onetrust | FUTURE OF PRIVACY FORUM | AI

Conformity Assessments under the EU AI Act: A step-by step guide

Updated with the final text of the AI Act

APRIL 2025

Authors: Andreea Serban, Vasileios Rovilos,* Katerina Demetzou*

Editors: Lee Matheson, Gabriela Zanfir-Fortuna, Rob van Eijk, Gulam Chagani, Anastasia Konova, Francesco Saturnino, Alexis Kateifides

Copyeditor: Alexander Thompson

The authors thank Chuma Akana for his research and contribution to the Guide.

* These authors worked on the Guide, including its first edition in 2023, during their time at the Future of Privacy Forum. Nothing from the Guide represents the position or opinion of their current employers.

Table of Contents

Table of Contents

I. Introduction to the Guide

This Guide explains Conformity Assessments (CAs) under the EU Artificial Intelligence Act (EU AIA or AIA) and provides a theoretical roadmap for conducting one. CAs are a key component of demonstrating compliance with product safety legislation and, consequently, an overarching accountability tool introduced by the AIA for high-risk artificial intelligence (AI) systems. They are expected to play a significant role in the implementation of the AIA in the EU. The content of this Guide is meant to assist with understanding the EU AIA and its processes but does not constitute legal advice for any specific compliance situation.

This Guide examines the CA as set out under the EU AIA. It does not offer a comparative study with other existing assessment processes required under other European legal acts that correspond to sectoral rules (for instance, the EU Medical Devices Regulation). The EU AIA includes various documentation obligations that are key for demonstrating compliance with the set CA process. We will refer to those obligations only where necessary to highlight their differences from the CA or explain where those other documentation obligations play a role in the performance of the CA.

In this updated version of the Guide, we will seek to address the final framework as adopted in April 2024 by the EU co-legislators and finalized in June 2024 through the corrigendum process. We hope the Guide serves as an essential resource for those who want to prepare for compliance with the EU AIA. Section II provides a high-level description of the EU AIA and the negotiations that led to its final version. Section III details the CA obligation's purpose, structure, and function. It identifies the questions that must be answered for an actor to assess whether they fall under the obligation to conduct a CA. Section III also explains when and how a CA should be performed and elaborates on all the requirements that need to be met during the CA process. Lastly, Section IV discusses the role of standards and the presumption of compliance with the requirements offered through adherence to harmonized standards.

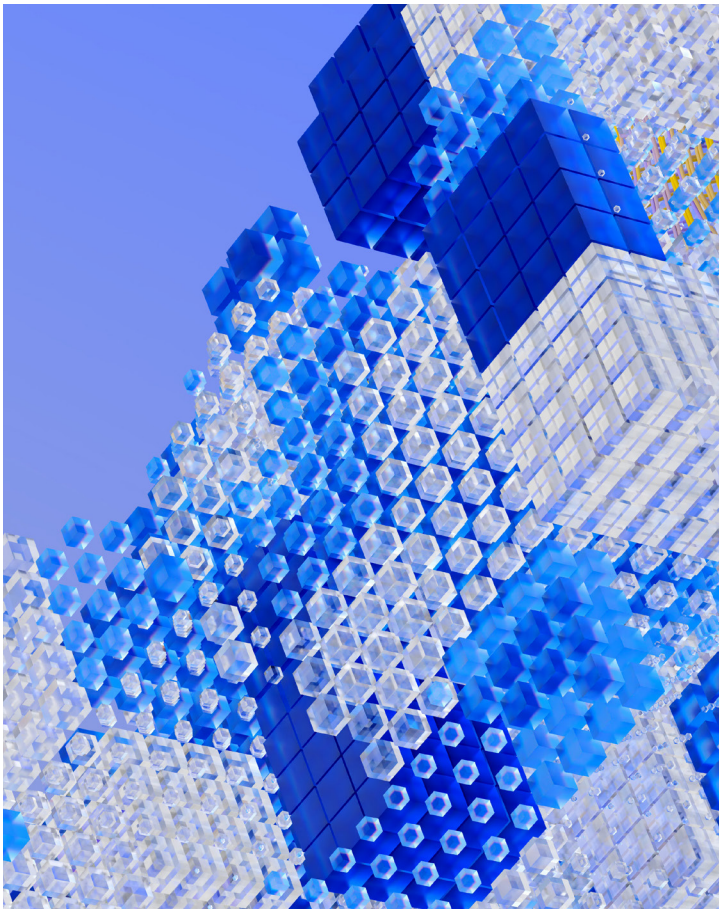

II. The EU Artificial Intelligence Act

A. The process leading to the final EU AIA

The European Commission is the main European institution that initiates legislation in the EU. In April 2021, it published the legislative proposal for a Regulation laying down harmonized rules on AI, the proposed EU AIA (COM(2021)206) (the Regulation). The Regulation was adopted in April 2024 by the European Parliament and the Council of the EU as co-legislators. The EU AIA was published in the Official Journal of the EU on July 12, 2024, entered into force on August 1, 2024, and will become fully applicable from August 2, 2026. Legally binding to all Member States, the provisions of the EU AIA will become applicable at different times over the next few years. Some of the significant milestones include the following.

- February 2, 2025: Prohibitions on certain AI systems began to apply (Chapter 1 and Chapter 2). The obligation to ensure AI literacy kicks off.
- August 2, 2025: The following rules start to apply:
    - notified bodies (Chapter III, Section 4);
    - general-purpose AI (GPAI)models (Chapter V);
    - governance (Chapter VII);
    - confidentiality (Article 78); and
    - penalties (Articles 99 and 100).
- August 2, 2026: The remainder of the AIA starts to apply, except Article 6(1) (classification as high risk of some AI systems which are safety components of products).
- August 2, 2027: Article 6(1) and the corresponding obligations in the Regulation begin to apply.

For a comprehensive overview of the EU AIA's implementation and compliance timeline, you can consult FPF's dedicated resource here.

B. The EU AIA is a risk-based regulation with enhanced obligations for high-risk AI systems

The EU AIA is structured on the basis of a precautionary and risk-based approach.

The EU AIA regulates AI technologies on the basis of the risks to the health, safety, and fundamental rights of individuals raised by their contextual use. The EU AIA prohibits outright certain uses of AI systems that raise unacceptable risks[1], and sets rules on the development and deployment of all other AI systems depending on whether they qualify as high, low, or minimum risk. The final version of the AI Act also addresses the systemic risk posed by GPAI, as illustrated in the AIA Risk Pyramid graphic below. This Guide focuses on high-risk AI systems. The CA obligation only applies to high-risk AI systems. The determination of whether an AI system qualifies as 'high-risk' is discussed under subsection Q2: Classification of the AI system as 'high-risk."

1 Pursuant to Article 5 of the EU AI Act.

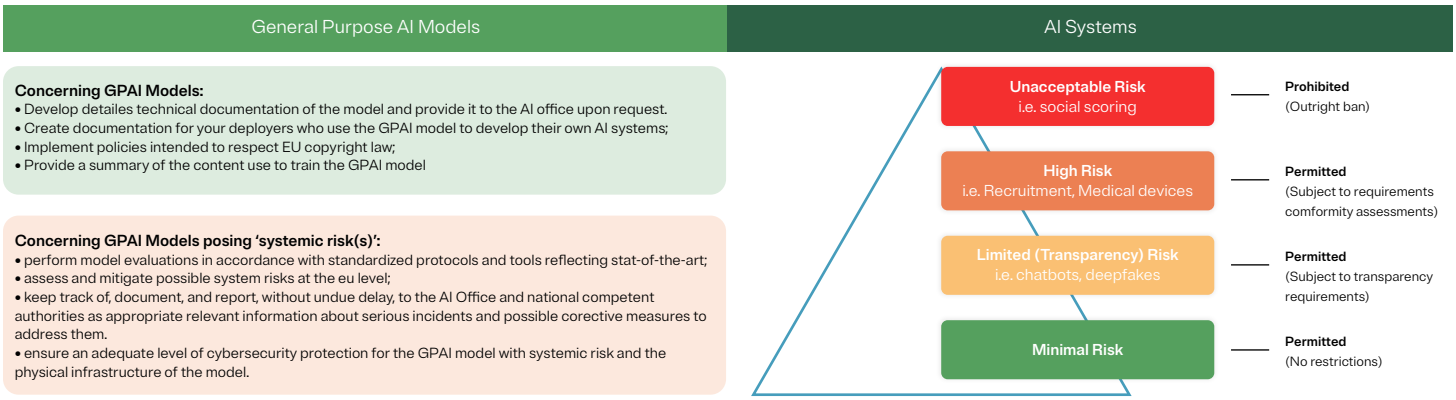

Illustration 1. Illustration of the pyramid of risk of AI systems, complemented by the obligations vested on GPAI Models (including the GPAI models with high impact capabilities that are considered to present 'systemic risks').

*The EU AIA was initially conceived as a '**product safety legislation.**'*

The Regulation aims to align with the processes and requirements found in laws that fall under the New Legislative Framework (NLF) in order to 'minimize the burden on operators and avoid any possible duplication' (Recital 124 of the AIA).[2] The EU AIA CA obligation is not a novelty in the broader EU context, as CAs are a tool for enhancing consistency within the EU Market under the NLF framework. CAs are also part of several EU product safety laws, such as the General Product Safety Regulation (GPSR),[3] the Machinery Regulation,[4] and the *In Vitro* Diagnostic Medical Devices Regulation.[5] It is possible an AI system that is a safety component of a product that falls under the scope of NLF laws will have been previously subject to a different, previously performed CA. Any AI system provider should consider this when determining their EU AIA CA obligations compliance strategy.

The EU AIA will apply without prejudice to other laws.

The application of the EU AIA is intended to be complementary and without prejudice to existing EU law, particularly on data protection, consumer protection, fundamental rights, employment, and protection of workers, and product safety, as mentioned in the Recital 9 of the AIA. One prominent example is the General Data Protection Regulation (GDPR).[8] While the EU AIA and the GDPR differ in material scope, if an AI system or model involves the processing of personal data, the legal obligations deriving from both Regulations may apply. For example, under Article 35 of the GDPR, the controller shall carry out a Data Protection Impact Assessment (DPIA); should the provider of a high-risk AI system qualify as a controller as defined by the GDPR, the provider will have to carry out both a DPIA as provided by the GDPR and a CA as provided by the EU AIA.

[2] An example of this can be found under the Risk Management System requirement in Step 4.1.

[3] Regulation (EU) 2023/988 of the European Parliament and of the Council of 10 May 2023 on general product safety, amending Regulation (EU) No 1025/2012 of the European Parliament and of the Council and Directive (EU) 2020/1828 of the European Parliament and the Council, and repealing Directive 2001/95/EC of the European Parliament and of the Council and Council Directive 87/357/EEC (Text with EEA relevance).

[4] Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on machinery and repealing Directive 2006/42/EC of the European Parliament and of the Council and Council Directive 73/361/EEC (Text with EEA relevance).

[5] Regulation (EU) 2017/746 of the European Parliament and of the Council of 5 April 2017 on in vitro diagnostic medical devices and repealing Directive 98/79/EC and Commission Decision 2010/227/EU (Text with EEA relevance)

[6] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance).

III. The Conformity Assessment obligation includes an overarching accountability framework for high-risk AI systems

'Conformity Assessment' is defined under Article 3(20) of the AIA as the process of demonstrating that a high-risk AI system complies with the requirements enumerated under Chapter III, Section 2 of the AIA. These requirements, which will be further elaborated under Step 4 of this Guide, are:

- risk management system (Article 9 of the AIA);
- data and data governance (Article 10 of the AIA);
- technical documentation (Article 11 of the AIA);
- record-keeping (Article 12 of the AIA);
- transparency and provision of information to deployers (Article 13 of the AIA);
- human oversight (Article 14 of the AIA); and
- accuracy, robustness, and cybersecurity (Article 15 of the AIA).

The CA consists of the assessment of whether the AI system qualifies as high-risk[7] and the assessment of risks that are part of the risk management system. The CA process additionally includes the assessment of requirements that must be built-in to any high-risk AI system (the assessment of the datasets used for training, validation, and testing, automatic recording of events, transparent operation of the AI system, human oversight capacity, AI system accuracy and robustness) as well as documentation obligations (technical documentation). The CA serves as a framework of assessments (technical and non-technical), requirements, and documentation obligations. This Guide follows four steps to determine if the obligation of performing a CA is applicable.

Step 1: Is a Conformity Assessment required?

The first step is to determine whether an organization is subject to the EU AIA CA legal obligation. The following flowchart outlines the key questions an organization should answer to determine whether a CA is required.

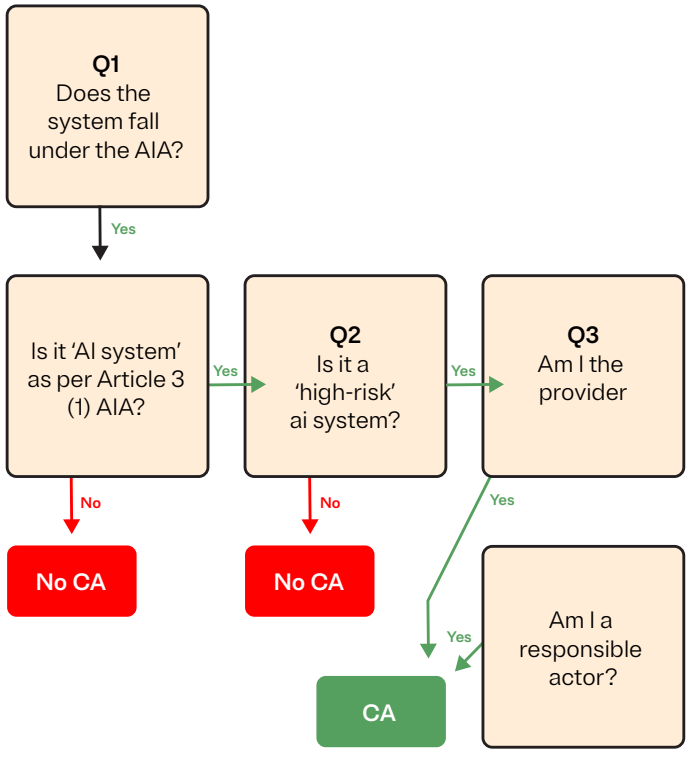

[7] The assessment of whether an AI system qualifies as high-risk and, therefore, requires a CA is a preliminary step, but it is necessary for the provider to determine if they are obligated to conduct a CA under the law.

Q1: Does the system fall under the AIA?

To determine whether a system is subject to the AIA, an organization must assess two key factors:

1. Does the system qualify as an AI system?

According to Article 3(1) of the AIA, 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.'

This definition aims to be technology-neutral and future-proof, ensuring that all AI systems with potential risks are covered under the AIA.[8]

2. Does the organization fall within the AIA's scope?

The AIA applies to both providers and deployers of AI systems, as defined by the provisions of Article 3(3)-(4) of the AIA:

- '*provider* means a natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge;' and
- '*deployer* means a natural or legal person, public authority, agency, or other body using an AI system under its authority, except where the AI system is used in the course of a personal non-professional activity.'

The organization should identify whether it falls under the AIA's material scope. Article 2(1) of the AIA sets out its scope, detailing which entities are covered by its provisions, specifically:

- providers placing on the market or putting into service AI systems or placing GPAI models on the EU market, regardless of where the provider is established or located (Article 2(1)(a));
- deployers that are located or established within the EU (Article 2(1)(b));
- providers or deployers of AI systems located in a third country, where the output produced by the AI systems is used in the EU (Article 2(1)(c));
- importers or distributors of AI systems (Article 2(1)(d));
- product manufacturers placing on the market or putting into service an AI system together with a product or under their own name or trademark (Article 2(1)(e));
- authorized representatives of providers not established in the EU (Article 2(1)(f)); and
- affected persons located in the EU (Article 2(1)(g)).

[8] See Guidelines on the Definition of an Artificial Intelligence System established by AI Act, AI Office, February 6, 2025.

Article 2 of the AIA also enumerates the cases that are not covered by the provisions of AIA:

- AI systems placed on the market, put into service, or used with or without modification exclusively for military, defense, or national security purposes (Article 2(3));
- public authorities in a third country or international organizations that use AI systems in the framework of international cooperation or agreements for law enforcement and judicial cooperation with the EU or member states (Article 2(4));
- AI systems or models, including their output, specifically developed or put into service for the sole purpose of scientific research and development (Article 2(6));
- AI systems in the research, testing, or development phase or AI models prior to being placed on the market or put into service, except for testing in real-world conditions (Article 2(8));
- obligations of deployers who are natural persons using AI systems in the course of purely personal non-professional activity (Article 2(10)); and
- free and open-source AI systems, unless they are high-risk or fall under specific prohibited practices (Article 2(12)).

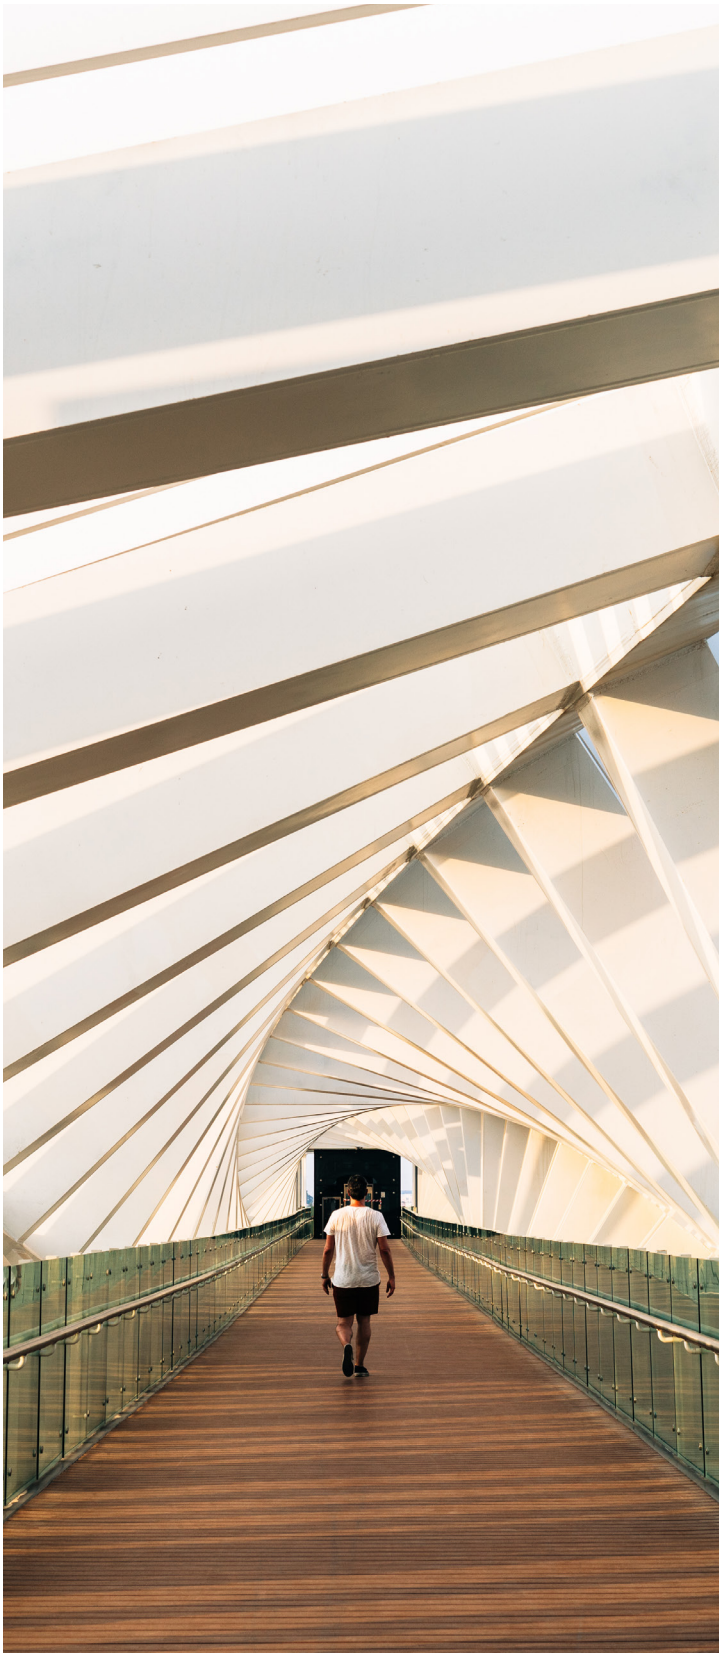

Q2: Is the AI system high risk?

Once the AIA's applicability has been established, it must be determined whether the AI system qualifies as 'high-risk.' As previously described in Section II B, the AIA classifies AI systems based on the following risk categories.

- **Unacceptable Risk:** Prohibited AI practices cover all AI systems considered a clear threat to people's health, safety, and rights, from social scoring by governments to real-time biometric identification systems in public systems.[9]

- **High-Risk:** AI Systems that pose a high risk to the health, safety, and rights of people, subject to strict obligations prior to being placed on the market.

- **Limited risk:** AI systems that pose risks associated with a lack of transparency in AI usage, such as chatbots or digital assistants.

- **Minimal or no risk:** AI systems that do not fall in any of the categories above and have no requirements to meet any obligations under the AIA.

- **GPAI:** the AIA provides specific rules for GPAI models that pose systemic risks.

    - According to Article 3(65) of the AIA, a 'systemic risk' refers to 'a risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain.'

The CA obligation applies only to high-risk AI systems. Article 6 of the AIA provides two categories of 'high-risk' AI systems:

1. **AI systems that are either safety components or standalone products already regulated by EU legislation in Annex I (Article 6(1) of the AIA):** AI systems intended to be used as a safety component in certain regulated products *and* where the AI system itself is a regulated product required to undergo a third-party conformity assessment in other EU harmonization legislation listed in Annex I.

2. **AI systems identified as particularly relevant or sensitive in Annex III (Article 6(2) of the AIA):** AI systems used for certain areas (such as biometrics, education, employment, financial services, critical infrastructure, access to essential services, law enforcement, migration, and asylum, administration of justice, or electoral processes). The only exceptions to this rule are considered in Article 6(3) of the AIA: AI systems in Annex III are not high-risk if they do not pose significant risks to health, safety, or fundamental rights. This applies when the AI system performs a narrow procedural task, improves results of prior human activities, detects decision-making patterns without influencing human assessment, or prepares for an assessment relevant to Annex III use cases.

[9] See 'Guidelines on Prohibited Artificial Intelligence practices established by Regulation (EU) 2024/1689 (AI Act)', AI Office, February 4, 2025.

Table 1: Classification of high-risk AI systems under the AIA

1st Category - Annex I AIA: product categories already covered by EU legislation	
The AI system is intended to be used as a safety component of a product or is itself a product, covered by the Union harmonization listed in Annex I AND The product is required to undergo a third-party conformity assessment under that legislation. ** Irrespective of whether an AI system is placed on the market or put into service independently from the product.*	• Machinery; • Safety of toys; • Recreational craft and personal watercraft; • Lifts and safety components of lifts; • Equipment and protective systems for use in explosive atmospheres; • Market of radio equipment; • Marker of pressure equipment; • Cableway installations; • Personal protective equipment; • Appliances burning gaseous fuels; • Medical devices (and in vitro diagnostic medical devices); • Civil aviation security; • Vehicles; • Marine equipment; • Interoperability of the rail system.
2nd Category - Sensitive AI systems categories listed in Annex III AIA	
AI system that falls under one or more of the eight critical areas and use cases referred to in Annex III.	1. Biometrics (and biometrics-based systems); 2. (Management and operations of) Critical infrastructure; 3. Education and vocational training; 4. Employment, workers management, and access to self-employment; 5. Access to and enjoyment of essential private services, public services, and benefits; 6. Law enforcement; 7. Migration, asylum, and border control management; 8. Administration of justice and democratic processes.

Under Article 7 of the AIA, the Commission has the power to amend Annex III through delegated acts to update the designation of high-risk AI systems listed under each of the eight purposes given. This ensures that the AIA remains flexible and future-proof. New AI systems can only be added to an Annex III purpose if two specific conditions are met.

1. The AI system is used in the context of any of the eight areas listed in Annex III (i.e., biometrics; critical infrastructure; education and vocational training; employment, workers' management and access to self-employment; access to and enjoyment of essential private services and essential public services and benefits; law enforcement; migration, asylum and border control management; administration of justice and democratic processes); and
2. The new system type presents risks to health, safety, or fundamental rights that are equivalent to or greater than those of previously-designated high-risk AI systems.

Article 7(2) of the AIA sets out eleven criteria[10] that the Commission must take into account when assessing the potential risk posed by an AI system use case. Developers and providers should consider these criteria closely before determining an AI system's high-risk status.

There are also more limited obligations imposed on providers of non-high-risk AI systems. Under Article 6(4) of the AIA, a provider who determines that an AI system is not high-risk shall document its assessment before that system is placed on the market or put into service. The provider shall provide the documentation of the assessment upon request of national competent authorities. Such a provider would also be obliged to register the AI system in the EU database established under Article 71 of the AIA, as per Article 80 of the AIA (read in conjunction with Article 71 of the AIA).

[10] According to Article 7(2) of the AIA, when assessing the potential risk posed by an AI system, the European Commission must take into account the following criteria: '(a) the intended purpose of the AI system; (b) the extent to which an AI system has been used or is likely to be used; (c) the nature and amount of the data processed and used by the AI system, in particular whether special categories of personal data are processes; (d) the extent to which the AI system acts autonomously and the possibility for a human to override a decision or recommendations that may lead to potential harm; (e) the extent to which the use of an AI system has already caused harm to health and safety, has had an adverse impact on fundamental rights or has given rise to significant concerns in relation to the likelihood of such harm or adverse impact, as demonstrated, for example, by reports or documented allegations submitted to national competent authorities or by other reports, as appropriate; (f) the potential extent of such harm or such adverse impact, in particular in terms of its intensity and its ability to affect multiple persons or to disproportionately affect a particular group of persons; (g) the extent to which persons who are potentially harmed or suffer an adverse impact are dependent on the outcome produced with an AI system, in particular because for practical or legal reasons it is not reasonably possible to opt-out from that outcome; (h) the extent to which there is an imbalance of power, or the persons who are potentially harmed or suffer an adverse impact are in a vulnerable position in relation to the deployer of an AI system, in particular due to status, authority, knowledge, economic or social circumstances, or age; (i) the extent to which the outcome produced involving an AI system is easily corrigible or reversible, taking into account the technical solutions available to correct or reverse it, whereby outcomes having an adverse impact on health, safety or fundamental rights, shall not be considered to be easily corrigible or reversible; (j) the magnitude and likelihood of benefit of the deployment of the AI system for individuals, groups, or society at large, including possible improvements in product safety; (k) the extent to which existing Union law provides for: (i) effective measures of redress in relation to the risks posed by an AI system, with the exclusion of claims for damages; (ii) effective measures to prevent or substantially minimise those risks.'

Q3: Who is responsible for performing the CA?

After having classified an AI system as 'high-risk,' the next task is identifying who is responsible for performing the CA. Typically, the provider is the primary responsible actor for conducting a CA. However, in exceptional circumstances, the obligation might fall on another actor.

Under Article 16 of the AIA, the actor responsible for performing the CA is the provider of the high-risk AI system. Article 3(3) of the AIA defines 'provider' as 'a natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model or that has an AI system or a general purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.' Even if the provider is not the direct designer/developer of the system, it still must ensure that requirements are embedded in the system prior to placing the system on the market or putting it into service.

Article 25 of the AIA determines that any distributor, importer, deployer, or other third party shall be considered to be a provider of a high-risk AI system (and be subject to the obligations under Article 16 of the AIA) if they fall under any of the following circumstances:

(i) they put their name or trademark on a high-risk AI system already placed on the market or put into service;

(ii) they make a substantial modification to a high-risk AI system that has already been placed on the market or has already been put into service in such a way that it remains a high-risk AI system; or

(iii) they modify the intended purpose of an AI system, including a general-purpose AI system, which has not been classified as high-risk and has already been placed on the market or put into service in such a way that the AI system concerned becomes a high-risk AI system.[11]

That initial provider shall closely cooperate with new providers and shall make available the necessary information and provide the reasonably expected technical access and other assistance that are required for the fulfillment of the obligations set out in the AIA, in particular regarding compliance with the conformity assessment of high-risk AI systems. This will not be the case where the initial provider has 'clearly specified that its AI system is not to be changed into a high-risk AI system and therefore does not fall under the obligations to hand over the documentation.'

[11] Article 25(1) of the EU AIA. However, pursuant to Recital 128, 'changes occurring to the algorithm and the performance of AI systems which continue to 'learn' after being placed on the market or put into service, namely automatically adapting how functions are carried out, should not constitute a substantial modification, provided that those changes have been pre-determined by the provider and assessed at the moment of the conformity assessment.'

Article 25(1)(b) and (c) of the AIA describes exceptional cases where a CA may need to be performed by a distributor[12], importer[13], deployer[14], downstream provider[15], or third party. Under Article 25(1)(b), one of these parties shall be considered a provider (and thus obligated to conduct a CA) when they have made substantial modifications to a high-risk AI system already placed on the market. Under Article 25(1)(c), a party's modifications of a non-high-risk AI system can also cause it to be considered a provider if the modifications cause the system concerned to become a high-risk system. Article 16(f) of the AIA obligates any "provider" to ensure that high-risk AI systems 'undergo the relevant conformity assessment procedure prior to … being placed on the market or put into service.'

The exact legal conditions that must be met for an actor other than the provider to be obliged to perform the original CA have yet to be determined beyond the examples given in the text of Article 25 of the AIA itself. However, the distributor, importer, deployer, downstream provider, or any other third party would, as a rule, be obliged to conduct a CA if they put their name or trademark on a high-risk AI system already placed on the market or put into service, or if they make a substantial modification to a high-risk AI system.

In some cases, Article 25(3) of the AIA designates the product manufacturer[16] as a 'provider' when the high-risk AI system is a safety component of a product covered by the Union harmonization legislation listed in Section A of Annex I of the AIA. Specifically, product manufacturers[17] may be subject to the obligations of Article 16 and can be responsible for a CA if, cumulatively:

(i) the high-risk AI system relates to products for which the laws in Annex I Section A apply;

(ii) the system is placed on the market or put into service together with the product; AND

(iii) under the name or trademark of the product manufacturer.

[12] According to Article 3(7) of the EU AIA, 'distributor' means 'any natural or legal person in the supply chain, other than the provider or the importer, that makes an AI system available on the Union market.

[13] According to Article 3(6) of the EU AIA, 'importer' means 'any natural or legal person located or established in the Union that places on the market an AI system that bears the name or trademark of a natural or legal person established in a third country.'

[14] According to Article 3(4) of the EU AIA, deployer means 'a natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.'

[15] According to Article 3(68) of the EU AIA, downstream provider means 'provider of an AI system, which integrates an AI model, regardless of whether the model is provided by themselves and vertically integrated or provided by another entity based on contractual relations.'

[16] The term 'product manufacturer' is used in the EU AIA but no definition was added. This was defined only under the Council proposal for the EU AIA as 'manufacturer within the meaning of any of the Union harmonization legislation listed in Annex II.' However, despite the definition's omission from the final AIA, Recital 87 refers to the 'relevant New Legislative Framework legislation' for the definition of 'product manufacturer.'

[17] Pursuant to Article 25(3) of the EU AIA.

Step 2: When should a CA be conducted?

Once a legal obligation to conduct a CA has been identified, it should be performed promptly, as the assessment must be completed before the high-risk system is placed on the market or put into service.

There are generally two circumstances in which a CA is required to be conducted.

1. **Before a high-risk system reaches the market:** A CA must be performed prior to placing an AI system on the EU market, which means prior to making it available (i.e., supplying for distribution or use) or prior to putting an AI system into service, which means prior to its first use in the EU market, either by the system's user or for (the provider's) own use.

2. **After an AI system reaches the market, if it has been substantially modified:** A CA will be required after the high-risk AI system has been placed on the market or put into service, in case the AI system is substantially modified. Substantial modification is considered any change, not foreseen or planned in the initial CA, that affects a system's compliance with the requirements for high-risk AI systems or results in a modification to the AI system's intended purpose.[18] For example, it will not be considered a 'substantial modification' when a high-risk AI system continues to learn after being placed on the market or put into service as long as these changes are pre-determined at the moment of the initial CA and are described in the initial technical documentation.[19] A CA would also be required if substantial modification of a system initially assessed as non-high-risk changes its designation to 'high-risk.'

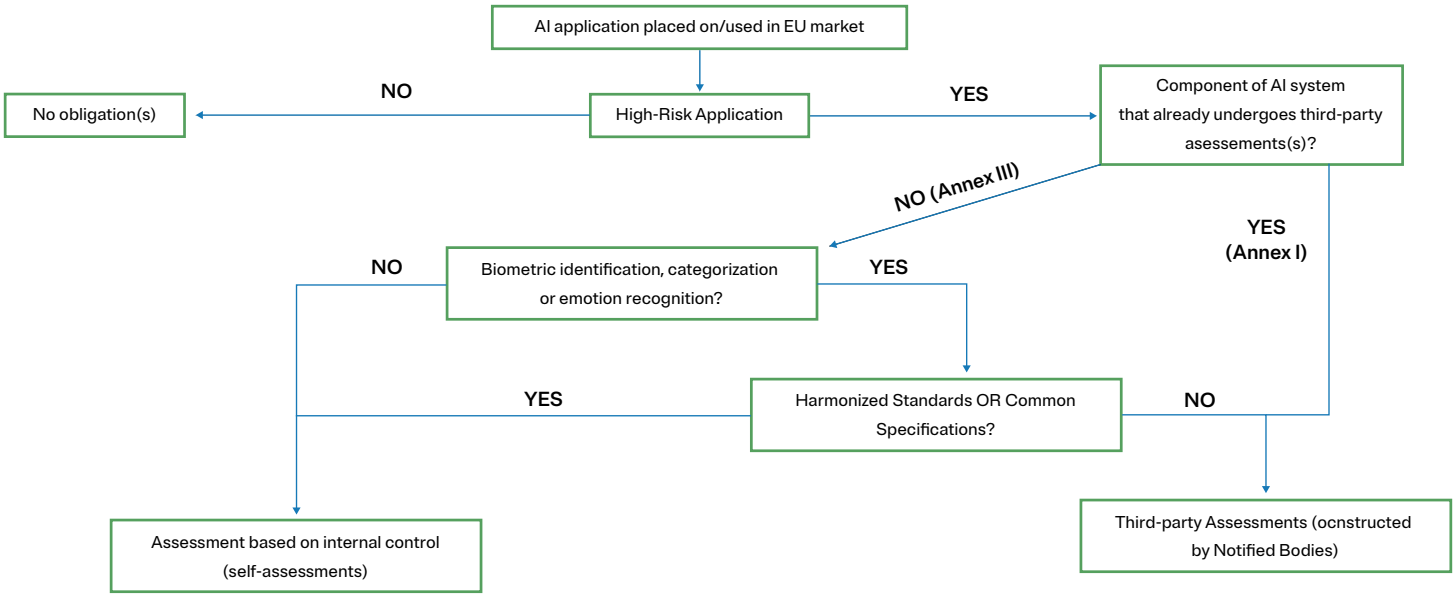

[18] Article 3(23) of the EU AIA defines 'substantial modification' as a 'change to an AI system after its placing on the market or putting into service which is not foreseen or planned in the initial conformity assessment carried out by the provider and as a result of which the compliance of the AI system with the requirements set out in Chapter III, Section 2 is affected or results in a modification to the intended purpose for which the AI system has been assessed.' In any case, a 'substantial modification' leads to a 'new' AI system, for which a new CA has to be conducted.

[19] Article 43(4) of the EU AIA. The inclusion of the possible changes in the high-risk AI system in the technical documentation (see Step 4.3) is a legal requirement for the change to qualify as 'non substantial modification.' It is an additional obligation of the provider to give, in the context of the transparency requirements (see Step 4.5), information about, inter alia, the 'characteristics, capabilities, and limitations of performance of the high-risk AI system.' Part of this information shall be any 'predetermined changes to the performance of the system.

Step 3: Who should conduct a CA?

A CA can be conducted internally or by a third party. In the internal process, the provider (or any other responsible actor)[20] may perform the CA, while the third-party process requires an assessment by an external independent third party, a 'notified body.' Article 43(2) of the AIA specifies that CAs for systems designated 'high risk' under points 2-8 of Annex III (see above) must follow the procedures for internal control.

Table 2 below details the type of CA process that should be followed in each case to assess a high-risk AI system.

Internal conformity assessment (Annex VI of the AIA)

The process for an internal CA is described in Annex VI of the AIA. The provider is required to verify:

- an established quality management system exists in compliance with the requirements detailed in Article 17 of the AIA;
- The information contained in the technical documentation is used to assess the compliance of the AI system with the relevant essential requirements set out in Chapter III, Section 2 of the AIA; and
- the design and development process of the AI system and its post-market monitoring, as referred to in Article 72 of the AIA, is consistent with the technical documentation.

Quality management system (Article 17 of the AIA)

The provider's obligation to have a Quality Management System (QMS) in place is set out in Article 17 of the AIA. The QMS must be clearly documented through written policies, procedures, and instructions organized in a systematic and orderly manner. The AIA further specifies that the QMS should encompass the following key elements.

- Regulatory compliance strategy, including compliance with conformity assessment procedures and procedures for managing any modifications made to the high-risk AI system.
- Design control procedures, which encompass techniques, procedures, and systematic actions to guide the design, design control, and quality assurance of the high-risk AI system.
- Development and quality assurance, which regards techniques, procedures, and systematic actions to ensure the quality control and quality assurance of the high-risk AI system.
- Examination, testing, and validation procedures carried out before, during, and after the development of the high-risk AI system, along with the specified frequency of these activities.
- Technical specifications and standards to be applied, and if harmonized standards are not fully applicable or do not cover all requirements, the means that should be used to ensure compliance with the required standards.
- Data management procedures (data acquisition, data collection, data analysis, data labeling, data storage, data filtration, data mining, data aggregation, data retention, and other data operations) performed before

[20] See Q3: Am I the responsible actor?

and for the purpose of placing the high-risk AI system on the market or putting it into service.

- Risk management system.
- The setup, implementation, and maintenance of a post-market monitoring system.
- Incident reporting procedures.
- Procedures for communication with national competent authorities, relevant regulatory bodies, other operators, customers, or interested parties.
- Record-keeping systems and procedures for relevant documentation and information;
- Resource management, including security-of-supply related measures.
- An accountability framework that outlines the responsibilities of the management and other staff with regard to the aspects listed above.

Technical documentation (Article 11 of the AIA)
Providers of high-risk AI systems must prepare and maintain technical documentation demonstrating compliance. This documentation should include general descriptions, specifications, development and training data, risk management documentation, testing and validation reports, and a post-market monitoring plan.

Design and development process and post-market monitoring (Article 72 of the AIA)
Providers must establish and maintain a system for collecting, documenting, and analyzing data on AI system performance throughout its lifecycle. This includes data collection, analysis, corrective actions, and continuous improvements based on monitoring outcomes.

EU declaration of conformity (Article 47 of the AIA)
When an internal CA is concluded, the provider must draw up an 'EU declaration of conformity' (Article 47 of the AIA) (the Declaration) and keep it for national competent authorities for 10 years after the AI system has been placed on the market or put into service. Annex V of the AIA specifies that the Declaration must include:

- AI system identification: name, type, and any additional references that allow for the identification and traceability of the AI system;
- provider information: name and address, and where applicable, the information of their authorized representative;
- a statement that the EU declaration of conformity is issued under the sole responsibility of the provider;
- conformity assurance: a statement confirming the AI system's compliance with relevant EU laws;
- compliance with the GDPR, Regulation 2018/1725, and Directive 2016/680, where the system involves the processing of personal data;
- reference to standards and specifications used in relation to the Declaration;

- information on the notified body, as well as a description of the CA procedure performed, and identification of the certificate issued; and

- the place and date of issue of the Declaration, as well as the identification details and the role of the signatory.

This structured Declaration ensures transparency and accountability, affirming that the AI system meets all regulatory requirements. In particular, the statement of compliance with the GDPR if the AI system processes personal data and the confirmation that the QMS, technical documentation, and design and development process align with AI requirements. This is relevant, given that the performance of a DPIA (if the legal conditions are met) is likely to be part of this declaration.

CE marking (Article 48 AIA)

To publicly display that the AI system has cleared an internal conformity assessment, pursuant to Article 48 of the AIA, the provider (or other responsible entity) is required to affix the CE marking of conformity[21] in a visible, legible, and indelible manner. For high-risk systems provided digitally, this may include the use of a digital CE marking if it can be easily accessed using the interface from which the system is accessed. Additionally, when a notified body is involved in the CA process, the identification number of the notified body responsible for the CA must be affixed next to the CE marking. The identification number shall be affixed by the notified body itself or, under its instructions, by the provider or its authorized representative.

Third-party conformity assessment (Annex VII of the AIA)

Alternatively, in the case of a third-party CA, an assessment by an independent CA body is required. Annex VII of the AIA describes the conformity assessment procedure that combines an evaluation of the provider's quality management system (QMS) and a review of technical documentation of high-risk AI systems by a 'notified body,' which performs conformity assessment activities, such as testing, certification and inspection of high-risk AI systems before they can be placed on the market.

This procedure is required for high-risk AI systems listed in point 1 of Annex III (biometrics) when harmonized standards (Article 40 of the AIA) or common specifications (Article 41 of the AIA) are unavailable or not fully applied. Additionally, certain high-risk AI systems, particularly those used by law enforcement, immigration, or asylum authorities, must undergo this assessment. The notified body also conducts periodic audits to confirm continued compliance, ensuring the AI system remains in line with regulatory standards throughout its lifecycle.

A 'notified body' is an independent CA body designated by an EU Member State to assess whether certain products, including high-risk AI systems under the AI Act, comply with EU regulations before they can be placed on the market. In order to become a notified body, a CA body may submit an application for notification to the notifying authority of the Member State in which it is established. Notifying authorities[22] may only notify CA bodies that satisfy the requirements laid down in Article 31 of the AIA. More information can be found in Chapter III, Section 4 of the AIA.

[21] The CE marking of conformity is subject to the general principles set out in Article 30 of Regulation (EC) No 765/2008.

[22] According to Article 28(1) AIA, each Member State shall designate or establish a notifying authority responsible for setting up and carrying out the necessary procedures for the assessment, designation and notification of conformity assessment bodies and for their monitoring.

[23] The provider may choose any of the notified bodies unless the system is intended to be put into service by law enforcement, immigration or asylum authorities as well as EU institutions, bodies or agencies. In that case, it is the market surveillance authority referred to in Article 63(5) or (6), that shall act as a notified body.

Following Article 43(1) of the AIA and in accordance with Annex VII, 4.1. of the AIA, it is understood that the provider shall submit two applications to the notified body of their choice[23] - one for the quality management system and one for the technical documentation. Annex VII of the AIA specifies the information that must be included in each application and outlines the criteria against which they should be assessed. According to Annex VII, 5 of the AIA, once the QMS is approved, the notified body conducts ongoing surveillance to ensure that the provider continues to fulfill the terms and conditions of the approved QMS. For instance, the notified body shall carry out periodic audits to ensure that the provider maintains and applies the QMS and shall provide the provider with an audit report.

Future updates (Article 43(5) of the AIA)
The European Commission may update these processes without the need to review the AIA: Article 43(5) of the AIA gives the Commission the power to adopt delegated acts in accordance with Article 97 for updating Annex VI ('CA Procedure based on Internal Control') and Annex VII ('Conformity Based on Assessment of Quality Management System and Assessment of Technical Documentation') in light of technical progress.

Certification decision by notified bodies (Article 44 of the AIA)
In accordance with Annex VII of the AIA, the notified body shall communicate to the provider the conclusions of the assessment of the QMS and the technical documentation (the documents submitted), as well as the reasoned assessment decision. The decision could have two outcomes, a positive decision of conformity or a negative decision of non-conformity:

A. In conformity:
Should the notified body find that the high-risk AI system is in conformity with the requirements of the AIA, it will issue an EU technical documentation assessment certificate (also see Article 44 of theAIA). The certificate has limited time validity and can be suspended or withdrawn by the notified body.

Responsibility remains with the provider to draw up an EU declaration of conformity and affix a CE marking of conformity, as with an internal CA, as described above.

Any change to the AI system that could affect the compliance of the AI system with the requirements of the AIA or change its intended purpose shall be approved by the notified body that issued the EU technical documentation assessment certificate.

Additionally, the notified body shall have the right to make periodic audits of the approved quality management system in order to make sure that the provider duly fulfills the terms and conditions of the approved quality management system (see Annex VII, 5).

B. Not in conformity:
Under Article 44(3) of the AIA, notified bodies may determine that a previously-certified AI system no longer meets the requirements set forth in Section 2 of the AIA. In such cases, the body may 'suspend or withdraw the certificate issued or impose restrictions on it,' unless compliance is 'ensured by appropriate corrective action taken by the provider ... within an appropriate deadline set by the notified body.' If a notified body withdraws or suspends a certification, it must provide reasons for its decisions, and an appeals procedure must be available.

The language of Article 43(3) of the AIA, which requires that '[a]n appeal procedure against decisions of the notified bodies, including on conformity certificates issued, shall be available' may also afford providers a right to appeal against an initial decision of the notified body to withhold certification, though the requirement is given at the end of a subsection describing the procedure for withdrawing or suspending previously granted certifications (Article 44(3) of the AIA).

Table 2. Internal or third-party CA according to the high-risk AI system

<table>
<tr><th colspan="3">AI systems as safety component of a product/product covered by other EU harmonization legislation (Annex I AIA)</th></tr>
<tr><th colspan="3">Type of CA</th></tr>
<tr><td colspan="2">Machinery
Safety of toys
Recreational craft and personal watercraft
Lifts and safety components of lifts
Equipment and protective systems for use in explosive atmospheres
Market of radio equipment
Marker of pressure equipment
Cableway installations
Personal protective equipment
Appliances burning gaseous fuels
Medical devices (and in vitro diagnostic medical devices)
Civil aviation security
Vehicles
Marine equipment
Interoperability of the rail system</td><td>Type of Conformity Assessment required under the respective EU legislation in Annex I (Article 43(3)).</td></tr>
<tr><th colspan="3">AI Systems expressly listed in Annex III</th></tr>
<tr><th colspan="3">Type of CA</th></tr>
<tr><td rowspan="2">1. Biometrics (and biometrics-based systems)</td><td>Provider has applied harmonized standards[24] or common specifications[25]?</td><td>Internal CA or Third-Party CA</td></tr>
<tr><td>Provider has not applied harmonized standards/ common specifications or has applied them only in part?</td><td>Third-Party CA</td></tr>
</table>

[24] According to Article 3(27) AIA, harmonized standard means 'a harmonized standard as defined in Article 2(1), point (c), of Regulation (EU) No 1025/2012.'

[25] According to Article 3(28) AIA, common specification means 'a set of technical specifications as defined in Article 2, point (4) of Regulation (EU) No 1025/2012, providing means to comply with certain requirements established under this Regulation.'

AI systems as safety component of a product/product covered by other EU harmonization legislation (Annex I AIA)	
	Type of CA
2. Critical infrastructure 3. Education and vocational training 4. Employment, workers management, and access to self-employment 5. Access to and enjoyment of essential private services and public services and benefits 6. Law enforcement 7. Migration, asylum, and border control management 8. Administration of justice and democratic processes	**Internal CA** **(The Commission may amend this rule and require third-party CA through delegated acts.)**

Ongoing requirements: post-market monitoring system

A CA is not a one-off exercise. Regardless of whether it is an internal CA or a third-party CA, the provider is required to establish a monitoring system that enables them to verify that the essential requirements are being complied with throughout the lifecycle of the high-risk AI system. For that, Article 72 of the AIA requires providers to establish a post-market monitoring system, which will form part of the 'quality management system' of Article 17 of the AIA. This post-market monitoring system must be proportionate to the nature of the AI technologies used and the risks of the high-risk system. Since the monitoring takes place after the AI system has entered the market, the user/deployer is also responsible for informing the provider regarding the AI system's performance. The AIA sets the conditions for effective communication and sharing of relevant information between the provider and the user/deployer of the high-risk AI system laid down in the provisions with regard to the transparency and provision of information to deployers (Article 13 of the AIA), the risk management system (Article 9 of the AIA).

A post-market monitoring system is defined under Article 3(25) of the AIA as 'all activities carried out by providers of AI systems to collect and review experience gained from the use of AI systems they place on the market or put into service for the purpose of identifying any need to immediately apply any necessary corrective or preventive actions.'

The AIA establishes specific conditions for effective communication and information sharing between the provider and the users/deployers of the high-risk AI system, including:

- incident reporting: users must promptly report malfunctions, failures, or other system irregularities to the provider;
- performance feedback: providers must actively collect feedback on system performance from deployers and end-users;
- risk updates: providers must share updates regarding newly identified risks or necessary adjustments with deployers; and
- regulatory notifications: if a significant safety or compliance issue arises, providers must notify national competent authorities and collaborate on mitigation measures. This ensures that providers can detect compliance issues early and take necessary preventive measures.

When is a CA not required?

Derogation for exceptional cases (Recital 130 and Article 46 of the AIA)

The AIA introduces exceptional cases under Article 46 of the AIA, where there can be a derogation from the normal CA process. Only for exceptional reasons of public security or the protection of life and health of persons, environmental protection, and the protection of key industrial and infrastructural assets can a high-risk AI system be authorized to be placed on the market or put into service by a market surveillance authority (within the territory of the Member State concerned) while the CA has not been concluded. Such authorization should be only for a limited period while the necessary CA procedures are being carried out, taking into account the exceptional reasons justifying the derogation.

Corrective actions

When a provider considers or has reason to consider that a high-risk AI system in use is not in conformity with the AIA, they shall immediately (1) inform the relevant actors (e.g., distributors, importers, user/deployer, national competent authority, etc.) and (2) take corrective actions, as required under Article 20 of the AIA. Corrective actions might range from bringing the system back to conformity to withdrawing or recalling the system from the market.

Additional requirements for high-risk systems presenting risks (Article 20(2) of the AIA)

If the non-compliant AI system presents risks within the meaning of Article 79(1) of the AIA, including risks to health, safety, or fundamental rights, the provider must: (i) investigate the causes, in collaboration with the reporting deployer, where applicable, and conduct an assessment to determine the extent of the threat posed by the AI system, and (ii) inform the market surveillance authorities.

[23] The provider may choose any of the notified bodies unless the system is intended to be put into service by law enforcement, immigration or asylum authorities as well as EU institutions, bodies or agencies. In that case, it is the market surveillance authority referred to in Article 63(5) or (6), that shall act as a notified body.

Step 4: Assess conformity with all requirements for high-risk AI systems

All high-risk AI systems must undergo the CA process, which ensures compliance with the requirements set out in Chapter III, Section 2 of the AIA. This section outlines these requirements, their significance, and the phase of the AI system's lifecycle at which they must be met.

All requirements must be fulfilled before a high-risk AI system is placed on the market or put into service unless otherwise specified by the AIA. The provider has the primary responsibility for ensuring compliance with these requirements. However, under certain circumstances, especially when substantial modifications are made, the deployer may also assume the role of a provider and take on compliance obligations (Article 43 and Recital 128 of the AIA).

- CA must verify that the AI system aligns with the generally acknowledged state-of-the-art, including harmonized standards and common specifications, as referred to in Articles 40 and 41 of the AIA, or those set out in Union harmonization law.
- Compliance with the requirements is not a one-time obligation. The system must be assessed for conformity throughout its lifecycle.
- The CA must evaluate whether the AI system is designed and used according to its intended purpose while also accounting for reasonably foreseeable misuse. The provider must establish a risk management system that proactively identifies and mitigates potential risks (Recital 64 of the AIA).

Requirements for High-Risk AI Systems

4.1	Risk Management (Article 9 & Recital 65)
4.2	Data & Data Governance (Article 10 & Recitals 67-69)
4.3	Technical Documentation (Article 11, Recital 71 & Annexes IV and VII)
4.4	Record Keeping (Articles 12 & Recitals 66 and 91)
4.5	Transparency Obligations (Article 13 & Recital 72)
4.6	Human Oversight (Article 14 & Recital 73)
4.7	Accuracy, Robustness & Cybersecurity (Article 15 & Recitals 74-76)

4.1. Is there a risk management system (RMS) in place?

Risk Management System (RMS)

Providers should establish, implement, document, and maintain a Risk Management System throughout the life-cycle of the high-risk AI system.

Elements to be included in a RMS:

1. Identification & assessment of risks (known and reasonably foreseeable);
2. Evaluation of other possibly arising risks (see ' post-market monitoring' and the requirement of 'automatic recording of events');
3. Adoption of Risk management measures (during the design and development phase of the AI system); and
4. Testing of the high-risk AI system.

Under Article 9 of the AIA, providers of high-risk AI systems must establish, implement, document, and maintain an RMS that runs throughout the entire lifecycle of the high-risk AI system. The RMS must be maintained and monitored even after the AI system has been placed on the market.

The RMS is a continuous and iterative process that should be monitored, reviewed, and updated regularly to ensure that it remains relevant and effective. The AIA also requires that the provider and the user/deployer of the high-risk AI system maintain good communication and share relevant risk-related information with one another to maintain the AI system's safety and compliance.

A detailed description of the RMS is required as part of both the technical documentation required by Article 11(1) of the AIA and the quality management system required by Article 17(g) of the AIA.[26] Providers operating within regulated sectors (e.g., credit institutions under Directive 2013/36/EU) must ensure that their existing risk management frameworks integrate the obligations set out in Article 9 of the AIA.[27]

Article 9 of AIA presents the following elements that should be part of an RMS.

1. Identification and assessment of risks

Providers shall identify and evaluate (a) known risks and (b) (reasonably) foreseeable risks that the AI system might pose to the health, safety, and fundamental rights of natural persons. The assessment shall be performed on the basis of the intended purpose[28] of the AI system as well as its reasonably foreseeable misuse.

2. Evaluation of other possibly arising risks

Providers must analyze data gathered during the post-marketing monitoring phase and, on the basis of that analysis, evaluate risks that may arise during the AI system's use. This part of the RMS is closely related to the requirement of record-keeping (Article 12 of the AIA). As will be further discussed under subsection 4.4, the automatic recording of events (logs) while an AI system is operating ensures a level of traceability regarding an AI system's functioning throughout its lifecycle. This enables monitoring of AI systems to identify situations where an AI system may present risks. Regular updates to the RMS are necessary to ensure that providers meet obligations to conduct post-market monitoring and evaluate future risks.

3. Adoption of risk management measures

The AIA does not provide examples of potential risk management measures. However, it identifies the criteria that providers should consider when deciding on the most appropriate risk management measures. Article 9 of the AIA mandates that providers establish an RMS that is continuously and systematically updated throughout the AI system's lifecycle. This should involve identifying possible risks associated with the AI system concerning health, safety, and fundamental rights; evaluating the likely impact; adopting suitable measures to mitigate identified risks; testing the AI system to assess its performance against potential risks; implementing monitoring measures to ensure that the AI system operates within acceptable risk levels during its deployment.

[26] See Annex VII, 4.2 (c).

[27] See Section II.B , point 2 'The EU AIA was initially conceived as a 'safety product legislation,' whereby the intention of the European legislature to avoid duplication of processes, is discussed.

[28] According to Article 3(12) of the AIA, intended purpose means 'the use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials and statements, as well as in the technical documentation.'

While the interaction between risk management measures and other requirements for AI systems is not explicitly provided for in the AIA, it should be understood as an integrated process. Risk management measures should be aligned with data governance (Article 10 of the AIA), record keeping (Article 12 of the AIA), transparency (Article 13 of the AIA), human oversight (Article 14 of the AIA), and cybersecurity (Article 15 of the AIA).

Any residual risk associated with each identified hazard, as well as the overall residual risks of the AI system as a whole, must be reduced to acceptable levels[29]; providers should aim to eliminate or mitigate risks as far as technically feasible (Article 9(5)(a)) of the AIA, and if risk elimination is not possible, they must implement appropriate measures for mitigation and control.

Regarding risks that may appear during the use of the AI system, when adopting the risk management measures, the provider shall take into account:

> (a) the technical knowledge, experience, education, and training to be expected by the user/deployer. The provider shall share all relevant information with the user/deployer and, where appropriate, train them; and
>
> (b) the environment in which the system is intended to be used (context of use).

4. Testing of the high-risk AI system

Part of the risk management system requires the testing of the high-risk AI system in order to make sure that the AI system performs in a manner that is consistent with its intended purpose and that the AI system is in compliance with the essential requirements for high-risk AI systems.

- The testing should take place during the development phase of the AI system and, in any case, before placing the AI system on the market.

- Testing shall be conducted against preliminarily defined metrics and probabilistic thresholds appropriate to the high-risk AI system's intended purpose.

The AIA highlights that while implementing the RMS, the provider should pay particular attention to 'persons under the age of 18 and, as appropriate, other vulnerable groups' that might interact with or be adversely impacted by the high-risk AI system.

[29] The term 'acceptable' refers to the level of residual risk that remains after all reasonable and feasible measures to eliminate or mitigate risks have been implemented, which is determined through risk minimization to the extent technically feasible; balancing against the intended purpose and benefits of the high-risk AI system; and alignment with the relevant legal and regulatory requirements.

4.2. Were high-quality datasets used for training, validation, and testing? (Article 10 of the AIA)

Data Governance	
If a high-risk AI system makes use of techniques involving the training of models with data, providers must ensure that they use high-quality datasets for training, validation, and testing.	
High-quality datasets: • relevant • representative • appropriately vetted for errors • as complete as possible • appropriate statistical properties • mitigation of bias	Data governance practices must cover the following: • relevant design choices • transparency as to the original purpose of data collection • data collection processes • data preparation processing operations • formulation of relevant assumptions • prior assessment of the availability, quantity, and suitability of the data sets that are needed • examination in view of possible biases • identification of any possible data gaps or shortcomings.
Presumption of conformity: if an AI system is trained and tested on data reflecting the specific geographical, behavioral, or functional setting within which the AI system is intended to be used.	

The AIA establishes in Article 10 that the training, validation, and testing of data sets should be subject to appropriate governance and management practices in line with the AI's intended purpose.

High-quality datasets are vital to building safe AI systems that perform as intended and do not lead to discriminatory outputs. High-risk AI systems that make use of techniques involving the training of models with data shall be developed on the basis of training, validation, and testing[30] data sets that meet specific quality criteria.

According to Annex IV(2)(d) of the AIA, detailed information about the datasets used for training, validation, and testing, such as their provenance, scope, and main characteristics, should be part of the technical documentation required by Article 11(1) of the AIA.

Criteria for high-quality datasets

High-quality datasets are datasets that are sufficiently relevant, representative, appropriately vetted for errors, and as complete as possible in view of the intended purpose of the AI system. The AIA does not explicitly define 'appropriate statistical properties,' yet it generally refers to ensuring datasets reflect the demographic, behavioral, and functional characteristics of the affected population. Providers should consider best practices from industry standards when determining appropriate statistical properties. Specific attention should be given to mitigating possible biases in the datasets, which might create risks to fundamental rights or discriminatory outcomes for the persons affected by the high-risk AI system.

Article 10 of the AIA enumerates the governance practices that providers should adhere to in further detail. Those practices include establishing transparent data collection processes and data-preparation processing operations,

[30] The AIA provides definitions on 'training data,' 'validation data,' 'validation data set,' and 'testing data' under Article 3 (29), (30), (31) and (32) respectively.

assessment of the availability, quantity, and suitability of the datasets, and implementing appropriate measures to detect, prevent, and mitigate possible biases likely to affect the health, safety, and fundamental rights, and ensuring that datasets account for geographical, behavioral and functional factors relevant to the deployment of AI systems.

The presumption of conformity established in Article 42 of the AIA applies where providers have trained and tested their high-risk AI systems on data reflecting the specific geographical, behavioral, or functional setting within which the AI system is intended to be used. However, this presumption primarily relates to compliance verification rather than data governance processes. It serves as an indicator that if datasets adequately reflect deployment conditions, the AI system is more likely to comply with data governance requirements under Article 10 of the AIA.

Processing of special categories of personal data (Article 10(5) of the AIA)

Article 10(5) of the AIA allows the processing of special categories of personal data as defined by Article 9 of the GDPR for the purposes of bias monitoring, detection, and correction. This exceptional processing shall be subject to appropriate safeguards for the fundamental rights and freedoms of natural persons.

In addition to the GDPR requirements, Article 10(5) of the AIA further established conditions that must be met before this processing can take place.

- The detection and correction of bias must not be possible using other types of data, including synthetic or anonymized data.
- Special category data are subject to state-of-the-art security and privacy-preserving measures, including pseudonymization, as well as technical limitations on re-use.
- The data must be secured and protected, subject to appropriate safeguards, including strict access controls, documentation of access, and confidentiality obligations for authorized persons to prevent misuse.
- Special category data must not be transmitted, transferred, or accessed by unauthorized third parties.
- Special category data must be deleted as soon as the bias is corrected or when the retention period expires, whichever comes first.
- Detailed records of processing activities must be maintained. This should include why the processing was strictly necessary to detect and correct bias and why the objective could not have been achieved by processing other data.

Annex VII of the AIA further reinforces data governance obligations related to third-party CAs. Where a negative CA has resulted from the assessment of a non-compliant AI system, the system must be retrained before applying for certification. In such cases, the reasoned assessment decision of the notified body refusing to issue the EU technical documentation assessment certificate must contain specific considerations on the quality of data used to train the AI system, including the reasons it is non-compliant. Furthermore, providers must ensure that all high-quality dataset requirements under the AIA remain aligned with GDPR principles, including lawfulness and fairness, purpose limitation, accuracy, and data minimization and necessity.

4.3. Has technical documentation been drawn up? (Article 11 of the AIA)

Technical Documentation
Providers should draw up technical documentation and keep it up-to-date.

Elements to be included in the TD (Annex IV):

1. a general description of the AI system;
2. a detailed description of the elements of the AI system and the process for its development;
3. detailed information about the monitoring, functioning, and control of the AI system;
4. a description of the appropriateness of the performance metrics for the specific AI system;
5. a detailed description of the risk management system (RMS);
6. a description of relevant changes made by the provider to the system through its lifecycle;
7. a list of harmonized standards applied in full or in part;
8. a copy of the EU Declaration of Conformity; and
9. a detailed description of the system in place to evaluate the AI system performance in the post-market phase, including the post-market monitoring plan.

The technical documentation shall be drawn up by the provider before the high-risk AI system is placed on the market or put into service and must be kept up-to-date.

With regard to the content, Annex IV of the AIA sets out the minimum elements to be included in the technical documentation:

- a general description of the AI system (e.g., intended purpose, nature of data processed, description of hardware and software, and the interaction with the AI system, whether the AI system is a component of a product, the system's expected output, scenarios of non-use of the AI system, etc.);
- a detailed description of the elements of the AI system and the process for its development;
- detailed information about the monitoring, functioning, and control of the AI system;
- a description of the appropriateness of the performance metrics for the specific AI system;
- a detailed description of the risk management system in accordance with Article 9;
- a description of relevant changes made by the provider to the system through its lifecycle;
- a list of the harmonized standards applied in full or in part. Where no such harmonized standards have been applied, a detailed description of the solutions adopted to meet the requirements set out in Chapter III, Section 2, including a list of other relevant standards and technical specifications applied;

- a copy of the EU declaration of conformity (which will be issued after the CA is successfully performed and hence cannot, in fact, be part of the technical documentation assessed as part of the CA process); and

- a detailed description of the system in place to evaluate the AI system performance in the post-market phase, including the post-market monitoring plan.

Pursuant to Article 18 of the AIA, providers should keep the technical documentation available for national competent authorities for a period of 10 years after a high-risk AI system has been placed on the market or put into service.

Where a national supervisory authority of a Member State finds that technical documentation is not available, incomplete, or not up to date, it shall require a provider to act and remedy the non-compliance. Article 83 of the AIA specifies that in cases of non-compliance with technical documentation obligations, corrective action may be taken by authorities, including requesting the provider to complete or update the documentation within a specified period of time or recalling or withdrawing the AI system from the market without delay.

4.4. Is the automatic recording of events ('logs') possible? (Article 12 of the AIA)

Article 12 of the AIA requires high-risk AI systems to be technically capable of automatically recording events (logs) throughout their lifetime. This requirement ensures traceability, transparency, and compliance, enabling responsible actors to monitor and assess the system throughout its lifecycle.

High-risk AI systems must be technically capable of logging events while in operation. Providers are responsible for implementing logging mechanisms that automatically record system events, ensuring that these records remain available for compliance verification and oversight (AI system traceability).

The aim of automatic event recording is to enable continuous monitoring of the system's operation, allowing responsible actors to identify and assess risks that might arise during the use of any substantial modifications. It also facilitates post-market monitoring and helps the user comply with their monitoring obligations (Article 12(2) of the AIA).

Article 12(3) of the AIA especially states that in the case of remote biometric identification systems, the logging capabilities shall provide, at a minimum:

- recording of the period of each use of the system;

- the reference database against which input data has been checked by the system;

- the input data for which the search has led to a match; and

- the identification of the natural persons involved in the verification of the results.

According to Article 19 of the AIA, providers of high-risk AI systems must retain logs for a period appropriate to the system's intended purpose, with a minimum retention period of six months.

High-risk AI systems must be designed and developed with capabilities that enable the automatic recording of events ('logs') while the AI system is operating. Logs include, for example, output data, start date, and time, and should be kept for a period that is appropriate to enable the responsible actors to fulfill their obligations (Recitals 71 and 73 of the AIA). Article 19 of the AIA requires that logs be maintained for at least six months. Providers should be aware that logs may qualify as personal data under the GDPR and, when retained, could constitute a 'processing operation,' thus requiring compliance with the GDPR provisions.

The logging capabilities must ensure a level of traceability of the AI system's functioning throughout its lifecycle that is appropriate to the intended purpose of the system. The record-keeping obligations are meant to monitor the AI system for the identification of situations that may result in the AI system presenting a risk,[31] for any substantial modifications and, subsequently, to facilitate the post-market monitoring, as required by Article 72 of the AIA, as well as the monitoring of the system's operation by the deployer by Article 26(5) of the AIA.

Providers are also bound by a duty of cooperation with the competent authorities: Article 21(2) of the AIA requires providers who receive a reasoned request from a national competent authority to give access to the logs to the extent such logs are under the provider's control, stressing that any information obtained by authorities should be treated in compliance with the confidentiality obligations set out in Article 78 of the AIA.

[31] See Step 4.1 on the requirement to have a Risk Management System in place. Part of this RMS is the 'Evaluation of the possibly arising risks.'

4.5. Is the AI system's operation sufficiently transparent? (Article 13 of the AIA)

Transparency and provision of information to deployers Providers should design and develop high-risk AI systems to ensure that their operation is sufficiently transparent and that the system's output is interpretable.	
A. Transparent operation of the AI system - Explainability of the system's decisions (interpretable output)	B. Instructions for use Information provided should be: • concise; • complete; • correct; • clear; • relevant; • accessible; and • comprehensible to the users. The information must include: • identity and contact details of the provider; and • characteristics, capabilities, and limitations of performance of the high-risk AI system.

Transparency is one of the core values of the AIA for high-risk systems. Article 13 of the AIA provides that such systems must be designed and developed to ensure sufficient transparency, enabling both providers and users/deployers to interpret the system's outputs and understand its functioning. Transparency is particularly relevant in the CA process, as compliance with these obligations must be demonstrated before a high-risk AI system is placed on the market.

To verify compliance, the technical documentation (Article 11 of the AIA) and the instructions for use (Article 13(2) of the AIA) should include detailed information that allows deployers to understand and appropriately use the AI system. The CA process must assess whether the transparency obligations outlined in Article 13 of the AIA are met through documentation, testing, and user instructions.

Article 13 of the AIA requires that high-risk AI systems be designed and developed to ensure that their operation is sufficiently transparent and their outputs are interpretable by both providers and users/deployers.[32]

[32] Read also Recital 72 of the AIA.

High-risk AI systems shall be accompanied by instructions for use in an appropriate digital format or otherwise (Article 13(2) of the AIA). All information provided should support informed decision-making by users/deployers and should be concise, complete, correct, clear, relevant, accessible, and comprehensible to the users/deployers. The information should concern the following (Article 13(3) AIA).

1. Identity and contact details of the provider (or authorized representative of the provider).

2. Characteristics, capabilities, and limitations of performance of the high-risk AI system. These instructions should contain, at the very least, the intended purpose of the AI system, the level of accuracy, robustness, and cybersecurity,[33] and any known or foreseeable circumstance/misuse that may lead to risks to the health and safety, fundamental rights, explainability of the AI system, performance of the system as regards the persons or groups of persons on which the system is intended to be used, specifications for the input data, or any other relevant information in terms of the training, validation, and testing data sets use, and where applicable, information to enable deployers to interpret the output of the high-risk AI system and use it appropriately.

3. Changes to the high-risk AI system and its performance, which have been pre-determined by the provider at the moment of the initial conformity assessment, if any[34].

4. Human oversight measures referred to in Article 14, including the technical measures put in place to facilitate the interpretation of the outputs of the high-risk AI systems by the deployers.

5. Computational and hardware resources needed, the expected lifetime of the high-risk AI system, and any necessary maintenance and care measures, including their frequency, to ensure the proper functioning of that AI system, including regarding software updates.

6. Description of the mechanisms included within the high-risk AI system (where relevant) that allow deployers to properly collect, store, and interpret the logs in accordance with Article 12 of the AIA.

The AIA also introduces a 'right to explanation' of decisions taken by a deployer on the basis of the output from a high-risk AI system: Article 86 of the AIA stresses that the user/deployer of the AI system should be able to explain the decisions taken by the AI system in order to satisfy the right to an explanation of individual decision-making when these decisions 'affect that person in a way that they consider to have an adverse impact on their health, safety or fundamental rights.' The level of transparency required must be appropriate to the provider's and deployer's respective obligations under the AIA. Of note, a similar right to explanation is granted by the GDPR at least in the case of solely automated decision-making that has a legal or similarly significant effect on individuals, as decided by the Court of Justice of the EU in 2025.[35]

[33] See also Recital 74 of the AIA.

[34] See Step 2 'When to conduct a CA?,' bullet point 'After an AI system reaches the market.'

[35] See Case C-203/22 Dun & Bradstreet, Judgement of February 27, 2025, interpreting Articles 13(2)(f), 14(2)(g) and 15(1)(h) of the GDPR.

It is important to highlight that when an AI system processes personal data, the transparency obligations under the GDPR[36] and the AIA must be aligned. This is particularly relevant when personal data is processed as part of a solely automated decision-making system that amounts to a high-risk AI system and is subject to enhanced transparency under the GDPR[37]. It is important to note that under the GDPR, transparency is due toward data subjects (an 'identified or identifiable natural person,' per Article 4(1) of the GDPR), while transparency under the AIA is due towards the user/deployer.

This is relevant given that Article 86 of the AIA established a 'right to explanation of individual decision-making' for any affected person subject to a deployer decision on the basis of the output from a high-risk AI system listed in Annex III (with the exception of 'critical infrastructure') which produces legal effects or similarly significantly affects that person in a way that they consider to have an adverse impact on their health, safety or fundamental rights. Affected persons should have the 'right to obtain from the deployer clear and meaningful explanations of the role of the AI system in the decision-making procedure and the main elements of the decision taken.' This provision established another form of transparency, that of the user/deployer towards the affected person, which nonetheless follows from the system's design and compliance with the transparency requirement of Article 13 of the AIA. In other words, the right to an explanation of Article 86 AIA will depend on, inter alia, whether the high-risk AI system has been designed and developed on the basis of the transparency requirements.

[36] Transparency in the GDPR takes the form of a general principle (Article 5(1)(a) of the GDPR 'Personal data shall be processed lawfully, fairly and in a transparent manner in relation to the data subject'), but also the form of legal obligations that fall on the data controller (e.g. Articles 12-14 of the GDPR) and data subject rights (e.g. Articles 15 and 22 of the GDPR).

[37] Article 22 of the GDPR reads: '(1) The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her; [...] (3) [...] the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller , to express his or her point of view and to contest the decision.' For more information, see the FPF Report, 'Automated Decision-Making under the GDPR - A Comprehensive Case-Law Analysis.'

4.6. Is there human oversight of the AI system? (Article 14 of the AIA)

Human Oversight	
Providers should design and develop high-risk AI systems in a way that they can be effectively overseen by natural persons during the period in which the AI system is in use. (Article 14, Recital 73)	
A. Effective human oversight AI systems designed such that the user can effectively oversee their operations and intervene where necessary (e.g., human-machine interface tools).	B. Effective human intervention Built-in operational constraints cannot be overridden by the system itself, and the system is responsive to the human operator. Measures must be implemented by the provider OR identified by the provider and implemented by the user.

• Oversight & intervention while the AI system is in use.
• Natural person that oversees the system:
  ◦ Competent;
  ◦ Trained; and
  ◦ With the authority to oversee and intervene.
• The provider should inform the user, and the user should follow the provider's instructions.	

Article 14 of the AIA requires that high-risk AI systems should be designed and developed in such a way that they can be 'effectively overseen by natural persons during the period in which the AI system is in use.' The goal of human oversight is to prevent or minimize the risks to health, safety, or fundamental rights from the use of a high-risk AI system either 'in accordance with its intended purpose or under conditions of reasonably foreseeable misuse.'

Human oversight under the AIA has two key dimensions that must be demonstrated in a CA:

1. AI system level: the provider should identify appropriate human oversight measures that guarantee that the AI system is subject to built-in operational constraints that cannot be overridden by the system and is responsive to a human operator.[35] Measures that enable human oversight could either be identified and built into the AI system by the provider before it is placed on the market or put into service when technically feasible, or could be first identified by the provider and then implemented by the deployer.

[38] Pursuant to Recital 73 of the AIA.

2. Natural person responsible for the oversight: a natural person responsible for overseeing the system's function. A natural person assigned to oversee an AI system shall have the necessary competence, training, and authority to carry out that role.[39] This dimension raises obligations for both providers and users/deployers. The former shall provide the high-risk AI system to the user/deployer in such a way that enables it to:

- properly understand the capacities and limitations of the system to monitor its operation, including to detect signs of anomalies, dysfunctions, and unexpected performance;
- remain aware of automation bias;
- correctly interpret the system's output;
- decide to override, reverse, or not use the system's output; and
- intervene in the operation of the system or interrupt it through a 'stop' button or a similar procedure.

Under Article 26(1) and (2) of the AIA, the user/deployer of the high-risk AI system also has an obligation to use the system in accordance with the instructions made available by the provider but also to assign human oversight to a person who is competent, properly qualified, and trained, and has the necessary resources in order to ensure the effective supervision of the AI system.

Detailed information, as well as an assessment of the human oversight measures, including an assessment of the technical measures needed to facilitate the interpretation of the outputs of AI systems by the users, in accordance with Articles 13(3)(d) of the AIA, should be part of the technical documentation, which must be included in a CA.[40]

[39] Pursuant to Recital 73 of the AIA.

[40] Annex IV (Technical Documentation), 2(e) of the AIA: the AIA requires the provider to include information on human oversight measures under the 'detailed description of the elements of the AI system and of the process for its development' part of the technical documentation.

4.7. Is the AI system accurate and robust? Are there cybersecurity measures in place? (Article 15 AIA)

Accuracy, Robustness, Cybersecurity
Providers should design and develop high-risk AI systems in a way that they achieve, in light of their intended purpose, an appropriate level of accuracy, robustness, and cybersecurity.
Key aspects • Consistent performance of the AI system throughout its lifecycle. • Accuracy metrics and level of accuracy to be communicated to the user. • Resilience against errors in the system or interaction with the environment (technical redundancy solutions); • Resilience against attempts of unauthorized parties, to alter the system's use, behavior, outputs, or performance by exploiting the system vulnerabilities.
Presumption of conformity An AI system which is certified under a cybersecurity scheme or declaration of conformity may not need to demonstrate conformity with this section separately if the previous certification or statement was issued pursuant to a cybersecurity scheme issued under Article 54(3) of the EU Cybersecurity Act.

The AIA establishes the principles of accuracy, robustness, and cybersecurity as key aspects that high-risk AI systems should observe. Article 15 of the AIA requires these systems to be 'designed and developed in such a way that they achieve an appropriate level[41] of accuracy, robustness, and cybersecurity, and that they perform consistently in those respects throughout their lifecycle.'

The level of accuracy and accuracy metrics should be declared in the accompanying instructions of use.[42]

With regard to robustness, high-risk AI systems should be as resilient as possible regarding errors, faults, or inconsistencies that may occur within the system or the environment in which the system operates. In order to achieve this higher level of resilience, providers should adopt 'technical and organizational measures.' The robustness of high-risk AI systems may be achieved through technical redundancy solutions, 'which may include backup or failsafe plans.' Article 15(4) of the AIA also provides the high-risk AI systems that continue to learn after being placed on the market or put into service should be 'developed to ensure that possibly biased outputs due to outputs used as an input for future operations ('feedback loops') are duly addressed with appropriate mitigation measures.'

The aspects of accuracy and robustness, as considered by the AIA, are particularly important when there is an interface between the AI system and its user (or any other natural person).

[41] Article 15(2) of AIA states that the European Commission should address the technical aspects of how to measure the 'appropriate level of accuracy and robustness' by 'encouraging the development of benchmarks and measurement methodologies' in cooperation with relevant stakeholders.

[42] Article 15(3) and Recital 74 of the AIA, and Section 4.5 of this Guide.

In Article 15(5) of the AIA, the AIA specifies that the AI system should also be resilient against attempts by unauthorized third parties to 'alter their use, outputs, or performance by exploiting the system vulnerabilities'[43]. The technical solutions aimed at ensuring cybersecurity to address AI-specific vulnerabilities should include, 'where appropriate, measures to prevent, detect, respond to, resolve and control for attacks trying to manipulate the training dataset ('data poisoning'), or pre-trained components used in training (model poisoning), inputs designed to cause the AI model to make a mistake ('adversarial examples'), confidentiality attacks, or model flaws.'

Regarding proof of compliance with these requirements, providers may benefit from a presumption of compliance with the cybersecurity requirement, which means that no additional cybersecurity testing or documentation may be required during the CA process. Providers of high-risk AI systems may not need to demonstrate cybersecurity compliance separately if their system has already been certified or for which a statement of conformity has been issued under a recognized cybersecurity scheme[44] under Article 54(3) of the EU Cybersecurity Act (CSA)[45].

For example, the EU Cyber Resilience Act[46] (CRA) from November 2024 sets horizontal cybersecurity requirements for digital products. High-risk AI systems under the AIA, which also fall within the scope of the CRA must comply with the essential cybersecurity requirements set out in the CRA. If they meet these requirements, they are presumed to comply with the cybersecurity obligations under Article 15 of the AIA, provided that the EU declaration of conformity under the CRA covers these aspects[47].

[43] Article 15(5), Recital 76 AIA.

[44] Article 42(2), Recital 122 AIA.

[45] Regulation (EU) 2019/881 of the European Parliament and of the Council of 17 April 2019 on ENISA (the European Union Agency for Cybersecurity) and on information and communications technology cybersecurity certification and repealing Regulation (EU) No 526/2013 (Cybersecurity Act).

[46] Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements and amending Regulations (EU) No 168/2013 and (EU) 2019/1020 and Directive (EU) 2020/1828 (Cyber Resilience Act).

[47] Article 12, Recital 51 CRA.

IV. Standards and presumption of conformity

The AIA emphasizes the role of harmonized standards in ensuring compliance with legal requirements. Recital 121 highlights that standardization should play a key role in providing technical solutions to providers to ensure compliance with the AIA. AI systems that comply with these harmonized standards will be presumed to conform to the legal requirements of the AIA, providing a strong incentive for AI companies to adopt standards to demonstrate compliance.

Article 40 of the AIA establishes that a high-risk AI system or general-purpose AI model is presumed to comply with AIA requirements if it aligns with relevant harmonized standards. Article 41 of the AIA allows the Commission to adopt common specifications when harmonized standards are unavailable or insufficient. Compliance with these specifications also grants a presumption of conformity.

Ongoing standardization efforts

In May 2023, the Commission issued a standardization request to the European Committee for Standardization (CEN) and the European Committee for Electrotechnical Standardization (CENELEC) to develop AI-specific harmonized standards. Draft European AI Standards are expected by April 30, 2025. Once finalized, the Commission will assess their compliance with AIA requirements. If approved, references to these harmonized standards will be published in the Official Journal of the European Union, making them legally recognized.

Standardization in the CA process

The Commission is also working on specific standards related to conformity assessments. These include procedures for CA of AI systems and quality management systems of AI providers. The Commission is also interested in standardizing competency criteria for individuals conducting conformity assessments (both internal and third-party assessments).

These standards will provide operational guidance, ensuring uniformity and reliability in the CA process under the AIA.

Presumption of conformity through regulatory sandboxes

In addition to harmonized standards, the AIA introduces regulatory sandboxes as a means of demonstrating compliance (Article 57 of the AIA). Regulatory sandboxes allow AI providers to develop, test, and validate high-risk AI systems under the supervision of competent authorities before market deployment.[48] AI systems successfully tested in a sandbox benefit from a presumption of conformity when exiting the program. Article 57(7) of the AIA states that written proof and exit reports issued by competent authorities are positively considered by market surveillance authorities and notified bodies and can serve as documentary evidence to streamline conformity assessments.

48 The definition of 'AI regulatory sandbox' is set out in Article 3(55) AIA: 'a controlled framework set up by a competent authority which offers providers or prospective providers of AI systems the possibility to develop, train, validate and test, where appropriate in real-world conditions, an innovative AI system, pursuant to a sandbox plan for a limited time under regulatory supervision.'

V. Key takeaways

1. A CA is the process of demonstrating that a 'high-risk AI system' complies with the requirements enumerated under Chapter III, Section 2 of the AIA. Those requirements include a risk management framework, system data governance, technical documentation, record-keeping, transparency, and provision of information to deployers, human oversight, accuracy, robustness, and cybersecurity.

2. The CA should be understood as a framework of assessments, (technical and non-technical) requirements, and documentation obligations. The provider should assess whether the AI system qualifies as high-risk and assess known or potential risks as part of the risk management system. The provider should additionally ensure that certain requirements are built into the high-risk AI system (e.g., automatic recording of events, human oversight capacity, transparent operation of the AI system) as well as whether documentation obligations (e.g., technical documentation) are met.

3. All requirements should be met before the high-risk AI system enters the market or is put into service (unless otherwise specified). Compliance with the requirements should, however, be ensured throughout the lifecycle of the system and until the AI system's withdrawal. For that, all actors involved in an AI system's supply chain should share information among themselves and should cooperate in a way that ensures compliance with the requirements.

4. Standardization is expected to play a key role in providing technical solutions to providers to ensure compliance with the Regulation. The AIA establishes a presumption of compliance with certain requirements for high-risk AI systems (e.g., cybersecurity requirements, high-quality datasets) as well as in the case where the AI system is developed in the context of a regulatory sandbox.

5. With the AIA's entry into force, an operational timeline for CAs will soon be in place. The infrastructure related to governance and the conformity assessment process should be operational before August 2, 2026. Therefore, the provisions on notified bodies and governance structure should apply from August 2, 2025. You can retrieve FPF's AI Act implementation and compliance timeline here.[49]

[49] FPF AI Act implementation and compliance timeline, https://fpf.org/fpf-resources-on-the-eu-ai-act/.

onetrust

FUTURE OF PRIVACY FORUM | AI

FPF is a global non-profit organization that brings together academics, civil society, government officials, and industry to evaluate the societal, policy, and legal implications of data use, identify the risks, and develop appropriate protections. FPF believes technology and data can benefit society and improve lives if the right laws, policies, and rules are in place. FPF has offices in Washington D.C., Brussels, Singapore, and Tel Aviv. For more information, visit www.fpf.org

OneTrust's mission is to enable the responsible use of data and AI. Our platform simplifies the collection of data with consent and preferences, automates the governance of data with integrated risk management across privacy, security, IT/tech, third-party, and AI risk, and activates the responsible use of data by applying and enforcing data policies across the entire data estate and lifecycle. OneTrust supports seamless collaboration between data teams and risk teams to drive rapid and trusted innovation. Recognized as a market pioneer and leader, OneTrust boasts over 300 patents and serves more than 14,000 customers globally, ranging from industry giants to small businesses. For more information, visit www.onetrust.com.