposed in enforcement scenarios. Four practices, ordered from quickest win to highest stakes.

01

AI fact sheet

A concise one-pager: which AI tools you use, for what purposes, and what safeguards are in place.

Why · Simplifies disclosure, builds confidence. Useful in sales, onboarding and support.

02

AI disclosures in contracts

Explicit statements in service agreements and FAQs that AI tools are used in service delivery.

Why · Manages expectations, reduces dispute risk, satisfies Article 13 transparency where it applies.

03

Opt-out mechanisms

Where technically feasible and non-central to the service, let customers request human-only handling.

Why · Builds trust, aligns with GDPR rights, reduces complaints to national authorities.

04

Incident communication

A clear protocol for notifying customers when an AI malfunction or incident may have affected them.

Why · Required for serious incidents (high-risk AI). Proactive communicatin reduces reputational and regulatory risk.

SME TIP

Build a one-page AI fact sheet for your business. Make it part of your standard sales, onboarding and support materials. Plain language only - customers do not need the technical architecture, just what AI does in their interaction with you.

SECTION 5.4 · ARTICLES 16, 26, 25

Transparency is shared. Liability is role-specific.

Transparency and compliance flow across the AI value chain - but liability is **role-specific**. Understanding which role your organisation holds for each system you develop, use, import or distribute is essential for accurate liability management.

ROLE	WHO YOU ARE	CORE TRANSPARENCY OBLIGATIONS
Provider	Develops an AI system (or has it developed) and places it on the market under its own name or trademark.	Implement a QMS. Maintain detailed technical documentation. Conformity assessment before CE marking. Supply clear instructions to deployers.
Deployer	Uses an AI system in a professional capacity, under its own authority.	Ensure human oversight. Monitor for and report serious incidents. Conduct FRIA for certain High-Risk systems used by public bodies or high-impact sectors.
Importer & Distributor	Brings AI systems into the EU or makes them available without modification.	Verify the provider's compliance (CE marking, documentation, instructions). Critical duty: inform the provider AND national MSAs if the system may present a risk to health, safety or fundamental rights.
Product manufacturer	Integrates a high-risk AI system into a regulated product (Annex I) under its own name - treated as a Provider.	Full provider obligations, incl. conformity assessment and CE marking under both the AI Act and the product-safety regulation.

STRATEGIC WARNING - WHEN YOUR ROLE CHANGES

An importer, distributor or deployer is legally reclassified as a Provider - and assumes the associated liability - in either situation:

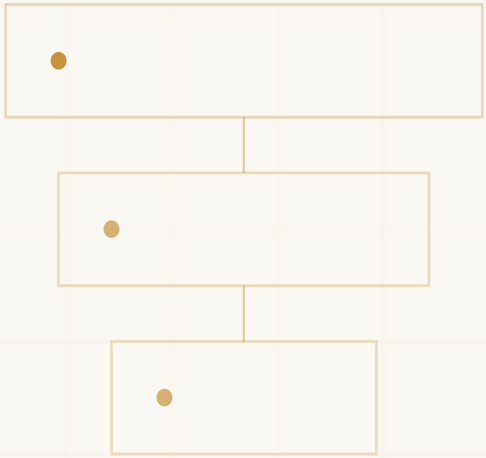
- × They place a high-risk AI system on the market under their own name or brand (rebranding).
- × They substantially modify a High-Risk AI system in a way not foreseen in the original conformity assessment.

Assess this explicitly and in writing before modifying, rebranding or significantly customising any third-party AI system.

06

CHAPTER SIX

General-Purpose AI Models *(GPAI).*



A separate regulatory layer for foundation models like GPT, Claude or Gemini. For most SMEs the bottom line is simple: you are almost certainly not the model's provider - those duties sit with your vendor. What is yours is the risk tier: the same tool can be minimal, limited or high-risk depending on how you use it.

SECTIONS
6.1 What is GPAI
6.2 Layer structure · 6.3 Use-case lens

READ TIME
15 minutes

KEY ARTICLES
Art. 3(63) · Art. 51–55

10^{25} FLOPs threshold

06 6.1 GPAI

You use GPAI. You almost certainly don't provide it.

GPT, Claude and Gemini are GPAI models; Copilot is a product built on one. Whichever your team uses, the organisation that trained the underlying model carries the model-provider duties - not you.

ONE MODEL, MANY JOBS

It can draft, translate, code and analyse.

Assess the task, not the model name.

BUILT INTO ANOTHER PRODUCT

You access it through an app, plug-in or API.

Your vendor is usually the provider.

TRAINED BY SOMEONE ELSE

Its training data and compute were controlled by the model developer.

Keep the vendor documentation on file.

Not all model providers carry the same load.
Every GPAI model provider owes a baseline set of duties. A small top tier of the most capable models owes more. The Act draws that line by training scale - models trained beyond 10^{25} FLOPs (a measure of computing power) are presumed to carry "systemic risk". Only frontier labs reach it, and either way these duties sit with the model's provider, above you.

FOR REFERENCE - THESE ARE YOUR VENDOR'S DUTIES, LIVE SINCE 2 AUGUST 2025
Baseline (Art. 53): a publicly available training-data summary, a copyright policy, and technical documentation made available to authorities on request. Systemic-risk models (Art. 55) add: model evaluation, adversarial testing, and incident reporting.

SECTION 6.2 · ARTICLES 25, 51–56

Three layers. *Three different owners* of responsibility.

The Act creates a clear layered responsibility framework. Understanding which layer your organisation occupies for each GPAI interaction determines your compliance obligations precisely.

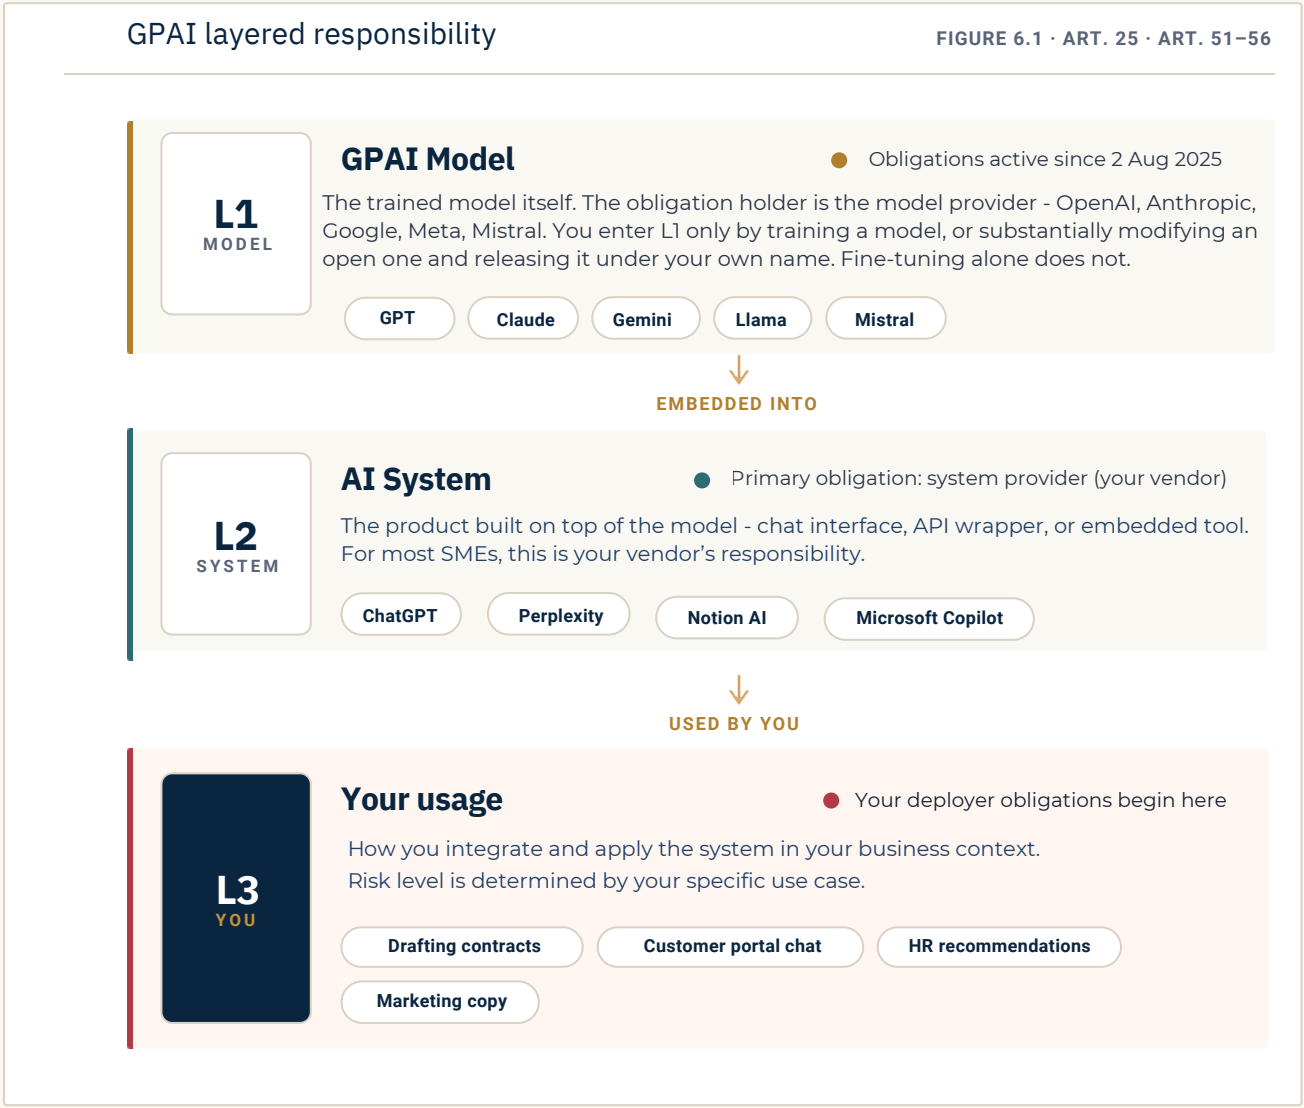


FIGURE 6.1 *The same model can sit beneath dozens of systems and hundreds of use cases. Each layer carries its own obligations - and your layer (L3) is almost always where your responsibility sits.*

SECTION 6.3 · USE-CASE LENS

The same AI tool can be minimal risk *or* high risk.

The risk framework from Chapter 3 applies to GPAI tools exactly as to any other AI system. GPAI adds one twist: take a single assistant your team already licenses - Copilot, ChatGPT, Gemini, whichever - and point it at four jobs. The tool never changes. The obligations do, because risk follows the task.

MINIMAL – LIMITED RISK

JOB 1 · Drafting internal documents

Minimal risk. AI literacy for staff (since 2 Feb 2025). Transparency (Article 50) if AI content is published externally. No formal assessment.

LIMITED RISK

JOB 2 · Answering customers in your chat

Your obligations - Transparency (Article 50): label AI-generated content, inform users they are interacting with AI. Staff AI literacy required.

HIGH RISK · ANNEX III — HIGH RISK

JOB 3 · Screening job applicants

High risk - Annex III. The same tool, now deciding about people: conformity assessment, technical documentation, logging, human oversight, FRA where applicable. Deadline - 2 December 2027.

POTENTIAL PROVIDER OBLIGATIONS — PROVIDER

JOB 4 · Built into a product you sell

Rebrand AI capability as your own product and you may become a provider under the Act. This is where the single-tool example ends - here you'd use an API or open model, not a licensed app. Full treatment in 6.4.

The tier attaches to your use, not to the vendor's product. Any comparable assistant behaves the same way.

CHAPTER SEVEN

07

Other regulations influencing *the EU AI Act.*

The AI Act does not stand alone. It sits inside a wider EU digital rulebook – GDPR, the Cyber Resilience Act, the Data Act, the Digital Services Act, NIS2 and sector-specific laws. For SMEs the practical insight is the same in every chapter: build one integrated compliance programme, not six silos.

SECTIONS
7.1 The wider ecosystem · 7.2 GDPR · 7.3 Product safety· 7.4 Platforms & critical infrastructure · 7.5 Sandboxes

READ TIME
16 minutes

KEY ACTS
GDPR · MDR · IVDR

CRA · Data Act · DSA · NIS2

07 7.1 THE WIDER ECOSYSTEM

One ecosystem. One compliance plan.

The EU AI Act sits inside a fast-evolving digital regulatory landscape encompassing data protection, product safety, cybersecurity, platform governance, and sector-specific requirements. For SMEs, the practical insight is this - compliance with one regulation typically supports compliance with another. A single integrated programme is both more efficient and more robust than treating each law as a separate workstream.

The EU Regulatory Ecosystem around the AI Act

FIGURE 7.1 · AI LEGISLATION · DIGITAL ECONOMY LAWS · SECTORAL REGULATION

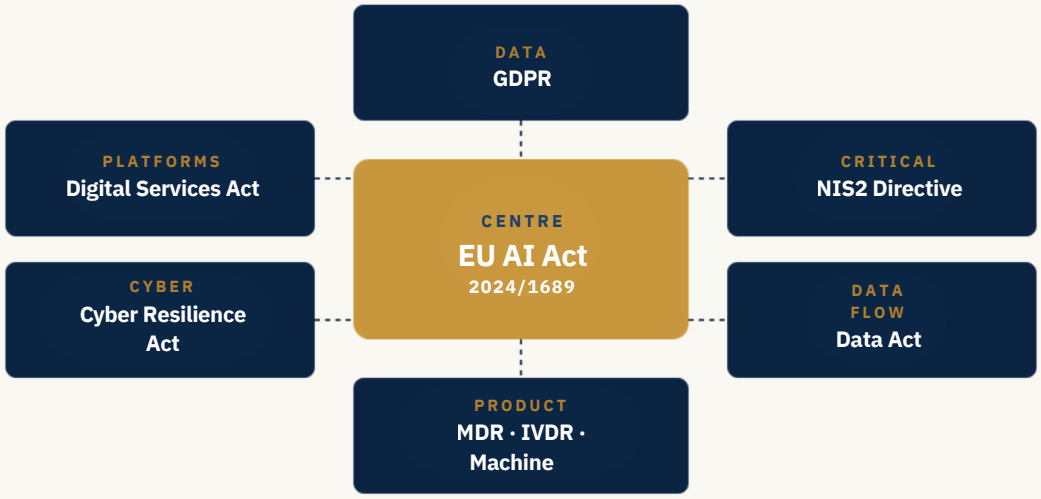


FIGURE 7.1 *Six EU regulations intersect with the AI Act, each covering a distinct regulatory angle - from GDPR's data rules to the Data Act's data flows.*

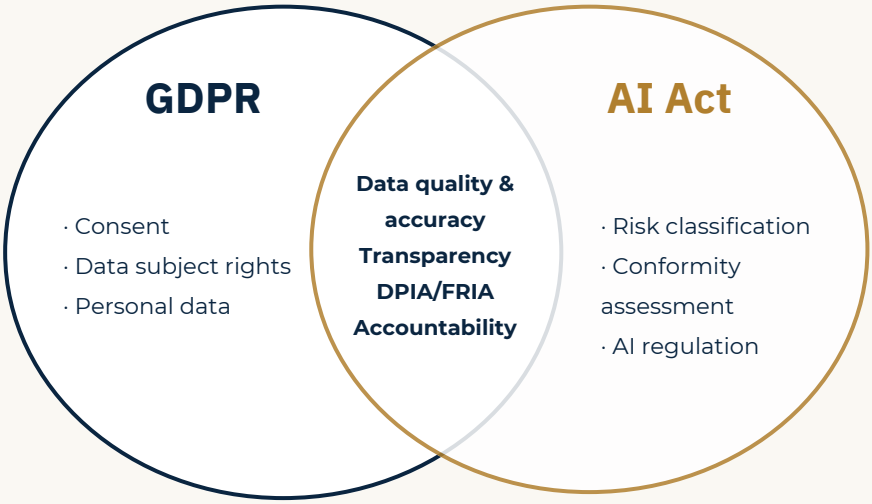
ONE PROGRAMME - MANY REGULATIONS

Most national authorities and EU bodies actively encourage a joint compliance approach. Technical documentation, risk and impact assessments, staff training, and audit trails can typically be designed to satisfy multiple regulatory requirements at once. *Do not build separate silos.*

SECTION 7.2 · REGULATION (EU) 2016/679

GDPR & the AI Act - two regulations, *one data programme*.

The GDPR governs how personal data is collected, processed, stored and protected. The AI Act adds AI-specific requirements on top of - but does not replace - GDPR. For any AI system that processes personal data (which is most of them, beyond pure internal productivity), both regimes must be addressed together.



INTERPLAY AREA	DETAIL	JOINT COMPLIANCE ACTION
Data quality **High-risk AI & personal data**	High-risk AI commonly processes personal data. AI Act Art. 10 data-quality requirements map directly onto GDPR principles of accuracy, minimisation and fairness.	Map all data flows for your AI systems. Identify where personal data is processed and ensure both regimes are covered in one data-management policy.
Impact **DPIAs & FRIAs**	High-risk AI deployers may need to run a GDPR DPIA and an AI Act FRIA. The two overlap substantially in scope and methodology.	Use a combined DPIA/FRIA template - several national DPAs publish one. One assessment covers both regimes; saves time and ensures consistency.
Penalties **Cumulative fine exposure**	GDPR: up to 4% of global annual turnover. AI Act: up to 7%. In a data-related AI incident, both regulators may act in parallel.	Do not silo GDPR and AI Act compliance. One responsible person or team for integrated AI & data compliance.

SECTION 7.3 · MACHINERY · MDR · IVDR

If your AI sits inside a regulated *physical product*.

AI embedded in physical products must comply with existing EU product-safety legislation in addition to the AI Act. The key regulations: the Machinery Regulation, the Medical Devices Regulation (MDR), and the In Vitro Diagnostic Regulation (IVDR). The AI Act adds cybersecurity and oversight requirements on top of pre-existing certification regimes - and a single CE marking can cover both.

REGULATION	RELEVANCE TO AI	SME ACTION
Machinery Regulation	Applies to AI used as a safety component in machinery. Category A (Annex I) High-Risk AI classification triggered.	Check whether your AI-enabled machinery needs updated technical documentation under both regimes. Engage your notified body early.
Medical Devices Regulation (MDR)	Applies to AI in medical devices. Category A High-Risk AI. Pre-existing MDR certification is the starting point.	Extend your existing MDR technical file to cover AI Act cybersecurity & human oversight. The body can typically run both assessments.
In Vitro Diagnostic Regulation (IVDR)	Applies to AI in in-vitro diagnostic medical devices. Same principle as MDR.	Audit your AI components against IVDR and AI Act simultaneously. Joint conformity assessment is the most efficient route.

SME TIP

If your AI qualifies as a "product" under any of these regulations, update your technical documentation for dual compliance. Notified-body fees for SMEs are subject to proportionate reductions - ask your notified body about SME-specific assessment packages.

ONE CE MARKING COVERS BOTH - WHEN DESIGNED IN FROM THE START

The CE marking system is the EU's shared compliance language across product safety, AI, cybersecurity and more. Designing for combined CE marking from day one is materially cheaper than two parallel certifications retro-fitted to a shipped product.

SECTION 7.4 · PLATFORMS & CRITICAL INFRASTRUCTURE

If you run a platform - or sit in critical infrastructure.

Two further regulatory frameworks have direct implications for SMEs operating AI in online platforms or in critical-service sectors. The DSA reinforces transparency duties on AI-generated content. NIS2 reinforces cybersecurity and incident-reporting duties. Both can usually be addressed in a single policy if planned for early.

LAW	WHO IT APPLIES TO	INTERACTION WITH AI ACT	SME ACTION
Digital Services Act (DSA)	Online platforms & intermediary services. Relevant if you operate a platform hosting user-generated or AI-generated content.	Requires platforms to label AI-generated content - directly aligned with AI Act Article 50 transparency. A single content policy can address both.	If you operate a platform, design your AI content labelling and moderation policy to satisfy both the Digital Services Act (DSA) and Article 50 requirements at the same time.
NIS2 Directive	Operators of essential services & critical infrastructure. Expanded scope vs. NIS1 - now also waste management, postal services, and manufacturing above certain thresholds.	Strengthens cybersecurity & incident-reporting for AI in essential-service sectors. NIS2 timelines and AI Act serious-incident reporting must be aligned.	If your AI sits in any NIS2-covered sector, align your incident-response & reporting to satisfy both regimes in a single protocol.

Where DSA / NIS2 / AI Act overlap in practice

DSA

- Content labelling on platforms
- Moderation transparency reports
- Algorithmic recommender transparency

SHARED WITH AI ACT

- Article 50 labelling
- Serious-incident reporting
- User & affected-person rights

NIS2

- Cybersecurity risk management
- 24-hour & 72-hour incident windows
- Supply-chain security duties

FIGURE 7.3 *The overlap zone is where joint compliance pays off. Build one labelling policy, one incident-reporting playbook, one supply-chain assurance set.*

SECTION 7.5 · ARTICLES 57–63

Regulatory sandboxes - *your SME advantage*.

Regulatory sandboxes - supervised testing environments provided by national competent authorities - are one of the most valuable tools available to SMEs navigating overlapping regulatory requirements. Member States must establish AI regulatory sandboxes by 2 August 2026 (see also Digital Omnibus 2026 changes). They let organisations test AI systems against multiple regulatory frameworks (AI Act, GDPR, CRA, sector laws) in a controlled environment, at low or no cost for SMEs.

01

Test multi-regime

Validate AI Act, GDPR, CRA and sector-law compliance simultaneously - with the regulator in the room - before market launch.

02

SME & startup priority

Priority access is granted to SMEs and startups by law. You do not need to compete with larger enterprises for places.

03

Practical regulator guidance

Direct interpretive guidance reduces compliance uncertainty and de-risks novel use cases that the regulation doesn't anticipate cleanly.

04

Faster time-to-market

Sandbox participation can accelerate go-to- market for compliant AI by surfacing regulatory issues early - and demonstrating good faith to the supervisor.

THREE STEPS TO TAKE THIS WEEK

1. Check your national competent authority's website for sandbox availability and the application calendar.
2. Map which adjacent regimes apply to your AI use case (GDPR, CRA, MDR/IVDR, DSA, NIS2).
3. Inventory the documentation you already have - most sandboxes accept your existing technical file as a starting point.

CLOSING THIS REPORT

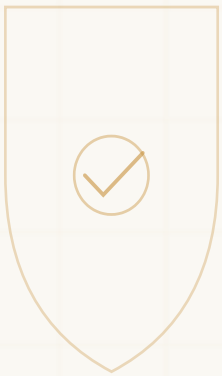
From rules to a single integrated programme.

Chapters 1–7 have walked the EU AI Act end to end: the definitions, scope and roles, the risk pyramid, AI literacy, transparency, GPAI, and the surrounding ecosystem. What remains is execution. The back cover summarises the through-lines - and the one principle that ties them together.

08

CHAPTER EIGHT

Practical compliance for SME *Deployers.*



Manage able obligations for SMEs using low-and limited-risk AI tools. Three duties carry the weight - human oversight, input data, and incident reporting - plus a vendor due-diligence checklist that protects you commercially as well as legally.

SECTIONS
8.1 Core Deployer Obligations

8.2 Vendor Due Diligence

WHO IT'S FOR
SMEs deploying ChatGPT,

Claude, Copilot & similar tools

READ TIME
7 minutes

SECTION 8.1 · FOR DEPLOYERS OF GPAI TOOLS

Three duties carry the weight.

The deployer obligations from Chapter 2 apply to GPAI tool use in the same way they apply to any AI system deployment. In the GPAI context - where the vendor controls the model but you control how and where it is used - three obligations carry particular weight.

OBLIGATION 01

Human Oversight

GPAI models can generate plausible but factually incorrect outputs ("hallucinations"). Staff must identify errors, question outputs, and retain authority over final decisions.

YOUR MOVE

Train staff to verify outputs before acting. Document who has override authority and when they must intervene.

OBLIGATION 02

Input Data Management

You are responsible for what you feed in. Customer data, personal data, and confidential documents sent as prompts are *your* compliance exposure - not the vendor's.

YOUR MOVE

Publish a clear policy on what data may and may not be entered into AI tools. This is also a live GDPR exposure.

OBLIGATION 03

Incident reporting

If a GPAI tool causes serious harm to health, safety or fundamental rights, you must report without undue delay to the provider and/or national authorities.

YOUR MOVE

Your vendor contract must support your incident reporting. If it doesn't, close that gap before deployment continues.

THE FINE-TUNING BOUNDARY - WHEN YOU BECOME A PROVIDER

If you significantly modify a GPAI model - by changing its performance, intended purpose or risk profile, including through fine-tuning or prompt engineering at scale - you may move from Deployer into Provider territory. Provider obligations are substantially heavier, including conformity assessment and CE marking for High-Risk systems. Seek specialist legal review before proceeding.

SME TIP · MAKE THIS A ONE-PAGE RECORD

For every AI tool in use, write a single page that names the owner, the override authority, the data input policy and the incident path. File it as a compliance record. Review when staff or vendors change.

SECTION 8.2 · BEFORE YOU SIGN OR RENEW

Eight questions your AI vendor should be able to answer

Work through this check list before signing or renewing any AI tool contract - GPAI or otherwise. These questions protect you commercially as well as legally: they give you contractual grounds to act if something goes wrong. Document responses and file as part of your compliance records.

01 Is the tool's intended purpose defined in writing?
Without it, you cannot classify the system's risk level. No written definition = no defensible compliance position. Ask for a system card or product sheet.

02 Does the vendor provide a system card or model documentation?
You need to understand what the system can and cannot do. Absent documentation = you cannot demonstrate due diligence to a regulator or auditor.

03 Will the vendor use our data or prompts to train their model?
If yes, live confidentiality + GDPR exposure. For business use, must be contractually excluded. Obtain written confirmation.

04 Where is data processed and stored?
Data leaving the EU without safe guards may breach GDPR. Non-EU processing requires documented transfer mechanisms (e.g. SCCs).

05 Can we export logs of AI activity for audit purposes?
For High-Risk use, activity logs are mandatory. If the vendor cannot provide them, you cannot meet your incident reporting obligations under the Act.

06 Will we be notified before major model updates or changes?
Un announced changes can shift the system's risk profile overnight. Advance notice lets you re-assess classification before update goes live.

07 Does the contract include a compliance cooperation clause?
If a regulator investigates, you need contractual grounds to compel vendor cooperation. Without this clause, you have no enforcement mechanism.

08 Does the vendor explicitly acknowledge the EU AI Act?
Vendors that cannot confirm awareness of the Act are a governance risk for any EU-market deployment. This is a basic due-diligence expectation.

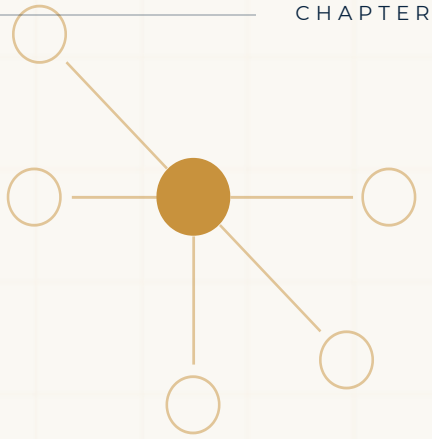


IF YOUR VENDOR CANNOT ANSWER Q3, Q4, Q5 OR Q7

Q3 & Q4 - Live GDPR exposure today. Pause use of the tool for personal or client data until resolved. | Q5 - Cannot meet incident reporting obligations for any High-Risk use case. Do not deploy at High Risk until closed. | Q7 - No contractual enforcement mechanism. Escalate to legal before renewal.

09

Practical compliance for providers & *high-risk deployers.*



Six steps to baseline compliance, a documentation toolkit you can build without a legal department, and a proportionate risk-management process built for small teams.

SECTIONS
9.1 Provider obligations · 9.2 Technical documentation · 9.3 Risk management · 9.4 Oversight · 9.5 Data governance

WHO IT'S FOR
High-Risk Deployers and all Providers

KEY ARTICLES
Art. 9 Risk · Art. 10 Data Art. 11 Docs · Art. 14 Oversight

SECTION 9.1 · THE PROVIDER'S PATH

Six steps to baseline compliance.

The Act itself implies a logical sequence-confirm what you have, assign accountability, document it, manage risk actively, establish oversight, and monitor ongoing. Start now with the minimum viable documentation set; expand as official guidance is published.

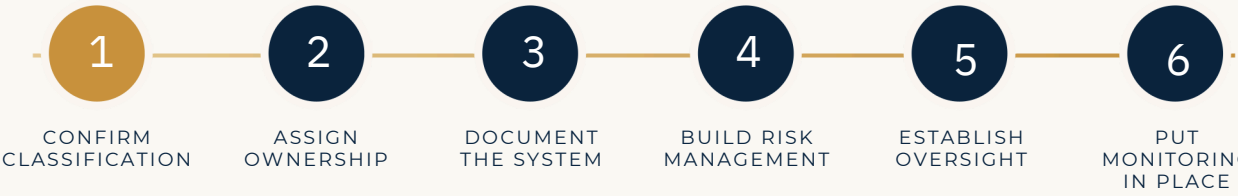


FIGURE 9.1 *The six-step implementation roadmap. Each builds on the last - without named ownership in step 2, nothing else holds.*

1 **CONFIRM YOUR CLASSIFICATION**
Document in writing that your system is High-Risk and record your role (Provider or Deployer).
An auditor's first question will be: how did you classify this system, and why?
Your answer must be written, dated, signed.

2 **ASSIGN OWNERSHIP**
Designate one named person responsible for AI compliance. In a small team this need not be a new hire - but someone must have the authority to pause deployment if something is wrong.

3 **DOCUMENT THE SYSTEM**
Technical documentation covering: what the system does, what data it uses, what it is designed to decide or recommend, and its known limitations. Start with a one-page system description; expand as needed (Art. 11).

4 **BUILD YOUR RISK MANAGEMENT PROCESS**
Risk management under the Act is a live process that runs the lifetime of the system - not a one-time exercise. Identify what can go wrong, document it, and review when the system changes or incidents occur.

5 **ESTABLISH HUMAN OVERSIGHT**
Before deployment, define who reviews AI outputs, under what conditions outputs can be overridden, and how the decision is recorded. Documented, not just practised. Staff training is mandatory under AI literacy (Art. 4).

6 **PUT MONITORING IN PLACE**
Log AI activity, review periodically, set an incident-reporting threshold and escalation path. Deployers retain logs ≥ 6 months. Providers retain longer per technical documentation.

IMPLEMENTATION TIMELINE

Most High-Risk obligations apply from 2 August 2026 under the current AI Act text - pending the Omnibus amendments described in Chapter 11. GPAI providers have been subject to obligations since August 2025; AI Literacy applies since 2 February 2025. *Do not wait for Commission templates before starting.*

SECTION 9.2 · ARTICLE 11 SIMPLIFIED

You do not need a legal department to satisfy Article 11.

The six documents below cover the most likely areas of regulatory scrutiny and are achievable for any small team. Start with the first four - they provide a practical starting point for compliance prepartaion

DOCUMENT	WHAT IT MUST COVER	WHO NEEDS IT	PRIORITY
System Description	Purpose, inputs, outputs, known limitations, and intended users of the AI system.	Providers & Deployers	START HERE
Risk Register	Identified risks, likelihood assessment, mitigation steps, and scheduled review dates.	Providers & Deployers	START HERE
Human Oversight Protocol	Who reviews AI outputs, the escalation process, authority and conditions for overriding output.	Providers & Deployers	START HERE
Incident Log	Date, description, action taken, outcome, and notification made (if any).	Providers & Deployers	START HERE
Data Governance Record	Data sources, quality checks performed, bias assessment methodology and findings.	Providers only	ADD NEXT
Conformity Assessment	Evidence the system meets all high-risk requirements before it is placed on the market.	Providers only	PRE-MARKET

MINIMUM VIABLE COMPLIANCE SET

The first four documents - System Description, Risk Register, Oversight Protocol, Incident Log - are achievable in one focused week for a small team. Migrate to the Commission's simplified templates (forthcoming under Art. 11) when published. Do not wait for them.

No official SME implementation framework had been published by the European Commission as of early 2026. The steps and documents on this page are built directly from the Act's requirements. Check the AI Act Service Desk (ai-act-service-desk.ec.europa.eu) regularly as official guidance emerges.

SECTION 9.3 · ARTICLE 9 IN PRACTICE

A risk management system. Not a risk management department.

Article 9 requires a risk management *system*, not a risk management *department*. For a small team, this means four ongoing practices - not a one-time assessment.

Identify before deployment

For each use case: what happens if the output is wrong? Who is affected? How seriously? Write the answers down - a one-page pre-deployment checklist reviewed by the named compliance owner constitutes a documented risk identification.

Mitigate what you can

Change the system, add human review steps, restrict the use case, or accept the risk with documented justification. Undocumented risk acceptance is *not* a valid compliance position.

Review when things change

A model update, a new use case, a new data source, or a reported error all trigger a fresh review. Build the triggers explicitly into your risk register - note "last reviewed" and "next review trigger".

Keep it proportionate

The Act requires regulators to consider SME scale when assessing compliance. A two-page risk register kept current is more defensible than a fifty-page document never opened.

PROPORTIONALITY IS ON THE PAGE

Compliance does not scale linearly with team size. The Act expects you to identify, mitigate, review and document risks proportionate to your operation. A small team's defensible risk file is short, current, and traceable to decisions - not long, comprehensive, and abandoned.

What "live process" looks like in practice

Risk management is not a binder you complete and shelve. The defensible position is a register that shows movement: dates of review, risks closed, new risks added when a model changes, justifications when a risk is accepted. Auditors will look for traceability - who decided, when, on what evidence - far more than they will look for comprehensiveness.

SECTION 9.4 · ARTICLE 14 IN PRACTICE

Three concrete oversight requirements.

Article 14 requires that high-risk AI systems enable humans to understand, monitor, and override outputs. For SMEs, this translates into three concrete requirements that must be in place before deployment.

Art. 14 · 01

TRAIN

Staff competence

Staff who use the system must be trained to recognise when outputs are unreliable. AI Literacy is a mandatory obligation - not optional guidance.

HOW TO IMPLEMENT

AI-specific training in onboarding and annual refresher cycles. Document completion. For High-Risk AI, training must be role-specific and cover the system's particular failure modes.

Art. 14 · 02

AUTHORISE

Override authority

Someone must have explicit authority and instruction to override an AI output. Documented - who the person is, and under what conditions they must intervene.

HOW TO IMPLEMENT

Name the person(s) in your Human Oversight Protocol. Define triggers (e.g. output confidence below threshold; output adversely affects an individual).

NO AUTONOMOUS FINAL DECISIONS

The system must not make final decisions autonomously in use cases where the Act requires human review. Audit your current AI-assisted workflows now. Any where AI output flows directly to a consequential outcome without human review must be redesigned before deployment.

SECTION 9.5 · ARTICLE 10 IN PRACTICE

Data governance - what you keep on file.

ROLE	TYPE	CORE OBLIGATION	MINIMUM COMPLIANCE POSITION
Deployer	**3rd-party AI**	Know what data you feed into the system. Personal data, client data and confidential information sent as inputs are your responsibility - not the vendor's.	Documented data input policy: what may and may not be used as AI inputs. Align with your GDPR data processing register. Brief all staff.
Deployer	**GDPR overlap**	Do not feed in data you are not permitted to use under GDPR. The AI Act sits alongside data protection law - it does not override it.	Review GDPR data processing records. Confirm AI tool use is covered by a lawful basis and, where applicable, a DPIA. A live gap for many SMEs.
Provider	**Own AI**	Document data sources, steps taken to assess data quality, and how bias was checked in training, validation and test datasets.	No data science team required. Honest documentation of what you checked, what you found and what you did satisfies the requirement.

10

CHAPTER TEN

Sector-specific *guidance.*

High-risk classification is not abstract - although the Act does not provide sector-specific guidance, it targets specific sectors where AI decisions affect fundamental rights. More prominent examples include human resources, finance, education, healthcare: **check your use cases against each sector and risk class**.

SECTIONS

10.1 High-impact sectors

10.2 Sector-specific priorities

ANNEXES

Annex III · Section 3, 4, 5

Annex I · Safety components

APPLY TO

All SMEs regardless of size if you operate in these sectors

SECTION 10.1 · FOUR SECTORS. FULL HIGH-RISK REGIME.

Sectors where the Act lands hardest.

If your organization deploys AI for applications specified in Annex I (AI systems integrated into products already subject to EU safety legislation, such as medical devices or vehicles) or Annex III (stand-alone AI systems used in critical areas, high-risk compliance requirements apply regardless of company size. However, classification depends strictly on the specific use case. Check your use cases against each use case example description - not every AI use in these example sectors make you subject to high-risk regime and its obligations .

ANNEX III · SECTION 4 · HIGH RISK

HR & Recruitment

AI used for CV screening, candidate evaluation, or employee monitoring. Applies regardless of company size. Focus: data quality to prevent discriminatory bias; human override authority at every decision point.

ANNEX III · SECTION 5 · HIGH RISK

Financial Services

AI for credit scoring or evaluating access to essential services (insurance, banking). Transparency, human oversight and record-keeping are key compliance requirements. Providers must maintain technical documentation.

ANNEX III · SECTION 3 · HIGH RISK

Education

AI for educational admissions or vocational training evaluation. Focus: robust human oversight to prevent the institutionalisation of systemic unfairness. Automated scoring or ranking systems require documented human review at consequential decisions.

ANNEX I · SAFETY COMPONENTS · HIGH RISK

Healthcare

AI integrated into medical devices or safety components covered by Annex I requires full conformity assessments and CE marking - in addition to compliance with MDR or IVDR. Seek specialist regulatory counsel.

MARKETING CLAIMS AND RISK CLASSIFICATION

Product descriptions, marketing materials and sales claims help define an AI system's intended purpose under the EU AI Act. If those materials describe the system as making or materially influencing decisions in areas covered by Annex I or Annex III, they may indicate that the system falls within the high-risk regime. Before publication, ensure that marketing claims accurately reflect the system's intended purpose and capabilities and are consistent with its regulatory classification.

The table highlights the most commonly encountered high-risk AI use cases in practice; however, Annex III also includes additional categories such as biometrics, critical infrastructure, law enforcement, migration and border control management, and the administration of justice and democratic processes.

SECTION 10.2 · WHAT TO FOCUS ON, SECTOR BY SECTOR

Where to put your compliance effort - four use cases

SECTOR	HIGH-RISK CLASSIFICATION	KEY COMPLIANCE FOCUS	SPECIALIST ACTION REQUIRED
HR & Recruitment	Annex III §4 - Employment, workers management, access to self-employment	Eliminate discriminatory bias through data quality controls. Human authority over final hiring, promotion, and disciplinary decisions. Notify employees when AI tools are used in decisions affecting them.	Data quality assessment, human oversight protocol, employee notification procedure
Financial Services	Annex III §5 - Access to and enjoyment of essential private services	Explainable decision logic for any AI-assisted credit, insurance, or lending decision. Documentation enabling affected individuals to understand and challenge automated outcomes.	Decision explainability documentation, individual rights procedure (right to explanation)
Education	Annex III §3 - Education and vocational training	Human review of AI-assisted admissions or assessment outcomes. Prevention of automated bias reinforcement. Documentation of review procedures.	Human oversight protocol, bias assessment, review of assessment methodology
Healthcare (AI in medical devices)	Annex I - Safety components of regulated pr o duct s	Full conformity assessment under both the AI Act and the Medical Devices Regulation (MDR) or IVDR. CE marking covering both regimes. Cybersecurity requirements.	Specialist regulatory counsel required, engage MDR notified body to extend assessment scope

CLOSING THE SECTOR CHAPTER

The EU AI Act is a living regulation.

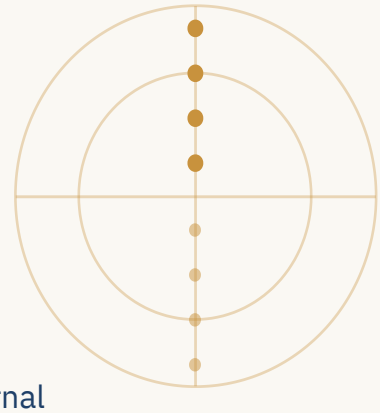
The Act is not a static compliance event - it is a framework designed to *evolve with the technology it governs*. Delegated acts, implementing regulations, harmonised technical standards, and Commission guidelines will continue to emerge through 2026, 2027, and beyond. Organisations that treat compliance as an ongoing commitment - rather than a one- time project - will be better positioned competitively and less exposed legally as the regulatory landscape matures.

For SMEs, the practical path forward is straightforward: start with the minimum viable compliance position outlined in Chapters 8 and 9, monitor the AI Act Service Desk and your national competent authority for emerging guidance, and treat AI governance as a standing agenda item - not one-time event or an annual exercise.

11

CHAPTER ELEVEN ADOPTED · 2026

The Digital Omnibus *on AI.*



The AI Digital Omnibus is now part of EU law. Adopted as Regulation (EU) 2026/1744 and published in the Official Journal of the European Union on 24 July 2026, it entered into force on **27 July 2026**.

The "Digital Omnibus on AI" is an amendment to the original AI Act - part of the EU's broader simplification agenda following the Letta and Draghi competitiveness reports. The Omnibus should be understood as a simplification package of AI regulation, not a full reset of the EU AI Act.
The amendments reduce administrative burdens, provide additional support for SMEs and eligible Small Mid-Caps (SMCs), extend certain compliance deadlines for high-Risk AI systems, and strengthen measures to foster innovation across the EU. The Commission's stated target: reduce administrative burdens by at least 25% for all businesses and at least 35% for SMEs by 2029.

Read in this Chapter what changed, what is now fixed in law, and what to do before each deadline arrives.

SECTIONS
11.1 Where things stand
11.2 What changes & when

STATUS
Adopted · 29 June 2026
Publication pending · OJ

KEY DATES
02.12.2026 · new prohibitions 02.12.2027 / 02.08.2028 · fixed

SECTION 11.1 - INTRODUCTION

Key practical changes introduced by *the AI Digital Omnibus.*

Digital Omnibus on the EU AI Act is not a rewrite: the risk-based approach and the substance of the rules stay as originally stated in the EU AI Act from 2024. What changes are **application dates**, two **new prohibitions**, and **real documentation relief** for smaller companies.

ADOPTED	OFFICIAL JOURNAL	IN FORCE	FIRST NEW DUTIES
29 June 2026	**24 July 2026**	**3 days after publication**	**2 December 2026**

01 Sandboxes & reinforced AI Office
The Omnibus establishes an EU-wide regulatory sandbox framework, strengthens the role of the AI Office, and expands access to regulatory sandboxes.

02 Two new prohibitions
From 02.12.2026: AI systems generating or facilitating the generation of non-consensual intimate imagery and AI generating child sexual abuse material. Penalty up to EUR 35 million or 7% of turnover.

03 Deadlines for High-Risk AI are postponed & fixed
High-Risk AI: 02.12.2027 (Annex III) and (Annex I) 02.08.2028
Parliament intends these deadlines to be final and not subject to further postponement.

04 SMEs and SMCs gain relief
Lighter documentation, a proportional quality-management system, and lower penalty ceilings - now for companies up to 750 employees and meeting financial criteria.

05 August 2026 stands
Transparency, governance, market surveillance and penalties take effect on schedule - no deferral for these.

06 Simplified Quality Management System (Art. 17) for all SMEs
The simplified QMS rules, previously limited to microenterprises, are extended to SMEs through the Omnibus.

SOURCES TO MONITOR FORMAL ADOPTION AND CHANGES

AI Act Service Desk – ai-act-service-desk.ec.europa.eu | European AI Office – ai-office.europa.eu | EDPB – edpb.europa.eu | ENISA – enisa.europa.eu
Your national competent authority for enforcement approach and sandbox applications.

WHAT CHANGES AND WHEN

2 August 2026: what takes effect - *and what doesn't.*

From 2 August: AI Act transparency rules for providers and deployers of certain AI systems apply (incl. marking of AI-generated content and labelling of deepfakes and certain publications)

APPLIES FROM 02.08.2026

Art. 50(1) - users must be told they are interacting with AI (chatbots, voice).

Art. 50(3) - disclosure of emotion recognition and biometric categorisation.

Art. 50(4) - deepfake and AI-generated text disclosure.

Art. 50(2) for systems newly placed on the market.

Governance, market-surveillance powers, and penalties up to EUR 35 million or 7% of turnover.

→

DELAYED BY THE OMNIBUS

Art. 50(2) labelling for systems already on the market: to 02.12.2026.

High-risk stand alone systems: to 02.12.2027.

High-risk AI embedded in products: to 02.08.2028.

National sandboxes: to 02.08.2027.

Delay is not repeal: what changes is the date, not the substance of the requirements.

High-risk deadlines: from condition to fixed date.

The most consequential structural change in the Omnibus - and the one SMEs are most likely to misread.

COMMISSION PROPOSAL (11.2025)

High-Risk rules would apply only after a Commission decision confirming that harmonised standards are available. The dates were conditional.

FINAL TEXT (06.2026)

Parliament removed that mechanism entirely. The dates now stand directly in the law: **02.12.2027** (Annex III) and **02.08.2028** (Annex I).

? **Why it was delayed**

Harmonised standards, conformity-assessment bodies and institutional capacity were not ready on the original schedule.

! **Why it will hold**

The same justification cannot be used twice. With the conditional mechanism gone, there is no legal route to a further extension.

-> **What SMEs should do**

Treat 2027/2028 as final. The extra time raises regulatory expectations - it does not lower them.

SECTION 11.2 - NEW PROHIBITIONS AND THE SAFE HARBOUR DEFENCE

Two new bans - and a defined way *to stay protected.*

From 2 December 2026, two new prohibitions join Article 5 - both at the Act's highest penalty tier: up to EUR 35 million or 7% of worldwide turnover. Both were added by the EU Parliament during the trilogue in response to real-world harms.

BAN 1 - NON-CONSENSUAL INTIMATE IMAGERY (NCII)
AI that generates or manipulates realistic intimate images of an identifiable person without their explicit consent. Targets "nudifier" apps and intimate deepfakes.

BAN 2 - CHILD SEXUAL ABUSE MATERIAL (CSAM)
AI that generates or manipulates child sexual abuse material (Directive 2011/93/EU). Applies to providers and deployers alike - and is absolute: no defence exists.

The “safe harbour” principle - how the NCII ban stays workable.

Legitimate image tools could, in theory, be misused to create intimate deepfakes. The “safe-harbour” principle keeps the NCII ban workable by linking the prohibition to a defined checklist of protective measures. Where the conditions set out in the Regulation are met and those safeguards can be demonstrated, providers may benefit from that protection.



THE DEADLINE IS THE PROOF, NOT THE PROMISE

Without documented, working measures by **2 December 2026,** you cannot claim the “safe harbour” - and are liable under Article 5 at the highest penalty tier.

WHAT PROVES SAFE HARBOUR IN YOUR COMPANY

Five ways to demonstrate appropriate safeguards - *evidence.*

If a national authority or the EU AI Office questions your claim, "trust us, it works" is not an answer. Documented evidence is one of the strongest ways to demonstrate that appropriate safeguards were implemented and maintained.

The Omnibus lists the possible evidence: data cleaning, refusal training, safe prompt design and output controls, runtime prompt guardrails, content classification and filtering mechanisms, usage restrictions, abuse detection mechanisms, and notice and action mechanisms.

The examples below illustrate practical evidence that may help you to comply*:

PILLAR	WHAT TO HOLD ON FILE
1 Technical architecture	Diagram and specification of input/output filters. Training data and safety measures used. Testing logs: false-positive and false-negative rates, sample blocks.
2 Monitoring logs	Incident reports and pattern analysis of circumvention attempts. A response protocol: who is notified, how fast.
3 Terms of use	Published policy in the local language, with a consent/acceptance log. Explicit NCII/CSAM prohibition language.
4 Staff training	Who understands the filters and why. Training records and an escalation path.
5 Audit trail	Contract with any third-party model provider. Proof of your own added filters and, where used, third-party stress-testing.

FIGURE 11.4 *The hard part: proving it worked - actual blocks, false positives, edge cases refined over time.*

Two roles, two burdens - and neither can hide behind the other.

PROVIDER / DEVELOPER HOLDS	
Filters built into the model or API itself. Training-data filtering and red-teaming records. A model card or technical report showing safeguards.	Your own added filters on top of the provider's model. Abuse monitoring and incident handling procedures. Your user policies or terms governing use of the system.



EACH ORGANISATIONAL ROLE HAS ITS DUTIES

The Omnibus expects both providers and deployers to take responsibility for preventing foreseeable misuse. A deployer cannot rely solely on the provider's safeguards, and a provider cannot assume deployers will add missing protections.

*The measures shown are recommendations derived from the safeguards described in Recital 12 of Regulation (EU) 2026/1744. They illustrate how providers can implement the "reasonable and adequate technical safety measures and other safeguards" required by Article 5(1a); they are not an exhaustive or mandatory checklist.

SECTION 11.3 - TRANSPARENCY, AI LITERACY AND PRODUCT RULES

The labelling deadline moves - Article 50 *itself does not.*

The one transparency change in the Omnibus is the labelling deadline. But Article 50 sits inside a bigger framework that every organisation using chatbots, generative tools or synthetic media needs to understand.

Where transparency sits in the EU AI Act risk-based system.

Unacceptable risk	prohibited - e.g. social scoring, manipulative biometric systems.	
High-risk	strict requirements - e.g. AI in employment, medical devices.	
Transparency risk	disclosure required - chatbots, deepfakes.	This is Article 50.
Minimal risk	no rules - e.g. spam filters.	

FIGURE 11.5 *A parallel track of transparency and risk-management duties applies to general-purpose AI models, which can sit as components inside any of these tiers.*

Article 50 at a glance - four obligations.

The main Objective of this obligation is to enable humans to **distinguish AI systems and AI-generated content** and foster trust and the integrity of the information ecosystem.

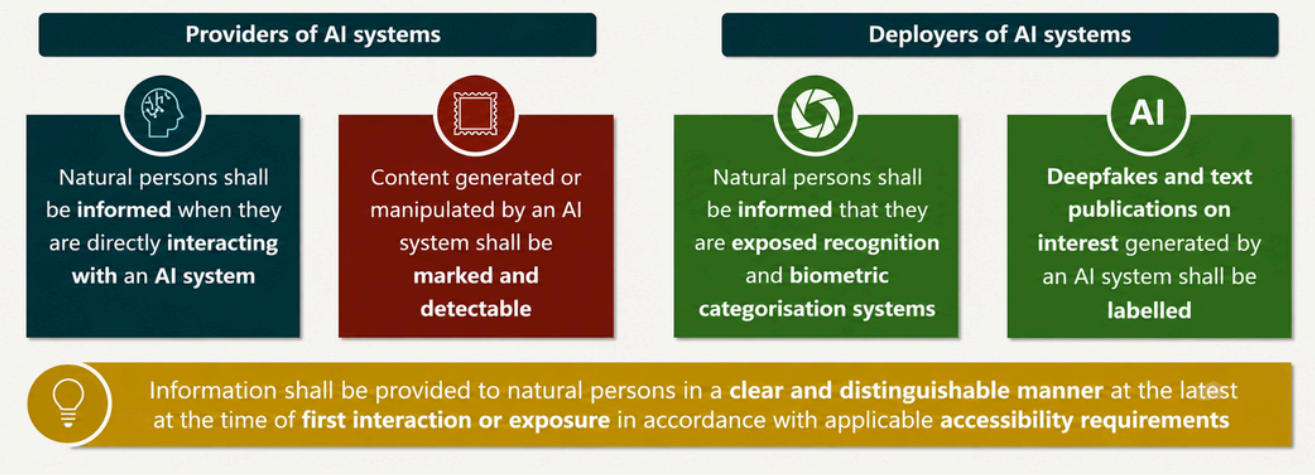


FIGURE 11.3 *Adopted from European Artificial Intelligence Office, 2026*

ENFORCEMENT AND THE GRANDFATHERING RULE

On the market before 02.08.2026: labelling under Art. 50(2) required by 02.12.2026 - a four-month grace period. On the market from 02.08.2026: all transparency rules apply from day one. Penalty: up to EUR 15 million or 3% of turnover, proportionate for SMEs and SMCs.

SECTION 11.4 - THE CODE OF PRACTICE ON TRANSPARENCY OF AI-GENERATED CONTENT

Code of Practice and a free set of *EU labelling icons.*

The Commission and the AI Board have assessed the Code of Practice as an adequate voluntary compliance tool. Six independent academic experts built it with broad input, to be practical, legally sound, and state-of-the-art.

The official EU icons for labelling AI-generated content.

Free and ready to use for the Article 50(4) labelling duty - in four variants (black, white, each also at 50% transparency), as SVG or PNG. Using them is optional; the labelling duty itself is not.

AI

AI was involved in creating the deepfake or published text - or a custom text label or interactive second layer is used with it.

AI GENERATED

The entire content is generated by AI with no human-created elements or editorial control, beyond prompting.

AI MODIFIED

Pre-existing,human-made content was partially modified by AI - e.g. a face swap in an authentic photograph.

FIGURE 11.6 *Source: European Commission, "EU Icons for labelling AI-generated content" - user-tested; performance improved when the icon was paired with a text label (e.g. "modified").*

How must a notice be designed?

ARTICLE 50(5) AI ACT

Clear, concise and readily available — at the latest at the first point of contact.

✓ CLEAR	No hidden links, no footnotes as the sole source of information.
✓ CONCISE	Must not be combined with other content — clearly separated from terms and conditions or other notices.
✓ ACCESSIBLE	Appropriate for the target audience, channel and accessibility requirements: text, audio, visual.
✓ TIMELY	At the latest at the first interaction, exposure or publication.
✓ VERIFIABLE	Document the version, text module, display location and the responsible point of contact.

Visual example



FIGURE 11.7 *Image generated with GenAI tool (Gemini) to illustrate the transparency principle in practice*

WHY THIS MATTERS TO SMES

Map where AI-generated content arises in your business - chatbots, marketing copy, product images - and ask each vendor to confirm their labelling and detection actually work. First-signatory deadline for the Code: 27 July 2026 (extended); list published 31 July 2026; signing stays open also afterwards.

SECTION 11.5 - RELIEF, GOVERNANCE AND ENFORCEMENT

A new company category and real *proportionality.*

NEW: SMC

Small mid-caps - fewer than *750 employees* and turnover up to EUR 150 million - now receive the reliefs previously reserved only for SMEs (under 250 employees, EUR 50 million).

THRESHOLD
< 750 staff

01 Simplified documentation
Lighter technical documentation - for SMCs as well as SMEs.

02 Proportional QMS
A simplified quality-management system for all SMEs and SMCs.

03 Lower penalty ceilings
Fines calibrated to company size - the lower of the two amounts applies.

04 Priority sandbox access
SMEs, start-ups and SMCs get priority; national sandboxes arrive by 02.08.2027, an EU-level sandbox sooner.

05 Simplified registration
The registration duty itself stays but with fewer AnnexVIII data points, thus, simplifying the procedure.

06 Proportional oversight
Post-market surveillance and monitoring applied proportionately to smaller companies.

Penalty tiers - architecture unchanged.

BREACH	MAXIMUM PENALTY
Prohibited AI practice (Article 5)	up to EUR 35 or 7% of worldwide annual million turnover
Other obligations (High-Risk systems, Article 50)	Up to EUR 15 million or 3% of worldwide annual
Misleading information to authorities	Up to EUR 7.5 million or 1% of worldwide annual

GDPR RELIEF IS NOT YET IN FORCE

Originally two digital omnibus proposals were published on 19 November 2025. As of today, only the AI Omnibus is adopted. The Data Omnibus - touching the GDPR, ePrivacy, the Data Act and NIS2 - remains in negotiation with no agreed timeline. Personal data processing in AI development stays fully subject to the GDPR as it stands today.

Sensitive data — permitted solely to *establish fairness.*

The Digital Omnibus introduces a targeted legal basis on which providers and deployers may process the most sensitive categories of personal data for a single, tightly defined purpose — the monitoring, detection and correction of bias — and only where a prescribed set of safeguards is satisfied.

The amendment does not create a general entitlement to process special-category data. It introduces a specific, purpose-limited derogation from the Article 9 prohibition (GDPR): where detection or correction of bias genuinely requires it, providers may lawfully process such data, beyond the confines of the High-Risk regime.

GDPR ART. 9 CATEGORIES

Racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership, genetic data, biometric identification, health data, sex life, sexual orientation.

Prohibited as a matter of principle; lawful only where an enumerated exception applies.

BEFORE THE OMNIBUS

Permitted only in the limited circumstances tied to High-Risk systems, leaving many providers unable to demonstrate fairness without the data needed to test it.

WHAT CHANGES

Providers may now process special-category data solely to detect, evaluate and correct bias, including outside the High-Risk regime, wherever the prescribed safeguards are observed.

NOT A GENERAL PERMISSION

The permission is *deliberately narrow*: available only where processing is genuinely necessary for bias detection or correction, and every safeguard below is satisfied in full.

Strict necessity • **Proportionality** • **Technical measures** • **Full GDPR compliance** • **Purpose limitation** • **Security**

✓ Credit scoring
A bank tests whether ethnicity affects scoring accuracy, under the safeguards.

✓ Diagnostic AI
A hospital checks diagnostic accuracy across sex and age groups.

✕ Marketing
A chatbot logs ethnicity with no bias-detection purpose — still prohibited.

Properly understood, the amendment doesn't dilute the Act — it renders responsible AI more workable.

SECTION 11.6 - KEY TAKEAWAYS, ACTION PLAN, SOURCES

Deadlines will not move again.
Use the time gained.

01	Build an AI inventory - every tool built or bought: purpose, provider, data, users	BY 08.2026
02	Reclassify risk against Annex III and the narrowed "safety component" definition	BY 08.2026
03	Ensure Article 50 transparency - chatbot disclosure, generative-content labelling	BY 02.08.2026
04	Audit generative tools against the new prohibitions - filters, monitoring, terms of use	BY 02.12.2026
05	Check your SME/SMC status - under 750 employees, up to EUR 150M turnover	IMMEDIATELY
06	Document AI literacy training, linked to the tools your teams actually use	ONGOING
07	Use the time gained - QMS, technical documentation, supplier contracts, oversight	THROUGH 2027

KEY TAKEAWAYS

Know your role - provider or deployer? Watch Article 25's reclassification traps.
Compliance is daily work , not a one-off audit
AI literacy is an obligation - documented, not just delivered.
The human factor - the "stop button" is your responsibility.

Sources: Commission proposal COM(2025) 836. Council compromise document 9247/26. European Parliament first reading, 16.06.2026. Council adoption, 29.06.2026. Directive (EU) 2024/2853 (product liability). AILD withdrawal, OJ C/2025/5423, 06.10.2025. European Commission, EU Icons for labelling AI-generated content (10.06.2026). European AI Office signatory communications, June-July 2026. Bird & Bird, Gibson Dunn, White & Case, Covington, Lewis Silkin, A&O Shearman. EPRS, 2026.

CLOSING

Build *one integrated programme*, not six silos.

The EU AI Act is the world's first comprehensive AI law - but it is not the only regulation that touches your AI portfolio. Risk classification is driven by intended purpose, not technology. AI literacy is a legal duty, not a nice-to-have. Penalties bite hardest where literacy is missing. Transparency, GPAI duties, sector rules, and the Omnibus simplifications all converge on one principle.

The SMEs that thrive will not be the ones with the longest binders. They will be the ones with traceable decisions: who owns the system, what risks were identified, what was done about them, and when. Build that record now, while the deadlines move and the standards land.

A REMINDER ON MATERIALITY

A two-person startup carries the same core legal obligations as a Fortune 500 enterprise. SMEs receive lighter procedural treatment - simplified docs, proportionate fines, sandbox access - not exemption from the substance. Plan accordingly.

01

Start small. Start now.

A two-page risk register kept current is more defensible than a fifty-page document never opened. Begin with the minimum viable set - system description, risk register, oversight protocol, incident log - and expand as standards land.

02

Name one owner.

In a small team, you do not need a new hire - you need one named person with the authority to pause deployment if something is wrong. Without owned accountability, nothing else in the programme holds.

03

Treat it as a living programme.

The Act will keep evolving - delegated acts, implementing regulations, harmonised standards, Commission guidelines, the Omnibus and the digital fitness check. Put AI governance on the standing agenda, not the annual one.

Regulation alone cannot protect us. The future requires AI literacy, critical thinking, and trustworthy AI that serves human dignity. Human agency is the driver.

This guide is provided for informational purposes only and does not constitute legal advice. References to Regulation (EU) 2024/1689 and to the Digital Omnibus on AI proposal reflect the position as of publication.

GLOSSARY

Glossary

This glossary provides plain-language explanations of key concepts used throughout this guide. These explanations are intended to support understanding and are not a substitute for the legal definitions contained in Regulation (EU) 2024/1689.

AI System
A machine-based system that generates predictions, recommendations, content, or decisions based on input data and can influence physical or virtual environments. Under the EU AI Act, an AI system is distinguished from traditional software by its ability to infer outputs rather than simply follow predefined rules.

AI Literacy
The skills, knowledge and understanding needed to use AI responsibly, recognise risks and limitations, and know when human judgement should override AI-generated outputs. AI literacy is a mandatory requirement under Article 4 of the EU AI Act.

AI Model
The underlying AI capability or engine that performs tasks such as generating text, images or predictions. An AI model becomes part of an AI system when integrated into a specific application or business process.

Provider
An organisation that develops an AI system or places it on the market under its own name or brand. Providers carry the most extensive compliance obligations under the AI Act.

Deployer
An organisation that uses an AI system in its professional activities. Most SMEs using AI tools such as ChatGPT or Microsoft Copilot act as deployers rather than providers.

General-Purpose AI (GPAI) Model
A foundation model capable of performing a wide range of different tasks and being integrated into many downstream applications. Examples include GPT-4, Gemini, Claude and Llama.

Intended Purpose
The specific use for which an AI system is designed, marketed or deployed. Under the AI Act, risk classification is determined primarily by intended purpose rather than by the technology itself.

High-Risk AI System
An AI system that can significantly affect health, safety or fundamental rights and is therefore subject to enhanced regulatory obligations, including documentation, human oversight and risk management requirements.

Limited-Risk AI
AI systems that are primarily subject to transparency requirements, such as informing users that they are interacting with AI or labelling AI-generated content.

GLOSSARY

Glossary

Minimal-risk AI
AI applications that do not trigger specific regulatory obligations under the AI Act, apart from the AI literacy requirement (productivity assistants, spam filters and recommendation engines).

Human Oversight
Measures ensuring that people remain able to monitor, review, intervene in, or override AI-generated outputs when necessary.

Risk Management System
A documented process for identifying, assessing, mitigating and monitoring risks throughout the lifecycle of an AI system.

Technical Documentation
The records and evidence describing how an AI system works, how it was developed, its limitations, risks and compliance measures.

Conformity Assessment
A formal process used to demonstrate that a High-Risk AI system complies with the requirements of the AI Act before being placed on the market.

CE Marking
The marking that indicates an AI system or product meets applicable European Union legal requirements and may be placed on the EU market.

Transparency Obligation
Requirements to inform users when AI is involved, disclose AI-generated content, and provide information necessary for trustworthy use of AI systems.

Fundamental Rights
Basic rights protected under EU law, including privacy, non-discrimination, human dignity, freedom of expression and fair treatment. The AI Act is designed to help protect these rights when AI is used.

Notified Body
An independent organisation authorised to assess whether certain High-Risk AI systems comply with EU requirements before they are placed on the market.

EU AI Office
The European Commission body responsible for oversight of General-Purpose AI models and coordination of AI Act implementation across the European Union.

Incident
An event where an AI system causes, or may cause, harm to health, safety, fundamental rights or other protected interests and may need to be reported to authorities.

AI Governance
The policies, processes, controls and responsibilities that ensure AI is used safely, ethically and in compliance with applicable regulations.

Build the programme. *Not the binder.*

A practical guide for organisations navigating the EU AI Act and the Digital Omnibus simplifications of 2026.

REGULATIONS

EU 2024/1689 — in force 1 Aug 2024.

Digital Omnibus Regulation on AI (EU) 2026/1744 in force from 27 July 2026.

USE THIS GUIDE

As a working reference for SMEs and organizations. Confirm specific obligations with qualified counsel or lawyer.

STAY CURRENT

ai-act-service-desk. ec.europa.eu

European AI Office · EDPB · ENISA

Some images and icons used in this guide are sourced from Canva and are used in accordance with Canva's license terms. AI tools were also used in the development of this guide, including in the preparation, editing, and processing of selected texts and images.