



TECHNOLOGY, MEDIA AND TELECOMMUNICATIONS

Classification of high-risk AI systems

Commission Guidelines

On 19 May 2026, the European Commission published draft Guidelines on the ‘Classification of high-risk AI systems under Article 6 of Regulation (EU) 2024/1689 (AI Act)’¹. The Commission’s objective is to assist in the classification of AI systems under the AI Act. Article 6(5) requires the Commission to “provide guidelines specifying the practical implementation of this Article in line with Article 96 together with a comprehensive list of practical examples of use cases of AI systems that are high-risk and not high-risk”. A consultation was launched in 2025 to gather input on the interpretation of Article 6, and this project is now open for consultation with interested parties until 23 July 2026.

These guidelines are not legally binding, but they reflect the Commission’s interpretation and will certainly be taken into account by supervisory authorities. They are intended to clarify the scope of the AI Act and have been published alongside other Commission guidelines on the subject².

Article 6(5) requires the Commission to “provide guidelines specifying the practical implementation of this Article in line with Article 96 together with a comprehensive list of practical examples of use cases of AI systems that are high-risk and not high-risk”.

1 Available [here](#).

2 Suggested footnote (ready to insert), with Portuguese titles and official Commission links: European Commission Guidelines under the AI Act: (i) Guidelines on the definition of ‘AI system’ (February 2025); (ii) Guidelines on prohibited AI practices (February 2025); and (iii) Guidelines for providers of general-purpose AI models (July 2025). Draft guidelines subject to public consultation: (iv) Draft Guidelines on the classification of high-risk AI systems under Article 6 of the AI Act. The consultation was initially open until 23 June 2026 but has been extended to 23 July 2026. (v) Draft Guidelines on transparency obligations under Article 50 of the AI Act. The consultation closed on 3 June 2026.

The AI Act adopts a risk-based approach and, in Article 6, which is not yet in force³, regulates the category of high-risk AI systems. Classification does not prohibit use but rather subjects the system to compliance with the specific obligations set out in the AI Act. A system may be included in this category via one of two classification pathways, which the Commission is now seeking to clarify.

Annex I (Article 6(1)) states that the system is considered high-risk where it constitutes a safety component of a product governed by Union harmonisation legislation (or is itself such a product) and is subject to third-party conformity assessment. The list of harmonisation legislation in Annex I is exhaustive and can only be updated by amending the scope of the relevant sector-specific legislation or Annex I of the AI Act.

The concept of ‘safety component’ should be interpreted independently in light of Article 3(14) of the AI Act. It should not be interpreted in accordance with the definitions in applicable sectoral legislation.

According to the Guidelines, the concept of ‘safety component’ should be interpreted independently in light of Article 3(14) of the AI Act. It should not be interpreted in accordance with the definitions in applicable sectoral legislation. The definition’s criteria are alternative: a system is a safety component if it fulfils a safety function in the AI product or system, or if its failure or malfunction endangers health or safety, or causes damage to property.

The Guidelines clarify the distinction between an AI system that constitutes a product component and a system that constitutes a product covered by EU harmonisation legislation in itself. They clarify the concept of a safety component by distinguishing safety functions from purely operational functionalities. They also set out the third-party conformity assessment requirements applicable

to placing such products on the market or putting them into service, providing examples to illustrate these requirements. The Guidelines also set out how to integrate the requirements of the AI Act into sector-specific technical documentation and procedures, thereby avoiding duplication. The Union’s harmonisation legislation covers a large number of systems and products. Accordingly, the Guidelines do not provide an exhaustive list of high-risk systems. Instead, they set out a classification methodology applicable to all sectors listed in Annex I, without prejudice to any sector-specific guidelines that may be developed in future.

Annex III (Article 6(2)) states that the system is classified as high-risk if the intended purpose falls within one of the areas listed in Annex III. These areas are biometrics, critical infrastructure, education and vocational training, employment, workforce management and access to self-employment, access to essential private services and essential public services and benefits, law enforcement, migration, asylum and border control management, and the administration of justice and democratic processes.

³ Under Article 113 of the AI Act, the obligations for high-risk AI systems classified under Article 6(2) (Annex III) will apply from 2 August 2026. Those classified under Article 6(1) (Annex I – regulated products) will be subject to these obligations from 2 August 2027. However, under the Digital Omnibus, the European Parliament and the Council reached an agreement on AI (based on the Commission’s proposal of 19 November 2025) to postpone these dates to 2 December 2027 (Annex III) and 2 August 2028 (Annex I), respectively.

The Guidelines clarify that the exemption criteria are alternative and exhaustive, and as they restrict fundamental rights, they must be interpreted strictly.

However, Article 6(3) provides a derogation, applicable only to systems classified under Article 6(2), for those that do not pose a significant risk of harm to the health, safety, or fundamental rights of natural persons, nor significantly influence decisions. Eligibility for this derogation depends on a self-assessment by the provider and applies where the system is intended to:

- Perform a limited procedural task.
- Improve the outcome of a previously completed human activity.
- Detect decision-making patterns or deviations from previous patterns without replacing or influencing a previously completed human assessment, and without prejudice to appropriate human verification.
- Perform a preparatory task in the context of an assessment relevant to the use cases set out in Annex III.

However, according to the provision, the derogation does not apply to systems that profile natural persons.

The Guidelines state that even if a system meets the conditions for exemption, it should not benefit from this mechanism if it is part of a complex system whose combined purpose or joint results influence individual decisions in high-risk use cases, or if it is integrated into interconnected systems such as agent-based AI systems. Conversely, functions that are genuinely separable and do not contribute to results that materially influence the analysis of individual cases remain eligible for the exemption. Autonomous components that do not contribute to a high-risk purpose fall outside the scope of that classification.

The Guidelines clarify that the exemption criteria are alternative and exhaustive, and as they restrict fundamental rights, they must be interpreted strictly. For example, a ‘limited procedural task’ excludes actions involving value judgements (useful/less useful), the assignment of scores or rankings, or suggestions for next steps. ‘Improvement’ means refining or verifying without altering the substantive result or providing materially different solutions. ‘Pattern detection’ covers only ex post comparative evaluation, without inferring criteria for new evaluations, replacing or influencing prior human decisions without proper review.

The eight areas set out in Annex III are systematised using classification matrices that include examples of cases that do or do not constitute high-risk use, as well as situations to which derogations may apply.

According to the Guidelines, classification is based on the system's intended purpose and the tasks it is designed to perform. The type and degree of human intervention during the system's implementation may also be relevant to the classification. However, this only applies in demonstrating that the tasks are strictly procedural, preparatory, or aimed at improving a human activity that has already been completed. A supplier cannot exempt a system from high-risk classification by adding a requirement for human intervention alone. Where a system is intended for use in various contexts and applications, all of these must be described in order to assess whether the intended purpose could lead to a high-risk classification.

General-purpose systems are designed for a wide range of uses and typically incorporate a general-purpose AI model. The Commission considers these systems to be high risk only in limited cases.

This applies where high-risk uses are not expressly excluded in relevant documentation, contractual terms and communications, and where those uses can reasonably be foreseen based on the system's functionalities and capabilities. This position acknowledges that AI systems, particularly those based on general-purpose models, are not static entities. Their behaviour may evolve after deployment and many capabilities only become apparent in real-world usage contexts. Accordingly, the adopted interpretation blurs the line between the intended purpose of a system and its reasonably foreseeable misuse, as defined in the AI Act.

The Guidelines also clarify that including a system as an example of a non-high-risk system does not make its use lawful. It must comply with all other applicable legislation, particularly the General Data Protection Regulation.

The Guidelines also clarify that including a system as an example of a non-high-risk system does not make its use lawful. It must comply with all other applicable legislation, particularly the General Data Protection Regulation.

In short, while these Guidelines are not binding, they reflect the Commission's interpretation and set out practical criteria for classifying high-risk AI systems. Suppliers and deployers should take them into account. ■