Standard Intelligence

White paper

EU AI Act Compliance Intelligence

A business and compliance guide to the whole regime

PUBLISHED

February 2026

WEB

standardintelligence.com

AUDIENCE

For boards and compliance leaders

CONTENTS

Summary 3
1. The Act reaches an organisation through its market or its output, not its location 4
2. Obligations arrive on a staged calendar, and the earliest are already enforceable 5
3. Classification decides which obligations apply, so it comes first 6
3.1 Prohibited practices are banned, with only narrow exceptions set in the Act 6
3.2 High-risk systems are identified by product legislation or by listed use 7
3.3 Transparency obligations attach to particular functions 8
3.4 Most systems carry no specific obligations beyond literacy 8
3.5 Whether the organisation is provider or deployer is a second classification 8
3.6 Classification has to be repeated as systems change 9
4. A high-risk system needs all eight pillars, evidenced together 9
5. General-purpose models carry their own duties, which reach those who build on them 12
6. Fines are tiered, and civil liability can exceed them 13
7. The Act has to be run with the regimes it overlaps 14
8. Sector rules add to the Act where it already bites hardest 15
9. Governance needs a lead with the authority to stop a deployment 17
Where to start, and what it costs 18
Standard Intelligence supports the inventory, the classification and the evidence 19
Basis and Standing 20
Sources 21

Summary

The EU AI Act (Regulation (EU) 2024/1689) is long, but for a board or a compliance leader it comes down to a short sequence of decisions. Its obligations differ in kind, not only in degree, from one class of system to the next, so the order in which those decisions are taken matters.

The Act reaches an organisation through its market or its output, not its location. It applies to providers placing AI systems or general-purpose AI models on the EU market, to deployers established or located in the EU, and to providers and deployers elsewhere whose systems' output is used in the EU.

The obligations arrive on a staged calendar, and the first are already enforceable. The prohibitions have applied since 2 February 2025 and carry the highest fines, and the duties on general-purpose AI models have applied since 2 August 2025. Under the Act as amended by Regulation (EU) 2026/1744 (the Digital Omnibus on AI), high-risk obligations apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I; systems already on the market by then are caught only if their design changes significantly, except those intended for public authorities, which must comply by 2 August 2030.

Classification is the most consequential compliance decision, so it comes first. Every system falls into one of four tiers (prohibited, high-risk, transparency obligations, minimal risk), and nothing else can be settled until the tier is known. A provider that classifies a high-risk system as not high-risk in order to circumvent the requirements can be fined, and an organisation can become a provider by rebranding or modifying a high-risk system someone else built.

A high-risk system needs all its requirements met and evidenced together. We group them into eight pillars (risk management, data governance, technical documentation, logging, transparency, human oversight, accuracy and robustness, and post-market monitoring), held together by the provider's quality management system. A weakness in one shows up in the others.

General-purpose models sit on a separate track that reaches those who build on them. Their providers owe documentation, copyright and training-content duties, with more for models with systemic risk, and an organisation building a system on such a model depends on the information its provider supplies.

Fines are tiered, and civil liability can exceed them. The ceilings are EUR 35 million or 7% of worldwide annual turnover for prohibited practices, EUR 15 million or 3% for most other operator obligations, and EUR 7.5 million or 1% for supplying misleading information. For products placed on the market after 8 December 2026, the revised Product Liability Directive presumes a product defective where it breaches mandatory product safety requirements, which may include AI Act requirements, and consumer representative actions can now be brought for AI Act infringements.

The Act has to be run with the regimes it overlaps. The GDPR, the Digital Services Act, the Cyber Resilience Act, the Medical Device Regulation and sectoral financial and employment rules reach the same systems, and separate programmes for each leave

gaps at the points where they meet. Healthcare, financial services, manufacturing and education each add their own sector rules to the Act's.

Governance needs a lead with the authority to stop a deployment. We recommend placing AI oversight within the board's existing risk and audit committees, appointing an AI governance lead whose reporting line bypasses the business units it oversees, and reaching the third of our five maturity levels before the Annex III date.

Five actions carry most of the value. Audit the portfolio against Article 5 now; build a complete AI inventory; classify every system and allocate governance accordingly; appoint governance leadership with real authority; and start building documentation, oversight, monitoring and incident-response capability well before the high-risk dates.

This paper states the position as at 11 October 2026. Basis and Standing, at the end, sets out which propositions are binding law, which are guidance, and which are our own recommended frameworks and controls.

1. The Act reaches an organisation through its market or its output, not its location

The AI Act entered into force on 1 August 2024, at the end of a legislative process that ran from the Commission's proposal of April 2021 to adoption in June 2024. The Commission describes it as "the first-ever comprehensive legal framework on AI worldwide". Its approach is risk-based, scaling requirements to the harm a particular use of AI can cause.

Its reach is set by Article 2(1), which applies the Act to:

- providers placing AI systems on the market or putting them into service in the EU, or placing general-purpose AI models on the EU market, wherever they are established;
- deployers of AI systems that are established or located in the EU;
- providers and deployers established outside the EU, where the output produced by the AI system is used in the EU.

A US firm selling credit decisioning tools to European lenders and an Indian software house supplying CV screening systems to German employers are both within scope. A UK healthcare AI provider is caught if it places systems on the EU market or its systems' output is used in the EU, but not merely because it serves NHS trusts in the United Kingdom. The Act's reach beyond the EU is comparable to that of the GDPR.

The Commission's 2021 proposal (COM(2021) 206, explanatory memorandum, section 1.1) set four specific objectives: that AI systems on the Union market are safe and respect fundamental rights and Union values; legal certainty to facilitate investment and innovation; better governance and effective enforcement of existing law on fundamental rights and safety as it applies to AI; and a single market for lawful, safe and trustworthy AI. The Act pursues them through obligations that arrive on a staged calendar and vary with each system's risk tier.

2. Obligations arrive on a staged calendar, and the earliest are already enforceable

Mapping obligations to dates is a governance task in its own right, because the first obligations to apply, the prohibitions, carry the highest fines. The calendar below follows Article 113 as amended by Regulation (EU) 2026/1744, in force since 27 July 2026.

Date	Milestone	Primary obligation	Maximum penalty
1 Aug 2024	Entry into force	Jurisdictional scope established	N/A
2 Feb 2025	Prohibitions and AI literacy apply	Prohibited practices (Art. 5); AI literacy (Art. 4)	Art. 5: EUR 35 million / 7% (penalty rules applied from 2 Aug 2025); Art. 4: no specific fine
2 Aug 2025	GPAI obligations and penalty rules apply	General-purpose AI model obligations; Chapter XII penalties	Commission fines on GPAI providers (Art. 101, EUR 15 million / 3%) from 2 Aug 2026
27 Jul 2026	Digital Omnibus on AI in force	Regulation (EU) 2026/1744 amendments apply, including the replaced Art. 4 and new Art. 4a; Arts. 102 to 110 apply (Art. 113(d))	N/A
2 Aug 2026	General date of application	Transparency (Art. 50) and remaining provisions; Commission GPAI enforcement	As applicable
2 Dec 2026	New prohibitions; Art. 50(2) for existing generative systems	Art. 5(1)(ba) and (bb); Art. 111(4)	EUR 35 million / 7% (Art. 5); EUR 15 million / 3% (Art. 50)
2 Dec 2027	High-risk AI systems (Annex III)	Chapter III high-risk requirements and obligations	EUR 15 million / 3%
2 Aug 2028	High-risk AI systems (Annex I)	As above, for Annex I products	EUR 15 million / 3%
2 Aug 2030	Existing high-risk systems used by public authorities	High-risk systems placed on the market or put into service before their high-risk date are caught only if their design changes significantly; those intended for use by public authorities must comply by this date (Art. 111(2))	EUR 15 million / 3%

> **The Article 5 audit is overdue.** The prohibitions have applied since 2 February 2025, and every day a prohibited system stays in use is a continuing infringement of the most heavily penalised provisions in the Act. An organisation that has scheduled its Article 5 audit as a medium-term task is already late.

3. Classification decides which obligations apply, so it comes first

The Act regulates uses of AI by their risk rather than applying one set of rules to all of it. A chess-playing application and a system assessing creditworthiness for a mortgage present different risks to the people they affect, and they carry different duties.

There are four tiers: prohibited practices, high-risk systems, systems with transparency obligations (often called limited risk) and minimal-risk systems. The obligations differ between the tiers in kind as well as in degree, which is why nothing else in a compliance programme can be settled until each system's tier is known.

3.1 Prohibited practices are banned, with only narrow exceptions set in the Act

Article 5 prohibits certain practices whatever the commercial justification or technical design, subject only to the narrow exceptions written into the Article itself. The Act treats them as "particularly harmful and abusive" because they contradict Union values and fundamental rights (Recital 28). The prohibited practices are:

- subliminal, purposefully manipulative or deceptive techniques that materially distort behaviour in ways that cause, or are likely to cause, significant harm;
- exploiting vulnerabilities arising from age, disability or a specific social or economic situation;
- social scoring, by public or private actors, that leads to detrimental treatment in unrelated contexts or to treatment that is unjustified or disproportionate;
- assessing the risk that a person will commit a criminal offence based solely on profiling or on personality traits;
- untargeted scraping of facial images to create or expand facial recognition databases;
- emotion recognition in the workplace and in education institutions, except for medical or safety reasons;
- biometric categorisation that infers race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation;
- real-time remote biometric identification in publicly accessible spaces for law enforcement, save for three narrowly defined objectives, each use requiring prior authorisation by a judicial or independent administrative authority or, in a duly justified urgency, authorisation requested at the latest within 24 hours;
- from 2 December 2026, AI that generates or manipulates realistic intimate images of an identifiable person without that person's explicit consent, or child sexual abuse material (Article 5(1)(ba) and (bb), inserted by Regulation (EU) 2026/1744).

The emotion recognition ban reaches only systems that infer emotions from biometric data, because that is how the Act defines an emotion recognition system (Article 3(39)). An employer that analyses staff messages for sentiment is outside this prohibition, though the processing still has to satisfy the GDPR.

No derogation equivalent to Article 6(3), which can take an Annex III system out of the high-risk tier, applies to Article 5. An organisation that finds prohibited functionality in a system it operates should stop the prohibited use without delay, even where removing

the feature cleanly from a larger platform will take longer. We expect regulators to look for documented urgency rather than a long planning cycle.

3.2 High-risk systems are identified by product legislation or by listed use

High-risk systems carry the Act's most extensive requirements. A system becomes high-risk by one of two routes.

Route A: product legislation (Article 6(1)). An AI system is high-risk where it is a safety component of a product, or is itself a product, covered by the Union harmonisation legislation listed in Annex I (such as medical devices, machinery, civil aviation and motor vehicles), and that product must undergo third-party conformity assessment. For high-risk AI systems related to products in Annex I Section B (including civil aviation, motor vehicles and, since Regulation (EU) 2026/1744, machinery), only Article 6(1), Article 60a and Articles 102 to 112 of the Act apply (Article 2(2)). The high-risk requirements reach the product through its own sectoral legislation.

Route B: listed use (Article 6(2)). An AI system is high-risk where it is intended for one of the uses listed in the eight areas of Annex III, unless the Article 6(3) derogation applies. The derogation covers systems that do not pose a significant risk of harm, for example because they perform only a narrow procedural or preparatory task, but it never applies to a system that profiles natural persons.

Annex III area	Examples of in-scope systems	Key risk driver
Biometrics	Remote biometric identification (facial, gait or voice), biometric categorisation by sensitive attributes, emotion recognition	Identity and privacy rights
Critical infrastructure	Safety components in the management and operation of critical digital infrastructure, road traffic, and water, gas, heating or electricity supply	Safety and societal continuity
Education and training	Admissions, assessment of learning outcomes, monitoring of students during tests	Access to opportunity
Employment and HR	CV screening, task allocation, performance monitoring	Economic rights and discrimination
Essential services	Credit scoring of natural persons; risk assessment and pricing in life and health insurance; eligibility for public assistance benefits (including housing benefits); emergency call triage	Financial exclusion
Law enforcement	Risk assessment of offending not based solely on profiling, polygraph-type tools, evaluation of the reliability of evidence, profiling in investigations (individual crime prediction based solely on profiling is prohibited under Article 5(1)(d))	Fundamental rights and liberty
Migration and border control	Assessment of asylum applications, security risk scoring	Vulnerability of affected persons

Annex III area	Examples of in-scope systems	Key risk driver
Administration of justice and democratic processes	Judicial decision support, alternative dispute resolution, AI intended to influence elections or referenda	Rule of law, due process and democratic integrity

Misclassification is costly. Where a market surveillance authority establishes that a provider misclassified a system as not high-risk in order to circumvent the Chapter III requirements, the provider is subject to fines under Article 99 (Article 80(7)). Article 80(7) does not name a tier, but the provider's obligations under Article 16 sit in Article 99(4): up to EUR 15 million or 3% of worldwide annual turnover, whichever is higher, and whichever is lower for SMEs and small mid-caps. Civil liability for harm caused by a non-compliant system can add to that exposure (section 6).

3.3 Transparency obligations attach to particular functions

Article 50 sets transparency duties for systems that perform certain functions, whatever their tier:

- providers must design systems that interact directly with people so that those people are informed they are dealing with AI;
- providers of systems that generate synthetic audio, image, video or text must mark the outputs in a machine-readable format so that they are detectable as artificially generated, with technical solutions that are "effective, interoperable, robust and reliable" as far as technically feasible;
- deployers of emotion recognition or biometric categorisation systems must inform the people exposed to them;
- deployers of systems that generate or manipulate deep fakes must disclose that the content has been artificially generated or manipulated.

Providers of generative systems placed on the market before 2 August 2026 have until 2 December 2026 to comply with the marking duty (Article 111(4)). In practice, marking means watermarking or provenance metadata that survives distribution, which for many providers is a new engineering capability. A system that triggers Article 50 may also be high-risk, in which case both sets of duties apply.

3.4 Most systems carry no specific obligations beyond literacy

The Commission states that "the vast majority of AI systems currently used in the EU" are minimal risk. They carry no specific obligations under the Act beyond the Article 4 duty on providers and deployers to take measures supporting AI literacy, although the Article 5 prohibitions apply to every system. For an organisation with a mixed portfolio, classification will therefore usually place most systems here, which concentrates compliance investment on the few that need it.

3.5 Whether the organisation is provider or deployer is a second classification

The Act distinguishes providers, who develop AI systems or have them developed and place them on the market under their own name, from deployers, who use them under their authority. Each carries distinct obligations, and the provider's are heavier. Under Article 25(1), a distributor, importer, deployer or other third party becomes the provider of a high-risk system where it:

(a) puts its name or trademark on a high-risk system already placed on the market or put into service;

(b) makes a substantial modification to such a system so that it remains high-risk;
(c) modifies the intended purpose of a system that was not high-risk, including a general-purpose AI system, so that it becomes high-risk.

An organisation offering white-labelled AI products, or adapting a licensed system to a new use, should assess its branding and modification practices against these tests with legal counsel. The provider's full obligations follow if any of them is met.

3.6 Classification has to be repeated as systems change

A system's tier can change after deployment. A customer service chatbot that is given medical advisory functions may move into the high-risk tier, and a system that continues to learn after deployment may change its behaviour in ways that alter its risk profile.

We recommend monitoring systems that sit near a classification boundary and reassessing them before a change of function or use takes them across it. Reassessment after the change leaves the organisation operating a misclassified system in the meantime.

4. A high-risk system needs all eight pillars, evidenced together

We group the high-risk requirements into eight pillars. The grouping is ours, not the Act's; the requirements for high-risk systems are Articles 8 to 15 (Chapter III, Section 2), and providers' obligations also include the quality management system in Article 17 and post-market monitoring in Article 72. The pillars apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I, subject to Article 111(2) for systems already on the market.

The pillars are not independent checkboxes; a weakness in data governance creates a gap in risk management, and inadequate documentation undermines conformity assessment. The Article 17 quality management system is what holds them together, and a high-risk system needs all eight, evidenced.

Pillar 1, risk management (Article 9), runs for the system's whole lifecycle. Providers must establish, implement, document and maintain a risk management system that identifies foreseeable risks to health, safety and fundamental rights and adopts measures to address them. It must test the system before it is placed on the market, judge residual risk acceptable and evaluate emerging risks using post-market monitoring data (Article 9(2), (5) and (8)).

The risk management system must be proportionate to the risk profile of the AI system it covers. In our experience, engaging the people affected (patients in healthcare, workers in manufacturing, applicants in financial services) surfaces risks that purely technical teams miss.

Pillar 2, data governance (Article 10), has to reconcile representativeness with data minimisation. Training, validation and test data must be subject to

governance practices covering relevance, representativeness, errors, completeness and the examination of possible biases, with measures to detect, prevent and mitigate them (Article 10(2) to (4)). The tension between the GDPR's data minimisation principle and the Act's demand for representative data is one of the most practically important overlaps in the regime, and an organisation should document how it resolves it.

Since 27 July 2026, Article 4a (which replaced Article 10(5)) has permitted providers of high-risk systems to process special categories of personal data where strictly necessary for bias detection and correction, under six cumulative conditions, including that the task cannot be effectively fulfilled by processing other data, such as synthetic or anonymised data. Article 4a(2) extends the permission, under the same conditions, to deployers and to other AI systems and models, and for those cases creates no obligation to carry out bias detection. It is a permission, not a legal basis; the processing still needs a lawful basis under Article 6 GDPR, and the condition under Article 9(2) GDPR is point (g), read with Article 4a.

In our experience, data governance works best with unambiguous roles: a governance committee with executive representation, data stewards managing day-to-day quality, engineers maintaining infrastructure and data scientists ensuring appropriate use in model development.

Pillar 3, technical documentation (Article 11), must exist before the system is placed on the market. Providers must draw it up beforehand and keep it up to date, so that it demonstrates compliance and gives national competent authorities and notified bodies the information they need to assess it (Article 11(1) and Annex IV). It covers the general description and intended purpose of the system, design choices and assumptions, training methodology and data governance, risk management measures, conformity assessment and the post-market monitoring plan.

Code written by AI tools raises a question the Act does not address. Annex IV asks for design choices and the rationale behind them, which are harder to document for a risk management module, data validation routine or safety filter that no engineer designed line by line. An organisation using AI coding tools in safety-critical development should document its review and testing of that code in unusual detail.

Pillar 4, logging (Article 12), must make the system's operation traceable. High-risk systems must technically allow for the automatic recording of events (logs) over their lifetime (Article 12(1)). The logs must enable the identification of situations that may result in a risk or a substantial modification, facilitate post-market monitoring and support the deployer's monitoring of the system's operation (Article 12(2)).

We recommend version control that goes beyond what suffices for conventional software, capturing the complete context of each model version: training data, configuration files, the computational environment and the random seeds that determine initialisation.

Pillar 5, transparency to deployers (Article 13), requires specific instructions. Systems must be designed so that deployers can interpret their output and use it appropriately. The instructions for use must state, among other things, the level of accuracy and its metrics (Article 13(3)(b)(ii)), "when appropriate, its performance regarding specific persons or groups of persons on which the system is intended to be used" (Article 13(3)(b)(v)) and the human oversight measures (Article 13(3)(d)).

We recommend instructions that say which categories of decision require human review, which indicators suggest the system is operating outside its reliable range and how to handle boundary cases. Performance claims also need context. A 95% accuracy rate stated in isolation may sound high, but if human experts in the same task achieve 99%, using the system in their place would lower the quality of decisions.

Pillar 6, human oversight (Article 14), places duties on the provider and on the deployer. The system must be designed and developed so that natural persons can effectively oversee it while it is in use (Article 14(1)), and the provider must supply it so that the people assigned to oversight can understand it, interpret its output and decide to disregard, override or stop it (Article 14(4)). The oversight measures must be commensurate with the risks, level of autonomy and context of use (Article 14(3)). The Act does not require a person to check every output, except that remote biometric identification must be verified separately by two people (Article 14(5)).

The deployer must then assign oversight to people with "the necessary competence, training and authority, as well as the necessary support" (Article 26(2)). In practice, oversight tends to take one of three forms: continuous monitoring of the system's behaviour in real time, periodic review of decisions and outcomes, or review of each output before it takes effect. The appropriate form depends on the stakes of individual decisions and the pace at which the system operates.

Meaningful oversight requires tooling, training, time and staff. A review stage whose reviewers lack the information, authority or time to exercise judgement is unlikely to meet the deployer's Article 26(2) duty, and where per-decision review is the measure adopted, a nominal review will not meet it.

Pillar 7, accuracy, robustness and cybersecurity (Article 15), must hold throughout the lifecycle. High-risk systems must achieve appropriate levels of all three and perform consistently over their lifetime. They must be resilient to errors, faults and inconsistencies, and to attempts by unauthorised third parties to alter their use, outputs or performance, including AI-specific attacks such as data poisoning and adversarial inputs.

Pillar 8, post-market monitoring (Article 72), must produce information someone acts on. Providers must establish and document a post-market monitoring system that actively and systematically collects and analyses data on the performance of deployed systems, so that they can evaluate continuing compliance. Serious incidents must also be reported under Article 73 once a system is high-risk under Chapter III, to the market surveillance authority or, for systems under the AI Office's exclusive competence, to the AI Office (Article 75(1a)).

We recommend defining retraining triggers during initial development rather than improvising them when performance degrades. A dashboard that nobody reviews does not monitor anything.

5. General-purpose models carry their own duties, which reach those who build on them

The provisions on general-purpose AI models are the Act's most novel territory. A general-purpose model "displays significant generality and is capable of competently performing a wide range of distinct tasks" (Article 3(63)), so it can be used across an open-ended range of applications, and the Act gives it a regime of its own in Chapter V.

A model has systemic risk if it has high-impact capabilities or the Commission designates it (Article 51(1)). High-impact capabilities are presumed where the cumulative compute used to train the model is greater than 10^25 floating point operations (Article 51(2)).

Every general-purpose model provider owes four duties. Under Article 53(1), it must:

(a) draw up and keep up to date technical documentation of the model, including at least the information in Annex XI (among it the model's known or estimated energy consumption);
(b) make information and documentation available to providers of AI systems who intend to integrate the model;
(c) put in place a policy to comply with Union copyright law, including identifying and complying with reservations of rights under Article 4(3) of Directive (EU) 2019/790;
(d) publish a sufficiently detailed summary of the content used for training, using the AI Office's template.

Marking AI-generated content is a duty on providers of AI systems under Article 50(2), not a general-purpose model obligation.

Models with systemic risk owe four more. Under Article 55(1), their providers must evaluate the model, including conducting and documenting adversarial testing; assess and mitigate possible systemic risks at Union level; track, document and report serious incidents to the AI Office and, as appropriate, national competent authorities without undue delay; and ensure an adequate level of cybersecurity protection.

The Code of Practice offers a recognised route to showing compliance. The General-Purpose AI Code of Practice was published on 10 July 2025, and the Commission and the AI Board have confirmed that it is "an adequate voluntary tool" (Commission opinion of 1 August 2025). Providers may rely on it to demonstrate compliance until a harmonised standard is published (Articles 53(4) and 55(2)); those who do not adhere to it "shall demonstrate alternative adequate means of compliance for assessment by the Commission", which is likely to require more evidence.

An organisation building on a general-purpose model inherits a dependency. The Article 53(1)(b) information is what a downstream provider uses to understand the model's capabilities and limitations and to meet its own obligations. If that information is thin, the downstream provider's documentation, risk management and instructions for use will be thin too, so the model provider's compliance posture becomes part of the downstream provider's.

6. Fines are tiered, and civil liability can exceed them

Enforcement is split between the AI Office and national authorities. The AI Office, within the Commission, has exclusive powers over providers of general-purpose AI models (Article 88). From 2 August 2026 it also has exclusive competence over AI systems based on a general-purpose model where the same provider develops both (with exceptions), and over AI systems that constitute or are integrated into designated very large online platforms or search engines (Article 75(1), as replaced by Regulation (EU) 2026/1744). It supports harmonised implementation and provides the secretariat of the AI Board, which has one representative per Member State, is chaired by one of them, and is attended by the AI Office without a vote (Article 65).

Each Member State must designate at least one notifying authority and one market surveillance authority to enforce the Act within its territory (Article 70(1)). For an organisation operating in several Member States, an enforcement action in one may prompt review in others.

The fines are set to be dissuasive. Article 99(1) requires penalties to be "effective, proportionate and dissuasive", and the maxima scale with turnover.

Tier	Violation category	Maximum penalty	Example at EUR 5 billion turnover
Tier 1	Prohibited practices (Article 5)	EUR 35 million or 7% of worldwide annual turnover	EUR 350 million
Tier 2	High-risk non-compliance and other operator obligations (Article 99(4))	EUR 15 million or 3% of worldwide annual turnover	EUR 150 million
Tier 3	Supplying incorrect, incomplete or misleading information to notified bodies or national competent authorities (Article 99(5))	EUR 7.5 million or 1% of worldwide annual turnover	EUR 50 million

All maxima are whichever is higher. For SMEs, including start-ups, the lower figure applies in every tier (Article 99(6)), and for small mid-caps the lower figure applies to Tiers 2 and 3 (Article 99(6a)). Providers of general-purpose AI models are fined separately by the Commission under Article 101, up to EUR 15 million or 3%, from 2 August 2026, and for the AI systems it supervises directly under Article 75(1) the AI Office can impose fines on the Article 99 scale and periodic penalty payments (Article 75c).

How an organisation responds affects the fine. In setting a fine, authorities must take account of, among other things, the nature, gravity and duration of the infringement and its consequences; the number of people affected and the damage they suffered; any action taken to mitigate that harm; the degree of cooperation with the authorities; whether the operator itself notified the infringement; and fines already imposed for the same conduct (Article 99(7)). An organisation that discovers an infringement, reports it promptly and corrects it at once puts several of these factors in its favour.

Civil liability can compound the fine. For products placed on the market or put into service after 8 December 2026, the revised Product Liability Directive (Directive (EU)

2024/2853) presumes a product defective where it does not comply with mandatory product safety requirements in Union or national law that are intended to protect against the risk of the damage suffered (Article 10(2)(b)), and AI Act requirements may be relied on for this purpose. The claimant must still prove the damage and the causal link (Article 10(1)) unless a further presumption applies, and every presumption is rebuttable (Article 10(5)). The Directive compensates natural persons only, excluding property used exclusively for professional purposes.

A regulatory fine, even at 3% of turnover, is capped and, in some jurisdictions, insurable. A presumption of defectiveness applied across every claim arising from a non-compliant system has no cap, and the cumulative liability could exceed any fine.

Since 27 July 2026, the AI Act has been listed in Annex I to the Representative Actions Directive (Directive (EU) 2020/1828), so qualified entities can bring representative actions on behalf of consumers harmed by infringements of it. An AI system that discriminates systematically in credit decisions could face such actions in several Member States at once. Meeting the substantive requirements that both the regulatory and the civil regimes demand is the most reliable way to limit that exposure.

No fine has yet been confirmed. As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source. The first significant enforcement actions will shape how regulators read the grey areas in classification and the practical scope of the prohibitions, and an organisation that has documented its decisions will be better placed to defend them.

7. The Act has to be run with the regimes it overlaps

The AI Act applies alongside the GDPR, the Digital Services Act, the Cyber Resilience Act and sectoral regimes such as the Medical Device Regulation and financial services regulation. Running a separate programme, with its own tools and team, for each regime costs more and leaves gaps at the points where they meet.

The GDPR reaches most of the same systems. Its lawful basis requirements, data minimisation principle, rights to meaningful information about the logic involved in automated decisions (Articles 13 to 15) and rules on automated decision-making interact with the Act across almost every pillar.

Article 22 GDPR prohibits in principle, or makes conditional, decisions based solely on automated processing that have legal or similarly significant effects (SCHUFA, C-634/21), and human involvement takes a decision outside it only where it is “meaningful, rather than just a token gesture” and carried out by someone with the authority and competence to change the decision (Article 29 Working Party, WP251rev.01). A deployer should design its oversight so that it meets its own duties under the Act (Article 26(1) and (2)) and Article 22 of the GDPR together, using the oversight measures the provider has built in under Article 14.

The Digital Services Act overlaps on risk assessment. Very large online platforms deploy recommender systems, content moderation and targeting tools that can engage both regimes, and the systemic risk assessments the DSA requires overlap substantially with the Act's risk management requirements. In our view a single process designed to satisfy both is feasible, and it should be the objective.

The Cyber Resilience Act can supply the Act's cybersecurity evidence. Many AI-enabled products with digital elements fall within both regimes, subject to the CRA's exclusions for sectors such as medical devices and motor vehicles.

A product classified as a high-risk AI system is deemed to comply with the cybersecurity requirements of Article 15 of the AI Act where it meets the CRA's essential cybersecurity requirements and its EU declaration of conformity demonstrates the required level of protection (CRA Article 12(1)). The CRA's reporting obligations have applied since 11 September 2026 and its other obligations apply from 11 December 2027, so an organisation should plan a single assessment strategy from the outset, using notified bodies competent in both regimes.

Medical devices are assessed once, against both regimes. For high-risk AI medical devices, the AI Act requirements are assessed within the MDR conformity assessment procedure by the MDR notified body (Article 43(3)), on a single set of technical documentation. The joint guidance of the AI Board and the Medical Device Coordination Group (AIB 2025-1 / MDCG 2025-6, June 2025) describes "a simultaneous and complementary application" of the two regimes rather than a general rule that one prevails, although where both Annex I and Annex III of the Act would apply, "Annex I alone should prevail".

The Commission's proposal of December 2025 to revise the MDR and IVDR (COM(2025) 1023) would move devices to Section B of Annex I, which would change this position; as at 11 October 2026 it had not been adopted.

> **Building to the EU standard is usually the efficient global strategy.** For a multinational organisation, building to the AI Act and supplementing for other jurisdictions will usually cost less than maintaining parallel systems for the United Kingdom, Brazil and other markets developing their own AI rules. ISO/IEC 42001 can support parts of an AI Act programme across several regimes at once, but it is not a harmonised standard and gives no presumption of conformity.

8. Sector rules add to the Act where it already bites hardest

Healthcare. Systems supporting clinical diagnosis, treatment recommendations, surgical planning and patient management are high-risk under the Act where they are medical devices requiring notified-body conformity assessment under the MDR (class IIa and above, and certain class I devices); the MDR itself uses risk classes rather than a "high-risk" category. Errors can directly harm patients, and systems trained on unrepresentative data can disadvantage whole patient populations. Compliance

therefore needs particular attention to explainability, so that clinicians can understand and challenge a recommendation, and to demographic bias.

Financial services. AI used to evaluate the creditworthiness of natural persons or establish their credit score (other than for fraud detection) is high-risk under Annex III point 5(b), as is risk assessment and pricing in life and health insurance under point 5(c), subject to the Article 6(3) derogation. Deployers of these systems must also carry out a fundamental rights impact assessment before deploying them (Article 27(1)).

Explainability matters particularly in credit. Consumer credit law already requires a lender that rejects an application on the basis of a database consultation to tell the consumer the result and the database consulted (Directive 2008/48/EC, Article 9(2)), and from 20 November 2026 the revised Consumer Credit Directive gives a consumer whose creditworthiness assessment involves automated processing the right to request human intervention, including an explanation, to express a view and to request a review (Directive (EU) 2023/2225, Article 18(8)).

Insurance risk assessment faces scrutiny for indirect discrimination through proxy variables. A system that charges higher premiums to residents of a predominantly minority neighbourhood may discriminate even if ethnicity is not an input, so an insurer should be able to show that its rating factors are actuarially justified and genuinely predictive.

Manufacturing. Defect detection, quality control, predictive maintenance and worker safety systems raise classification questions that turn on deployment context and integration with physical systems. An AI system that operates machinery capable of harming workers, or controls a safety-critical process, may be high-risk under Article 6(1) where it is a safety component requiring third-party assessment under the Machinery Regulation, which applies from 20 January 2027. Under Regulation (EU) 2026/1744, machinery now sits in Annex I Section B, and the AI requirements will reach it through delegated acts under the Machinery Regulation, which are to apply by 2 August 2028.

Worker-facing systems need extra care. AI that monitors employee performance or manages workflows requires works council consultation in many Member States before deployment, and for a high-risk system a deployer that is an employer must inform workers' representatives and the affected workers before putting it into service at the workplace (Article 26(7)).

Education. Student monitoring, remote proctoring, admissions and learning analytics systems face heightened scrutiny. The power imbalance between institutions and students, the formative nature of educational decisions and the long-term consequences of assessments all amplify the effect on fundamental rights. We recommend involving student representatives and pedagogical experts throughout design.

9. Governance needs a lead with the authority to stop a deployment

AI governance works across three layers. Strategic governance, at board and executive level, sets AI risk appetite, approves policies, allocates resources and aligns AI governance with corporate strategy. Operational governance, in business units and product teams, implements policies, manages day-to-day compliance and escalates issues; technical governance, at system and development level, builds compliance into engineering workflows and pipelines.

In our view, board oversight of AI belongs within the existing risk and audit committees rather than in a separate AI committee. Board members need not be technical experts, but they do need to understand AI risk in terms of fines of up to 7% of worldwide turnover, reputational damage and liability for AI decisions that harm individuals.

The AI governance lead is the role on which the rest depends. The lead maintains the AI inventory, coordinates risk assessments, manages conformity assessment, acts as the regulatory contact point and reports to the board. The role needs real authority, including standing to challenge business unit leaders and to halt a non-compliant deployment; a lead who cannot stop a product team deploying a system without a completed conformity assessment has the title without the function.

We recommend three supporting roles. AI risk officers conduct and maintain risk assessments for individual systems or clusters, monitor post-market performance and maintain risk registers; data governance officers own data quality and lineage across development pipelines; and technical compliance leads translate regulatory obligations into development standards and automated pipeline checks.

Maturity is a better measure than a binary pass. Our framework describes five levels:

Level	Designation	Characteristics	Compliance risk
1	Initial	Ad hoc; no formal processes; AI deployed without systematic oversight	Critical
2	Developing	Basic processes defined; inconsistent application across teams	High
3	Managed	Systematic implementation; consistent documentation; automated pipeline checks	Moderate
4	Advanced	Predictive governance; risk-driven rather than compliance-driven; industry engagement	Low
5	Optimised	Governance culture embedded; continuous improvement; standards contribution	Minimal

Most organisations aiming for initial compliance by 2 December 2027 (Annex III) should target Level 3. Moving from Level 2 to Level 3 requires automated pipeline checks that prevent non-compliant deployment, consistent documentation quality across engineering teams, systematic bias detection with documented outcomes, and oversight mechanisms that are themselves monitored.

Resource plans need to be realistic. Maturity cannot be bought through technology alone; it requires people, processes and cultural change, and an organisation that

underinvests in people will find its tools unused and its dashboards populated but not acted on.

Governance that lasts has four features. We recommend:

- principle-based frameworks with specific processes beneath them, so that when a requirement changes, as standards develop and enforcement clarifies grey areas, the organisation adjusts a process rather than redesigning from scratch;
- active tracking of legislative, standards and enforcement developments, including participation in consultations and codes of practice;
- compliance built into development workflows, with automated checks in the pipeline, so that a bias check run at model evaluation catches an issue before deployment rather than after;
- risk-driven rather than compliance-driven governance, with predictive indicators and independent assurance such as external red-teaming and algorithmic audit alongside internal validation.

An organisation with this capability will have the documentation, oversight records and compliance evidence to defend its decisions when an authority examines them. One that has built processes only on paper will find that examination far more difficult.

Where to start, and what it costs

The Act's decisions come in an order: what each system is, what tier it is in, who is responsible for it and what evidence it needs. The five actions below follow that order, dealing first with the most serious risks and then building the capability for the high-risk dates.

First, immediately: audit the portfolio against Article 5. Every AI system in operation should be assessed against the prohibited categories, and any prohibited use stopped without delay, with the removal of the underlying feature planned and documented as a matter of urgency. The 2 December 2026 additions in Article 5(1)(ba) and (bb), and the Article 50(2) marking deadline for existing generative systems on the same date, belong in the same review.

Second: build a complete and maintained AI inventory. Classification, risk assessment, governance and conformity assessment all depend on knowing which AI systems the organisation operates. The inventory should record each system's technical architecture, intended use, deployment context, user population and integration with other systems.

Third: classify every system and allocate governance accordingly. Apply the Act's tiers and the provider and deployer tests to each entry in the inventory, and concentrate governance investment on the highest-risk systems. Classification is not a one-time exercise; systems near a boundary need monitoring for changes of function or use.

Fourth: appoint governance leadership with real authority and a board-level mandate. The AI governance lead should report through a line that bypasses the

functions it oversees, typically to the general counsel or chief risk officer, with direct access to the board.

Then, starting now and ongoing: build documentation, oversight, monitoring and incident-response capability. The high-risk requirements apply from 2 December 2027 (Annex III) and 2 August 2028 (Annex I), subject to Article 111(2) for systems already on the market. The documentation architecture, human oversight arrangements and post-market monitoring in Articles 9 to 17 and 72, and the serious incident reporting in Article 73, take time to design, implement and integrate into existing workflows.

The first four actions cost mainly the time of people who know the organisation's systems. The fifth is the main build cost, and it does not shrink if started later; a late start makes it more expensive and more rushed, and more likely to produce documentation that satisfies the form of the requirements without their substance.

An organisation that has taken these steps can say, for each system, what tier it is in and why, who is responsible for it and what evidence supports its compliance. The Act will generate years of interpretive debate over classification, its overlaps with adjacent regimes and the practicalities of documentation and oversight. In our view its core logic is sound, however: AI systems that decide people's access to credit, healthcare, employment and justice should meet requirements for transparency, oversight, accuracy and accountability, and an organisation that governs its AI to that standard builds systems its customers, employees and communities can trust.

Standard Intelligence supports the inventory, the classification and the evidence

Understanding the whole regime is one task; operating against it across a portfolio is another. The Standard Intelligence platform is built to carry an organisation from understanding its obligations to evidencing them continuously.

Platform capability	How it supports whole-regime compliance
AI Regulatory Navigator	Turns a plain-language description of a system into a structured, cited assessment of which obligations apply across the regime.
AI Intelligence Centre	Holds the inventory of systems, their classifications and their documentation as a single source of truth.
AI Integrity Engine	Provides a continuous monitoring layer that helps keep high-risk systems within their obligations after deployment.
TRACE Maturity Model	Scores current capability across governance domains and produces a prioritised roadmap toward audit-ready conformity.

This paper is provided for educational and informational purposes only and does not constitute legal advice. It describes the AI Act, as amended by Regulation (EU) 2026/1744, and related law and guidance as at 11 October 2026. Organisations should seek qualified legal, regulatory and technical advice before taking compliance decisions.

Basis and Standing

Binding law. AI Act (Regulation (EU) 2024/1689), as amended by Regulation (EU) 2026/1744 (the Digital Omnibus on AI), adopted 8 July 2026 and in force 27 July 2026: Articles 2(1) and 2(2); 3(39) and 3(63); 4 (as replaced) and 4a (inserted, replacing Article 10(5)); 5, including points (ba) and (bb) of paragraph 1 from 2 December 2026; 6; 9 to 15; 17; 25(1); 26(2) and 26(7); 27(1); 43(3); 50; 51; 53; 55; 65; 70(1); 72; 73; 75(1) as replaced; 75(1a) (inserted); 75c; 80(7); 88; 99; 101; 111(2) and 111(4) as amended; 113 as amended, under which high-risk obligations apply from 2 December 2027 (Annex III) and 2 August 2028 (Annex I); Annexes I, III, IV and XI. Product Liability Directive (EU) 2024/2853, Articles 5(1), 6(1) and 10, applying to products placed on the market or put into service after 8 December 2026 (Article 2(1) as corrected on 7 May 2026). Representative Actions Directive (EU) 2020/1828, Annex I as amended by Article 110 of the AI Act. Cyber Resilience Act (Regulation (EU) 2024/2847), Articles 2(2), 12(1) and 71(2). MDR (Regulation (EU) 2017/745). Machinery Regulation (EU) 2023/1230, applying from 20 January 2027, as amended by Regulation (EU) 2026/1744. Directive 2008/48/EC, Article 9(2); Directive (EU) 2023/2225, Article 18(8), applying through national measures from 20 November 2026. GDPR (Regulation (EU) 2016/679), Articles 6, 9, 13 to 15 and 22, as interpreted in SCHUFA Holding (C-634/21). Digital Services Act (Regulation (EU) 2022/2065).

Proposals, not law. The Commission's proposal of December 2025 to revise the MDR and IVDR (COM(2025) 1023) had not been adopted as at 11 October 2026.

Guidance, not law. The General-Purpose AI Code of Practice (10 July 2025) and the Commission's opinion of 1 August 2025 on its adequacy; AIB 2025-1 / MDCG 2025-6 on the interplay of the medical devices regulations and the AI Act; the Article 29 Working Party's guidelines on automated decision-making (WP251rev.01), endorsed by the EDPB; the Commission's description of the Act on its regulatory framework page. The explanatory memorandum to COM(2021) 206 is cited for the proposal's objectives only. ISO/IEC 42001 is a voluntary management system standard, not a harmonised standard, and confers no presumption of conformity.

Not a regulatory metric. No instrument cited prescribes the eight pillars, the three forms of human oversight, the three governance layers, the AI governance lead and supporting roles, the five-level maturity framework or the five priority actions; these are this paper's own frameworks and recommended controls. The pillars group the substantive Chapter III, Section 2 requirements (Articles 9 to 15) with Articles 17 and 72; the oversight forms are a way of meeting Articles 14(3) and 26(2); and the governance lead's authority to halt a deployment supports, but is not required by, the provider's quality management system under Article 17. The version control, retraining triggers, instructions for use content and stakeholder engagement recommended in section 4 go beyond what the Act requires.

Research. The paper relies on no empirical research. Statements introduced with "in our experience" or "in our view" are the authors' professional judgement.

Enforcement position. As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source. Article 5 fines have been available since

2 August 2025, and Commission fines on general-purpose AI model providers under Article 101 since 2 August 2026.

Sources

All Article, Annex and Recital references are to Regulation (EU) 2024/1689 (the AI Act), as amended by Regulation (EU) 2026/1744, unless otherwise stated.

Primary legislation and case law

1. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng; consolidated text of 27 July 2026. https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng; European Commission, AI Act Service Desk (consolidated text by article). https://ai-act-service-desk.ec.europa.eu/en/ai-act
2. Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 (Digital Omnibus on AI), OJ L, 24.7.2026. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
3. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC. https://eur-lex.europa.eu/eli/dir/2024/2853/oj; corrigendum of 7 May 2026 (Article 2(1) date). https://eur-lex.europa.eu/eli/dir/2024/2853/corrigendum/2026-05-07/oj
4. Directive (EU) 2020/1828 on representative actions for the protection of the collective interests of consumers (Annex I as amended by Article 110 of the AI Act). https://eur-lex.europa.eu/eli/dir/2020/1828/oj
5. Regulation (EU) 2024/2847 (Cyber Resilience Act). https://eur-lex.europa.eu/eli/reg/2024/2847/oj
6. Regulation (EU) 2017/745 on medical devices (MDR). https://eur-lex.europa.eu/eli/reg/2017/745/oj
7. Regulation (EU) 2023/1230 on machinery (Machinery Regulation). https://eur-lex.europa.eu/eli/reg/2023/1230/oj
8. Directive 2008/48/EC on credit agreements for consumers. https://eur-lex.europa.eu/eli/dir/2008/48/oj
9. Directive (EU) 2023/2225 on credit agreements for consumers (national transposing measures apply from 20 November 2026). https://eur-lex.europa.eu/eli/dir/2023/2225/oj
10. Regulation (EU) 2016/679 (General Data Protection Regulation). https://eur-lex.europa.eu/eli/reg/2016/679/oj
11. Regulation (EU) 2022/2065 (Digital Services Act). https://eur-lex.europa.eu/eli/reg/2022/2065/oj
12. Case C-634/21, OQ v Land Hessen (SCHUFA Holding), judgment of 7 December 2023, ECLI:EU:C:2023:957. https://curia.europa.eu/juris/liste.jsf?num=C-634/21

Proposals and guidance

13. European Commission, Proposal for a Regulation laying down harmonised rules on artificial intelligence, COM(2021) 206 final, 21 April 2021 (explanatory memorandum, section 1.1). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206
14. European Commission, Proposal to simplify the medical devices and in vitro diagnostic medical devices regulations, COM(2025) 1023, December 2025 (not adopted as at 11 October 2026). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC1023
15. European Commission, Regulatory framework for AI (AI Act overview). https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
16. European Commission, General-Purpose AI Code of Practice (published 10 July 2025), and Commission opinion of 1 August 2025 on its assessment. https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai; https://digital-strategy.ec.europa.eu/en/library/commission-opinion-assessment-general-purpose-ai-code-practice
17. Artificial Intelligence Board and Medical Device Coordination Group, AIB 2025-1 / MDCG 2025-6, joint guidance on the interplay between the medical devices regulations and the AI Act (June 2025). https://health.ec.europa.eu/document/download/b78a17d7-e3cd-4943-851d-e02a2f22bbb4_en?filename=mdcg_2025-6_en.pdf
18. Article 29 Working Party, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, WP251rev.01, as last revised 6 February 2018, endorsed by the EDPB. https://ec.europa.eu/newsroom/article29/items/612053
19. ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system (subscription access). https://www.iso.org/standard/81230.html

Standard Intelligence

GET IN TOUCH

Turn the Act into five priority actions.

EMAIL

hello@standardintelligence.com

WEB

standardintelligence.com

LOCATION

Manchester, United Kingdom

AI Use: AI tools were used in research, editorial review and drafting support. All sources were retrieved and checked against the published record by the author, who holds editorial responsibility for the contents and for the conclusions drawn. This disclosure is made voluntarily. To the extent the Article 50(4) obligation in Regulation (EU) 2024/1689 applies to this paper as text published to inform the public on matters of public interest, the exception for content that has undergone human review and for which a person holds editorial responsibility applies.

STANDARD INTELLIGENCE LIMITED · REGISTERED IN ENGLAND AND WALES · COMPANY NO. 16737890