Standard Intelligence

White paper

Risk Classification

Classifying AI systems under Regulation (EU) 2024/1689

PUBLISHED

February 2026

WEB

standardintelligence.com

AUDIENCE

For AI governance, legal and procurement teams

CONTENTS

Summary 3
1. Every obligation in the Act attaches to a classification, so classification comes first 4
2. The four tiers differ in kind, not in degree 5
3. The obligations arrive on a staged calendar, and the first are already enforceable 6
4. Classification follows a fixed sequence, beginning with whether the system is an AI system at all 7
5. Article 5 comes first because a prohibited practice has no compliance route 8
6. A system is high-risk by one of two routes, and Annex III classifies the use, not the technology 9
7. Annex III lists eight areas, and the employment area reaches every employer 10
8. The Article 6(3) exception is narrow, never covers profiling, and must be documented and registered 11
9. General-purpose, mixed-purpose, agentic and embedded systems need the most care 12
10. Classification is also a role question, and a deployer can become a provider 13
11. Procurement is the cheapest point at which to check a classification 14
12. Sector context changes the classification more than the technology does 15
13. A classification stays defensible only if it is recorded and kept current 17
Where to start, and what it costs 18
Standard Intelligence makes each classification a recorded decision 19
Basis and Standing 19
Sources 20

Summary

Every obligation in the EU AI Act (Regulation (EU) 2024/1689) attaches to a classification. The Act imposes no general duty of care on AI and no universal transparency rule; it attaches distinct sets of obligations to distinct categories of system. Until an organisation knows how each system is classified, it cannot know what it owes.

Classification is therefore the first compliance decision, and the one from which cost, product design and legal exposure follow. The Act's four tiers (prohibited, high-risk, limited-risk and minimal-risk) differ in kind rather than degree. The gap that drives cost is between high-risk and limited-risk: a typical medium-complexity high-risk system can, in our experience, take six to eighteen months from initial compliance assessment to lawful market placement, while a limited-risk system mainly owes transparency to the people who use it or are exposed to it.

The obligations arrive on a staged calendar, and the first are already enforceable. The Article 5 prohibitions have applied since 2 February 2025, so an organisation that has not screened its portfolio against them may already be behind, as may a provider of general-purpose AI models on the obligations that applied from 2 August 2025. As amended by Regulation (EU) 2026/1744, the Act applies its high-risk requirements to Annex III systems from 2 December 2027 and to Annex I systems from 2 August 2028; a high-risk system already on the market by then is caught only on a significant change in its design, unless it is intended for use by public authorities, which must comply by 2 August 2030.

Classification follows a fixed sequence. Is the system an AI system within scope; is the practice prohibited by Article 5; is it high-risk under Article 6(1) or Annex III; does the Article 6(3) exception apply; do the Article 50 transparency duties apply; and in what role (provider, deployer, importer or distributor) does the organisation stand? Most of the answers turn on the use to which the system is put rather than on the technology, so the same model can be minimal-risk in one deployment and high-risk in another.

The Article 6(3) exception is narrow and must be evidenced. An Annex III system escapes high-risk status only where it poses no significant risk of harm and meets one of four conditions; an Annex III system that profiles natural persons is always high-risk. A provider relying on the exception must document its assessment before placing the system on the market, register the system in the EU database, and be ready to defend the assessment to a market surveillance authority.

Misclassification carries direct legal exposure. Where a market surveillance authority finds that a provider misclassified an Annex III system as not high-risk in order to circumvent the requirements, the provider is subject to fines of up to EUR 15 million or 3% of worldwide annual turnover. Non-compliance with Article 5 carries fines of up to EUR 35 million or 7%.

The hardest cases need assessment deployment by deployment. General-purpose models, systems used for several purposes, agentic systems and AI features added to commercial software all need it, and classification is often uncertain in practice: a 2023 study of 106 enterprise AI systems against the draft texts found that 40% could not be

clearly classified. A deployer that modifies or repurposes a system can also become its provider.

A classification is defensible only if it is recorded and kept current. We recommend screening every system against Article 5 now, classifying the whole portfolio against Annex III well before 2 December 2027, building classification into procurement and change management, and recording each determination with its reasoning and an accountable owner. Where a case remains genuinely uncertain after analysis and advice, early engagement with the competent authority is likely to serve better than a confident position later found to be wrong.

Basis and Standing, at the end, sets out which propositions are binding law, which are Commission guidance, and which are our own recommended controls. The position is stated as at 11 October 2026.

1. Every obligation in the Act attaches to a classification, so classification comes first

The AI Act entered into force on 1 August 2024 and applies a risk-tiered structure to AI systems with a connection to the EU market. Under Article 2(1) it reaches providers placing AI systems on the EU market or putting them into service, wherever they are established; deployers established or located in the EU; and providers and deployers in third countries where the output produced by the system is used in the EU.

The Act creates obligation sets, not a general duty. It imposes no general duty of care on AI systems, no universal transparency requirement and no blanket right of explanation for AI-assisted decisions. Instead it attaches specific obligations to specific categories of system, and classification decides which set applies: prohibition, the full high-risk regime, targeted transparency duties, or almost nothing beyond AI literacy.

For a high-risk system, the set is large. Before it can lawfully be placed on the EU market or put into service, the system must meet the requirements of Articles 9 to 15: a risk management system, data and data governance, technical documentation, record-keeping, transparency and information to deployers, human oversight, and accuracy, robustness and cybersecurity. The provider must also operate a quality management system, complete a conformity assessment, register an Annex III system in the EU database (Article 49, with point 2 systems registered at national level under Article 49(5)) and run post-market monitoring for as long as the system is in use (Article 72).

Figure	What it measures
4	Risk tiers, each with its own obligation set
6 to 18 months	High-risk lead time from initial compliance assessment to lawful market placement, in our experience
EUR 35 million or 7%	Top penalty tier, for non-compliance with Article 5

The gap between high-risk and limited-risk is the one that drives cost. For a typical medium-complexity high-risk system, the path from initial compliance assessment to lawful market placement can, in our experience, take six to eighteen months. It needs

a dedicated cross-functional team, a large legal and technical advisory budget and, for Annex I products and some biometric systems, assessment by a notified body, whereas a limited-risk system needs transparency measures. Starting the high-risk work before the classification is settled is, in our experience, a common and costly mistake.

Classification also carries direct legal exposure. Where a market surveillance authority evaluating a system finds that the provider misclassified it as not high-risk in order to circumvent the requirements of Chapter III, Section 2, the provider is subject to fines under Article 99 (Article 80(7)). Article 80(7) names no tier, but the provider obligations it concerns sit in the Article 99(4) tier: up to EUR 15 million or 3% of worldwide annual turnover, whichever is higher, or whichever is lower for SMEs (Article 99(6)) and small mid-cap enterprises (SMCs, Article 99(6a)). Non-compliance with the Article 5 prohibitions carries fines of up to EUR 35 million or 7% of worldwide annual turnover, whichever is higher (Article 99(3)), or whichever is lower for SMEs.

Classification is often uncertain in practice. appliedAI's practical study of 106 enterprise AI systems, classified in March 2023 against the then draft texts, found 18% high-risk, 42% low-risk and 40% that could not be clearly classified as high-risk or not, with the unclear cases concentrated in critical infrastructure, employment, law enforcement and product safety. A Cloud Security Alliance research note of 13 March 2026, citing a TechPinions article of 23 February 2026, reports that more than half of organisations still lack a systematic inventory of the AI systems they operate. Without one, classification cannot be done systematically.

2. The four tiers differ in kind, not in degree

The Act's obligations fall into four groups, usually called tiers. The label, which the Commission also uses, is not statutory wording, and the tiers do not form a smooth scale; each carries a distinct and largely self-contained regime.

Prohibited practices (Article 5) are banned. The legislator has judged these practices incompatible with EU values and fundamental rights, so no compliance pathway exists. Safeguards, human oversight and risk mitigation cannot make a prohibited practice lawful; the only lawful course is not to place it on the market, put it into service or use it in the EU.

High-risk systems (Article 6, with Annexes I and III) carry the Act's full substantive requirements. "High-risk" is a legal classification, not an assessment of operational risk. A system can be legally high-risk, with every associated obligation, even where the organisation using it considers the probability of harm low in practice.

Limited-risk systems are defined by the Article 50 transparency obligations. These cover, among others, systems that interact directly with people, systems that generate synthetic content, and the deployment of emotion recognition, biometric categorisation and deep fakes. Article 50 can also apply to a high-risk system, in which case both regimes apply together.

Minimal-risk systems owe nothing under the Act beyond the Article 4 AI literacy measures. Article 4, as replaced by Regulation (EU) 2026/1744, requires providers and deployers of all AI systems to "take measures to support the development of AI literacy" of their staff and others operating the systems on their behalf. The Commission describes the vast majority of AI systems currently used in the EU as falling into this tier, giving spam filters and AI-enabled video games as examples.

The absence of AI Act obligations does not leave minimal-risk systems unregulated, since the GDPR, the Digital Services Act, product liability law and sector law may still apply.

3. The obligations arrive on a staged calendar, and the first are already enforceable

The Act's obligations do not all apply from one date. The calendar below, under Articles 111 and 113 as amended by Regulation (EU) 2026/1744, should set the order of the work.

Date	Obligations applicable	Detail
1 August 2024	Act enters into force	
2 February 2025	Article 5 prohibitions and Article 4 AI literacy	Chapters I and II apply (Article 113, point (a)).
2 August 2025	General-purpose AI model obligations; governance; penalties	Chapter V, Chapter VII and Chapter XII, except Article 101 (Article 113, point (b)). Fines for Article 5 breaches are available from this date.
27 July 2026	Regulation (EU) 2026/1744 in force	Article 4 as replaced and the new Article 4a apply; Articles 102 to 110 apply (Article 113, point (d)).
2 August 2026	General date of application	Includes the Article 50 transparency obligations and Commission fines on general-purpose AI model providers (Article 101).
2 December 2026	New Article 5 prohibitions	Article 5(1)(ba) and (bb), with Article 5(1a) and (1b) (Article 113, point (a)); Article 50(2) marking for generative systems placed on the market before 2 August 2026 (Article 111(4)).
2 August 2027	Legacy general-purpose AI models	Models placed on the market before 2 August 2025 must comply (Article 111(3)).
2 December 2027	High-risk obligations for Annex III systems	The most commercially significant date for most organisations: requirements, provider and deployer obligations, the fundamental rights impact assessment under Article 27 and, in practice, EU database registration (Article 113, point (c)(i)).
2 August 2028	High-risk obligations for Annex I systems	Article 6(1) systems (Article 113, point (c)(ii)).
2 August 2030	Legacy public-authority systems	High-risk systems placed on the market or put into service before their high-risk

Date	Obligations applicable	Detail
		date are caught only if their design changes significantly, but those intended for use by public authorities must comply by this date (Article 111(2)). AI components of the Annex X large-scale IT systems placed on the market before 2 August 2027 must comply by 31 December 2030 (Article 111(1)).

The Article 111(2) rule matters for classification, though it does not shelter any system from Article 5. A high-risk system already on the market before 2 December 2027 (or 2 August 2028 for Annex I) becomes subject to the high-risk requirements only on a significant change in its design, unless it is intended for use by public authorities. The classification record therefore also needs to show when each system was placed on the market and what has changed since.

> *Organisations that have not yet completed a full classification of their AI portfolio may already be behind on Article 5 (and, if they provide general-purpose AI models, on the obligations for those models), and should prepare now for the 2 December 2027 Annex III date.*

4. Classification follows a fixed sequence, beginning with whether the system is an AI system at all

The questions have to be asked in order, because each answer can end the analysis or change the next question:

1. Is the system an "AI system", and is it within the Act's scope?
2. Is the practice prohibited by Article 5?
3. Is the system high-risk under Article 6(1), as a safety component of, or itself, a product in Annex I?
4. Is it listed in Annex III, and if so, does the Article 6(3) exception apply?
5. Do any Article 50 transparency obligations apply, alongside or instead of the above?
6. In what role does the organisation stand for this system: provider, deployer, importer or distributor?

The threshold question is the definition in Article 3(1). An AI system is "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". Its core criterion is inference. Recital 12 excludes systems "based on the rules defined solely by natural persons to automatically execute operations", and names both machine learning and "logic- and knowledge-based approaches" as techniques that enable inference.

Simpler statistical systems may fall outside. The Commission's guidelines on the definition (C(2025) 5053) treat some systems as outside it: systems for improving mathematical optimisation, basic data processing, classical heuristics, and simple prediction systems whose performance can be achieved with a basic statistical learning rule (paragraphs 41 to 51). The guidelines are not binding, and the boundary is decided system by system.

Some systems are excluded from scope entirely. These are AI systems used exclusively for military, defence or national security purposes (Article 2(3)); systems and models specifically developed and put into service for the sole purpose of scientific research and development (Article 2(6)); research, testing and development before placing on the market, though not testing in real-world conditions (Article 2(8)); and systems released under free and open-source licences, unless placed on the market or put into service as high-risk systems or as systems falling under Article 5 or 50 (Article 2(12)).

5. Article 5 comes first because a prohibited practice has no compliance route

The original eight prohibitions have applied since 2 February 2025, with fines available since 2 August 2025. Two more, inserted by Regulation (EU) 2026/1744, apply from 2 December 2026: Article 5(1)(ba), on generating or manipulating realistic intimate imagery of an identifiable person without consent, and Article 5(1)(bb), on generating child sexual abuse material. The eight are:

- subliminal, purposefully manipulative or deceptive techniques that materially distort behaviour and cause, or are reasonably likely to cause, significant harm (Article 5(1)(a));
- exploitation of vulnerabilities due to age, disability or a specific social or economic situation (Article 5(1)(b));
- social scoring leading to detrimental or unfavourable treatment in unrelated contexts, or treatment that is unjustified or disproportionate, whether by public or private actors (Article 5(1)(c));
- risk assessments predicting criminal offending based solely on profiling or on assessing personality traits and characteristics (Article 5(1)(d));
- untargeted scraping of facial images from the internet or CCTV footage to create or expand facial recognition databases (Article 5(1)(e));
- inferring emotions in the workplace and in education institutions, except for medical or safety reasons (Article 5(1)(f));
- biometric categorisation to deduce or infer race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation (Article 5(1)(g));
- real-time remote biometric identification in publicly accessible spaces for law enforcement, outside three narrowly defined objectives (Article 5(1)(h)).

The emotion recognition ban turns on biometric data. On the Commission's guidelines on prohibited practices (C(2025) 5052, paragraph 245), Article 5(1)(f) covers

only the inference of emotions from biometric data, such as facial expressions or voice characteristics. Sentiment inferred from text is outside the ban, although a biometric emotion recognition system that escapes it can still be high-risk under Annex III point 1(c).

The word "solely" marks the boundary between the predictive policing ban and the high-risk regime. A system that assesses or predicts a person's risk of committing a criminal offence based solely on profiling or personality traits is prohibited. Article 5(1)(d) does not apply to systems "used to support the human assessment of the involvement of a person in a criminal activity, which is already based on objective and verifiable facts directly linked to a criminal activity". Such a system falls instead to be assessed against the law enforcement uses in Annex III point 6.

6. A system is high-risk by one of two routes, and Annex III classifies the use, not the technology

Once a system is outside Article 5, the next question is whether it is high-risk. Article 6 provides two independent routes.

The first route, Article 6(1), runs through product safety law. A system is high-risk where it is intended to be used as a safety component of a product, or is itself a product, covered by the legislation in Annex I, and that product is required to undergo third-party conformity assessment. Since 27 July 2026, Article 6(1a) to (1c) narrow this: AI systems "solely used for non-safety related aspects of user assistance, performance optimisation, service efficiency, automation or convenience or quality control" are not safety components unless their failure or malfunctioning would endanger health and safety. A product requiring third-party assessment solely for risks other than health and safety does not meet the condition.

Annex I has two sections, and the second is treated differently. Section A covers, among others, toys, lifts, radio equipment, pressure equipment, personal protective equipment, medical devices and in vitro diagnostic devices. Section B covers motor vehicles, civil aviation, marine equipment, rail and, since 27 July 2026, machinery.

For Section B products, the high-risk requirements arrive through sectoral law. For high-risk systems related to Section B products, only Article 6(1), Article 60a and Articles 102 to 112 apply (Article 2(2)). The limitation reaches only systems that are high-risk under Article 6(1) and does not list Articles 4, 5 or 50, so whether those articles reach such systems is not settled; other AI in those products remains subject to them.

The second route, Article 6(2), runs through Annex III. AI systems referred to in Annex III are high-risk, subject to the exception in Article 6(3). Annex III describes uses ("AI systems intended to be used for..."), so classification turns on the intended purpose in a given deployment, not on the underlying technology.

The same model can sit in different tiers. A large language model drafting marketing emails is not high-risk. The same model embedded in a tool that analyses, filters

and ranks job applicants is high-risk under Annex III point 4(a), and, because ranking applicants involves profiling them, the Article 6(3) exception cannot take it out.

7. Annex III lists eight areas, and the employment area reaches every employer

Annex III organises the high-risk uses into eight numbered areas. They are referred to here by point number (point 4 is employment), with the Annex's own lettering for the uses within each.

Point 1, biometrics, covers remote biometric identification (but not verification whose sole purpose is to confirm that a person is who they claim to be), biometric categorisation according to sensitive or protected attributes, and emotion recognition, in each case in so far as the use is permitted by law. It is the only Annex III area for which a notified body can be required: a provider that has not applied harmonised standards or common specifications in full must use the Annex VII procedure with a notified body, and may otherwise choose internal control (Article 43(1)). Points 2 to 8 follow internal control (Article 43(2)).

Point 2, critical infrastructure, covers AI systems intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, or the supply of water, gas, heating or electricity.

Point 3, education and vocational training, covers admission and assignment to institutions, evaluating learning outcomes, assessing the appropriate level of education a person will receive, and monitoring prohibited behaviour during tests.

Point 4, employment, workers' management and access to self-employment, is one of the broadest areas and, because nearly every organisation is an employer, one of the most consequential. It covers recruitment and selection, including targeted job advertisements, analysing and filtering applications and evaluating candidates (point 4(a)), and decisions affecting terms of work, promotion and termination, allocating tasks based on individual behaviour or personal traits, and monitoring and evaluating performance and behaviour (point 4(b)). A tool that ranks applicants and produces a shortlist of five for a manager to choose from is influencing the hiring decision and is profiling the applicants, so it remains high-risk.

Point 5, essential private and public services and benefits, covers eligibility for public assistance benefits and services assessed by or for public authorities (5(a)); creditworthiness and credit scoring of natural persons, except systems used to detect financial fraud (5(b)); risk assessment and pricing in relation to natural persons in life and health insurance (5(c)); and evaluating emergency calls, dispatching emergency services and emergency patient triage (5(d)).

Point 6, law enforcement, covers, for use by or on behalf of law enforcement authorities, assessing a person's risk of becoming a victim, polygraphs and similar tools, evaluating the reliability of evidence, assessing the risk of offending or re-offending not

solely on the basis of profiling, and profiling in the course of detecting, investigating or prosecuting crime.

Point 7, migration, asylum and border control management, covers polygraphs, risk assessments of people entering or who have entered a Member State, examination of applications for asylum, visas and residence permits, and detecting, recognising or identifying people other than by verifying travel documents.

Point 8, administration of justice and democratic processes, covers systems intended to be used by a judicial authority or on its behalf to research and interpret facts and law and apply the law to facts, or used similarly in alternative dispute resolution (8(a)), and systems intended to influence the outcome of an election or referendum or voting behaviour (8(b)). Point 8(b) excludes tools whose output people are not directly exposed to, such as tools used to organise or optimise political campaigns from an administrative or logistical point of view. That boundary is hard to apply at the margins, and providers of political technology should take legal advice before the next electoral cycle in which they operate.

8. The Article 6(3) exception is narrow, never covers profiling, and must be documented and registered

Article 6(3) gives providers a self-assessed route out of high-risk status for a system listed in Annex III. In our experience, it is where over-optimism most often enters a classification.

The exception has a threshold and four conditions. A listed system "shall not be considered to be high-risk where it does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons, including by not materially influencing the outcome of decision making". That rule applies where the system is intended to do any one of the following:

(a) perform a narrow procedural task;
(b) improve the result of a previously completed human activity;
(c) detect decision-making patterns or deviations from prior patterns, without being meant to replace or influence the previously completed human assessment without proper human review; or
(d) perform a preparatory task to an assessment relevant to an Annex III use case.

Profiling always defeats the exception. Notwithstanding the conditions, "an AI system referred to in Annex III shall always be considered to be high-risk where the AI system performs profiling of natural persons". Profiling has its GDPR meaning (Article 3(52), referring to Article 4(4) of the GDPR): automated processing of personal data to evaluate personal aspects of a person, in particular to analyse or predict aspects such as performance at work, economic situation, health, preferences, reliability or behaviour, which many employment, credit and insurance systems do by design.

The provider must document, register and be ready to produce the assessment. A provider that considers an Annex III system not high-risk "shall document its assessment before that system is placed on the market or put into service" and must provide the documentation to national competent authorities on request (Article 6(4)). It is also subject to registration under Article 49(2), which must take place before the system is placed on the market or put into service.

A market surveillance authority can test the assessment. Where it has sufficient reason to consider that a system classified as not high-risk under Article 6(3) is high-risk, the authority evaluates it against the Article 6(3) conditions and the Commission guidelines (Article 80(1)). If it finds the system high-risk, it requires the provider to bring it into compliance, with fines under Article 99 if the provider does not (Article 80(2) and (4)), and the misclassification fine in section 1 where the classification was made to circumvent the requirements (Article 80(7)).

Guidance is still in draft. Article 6(5) required the Commission to publish guidelines on the practical implementation of Article 6, with practical examples of high-risk and non-high-risk use cases, by 2 February 2026. The Commission published draft guidelines on 19 May 2026 for a stakeholder consultation that closed on 23 July 2026, and as at 11 October 2026 they had not been formally adopted.

The Commission may also add, modify or delete Article 6(3) conditions by delegated act (Article 6(6) and (7)).

Whether a system materially influences a decision is a question about how it works in practice. In our view the relevant factors include how prominent the output is in the decision-maker's interface, whether the output is a score or rank that anchors the review, how much time the reviewer has for an independent assessment, and whether targets or performance pressure push reviewers to follow the output. We recommend applying the exception conservatively and treating it as a documented position the provider must be able to defend, not as a route around the regime.

Some vendors are designing for the exception. In our experience, several vendors in human resources, legal and financial technology are restructuring products so that narrow procedural functions can be separated from functions that influence decisions, in order to preserve the Article 6(3) exception for lower-risk product tiers. That separation helps only if it holds in the deployed product and the separated functions do not profile.

9. General-purpose, mixed-purpose, agentic and embedded systems need the most care

Four kinds of case arise often enough, and are hard enough, that they should be approached with particular care and in many cases with legal advice.

A general-purpose AI model is not high-risk merely because it could be used for a high-risk purpose. High-risk classification attaches to an AI system deployed for a specific intended purpose. Where a general-purpose model is integrated into a high-risk

system, the model provider owes information to downstream providers under Article 53, and the provider of the high-risk system and any third party supplying an AI model, tool or component used in it must specify by written agreement the information, technical access and assistance needed for compliance (Article 25(4)).

A system used for several purposes needs a classification for each. Many systems serve different users for different purposes, some listed in Annex III and some not. An organisation deploying a system for mixed purposes should assess each use separately, and the system is high-risk for any use that falls within Annex III and outside Article 6(3).

Agentic systems raise classification questions at several levels. A single agentic workflow may engage more than one Annex III area at once, and a system that takes consequential actions may be high-risk in ways its interface description does not reveal.

Article 14(3) requires oversight measures "commensurate with the risks, level of autonomy and context of use", which is harder to design for a system built to act without step-by-step review. Until the Commission's guidance addresses these systems, we recommend treating an agentic system as high-risk where any step in its intended workflow falls within Annex III, unless a documented Article 6(3) assessment supports a different conclusion.

AI features embedded in commercial software are often missed. Software, particularly human resources platforms, contract management systems and customer service tools, increasingly gains AI features after it was first assessed, and in our experience the classification of these features is often misunderstood.

A platform outside the high-risk regime before an AI performance-scoring feature was added may be high-risk after it, because the feature falls within Annex III point 4(b). Any update that adds or materially changes an AI feature should therefore trigger a fresh classification.

10. Classification is also a role question, and a deployer can become a provider

The Act allocates most high-risk obligations by role, so the classification is incomplete until the organisation knows which role it holds. A provider (Article 3(3)) is a natural or legal person, public authority, agency or other body that develops an AI system, or has one developed, and places it on the market or puts it into service under its own name or trademark, whether for payment or free of charge. A deployer (Article 3(4)) is a natural or legal person, public authority, agency or other body using an AI system under its authority, except in the course of a personal non-professional activity.

Under Article 25(1), a distributor, importer, deployer or other third party becomes the provider of a high-risk AI system in three circumstances:

(a) it puts its name or trademark on a high-risk AI system already placed on the market or put into service, without prejudice to contractual arrangements allocating the obligations otherwise;

(b) it makes a substantial modification to a high-risk AI system already placed on the market or put into service, such that the system remains high-risk;
(c) it modifies the intended purpose of an AI system, including a general-purpose AI system, that has not been classified as high-risk, so that it becomes high-risk.

A substantial modification is a change the original assessment did not foresee. Article 3(23) defines it as a change after placing on the market or putting into service that was not foreseen or planned in the provider's initial conformity assessment and that affects compliance with the Chapter III, Section 2 requirements or modifies the intended purpose. In our view, retraining a deployed model on new data, changing outputs in ways that affect their impact, and connecting the system to new data sources may each qualify.

Pre-determined changes are not substantial modifications. A high-risk system that is substantially modified must undergo a new conformity assessment, but changes to a continuously learning system that the provider pre-determined at the initial assessment and recorded in the technical documentation are not substantial modifications (Article 43(4)).

The new provider does not act alone. Article 25(2), as amended by Regulation (EU) 2026/1744, requires the initial provider to cooperate closely with the new provider, including by making available technical documentation sufficient to assess compliance, informing it of known limitations and failure modes, and giving targeted technical access for testing and validation. That duty does not apply where the initial provider has clearly specified that its system is not to be changed into a high-risk system.

A deployer that repurposes a system its provider has marked in that way takes on provider obligations without a statutory route to the documentation it needs. That makes the contractual allocation of intended purpose and change notification a commercial priority as well as a compliance one.

Article 25(1) sets out when an importer or distributor becomes the provider of a high-risk system. An importer or distributor does not become a provider merely because the original provider has failed to meet its obligations; each has its own obligations under Articles 23 and 24.

11. Procurement is the cheapest point at which to check a classification

A deployer buying a high-risk system depends on information only the provider holds, so procurement is where the classification, and the evidence behind it, should be checked. Classification review should be a mandatory checkpoint before a system is approved for deployment, not a task performed afterwards.

The Act gives deployers instructions for use, not the technical documentation. Under Article 13, a high-risk system must come with instructions for use covering, among other things, the system's characteristics, capabilities and limitations, its

expected accuracy, robustness and cybersecurity, the human oversight measures and, where appropriate, information about input data and the training, validation and testing data sets. The Act gives a deployer, as such, no right to the provider's technical documentation, so access to it is a matter for contract.

A non-EU provider needs an authorised representative. Providers established in third countries must appoint, by written mandate, an authorised representative established in the EU before making high-risk systems available on the EU market (Article 22(1)). An importer must verify that the representative has been appointed before placing the system on the market (Article 23(1)(d)), so due diligence on the compliance status of non-EU providers is a core part of procurement.

Procurement checkpoint	What to verify
Classification	The provider's classification and, for an Annex III system it treats as not high-risk, its Article 6(3) assessment and Article 49(2) registration.
Conformity assessment	Internal control (Annex III points 2 to 8, and point 1 where harmonised standards or common specifications were applied in full) or a notified body (point 1 otherwise). Request documentary evidence.
EU database registration	Annex III high-risk systems must be registered (Article 49(1); point 2 systems at national level, Article 49(5)). Verify the registration before sign-off; a public-authority deployer must not use an unregistered system (Article 26(8)).
Technical documentation	Drawn up by the provider under Annex IV; deployer access is a matter for contract. Review for completeness and recency where provided.
Instructions for use	Must include capabilities, limitations, performance metrics and human oversight measures (Article 13(3)).
EU authorised representative	Required for non-EU providers of high-risk systems; without one, an importer may not place the system on the market (Article 23(1)(d)).
Post-market monitoring plan	Providers must maintain one. Request evidence of active monitoring and incident reporting procedures.
Substantial modification policy	Contractual clarity on what counts as a substantial modification and on notice of system changes, so that re-assessment is triggered when needed.

The contract should settle what the Act leaves open. Procurement due diligence should include a preliminary classification, review of the vendor's compliance documentation and verification of any authorised representative, together with contract terms on the provider and deployer split of obligations, information rights, technical documentation access, audit rights, notice of system changes and incident reporting. In our experience, a vendor that can show a completed conformity assessment, maintained technical documentation and EU database registration has a clear advantage with enterprise buyers in regulated industries.

12. Sector context changes the classification more than the technology does

The same techniques sit in different tiers depending on who uses them and for what. The sectors below show the pattern.

Financial services. AI evaluating the creditworthiness or credit score of natural persons is high-risk under Annex III point 5(b), and because such systems ordinarily profile, the Article 6(3) exception will rarely be available. Fraud detection is excluded from point 5(b) but needs analysis against the other areas, while AI used for risk assessment and pricing in relation to natural persons in life and health insurance is high-risk under point 5(c). The Digital Operational Resilience Act (DORA, Regulation (EU) 2022/2554) runs in parallel for in-scope financial entities, and its third-party risk requirements bear directly on AI supply chains.

Healthcare and life sciences. AI that is medical device software requiring notified-body assessment under the Medical Devices Regulation (MDR, Regulation (EU) 2017/745) or the In Vitro Diagnostic Medical Devices Regulation (IVDR, Regulation (EU) 2017/746) is high-risk under Article 6(1), and from 2 August 2028 must meet the AI Act requirements, assessed within the MDR or IVDR conformity assessment (Article 43(3)). Article 2(13), inserted by Regulation (EU) 2026/1744, allows the Commission to limit some requirements for Section A products where sectoral law gives equivalent or higher protection, by delegated acts due by 2 August 2027. The Commission's proposal of 16 December 2025 to revise the MDR and IVDR (COM(2025) 1023) would move both to Annex I Section B; as at 11 October 2026 it remained a proposal, and Annex I as amended still lists them in Section A.

Health research and workplace wellbeing. AI used in drug discovery, clinical trial design and pre-clinical research may fall outside the Act under Article 2(6) or 2(8) while it is used for research and not placed on the market, and its classification should be revisited when it moves into clinical or commercial use. Mental health and wellbeing tools used on employees engage the Article 5(1)(f) prohibition if they infer emotions from biometric data, unless the medical or safety exception applies.

Legal services. AI legal research tools used by private practitioners are not high-risk under Annex III point 8(a), which covers systems used by or on behalf of a judicial authority, or in a similar way in alternative dispute resolution. The same platform used by or on behalf of a court to research and interpret facts and law is high-risk. AI used in associate performance assessment, partnership track decisions or redundancy selection within law firms is likely to be high-risk under point 4, subject to Article 6(3).

Public sector. Public bodies are among the operators most likely to deploy systems under points 6, 7 and 8, and the Article 5(1)(c) social scoring prohibition applies to them as it does to private actors. Deployers that are bodies governed by public law or private entities providing public services must carry out a fundamental rights impact assessment before deploying an Annex III high-risk system other than one under point 2, as must any deployer of a point 5(b) or 5(c) system (Article 27(1)). Public procurement frameworks should, in our view, make AI Act compliance a selection criterion, covering conformity assessment, technical documentation, EU database registration and post-market monitoring.

13. A classification stays defensible only if it is recorded and kept current

A classification is a judgement made on facts that change: the system's functions, its deployment, its users and the guidance that interprets the Act. It is defensible later only if the reasoning was written down when it was made and revisited when the facts moved.

Keep a register of every AI system. Every organisation placing AI systems on the EU market or deploying them should keep a documented AI system register, reviewed at least annually and updated whenever a system is modified or relevant guidance changes. Each entry should record:

- system name and version;
- vendor and internal reference;
- intended use as deployed (not as the vendor markets it, since the deployed use is the operative criterion for Annex III);
- the classification and the provisions on which it rests, including any Article 6(3) assessment;
- the date of assessment and the accountable reviewer;
- the date the system was placed on the market or put into service, for Article 111(2);
- current compliance status.

Set re-assessment triggers. A formal review of the classification should follow changes to the system's functions, changes to its deployment context or user population, connection to new data sources or downstream decision systems, and changes in regulatory interpretation. Where a change to a high-risk system is a substantial modification, a new conformity assessment is required (Article 43(4)).

Track the instruments that move the boundary. The Commission may amend Annex III by delegated act, adding or modifying uses within its eight areas or removing uses that no longer pose significant risks (Article 7), amend the Article 6(3) conditions (Article 6(6) and (7)), and amend the thresholds for general-purpose AI models with systemic risk (Article 51(3)). Regulatory monitoring should cover delegated and implementing acts, Commission guidelines under Articles 6(5) and 96, AI Board opinions, national authority guidance and data protection authorities' positions on the interaction with the GDPR. The Commission updates its Article 96 guidelines at the request of Member States or the AI Office (Article 96(2)).

Engage the authority where a case remains genuinely uncertain. Where classification is still ambiguous after thorough analysis and legal advice, early engagement with the relevant national competent authority is advisable, and in our experience some authorities welcome good-faith engagement on classification questions.

In deciding whether to fine and how much, authorities must weigh, among other factors, the operator's cooperation in remedying an infringement and whether the operator itself notified it (Article 99(7)(f) and (h)). An early, open account of an uncertain classification is therefore likely to be treated more favourably than a confident assertion of compliance later found to be wrong.

Where to start, and what it costs

The work comes in stages, beginning with the one that is already due.

Now: screen every system against Article 5. The prohibitions have applied since 2 February 2025, so screen every existing deployment against the eight categories, and against the two added from 2 December 2026, and withdraw or redesign without delay any system that falls within one. Document the screening even for systems clearly outside the prohibitions, since the record shows a systematic process.

Workforce tools that infer employee mood or engagement from facial or voice analysis have been prohibited in the workplace since 2 February 2025, unless the medical or safety exception applies (Article 5(1)(f)).

Next, well before 2 December 2027: inventory and classify the portfolio. Classify every AI system within the Act's reach under Article 2(1) against Article 6 and Annex III, and treat the classification as the precondition for every other workstream.

Set up a cross-functional working group (legal, compliance, data protection, engineering and the relevant business units) and name a senior accountable owner for each determination. Apply Article 6(3) conservatively, documenting and registering any reliance on it before the system is placed on the market, and give the cases in section 9 first call on the group's time.

Alongside: build classification into procurement and the development lifecycle. Make classification review a mandatory checkpoint before any AI system is approved, with the contract terms in section 11. Build re-assessment triggers into change management, so that any update that changes AI functionality, training data or deployment context prompts a review, and a substantial modification prompts a new conformity assessment (Article 43(4)).

Before deploying an in-scope system: the fundamental rights impact assessment. Public bodies, private entities providing public services and deployers of credit-scoring and life and health insurance systems should develop the method, accountability and process integration for the assessment required by Article 27 before 2 December 2027. It is not a GDPR data protection impact assessment, but the two are closely related and are best carried out together.

Ongoing: keep the register current. Review the register at least annually and on every trigger, track the delegated acts, guidelines and opinions in section 13, and engage the authority on cases that remain genuinely uncertain.

On timing and cost. Given the lead time set out in section 1, in our experience six to eighteen months for a typical medium-complexity high-risk system, a classification settled late compresses the compliance work against the 2 December 2027 and 2 August 2028 dates. A high-risk system already on the market by then is caught only on a significant change in its design, unless it is intended for use by public authorities, which

must comply by 2 August 2030 (Article 111(2)). As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source.

Classification remains a continuing discipline rather than a single exercise, driven by the staged calendar, the Commission's delegated powers and the spread of agentic and embedded AI. An organisation that has classified its portfolio, recorded the reasoning and kept it current can say what it owes for each system, and can show a market surveillance authority why.

Standard Intelligence makes each classification a recorded decision

Classifying a system determines its whole obligation set, and the classification has to be demonstrable. The Standard Intelligence platform is built to make classification structured, cited and recorded.

Platform capability	How it supports risk classification
AI Regulatory Navigator	Turns a plain-language description of a system into a structured classification assessment against Article 6 and Annex III, with citations attached.
AI Intelligence Centre	Records the classification, the applicable Annex III point and the rationale, with version history as the system changes.
Guided Discovery	Carries a newly classified high-risk system into the documentation obligations that follow from the classification.

Classification then becomes a recorded, evidenced decision rather than an undocumented judgement.

This paper is provided for information and guidance only and is not legal advice. It states the position as at 11 October 2026. Classification determinations with material commercial or regulatory consequences should be reviewed by qualified legal counsel with expertise in EU AI regulation.

Basis and Standing

Binding law. Regulation (EU) 2024/1689 (the AI Act), as amended by Regulation (EU) 2026/1744 (the Digital Omnibus on AI), in force 27 July 2026: Articles 2, 3(1), (3), (4), (23) and (52), 4, 5, 6, 7, 13, 14(3), 22, 23, 25, 26(8), 27, 43, 49, 50, 51(3), 53, 72, 80, 96, 99, 111 and 113, and Annexes I and III. High-risk obligations apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems (Article 113, point (c)), subject to Article 111(2) for systems already on the market. Article 6(3) and (4) were not amended by Regulation (EU) 2026/1744; Article 6(1a) to (1c), Article 2(2) and (13), Article 25(2), Article 99(6a) and the

move of machinery to Annex I Section B were. Profiling has the meaning in Article 4(4) of the GDPR (Regulation (EU) 2016/679). The MDR, IVDR and DORA are cited where used.

Commission guidance, not law. The guidelines on prohibited practices (C(2025) 5052) and on the definition of an AI system (C(2025) 5053) are non-binding. The guidelines on the classification of high-risk AI systems under Article 6(5) were published in draft on 19 May 2026 and had not been formally adopted as at 11 October 2026. The Commission's MDR and IVDR proposal (COM(2025) 1023) is a proposal, not law.

Our recommended controls, not regulatory requirements. The six-question classification sequence in section 4, the factors for material influence in section 8, the default treatment of agentic systems in section 9, the procurement checkpoints in section 11, the register fields and re-assessment triggers in section 13 and the staged actions in the closing section are our design. The register reflects Article 6(4) (documenting an Article 6(3) assessment), Article 49 (registration) and Article 111(2) (date of placing on the market); the procurement checkpoints reflect Articles 13, 22, 23, 26(8), 43 and 49.

Research and market evidence, and what it supports. The appliedAI study classified 106 enterprise AI systems in March 2023 against the draft texts of the Commission, Parliament and Council, not the adopted Act; it shows that classification is often uncertain in practice, not what share of systems is high-risk under the final text. The Cloud Security Alliance figure on AI inventories is reported at second hand, from a TechPinions article, and is indicative only. The six-to-eighteen-month lead time and the statements made "in our experience" rest on our own engagements, not on a published study.

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

Sources

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
2. Regulation (EU) 2024/1689, consolidated text of 27 July 2026: Articles 2, 3, 4, 5, 6, 7, 13, 14, 22, 23, 25, 26, 27, 43, 49, 50, 51, 53, 72, 80, 96, 99, 111 and 113; Recital 12; Annexes I and III. https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng
3. 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
4. Regulation (EU) 2016/679 (GDPR), Article 4(4). https://eur-lex.europa.eu/eli/reg/2016/679/oj
5. European Commission, Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689, C(2025) 5052 final, 29 July 2025 (non-binding), paragraph 245. https://ai-act-service-desk.ec.europa.eu/sites/default/files/2026-01/guide-prohibited_en.pdf

6. European Commission, Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689, C(2025) 5053 final, 29 July 2025 (non-binding), paragraphs 41 to 51. https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-ai-system-definition-facilitate-first-ai-acts-rules-application
7. European Commission, Draft Commission guidelines on the classification of high-risk AI systems, 19 May 2026 (consultation closed 23 July 2026; not adopted as at 11 October 2026). https://digital-strategy.ec.europa.eu/en/library/draft-commission-guidelines-classification-high-risk-ai-systems
8. European Commission, AI Act: regulatory framework for AI (risk tiers; minimal or no risk). https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
9. appliedAI, AI Act: Risk Classification of AI Systems from a Practical Perspective (27 March 2023; 106 enterprise AI systems against the draft texts; 18% high-risk, 42% low-risk, 40% unclear). https://www.appliedai.de/en/ai-resources/white-papers/ai-act-risk-classification-of-ai-systems-from-a-practical-perspective/
10. Cloud Security Alliance, research note: EU AI Act High-Risk Deadline, Enterprise Readiness Gap (13 March 2026), citing TechPinions, "Why the EU's AI Act Is About to Become Every Enterprise's Biggest Compliance Challenge" (23 February 2026). https://labs.cloudsecurityalliance.org/research/csa-research-note-eu-ai-act-high-risk-compliance-deadline-20/
11. Regulation (EU) 2017/745 (MDR) and Regulation (EU) 2017/746 (IVDR). https://eur-lex.europa.eu/eli/reg/2017/745/oj; https://eur-lex.europa.eu/eli/reg/2017/746/oj
12. European Commission, proposal to revise the MDR and IVDR, COM(2025) 1023 final, 16 December 2025 (not adopted). https://www.gleisslutz.com/en/know-how/radical-simplification-mdr-and-ivdr-and-broad-inapplicability-ai-act-medical-devices-and-ivds (secondary summary)
13. Regulation (EU) 2022/2554 (DORA). https://eur-lex.europa.eu/eli/reg/2022/2554/oj

Provisions are cited by article, paragraph or point in the text above.

Standard Intelligence

GET IN TOUCH

Get classification right, first.

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