AI & Partners

Amsterdam - London - Singapore

EU AI Act

Implementation Guide

Contents

1 Introduction

1.1 The Artificial Intelligence Act

1. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (the Artificial Intelligence Act)[1] entered into force on 1 August 2024. The Regulation aims to improve the functioning of the internal market and to promote the uptake of human-centric and trustworthy artificial intelligence, while ensuring a high level of protection of health, safety and the fundamental rights enshrined in the Charter of Fundamental Rights of the European Union, including democracy, the rule of law and environmental protection, against the harmful effects of AI systems in the Union.

2. The Artificial Intelligence Act (AI Act) follows a risk-based approach. It prohibits a limited number of AI practices, lays down requirements and obligations in respect of AI systems classified as high-risk, establishes transparency obligations for certain AI systems irrespective of their risk classification, and sets out a dedicated regime for general-purpose AI models. Chapter III of the AI Act is built upon the EU's New Legislative Framework (NLF) set out in Regulation (EC) No 765/2008 and Decision No 768/2008/EC[2], applying the familiar logic of essential requirements, conformity assessment, CE marking and market surveillance to high-risk AI systems. Chapters II, IV and V, by contrast, do not follow that logic and should not be interpreted by analogy with it.

Figure 1 — The four regulatory tiers and the parallel model layer

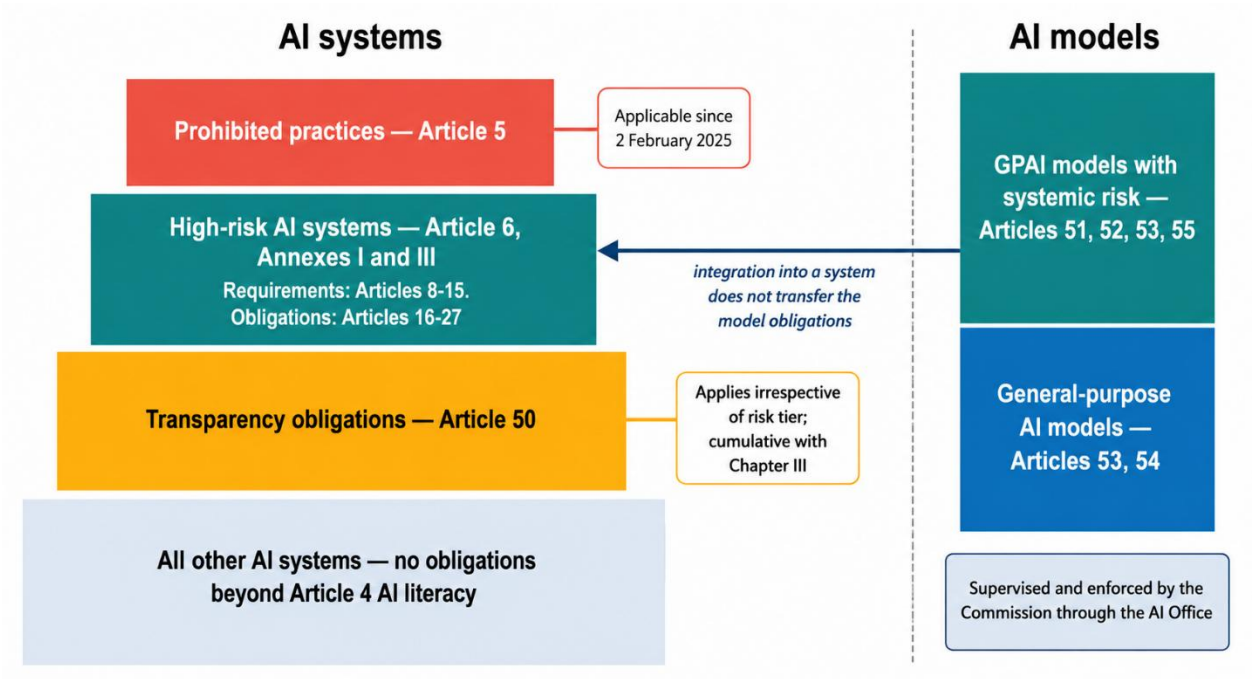

[1] OJ L, 2024/1689, 12.7.2024, ELI: http://data.europa.eu/eli/reg/2024/1689/oj.
[2] Regulation (EC) No 765/2008 of the European Parliament and of the Council of 9 July 2008 setting out the requirements for accreditation and repealing Regulation (EEC) No 339/93, and Decision No 768/2008/EC of the European Parliament and of the Council of 9 July 2008 on a common framework for the marketing of products, and repealing Council Decision 93/465/EEC, OJ L 218, 13.8.2008.

3. Market surveillance and enforcement in respect of AI systems is carried out by national market surveillance authorities, in accordance with Chapter IX of the AI Act and with Regulation (EU) 2019/1020 of the European Parliament and of the Council of 20 June 2019 on market surveillance and compliance of products. In respect of general-purpose AI models, supervision, investigation and enforcement are entrusted exclusively to the Commission, acting through the AI Office. The AI Office, the European Artificial Intelligence Board[3], the advisory forum and the scientific panel of independent experts support the Commission and the Member States in the application of the AI Act.

4. The AI Act applies from 2 August 2026, subject to the staggered dates of application set out in Article 113. Chapters I and II, containing the general provisions and the prohibited AI practices, have applied since 2 February 2025. Chapter III Section 4, Chapter V, Chapter VII and Chapter XII, as well as Article 78, have applied since 2 August 2025, with the exception of Article 101. Article 6(1) and the corresponding obligations, which concern AI systems that are themselves products, or safety components of products, covered by the Union harmonisation legislation listed in Section A of Annex I, apply from 2 August 2027.

Figure 2 — Staggered dates of application

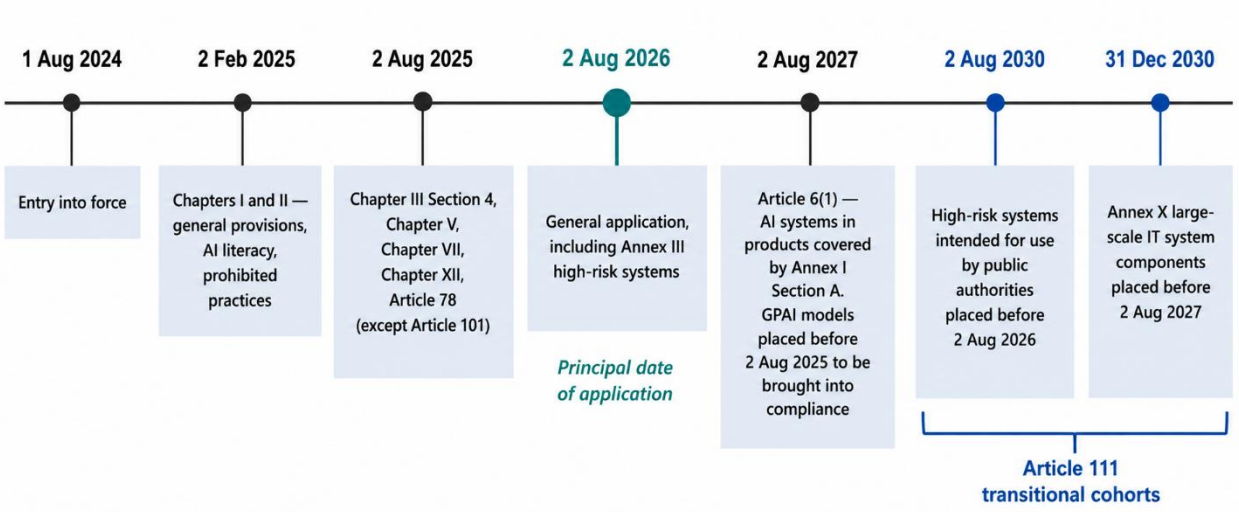

1.2 Purpose of the guidance

5. Article 96(1) of the AI Act requires the Commission to develop guidelines on the practical implementation of the Regulation. That provision sets out the minimum aspects that such guidelines should address. These include: (i) the application of the requirements and obligations referred to in Articles 8 to 15 and in Article 25; (ii) the prohibited practices referred to in Article 5; (iii) the practical implementation of the provisions related to substantial modification; (iv) the practical implementation of the transparency obligations laid down in Article 50; (v) detailed information on the relationship of the AI Act with the Union harmonisation legislation listed in Annex I as well as with other relevant Union law, including as regards consistency in their enforcement; and (vi) the application of the definition of an AI system as set out in Article 3(1).

[3]Established by Article 65 of the AI Act.

6. Article 96(2) requires the Commission, when issuing such guidelines, to pay particular attention to the needs of small and medium-sized enterprises (SMEs), including start-ups, of local public authorities and of the sectors most likely to be affected by the AI Act. This is complemented by Article 62, which requires Member States and the Commission to take a number of specific measures for providers and deployers that are SMEs, including priority access to regulatory sandboxes, dedicated awareness-raising and information activities, and simplified technical documentation.

7. The Commission has already published guidelines on the definition of an artificial intelligence system and guidelines on prohibited artificial intelligence practices[4], as well as guidelines on the scope of the obligations for providers of general-purpose AI models, accompanied by the General-Purpose AI Code of Practice and by the template for the public summary of training content[5]. The present guidance complements those documents; it does not replace them and should be read together with them.

8. This guidance is intended to help operators comply with the AI Act and to support the activities of market surveillance authorities, notifying authorities and notified bodies, with a view to ensuring the harmonised enforcement of the AI Act across the Union. This guidance is not intended to cover the AI Act in its entire scope, but rather to provide clarifications on the rationale of certain key provisions and on how they could be implemented in practice. This guidance concerns the AI Act and is not applicable to other EU laws.

9. Stakeholders were extensively consulted in the preparation of this guidance, including through the advisory forum referred to in Article 67, the European Artificial Intelligence Board, and a public consultation held between 9 February and 20 March 2026.

10. This guidance is not binding on operators or other actors subject to the AI Act. An authoritative interpretation of the AI Act may only be given by the Court of Justice of the European Union. Nevertheless, these guidelines set out the Commission's interpretation of the AI Act, with a view to supporting compliance and contributing to the effective implementation of the Regulation. Furthermore, the guidance contains a number of examples to illustrate the application of specific provisions. However, these examples are not intended to replace a case-by-case assessment, which will always be necessary to account for the specifics of each individual case.

11. In line with Article 96, the Commission may consider issuing further guidance, including guidance targeted at providers subject to the AI Act and to other Union harmonisation legislation or to other related Union legal acts. Such guidance may, for example, address questions on the interplay between the AI Act and Regulation (EU) 2016/679 (the General Data Protection Regulation), Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA), and Regulation (EU) 2024/2847 (the Cyber Resilience Act).

2 Scope

[4]Commission Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689, and Commission Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689, both published in February 2025 and available via the Commission's AI Act implementation website.

[5]Commission Guidelines on the scope of the obligations for providers of general-purpose AI models under the AI Act; General-Purpose AI Code of Practice; template for the public summary of training content, published in July 2025.

2.1 Placing on the market and putting into service

12. Article 2(1) sets out the personal and territorial scope of the AI Act. It applies, among others, to providers placing on the market or putting into service AI systems, or placing general-purpose AI models on the market, in the Union, irrespective of whether those providers are established or located in the Union or in a third country; to deployers of AI systems that have their place of establishment or are located within the Union; to providers and deployers of AI systems that have their place of establishment or are located in a third country, where the output produced by the AI system is used in the Union; and to importers, distributors, product manufacturers and authorised representatives.

13. Article 3(9) defines 'placing on the market' as the first making available of an AI system or a general-purpose AI model on the Union market. Article 3(10) defines 'making available on the market' as the supply of an AI system or a general-purpose AI model for distribution or use on the Union market in the course of a commercial activity, whether in return for payment or free of charge. Article 3(11) defines 'putting into service' as the supply of an AI system for first use directly to the deployer or for own use in the Union for its intended purpose.

14. Unlike most Union harmonisation legislation built on the NLF, the AI Act attaches obligations not only to the placing on the market of an AI system but also to its putting into service. A provider that develops an AI system and uses it exclusively for its own purposes within the Union puts that system into service and is subject to the obligations of a provider, even though the system is never supplied to any third party. In such cases, the same legal person is both provider and deployer, and is subject to both sets of obligations.

    Example 1: A bank develops in-house an AI system intended to be used to evaluate the creditworthiness of natural persons, and uses it exclusively within its own organisation. The bank has put a high-risk AI system into service under its own name. It is therefore the provider of that system for the purposes of the AI Act, and is subject to the obligations laid down in Article 16, as well as to the deployer obligations laid down in Article 26.

15. The latest edition of the 'Blue Guide on the implementation of EU product rules'[6] provides guidance facilitating the understanding of the EU product rules and their uniform application across the different sectoral legal frameworks aligned to the NLF. The Blue Guide indicates that the concepts of 'placing on the market' and of 'making available' are to be understood as referring to each individual product, not to a type of product, and whether it was manufactured as an individual unit or in series (Section 2.3).

[6]Commission notice – The 'Blue Guide' on the implementation of EU product rules 2022 (Text with EEA relevance) 2022/C 247/01, OJ C 247, 29.6.2022, pp. 1–152.

Figure 3 — Placing on the market and putting into service

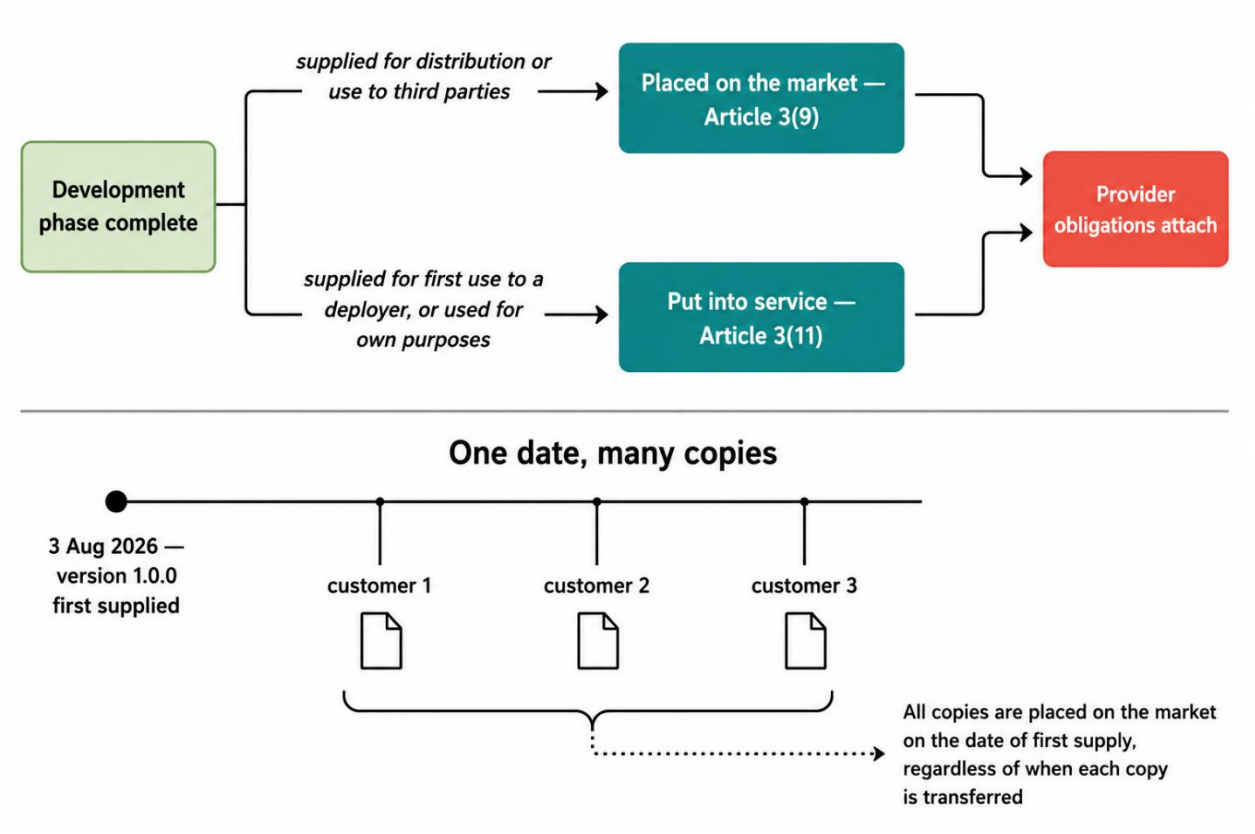

16. However, given the nature of AI systems, which are typically intangible, further guidance is useful on when such systems are considered to be placed on the market. Once the development phase of the AI system is complete and the system is offered to prospective users in the Union, its provider can be regarded as having produced multiple copies of the same AI system and as having supplied them for distribution or use. Unlike tangible products, software is not subject to physical production or stock limitations: each act of making the system available for download or distribution results in a new identical copy being created for the user. As long as that version is not modified in a way that affects compliance with the AI Act, the placing on the market is to be considered to have occurred at the moment of the first offering for distribution or use.

17. Therefore, an AI system should be considered to have been placed on the market when its development phase is complete and it is first supplied for distribution or use on the Union market in the course of a commercial activity. The provider should be considered to have placed on the market multiple copies of the same AI system at the same time, regardless of when possession or use of each individual copy is transferred to another natural or legal person. At the same time, where the provider makes an AI system available in different variants that differ in their included components, configurations or enabled functionalities, those variants cannot be regarded as multiple copies of the same AI system and should be treated as distinct AI systems for the purposes of placing on the market.

Example 2: On 3 August 2026, company A first supplies for distribution via its website version 1.0.0 of its AI system X. On the same day, customer 1 obtains a copy of version 1.0.0 of AI system X. On

17 August 2026, customer 2 obtains a copy of version 1.0.0 of AI system X. Both copies of version 1.0.0 of AI system X are placed on the market on 3 August 2026.

18. Given the iterative nature of software development, it should be further clarified that subsequent iterations of an AI system are considered as newly placed on the market where those iterations qualify as a 'substantial modification' within the meaning of Article 3(23). Iterations that do not qualify as substantial modifications do not require the provider to perform a new conformity assessment procedure under Article 43(4) and therefore do not modify that system's date of placing on the market. For more guidance on the concept of substantial modification, see Section 4.

Example 3: On 3 August 2026, company A first supplies for distribution via its website version 1.0.0 of AI system X. On 17 August 2026, company A issues an updated version 1.0.1 of AI system X that does not constitute a substantial modification. On 30 August 2026, customer 2 obtains a copy of version 1.0.1. Both copies are considered to be placed on the market on 3 August 2026.

19. The AI Act is not confined to AI systems that are supplied to the user for local execution. Where an AI system is made available exclusively through an interface operated by the provider, and is therefore accessed rather than obtained by the user, that supply nonetheless constitutes a making available on the market, since Article 3(10) refers to the supply of an AI system for distribution or use. Providers of AI systems offered as a service are therefore subject to the AI Act on the same footing as providers of systems distributed for installation. Furthermore, in accordance with Article 2(1), point (c), a provider established in a third country falls within the scope of the AI Act where the output produced by its system is used in the Union.

Figure 4 — The seven elements of the definition of an AI system

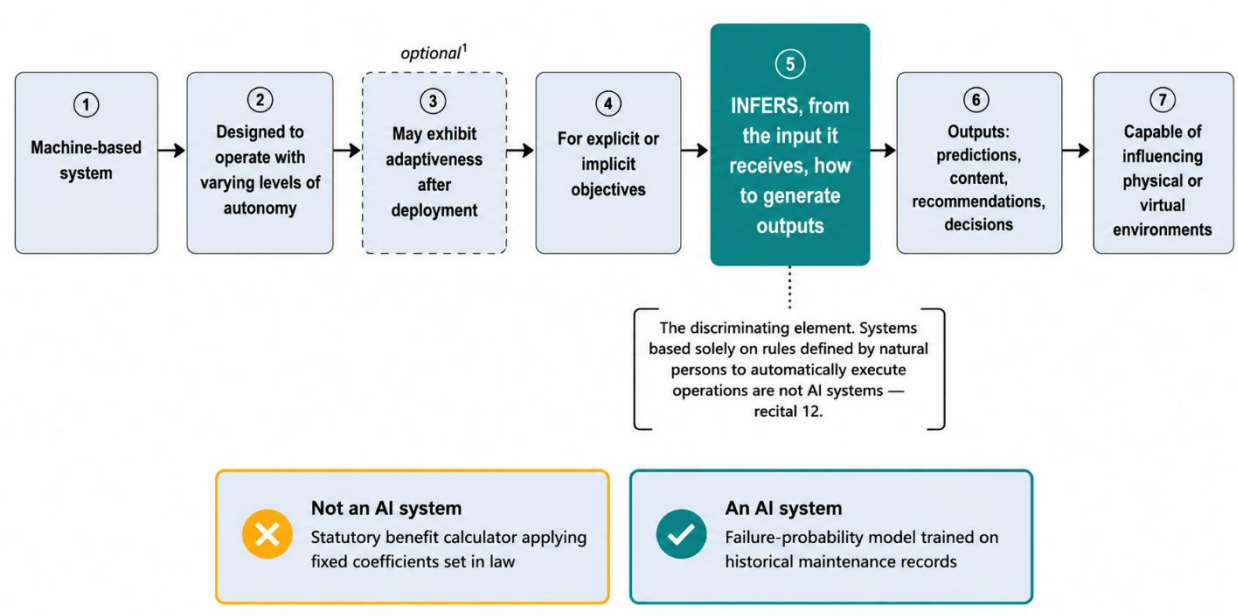

2.2 The notion of an AI system

20. Article 3(1) defines an AI system as 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. That definition is composed of seven elements, not all of which need to be present at all times throughout the lifecycle of the system.

21. The element that principally distinguishes an AI system from simpler traditional software is the capability to infer how to generate outputs. As explained in recital 12, that notion refers to the process of obtaining the outputs and to the capability of AI systems to derive models or algorithms, or both, from inputs or data. Systems that are based solely on rules defined by natural persons to automatically execute operations should not be considered AI systems, even where those rules are numerous, complex, or produce outputs that are difficult for a human to anticipate.

    Example 4: A spreadsheet application that applies a statutory formula fixed in law in order to calculate a social benefit entitlement, where all thresholds and coefficients are set by natural persons, does not infer how to generate its output. It is not an AI system.

    Example 5: A model trained on historical maintenance records to estimate the probability that a given machine will fail within a given period derives its output from patterns learned from data. It infers how to generate its output and is an AI system.

    Example 6: A descriptive analytics tool that computes averages, medians and standard deviations across a dataset and presents them in a dashboard is not an AI system, even where it is described as 'AI-powered' in the provider's promotional materials. The classification of a system depends on its objective technical characteristics and not on the way it is marketed.

22. The reference to 'varying levels of autonomy' indicates some degree of independence of action from human involvement and of capability to operate without human intervention. A system that requires a manual human input in order to generate an output may nevertheless operate with the degree of autonomy required by the definition, provided that it is not fully and exclusively determined by human instruction at every step.

Figure 5 — Models are components; systems are products

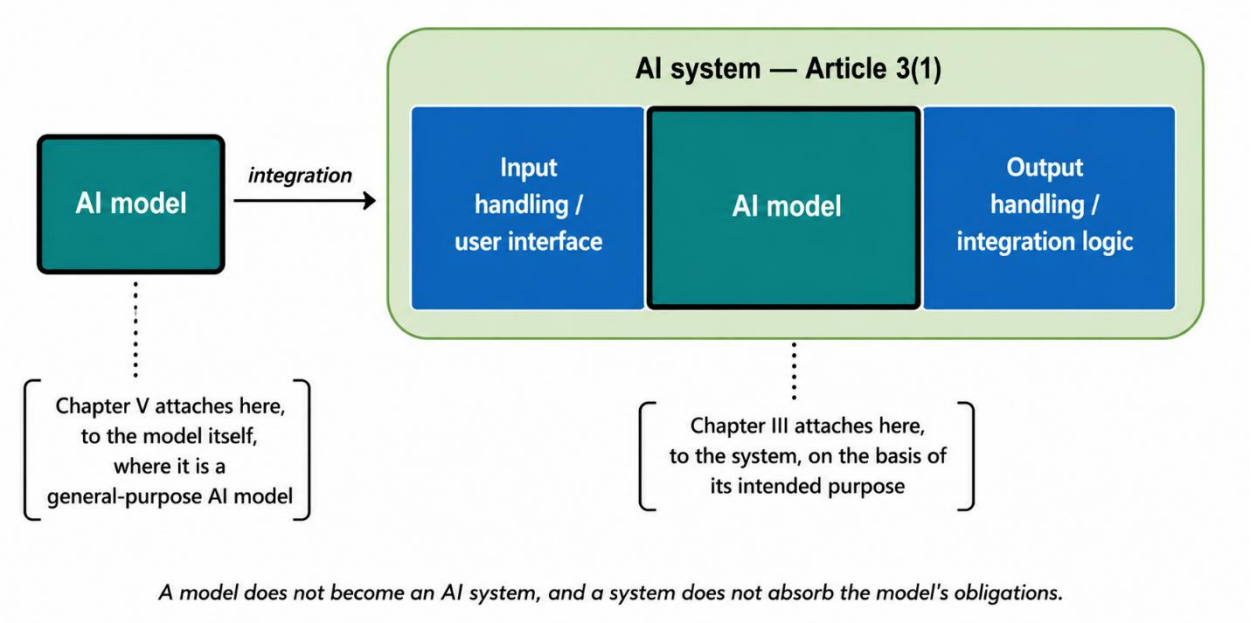

23. The reference to a system that 'may exhibit adaptiveness after deployment' is expressed in permissive terms. Self-learning behaviour, or the capacity of a system to change while in use, is therefore not a necessary condition for a system to qualify as an AI system. Conversely, where such adaptiveness is present, it has consequences for the provider's obligations, in particular under Articles 9, 15 and 43(4), as further discussed in Sections 4.2 and 7.

2.3 Models, computer code and AI components

24. The AI Act distinguishes AI models from AI systems. As explained in recital 97, AI models, including general-purpose AI models, are essential components of AI systems but do not by themselves constitute AI systems. AI models require the addition of further components, such as a user interface, in order to become AI systems. Chapter V accordingly establishes a distinct set of obligations attaching to general-purpose AI models as such, independently of any system into which they may subsequently be integrated.

25. It is also useful to clarify the treatment of computer code. Source code, training scripts, model weights and datasets, where supplied as such and not as an AI system or as a general-purpose AI model, are not themselves regulated by the AI Act. By contrast, a general-purpose AI model is subject to Chapter V irrespective of the way in which it is placed on the market, whether through an application programming interface, as a downloadable set of parameters, through a model repository, or as a component of an AI system supplied by the same provider.

26. In accordance with Article 2(8), the AI Act does not apply to any research, testing or development activity regarding AI systems or AI models prior to their being placed on the market or put into service. That exclusion does not extend to testing in real world conditions[7]. Unfinished code and intermediate model checkpoints shared during the design and development phase, for example for

[7] Testing in real world conditions is governed by Articles 60 and 61, and, for the purposes of AI regulatory sandboxes, by Article 57.

testing or peer review, are therefore not considered to be placed on the market. Similarly, sample or demonstration code provided as part of tutorials or training materials is not considered to be placed on the market for the purposes of the AI Act.

Example 7: A research team trains a model and shares intermediate checkpoints within a closed consortium exclusively for the purposes of evaluation, prior to any decision on release. That activity falls within Article 2(8) and the model is not placed on the market.

Example 8: Company A licenses to company B the source code and weights of a model together with an inference interface, for integration into company B's internal platform. Even if further adaptation and integration by company B is required before the system can be used, company A is placing an AI system, and potentially a general-purpose AI model, on the market. Company A is not responsible for the compliance of company B's subsequent adaptations.

2.4 Combination of hardware and AI systems

27. Whether an AI system forms part of a product should be determined not by how or when the software is delivered to the user, but by whether, in light of the product's intended purpose and reasonably foreseeable use, that software is necessary for the product to perform its intended functions. Where a hardware product is designed to work together with a specific AI system provided by the same provider in order to perform its functions, the hardware and that AI system together constitute the product placed on the market, even where the software is obtained through a separate channel.

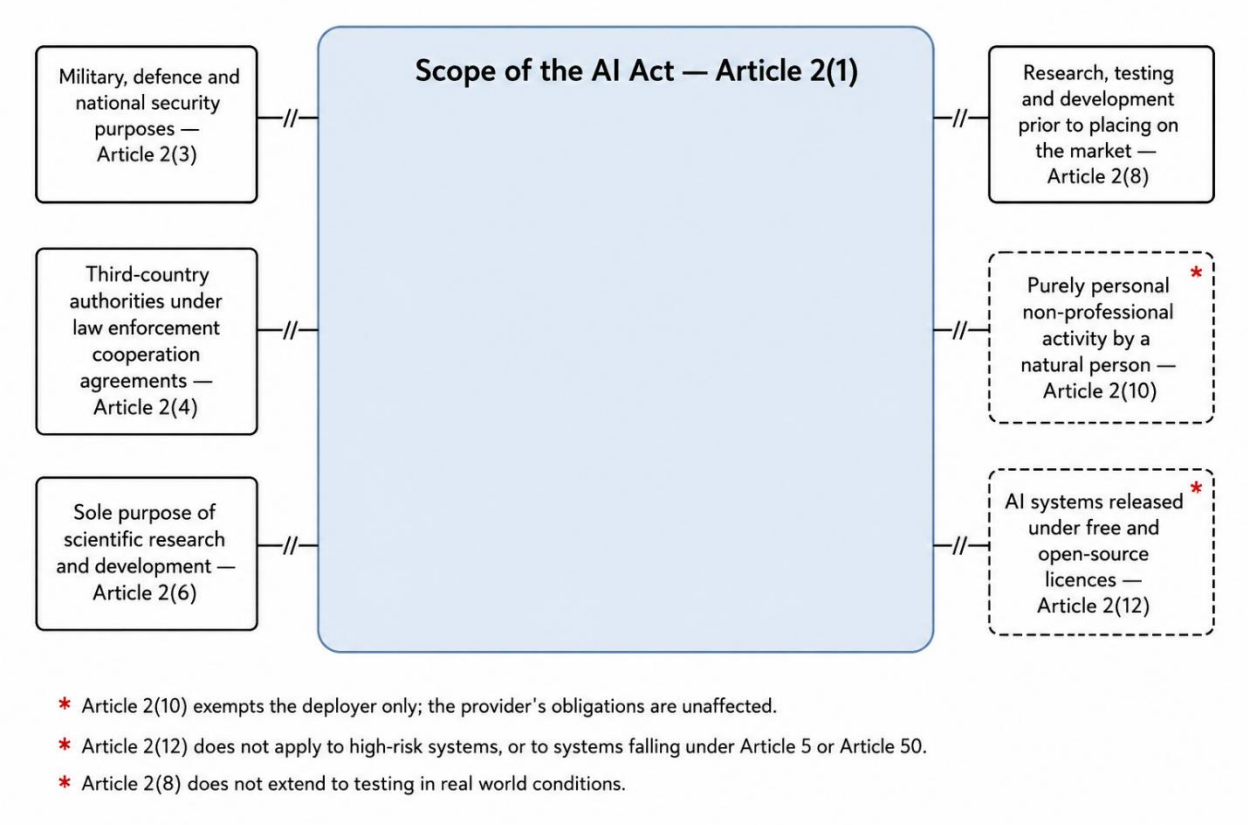

28. Where an AI system is a safety component of a product, or is itself a product, covered by the Union harmonisation legislation listed in Section A of Annex I, and where that product is required to undergo a third-party conformity assessment under that legislation, the AI system is classified as high-risk under Article 6(1). In accordance with Article 25(3), where a product manufacturer places on the market or puts into service a high-risk AI system together with the product under its own name or trademark, that manufacturer is considered to be the provider of the AI system. Article 8(2) allows the provider to integrate the testing, reporting and documentation obligations of the AI Act into the documentation and procedures already existing under that other legislation, in order to avoid duplication.

Example 9: A robotic lawnmower is placed on the market as a product, while the obstacle-recognition functionality that allows it to stop when a person or animal is detected is delivered through an application supplied by the same manufacturer. Although supplied through different channels, the mower and that functionality together constitute a single product, because the mower cannot fulfil its intended purpose safely without it.

Example 10: A surgical robot incorporates an AI-based navigation module that constitutes a safety component. The robot is a medical device subject to third-party conformity assessment under Regulation (EU) 2017/745. The navigation module is therefore a high-risk AI system under Article 6(1). In accordance with Article 43(3), compliance with the requirements of Chapter III, Section 2 is assessed as part of the conformity assessment procedure carried out under Regulation (EU) 2017/745, and a single EU declaration of conformity is drawn up in accordance with Article 47(1).

2.5 Activities excluded from the scope of the AI Act

29. Article 2 excludes a number of activities from the scope of the AI Act. These include AI systems and models placed on the market, put into service or used, with or without modification, exclusively for military, defence or national security purposes (Article 2(3)); public authorities in third countries and international organisations acting under international cooperation agreements for law enforcement and judicial cooperation (Article 2(4)); AI systems and models specifically developed and put into service for the sole purpose of scientific research and development (Article 2(6)); research, testing and development activity prior to placing on the market or putting into service (Article 2(8)); deployers who are natural persons using AI systems in the course of a purely personal non-professional activity (Article 2(10)); and AI systems released under free and open-source licences, subject to the limits set out in Article 2(12) and discussed in Section 3.

30. As exceptions to the general scope of the Regulation, these exclusions are to be interpreted strictly. In particular, the exclusion in Article 2(3) applies by reference to the purpose for which the system is placed on the market, put into service or used, and not by reference to the identity of the actor concerned. Where an AI system is developed for a military purpose but is also placed on the market or put into service for a civilian, humanitarian, law enforcement or public security purpose, the AI Act applies in respect of that placing on the market or putting into service.

Example 11: An individual uses a general-purpose AI chatbot to draft a letter to a friend. That individual is acting in the course of a purely personal non-professional activity and the deployer

obligations of the AI Act do not apply to them. This has no effect on the obligations of the provider of that chatbot, which continue to apply in full.

2.6 Complex systems

31. AI systems within the scope of the AI Act may consist not only of single models or software components, but also of complex systems, including systems composed of multiple hardware and software elements, several models, and orchestration layers that operate together to perform a certain function. Where such a system is placed on the market or put into service as a single AI system, it constitutes an AI system within the meaning of the AI Act.

32. Such complex systems are often characterised by long design and development cycles, with contracts that may have been signed before the AI Act became applicable, by extended operational lifetimes, and by a high degree of technical and organisational complexity. Such systems may rely on components placed on the market before the AI Act entered into application, on established architectures, or on widely used interoperability standards, including standards referred to in other Union legislation or in sector-specific frameworks. As a result, certain technical characteristics of those systems may be difficult or disproportionate to modify without affecting their intended purpose, safety, reliability or interoperability with existing infrastructure.

33. These characteristics do not in themselves exclude complex systems from the scope of the AI Act. Rather, they illustrate the application of the risk-based approach that underpins the Regulation. In accordance with Article 9(3), the risks to be addressed through the risk management system are only those which may be reasonably mitigated or eliminated through the development or design of the high-risk AI system, or through the provision of adequate technical information. Those characteristics form part of the system's intended purpose and operating context and are therefore relevant when assessing compliance with the requirements set out in Chapter III, Section 2.

34. Accordingly, where the intended purpose of the system requires it to interact with existing dependencies or to follow certain interoperability requirements, and where this constrains the implementation of a given requirement, providers should identify and document those specific constraints, assess the associated risks, and implement appropriate alternative or compensatory risk-mitigation measures, so as not to undermine the performance or safety of the system. For these purposes, both the technical documentation referred to in Article 11 and Annex IV and the instructions for use referred to in Article 13 play a key role in transparently describing the identified constraints, the associated risks and the mitigation measures implemented. In accordance with the obligation under Article 9(2) to maintain the risk management system as a continuous iterative process throughout the entire lifecycle of the system, providers should periodically reassess whether such constraints continue to exist, and update the system where they can be lifted or reduced over time.

Example 12: A provider places on the market a high-risk AI system that must exchange data with an existing rail signalling infrastructure using a legacy message format that does not support authenticated exchanges. On the basis of the risk management system, the provider determines that authenticated exchanges would be necessary to ensure the integrity of the data on which the

system relies. Where the intended purpose requires interoperability with the existing infrastructure, the provider may implement the legacy format, provided that the associated risks are identified and mitigated through other means, such as plausibility checks on incoming data, logging of anomalous inputs and defined fail-safe states. Where it is technically feasible for the system to support both the legacy and an authenticated format, the provider is expected to implement the latter and to enable its use by default.

Example 13: A provider of an industrial automation system procures a vision component before the AI Act enters into application. The provider carries out the risk management process for the automation system, which is to be placed on the market after that date. On the basis of that process, the provider does not identify risks associated with that component that cannot be addressed at the level of the system. The provider integrates the component and ensures that the system as a whole meets the requirements of Chapter III, Section 2.

2.7 AI systems placed on the market before the AI Act entered into application

35. Article 111 lays down a set of transitional rules. In accordance with Article 111(2), the AI Act applies to operators of high-risk AI systems, other than those referred to in Article 111(1), that have been placed on the market or put into service before 2 August 2026 only if, as from that date, those systems are subject to significant changes in their designs. In any event, providers and deployers of high-risk AI systems intended to be used by public authorities are required to take the necessary steps to comply by 2 August 2030. In accordance with Article 111(3), providers of general-purpose AI models that were placed on the market before 2 August 2025 are required to take the necessary steps to comply by 2 August 2027. In accordance with Article 111(1), AI systems that are components of the large-scale IT systems established by the legal acts listed in Annex X and that were placed on the market or put into service before 2 August 2027 are to be brought into compliance by 31 December 2030.

36. The notion of a 'significant change in the design' of a system used in Article 111(2) is distinct from, although closely related to, the notion of 'substantial modification' defined in Article 3(23). A change should be regarded as significant where it alters the design of the system in a manner that affects its compliance with the requirements set out in Chapter III, Section 2, or that results in a modification of the intended purpose. By contrast, routine maintenance, correction of defects, security patching, changes confined to the user interface, and retraining on data of the same nature and distribution within parameters already determined by the provider do not, in themselves, amount to significant changes in the design of the system.

37. In some cases, a provider will place on the market, after 2 August 2026, units of an AI system developed in accordance with a design finalised before that date. In such cases, compliance with the AI Act does not necessarily require that the system be redesigned. The provider is required to establish, implement, document and maintain the risk management system referred to in Article 9 in order to determine, on the basis of the intended purpose and reasonably foreseeable misuse of the system, which risks are relevant and how the requirements of Chapter III, Section 2 are implemented.

38. Where the outcome of that process demonstrates that the system already incorporates appropriate and effective measures addressing the relevant risks, the provider may rely on those existing measures to demonstrate compliance. The AI Act does not in itself impose an obligation to introduce new features or to redesign the system where this is not necessary to address the identified risks.

39. Nevertheless, the provider remains subject to all the obligations laid down in the AI Act. These include ensuring, before the system is placed on the market or put into service, that the applicable conformity assessment procedure has been carried out, that the technical documentation has been drawn up, that the EU declaration of conformity has been drawn up, that the CE marking has been affixed, and that the system has been registered in the EU database in accordance with Article 49. Compliance with those obligations is independent of whether the design of the system required modification as a result of the risk management process.

40. Where it is not possible for the provider to demonstrate how risk considerations were taken into account during the original design and development phase, the obligations of Article 9 should be understood as requiring the provider to carry out a current risk analysis and to demonstrate on that basis that the system incorporates adequate measures with a view to eliminating or reducing risks as far as technically feasible. Providers are not expected to recreate historical design or test documentation where this would not contribute to the performance or safety of the system. Where tests are nonetheless necessary, providers are not expected to provide evidence of tests carried out on all variants, but may group such tests across families of AI systems, as further discussed in Section 7.4.

Example 14: A provider places on the market, after 2 August 2026, an AI system for the evaluation of applications to educational institutions that was designed and validated in 2024. Before placing new instances on the market, the provider establishes the risk management system required by Article 9 and identifies the relevant risks to fundamental rights, in particular the risk of discriminatory outcomes across protected groups. The outcome of that analysis shows that the existing design, including its data governance measures and its bias testing regime, already addresses the identified risks. The provider concludes that no redesign is necessary. It documents a current risk analysis explaining how the existing design mitigates the identified risks, completes the conformity assessment procedure under Annex VI, draws up the EU declaration of conformity, affixes the CE marking and registers the system in the EU database.

3 Free and open-source AI

41. Article 2(12) provides that the AI Act does not apply to AI systems released under free and open-source licences, unless they are placed on the market or put into service as high-risk AI systems or as AI systems falling under Article 5 or Article 50. A partial and separately drafted exemption applies to general-purpose AI models under Article 53(2), as discussed in Section 3.3.1.

42. The structure of the exemption in the AI Act differs from that found in other Union legislation addressing free and open-source software. It is expressed as an exclusion from scope rather than

as a question of whether a supply takes place in the course of a commercial activity. Nevertheless, recital 103 makes clear that free and open-source AI components that are provided against a price or otherwise monetised should not benefit from the exemption. Monetisation is described as including the provision of technical support or other services in relation to the AI component, the provision of a software platform through which the provider monetises other services, and the use of personal data for reasons other than exclusively for improving the security, compatibility or interoperability of the software, with the exception of transactions between microenterprises[8].

43. As with other Union legislation in this field, the mere circumstances under which a free and open-source AI component has been developed, or the way in which its development has been financed, should not be taken into account when determining whether it is monetised. The distinction between the development phase and the supply phase should be preserved.

44. It is therefore useful to (i) clarify what free and open-source AI is; (ii) clarify when it is deemed to fall under the responsibility of a natural or legal person; (iii) provide examples to help stakeholders understand when a free and open-source AI component is monetised; and (iv) set out the limits of the exemption. It should be noted at the outset that, unlike Regulation (EU) 2024/2847, the AI Act does not establish a category of actor comparable to the 'open-source software steward'. A natural or legal person that publishes a free and open-source AI component without monetising it, and outside the categories listed in Article 2(12), bears no obligations under the AI Act in respect of that component.

45. Recital 102 describes free and open-source AI by reference to software and data, including models, released under a licence that allows them to be openly shared and where users can freely access, use, modify and redistribute them or modified versions thereof. It follows that only software that cumulatively fulfils two conditions should be regarded as free and open-source for the purposes of the AI Act: (i) it must be made available under a licence granting the full set of rights referred to in recital 102; and (ii) it must be openly shared, that is, made publicly available rather than provided on a restricted or conditional basis. Software distributed under a free and open-source licence but whose source code, or, in the case of a model, whose parameters, are shared only with paying customers or with a limited group of users is not to be considered free and open-source AI for the purposes of the AI Act.

[8] It should also be recalled that, while AI components are often monetised through the processing of personal data, the protection of personal data is a fundamental right and personal data cannot therefore be considered a commodity, as indicated in recital 24 of Directive (EU) 2019/770, OJ L 136, 22.5.2019, pp. 1–27.

Figure 7 — Decision sequence for free and open-source AI

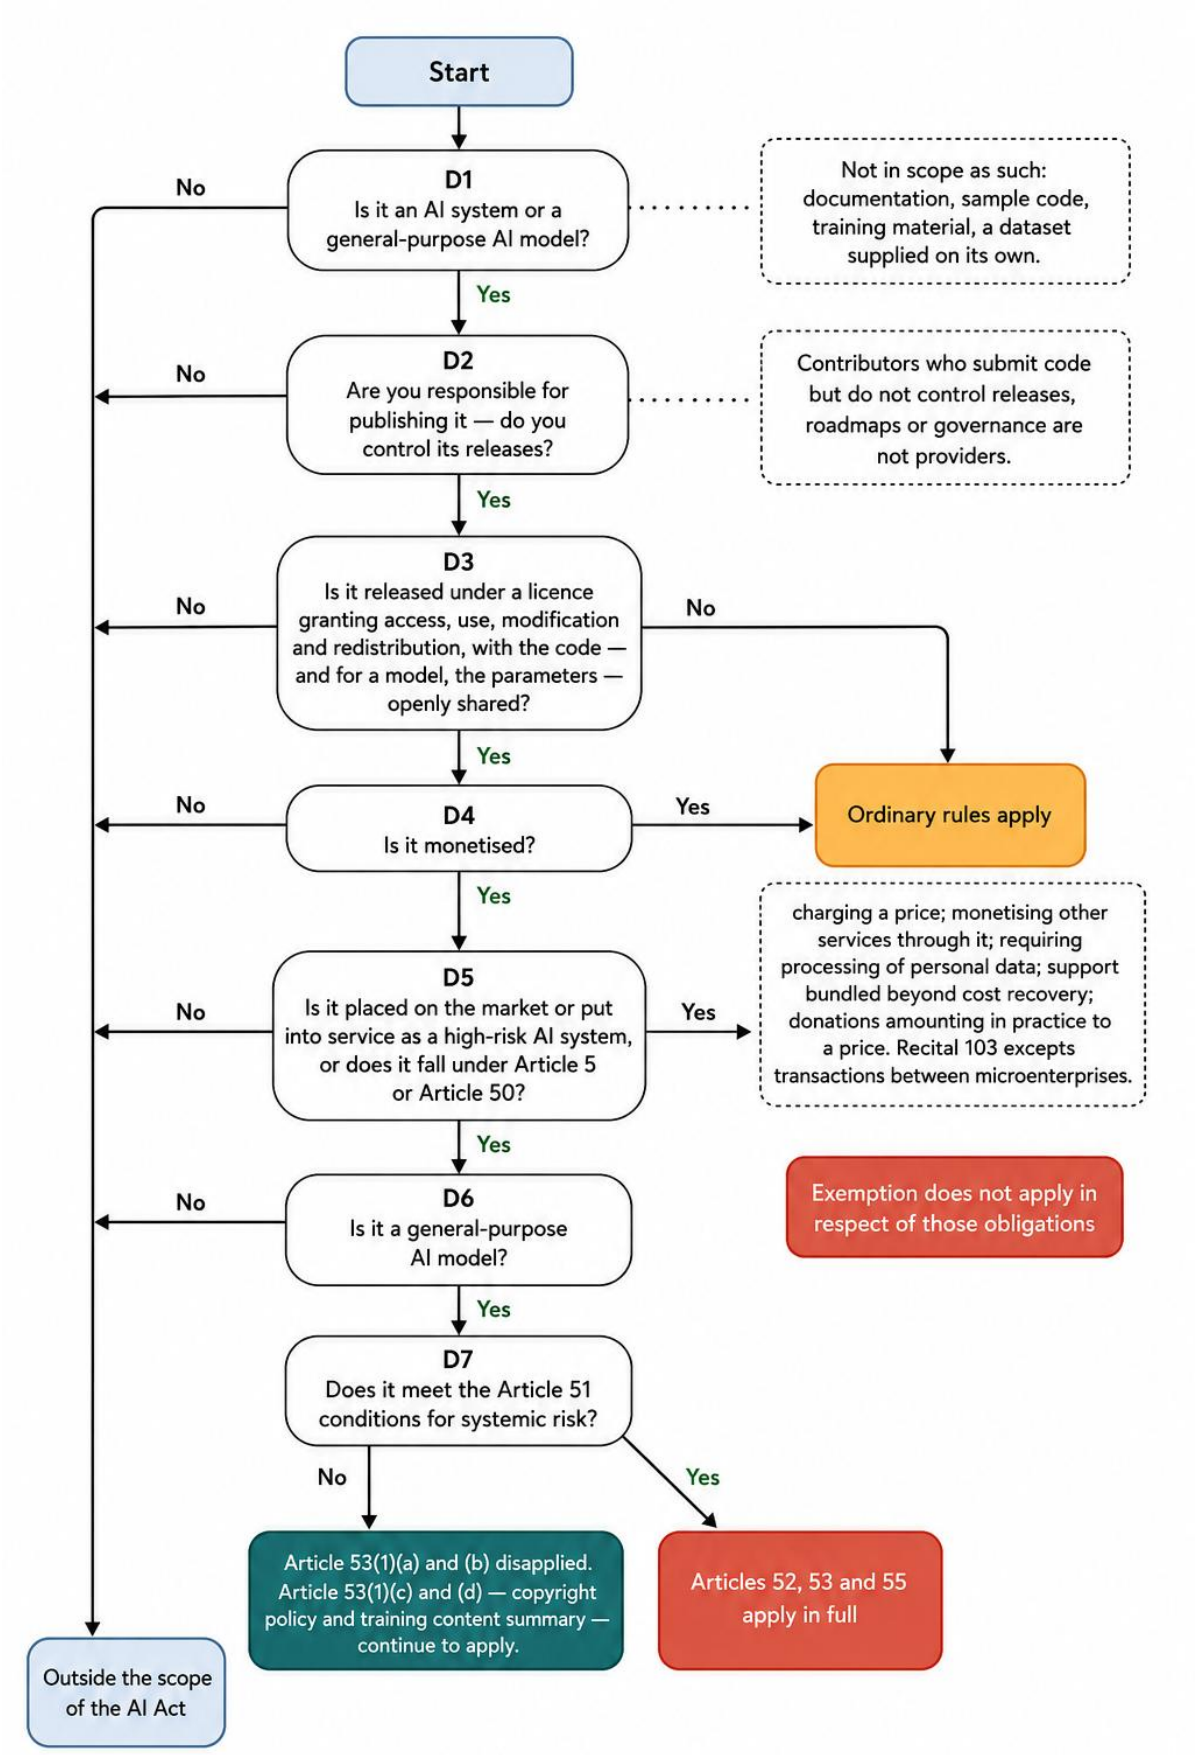

3.1 Determining if free and open-source AI is under one's responsibility

46. The AI Act places its key obligations on providers. Article 3(3) defines a 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.

47. To determine whether a natural or legal person is a provider in respect of a free and open-source AI component, it is first necessary to establish whether that person is actually supplying it. Given the specificities of the development of free and open-source software, which often involves multiple contributors, decentralised collaboration models and a separation between contribution and decision-making, it is useful to clarify that a component is considered to be under the responsibility of the natural or legal persons who publish it and who exercise primary control over its development, releases and distribution decisions, often referred to as 'maintainers'. Persons who contribute source code, training data or evaluation results but who do not control releases, roadmaps or governance decisions are 'contributors'; in such cases the component is not under their responsibility, even though they have contributed to it. The mere existence of technical permissions, such as commit access, is not sufficient to establish responsibility.

Example 15: An individual developer or an employee of a company submits a pull request containing a bug fix or a new evaluation harness to a free and open-source AI project. The project's maintainers review, accept and merge the contribution and subsequently include it in a new release. The person who submitted the pull request is a contributor and is not the provider of the resulting component.

3.2 Determining if free and open-source AI is monetised

48. Once it is established that a free and open-source AI component is under the responsibility of a natural or legal person, it must be established whether that person monetises it, in which case the exemption does not apply.

3.2.1 Charging a price

49. Where the natural or legal person that supplies a free and open-source AI component charges a price for the component itself, for example by charging for access to model weights or to a packaged distribution, that component is monetised and the exemption does not apply. That person is therefore a provider under the AI Act in respect of that component.

50. Providers frequently supply, alongside a paid offering, a version of the same component free of charge, often referred to as a 'community' version, whose codebase or model parameters are identical or nearly identical. Those are nevertheless distinct offerings: the paid version is monetised and therefore does not benefit from the exemption, whereas the version provided free of charge and without any of the other forms of monetisation described in this Section does benefit from it. This is also the case where the paid version is an enhanced commercial version that extends the free version, or that incorporates it into a broader offering, as in the case of the 'open-core' model.

Example 16: A legal person publishes a free and open-source retrieval framework under a licence granting the full set of rights referred to in recital 102, and separately offers an enterprise edition of the same framework under a commercial licence, with additional connectors and management features. The community edition benefits from the exemption in Article 2(12), subject to the limits set out in Section 3.3. The enterprise edition is monetised and does not benefit from it.

3.2.2 Monetisation of other services or requiring the processing of personal data

51. A natural or legal person publishing a free and open-source AI component may also be monetising it where the component is provided through a platform by means of which that person monetises other products or services, or where the use of the component is conditional on the processing of personal data for reasons other than exclusively for improving the security, compatibility or interoperability of the software.

Example 17: A natural or legal person publishes a free and open-source agent framework. The framework is available free of charge but is configured to route model calls, by default, through a paid inference service operated by the same person, from which that person derives revenue. The framework allows its publisher to monetise other services offered through it and is therefore monetised.

Example 18: A natural or legal person publishes a free and open-source personal assistant application. The application is available free of charge, but its use is conditional upon the processing of users' personal data for purposes such as targeted advertising or analytics unrelated to the security, compatibility or interoperability of the software. By requiring such data processing as a condition for use, the application is monetised.

3.2.3 Support services

52. The mere fact that a natural or legal person publishing a free and open-source AI component also offers paid support services related to it does not, as such, mean that the component is monetised. Support services may include activities such as consultancy, training or professional services related to the use, documentation, configuration and deployment of the component by a customer.

53. The decisive factor is whether access to the component itself, including its maintenance, is conditioned on remuneration, rather than the mere offering of professional services around a freely available component. Where the component can be obtained and used freely, and users can optionally choose to purchase professional services separately, the exemption is not lost.

54. By contrast, where access to a specific version of the component, including certain benefits such as technical assistance, guaranteed response times or performance optimisation, is conditioned on remuneration, that constitutes monetisation. This includes cases where a paid edition or enterprise version is made available under a commercial agreement, irrespective of whether functionally equivalent software is also available free of charge under a free and open-source licence.

Example 19: A natural or legal person publishes a free and open-source model-serving stack and separately offers optional consultancy services to train users and to support them in deploying it. Anyone can obtain and use the stack without purchasing those services. The stack is not monetised.

Example 20: A natural or legal person publishes a free and open-source model-serving stack and offers a paid edition of that stack that includes technical assistance and performance optimisation. The paid edition is monetised.

55. Particularly in the case of natural persons publishing free and open-source AI, offering technical assistance bundled with access to the component would still not amount to monetisation if the price charged serves only to recuperate actual costs. Such actual costs include a variety of costs related to the component's design, development and maintenance, including the person's reasonable living expenses. A natural person publishing a free and open-source AI component and offering technical assistance to cover their costs and obtain fair remuneration is not to be considered, on that basis alone, as monetising it.

56. A natural or legal person offering technical assistance in relation to a free and open-source AI component that is not under its responsibility is not, on that basis, a provider of that component, unless that person substantially modifies it and places it on the market or puts it into service under its own name or trademark.

3.2.4 Donations

57. Natural or legal persons publishing free and open-source AI projects routinely include ways for users to voluntarily donate money, both to thank the project's publishers and to ensure that the project remains actively maintained. Accepting donations without the intention of making a profit should not be considered monetisation. The mere fact of including a link to a donation platform or similar tools should not be viewed as an intention to make a profit, even where the amount collected exceeds the costs associated with the design, development and provision of the component. Such costs include reasonable compensation for contributors hired by a legal person and a natural person's reasonable living expenses. Donations, by their very nature, fluctuate over time, and a degree of flexibility is therefore to be exercised in this assessment. A free and open-source AI component supported only through donations is unlikely to be considered monetised.

Example 21: A natural or legal person publishes an AI component in a public online repository, allowing anyone to obtain, use, modify and redistribute it under a free and open-source licence. The publisher invites users to make voluntary donations to support the project's continued development. Access to the component, its source code, its parameters and its updates is not conditional on making a donation. The component is not monetised.

58. Nonetheless, there are instances where a component supported through donations is to be regarded as monetised. This is the case where, on an overall assessment of the circumstances, the donations are de facto equivalent to charging a price for access to the component or to certain of its functionalities. This may be the case in particular where access to the component, to essential functionalities or to updates is conditioned in practice on making a donation, or where donations are associated with contractual benefits or exclusive advantages that go beyond community perks.

Example 22: A natural or legal person publishes an AI component under a free and open-source licence, but provides the trained parameters and subsequent updates only to users who make a donation. Users who do not make a donation do not have access to a usable version of the component. The donations are de facto a condition for access and therefore amount to charging a price. The component is monetised.

3.2.5 Financing of free and open-source AI

59. The mere fact that a third party has paid for, sponsored or otherwise financed the development of a free and open-source AI component does not in itself determine whether that component is monetised. It is common practice for individual developers, foundations or communities to receive funding from companies, public bodies or other sponsors in order to work on specific features, fix defects, run evaluations or maintain critical components. That funding may take many forms, including grants, sponsorships, compute credits, service contracts or paid development work.

Figure 8 — What survives the Article 53(2) exemption

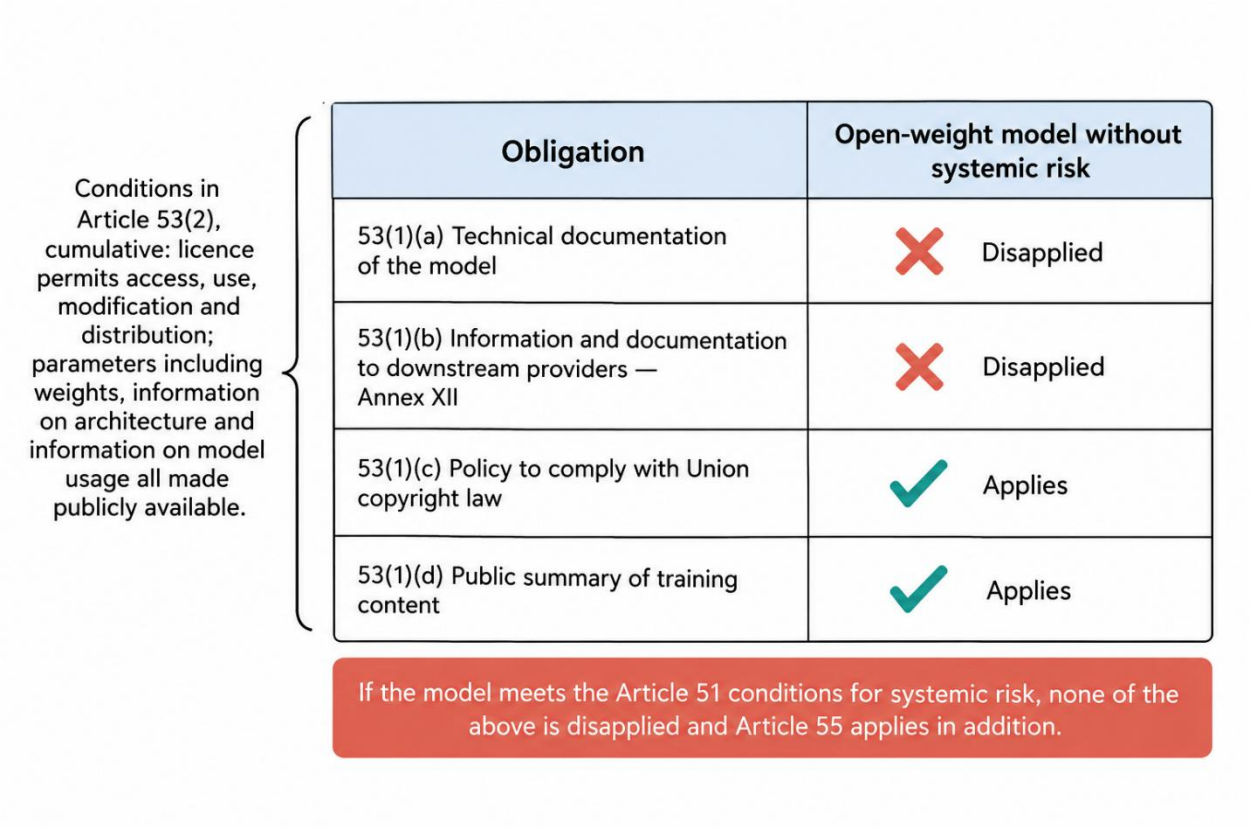

60. Where the component is openly shared and made freely available for all to access, use, modify and redistribute, the fact that a commercial entity may have paid for its development should not contribute to determining whether it is monetised. Where the entity that funded the development, or any other entity, integrates the component into a high-risk AI system, that entity is required to comply with the AI Act in respect of that system and, where applicable, with Article 25(4).

Example 23: An individual developer publishes an AI component and actively maintains it. Provider A requests that the developer add a specific capability and funds that development. The developer adds the capability, openly sharing it and making it freely available for all to access, use, modify and redistribute. The developer does not monetise the component. If provider A integrates the component into a high-risk AI system, provider A is responsible for the compliance of that system as a whole.

3.2.6 Not-for-profit entities set up to achieve not-for-profit objectives

61. Where a legal person publishing a free and open-source AI component is a not-for-profit organisation set up in such a way that all earnings after costs are used to achieve not-for-profit objectives, the component should not be regarded as monetised for the purposes of the exemption, even where the organisation receives revenue in connection with it. This reflects the absence of an intention to make a profit. The limits set out in Section 3.3 continue to apply.

3.2.7 Integration by other providers

62. In some cases, a free and open-source AI component is published by a clearly identifiable natural or legal person but is intended for integration by other providers into their own AI systems. In such cases, the component benefits from the exemption unless it is monetised by the person that publishes it, or unless it falls within one of the categories listed in Article 2(12).

63. The mere fact that providers integrate free and open-source components into their own AI systems has no effect on the status of those components under the AI Act. Maintainers of components that benefit from the exemption do not bear obligations in relation to other entities that may integrate those components into their systems. Whether the AI Act applies to a given component depends on whether the person that publishes it monetises it and on whether the component is placed on the market or put into service as a high-risk AI system, or as a system falling under Article 5 or Article 50.

Example 24: A legal person publishes, without monetising it, a free and open-source library for building conversational interfaces, and separately integrates that library into one of its own AI systems, which it places on the market. The library benefits from the exemption in Article 2(12). The AI system into which it is integrated does not benefit from that exemption and is subject to the AI Act in accordance with its classification.

3.3 Limits of the free and open-source exemption

64. The exemption in Article 2(12) is subject to three express limits. It does not apply where the AI system is placed on the market or put into service (i) as a high-risk AI system; (ii) as an AI system falling under Article 5, that is, an AI system whose placing on the market, putting into service or use constitutes a prohibited practice; or (iii) as an AI system falling under Article 50, that is, a system subject to the transparency obligations laid down in that Article.

65. The third limit is of considerable practical importance. Article 50 applies, among others, to AI systems intended to interact directly with natural persons and to AI systems, including general-

purpose AI systems, that generate synthetic audio, image, video or text content. A very large proportion of freely available generative AI systems therefore falls within Article 50, with the consequence that the transparency obligations of that Article apply to their providers notwithstanding the free and open-source character of the system. The exemption is lost only in respect of the obligations that give rise to it: a free and open-source system falling under Article 50 but not classified as high-risk is subject to Article 50 and not to Chapter III.

Example 25: A natural person publishes, under a free and open-source licence and without any form of monetisation, a system that generates synthetic images from text prompts. The system is not high-risk and does not implement any prohibited practice. It nonetheless falls under Article 50(2), and the publisher is required to ensure that its outputs are marked in a machine-readable format and detectable as artificially generated, in accordance with that provision.

Example 26: A legal person publishes, under a free and open-source licence and without monetisation, a system intended to be used for the screening and filtering of job applications. That intended purpose falls within point 4(a) of Annex III. The system is therefore placed on the market as a high-risk AI system, the exemption in Article 2(12) does not apply, and the publisher is subject to the obligations of a provider of a high-risk AI system.

Example 27: A legal person publishes, under a free and open-source licence, a general-purpose model and a reference implementation, and states in its documentation that the system is not intended for any use listed in Annex III. The system is not placed on the market as a high-risk AI system. Should a downstream actor modify its intended purpose so that it becomes a high-risk AI system, that actor becomes the provider of the resulting high-risk system in accordance with Article 25(1), point (c).

3.3.1 Free and open-source general-purpose AI models

66. A separate regime applies to general-purpose AI models. Article 53(2) provides that the obligations laid down in Article 53(1), points (a) and (b) — namely the obligation to draw up and keep up to date the technical documentation of the model, and the obligation to draw up and make available information and documentation to downstream providers — do not apply to providers of AI models released under a free and open-source licence that allows for the access, usage, modification and distribution of the model, and whose parameters, including the weights, the information on the model architecture and the information on model usage, are made publicly available. That exemption does not apply to general-purpose AI models with systemic risk.

67. Two consequences follow. First, the exemption in Article 53(2) is narrower than that in Article 2(12): it relieves the provider of two specific documentation obligations only. The obligations laid down in Article 53(1), points (c) and (d) — to put in place a policy to comply with Union law on copyright and related rights, in particular to identify and comply with a reservation of rights expressed pursuant to Article 4(3) of Directive (EU) 2019/790, and to draw up and make publicly available a sufficiently detailed summary about the content used for training, according to the template provided by the AI Office — continue to apply in full. Second, the conditions are cumulative and

specific: it is not sufficient that the licence be a free and open-source licence; the parameters, the information on the architecture and the information on model usage must all be publicly available[9].

68. Where a general-purpose AI model is classified as a model with systemic risk under Article 51, the exemption in Article 53(2) does not apply, and the provider is subject to the obligations laid down in Article 53(1) in their entirety as well as to those laid down in Article 55, irrespective of the licence under which the model is released.

Example 28: A legal person releases a general-purpose AI model under a licence permitting access, use, modification and distribution, and publishes the weights, a description of the architecture and documentation on model usage. The model does not meet the conditions for classification as a model with systemic risk. The provider is exempt from the obligations in Article 53(1), points (a) and (b). It remains subject to the copyright policy obligation and to the obligation to publish a sufficiently detailed summary of the training content.

Example 29: A legal person releases the weights of a general-purpose AI model that was trained using a cumulative amount of computation exceeding the threshold referred to in Article 51(2). The model is presumed to be a general-purpose AI model with systemic risk. The exemption in Article 53(2) does not apply. The provider is required to notify the Commission in accordance with Article 52, to comply with Article 53(1) in full, and to comply with Article 55, including as regards model evaluation, adversarial testing, the assessment and mitigation of systemic risks, the reporting of serious incidents and the ensuring of an adequate level of cybersecurity protection.

3.4 Contributors and downstream uses

69. Providers of AI systems that integrate free and open-source components into their own systems do not become responsible for those components' individual compliance with the AI Act, even where they contribute source code, data or evaluation results to their maintenance. They are, however, required to comply with the AI Act in respect of their own systems.

70. In accordance with Article 25(4), the provider of a high-risk AI system and the third party that supplies an AI system, tools, services, components or processes that are used or integrated into that high-risk AI system are required to specify, by written agreement, the information, capabilities, technical access and other assistance, based on the generally acknowledged state of the art, that are necessary to enable the provider of the high-risk AI system to comply fully with its obligations. That obligation does not apply to third parties making accessible to the public tools, services, processes or components, other than general-purpose AI models, under a free and open-source licence. The Commission is to develop and recommend voluntary model contractual terms for that purpose[10].

71. It is important to note that the carve-out in Article 25(4) expressly excludes general-purpose AI models. The obligation to conclude a written agreement therefore applies in respect of general-purpose AI models even where those models are made accessible to the public under a free and open-source licence. In practice, for open-weight models supplied without any contractual

[9] Recital 104 further clarifies that the exception should not apply where the model is monetised, and recital 103 excludes from the benefit of the free and open-source exceptions components provided against a price or otherwise monetised.
[10] In accordance with the third subparagraph of Article 25(4), the AI Office may develop and recommend voluntary model terms for contracts between providers of high-risk AI systems and third parties supplying tools, services, components or processes.

relationship, the provider of the high-risk AI system will need to assess whether the publicly available documentation, including that published pursuant to Article 53(1), point (b) and Annex XII where applicable, is sufficient to enable it to comply with its own obligations, and, where it is not, to select a different model or to generate the necessary evidence itself.

3.5 Illustrative scenarios

72. The scenarios set out below are entirely hypothetical and are only intended to illustrate the different cases explained in the preceding sections.

Example 30: Individual developer A has developed a free and open-source AI component and publishes it under their own name. Developer A does not charge a price for its use. The component is openly shared and freely available for all to access, use, modify and redistribute. Developer A includes a link to a platform to collect voluntary donations.

Companies B, C and D integrate that component into their own AI systems, some of which are high-risk. To support ongoing maintenance, companies B, C and D make voluntary donations to developer A.

The component is not monetised and benefits from the exemption in Article 2(12). Developer A has no obligations under the AI Act in respect of that component. Companies B, C and D are responsible for the compliance of their own systems, including as regards the risk management system under Article 9 and the accuracy, robustness and cybersecurity requirements of Article 15.

Example 31: Not-for-profit foundation F publishes a free and open-source model and evaluation suite intended for integration into other systems. Foundation F is set up in such a way that all earnings after costs are used to achieve not-for-profit objectives. Maintenance is financed through public research grants and membership fees, and developers employed by member organisations contribute code and compute.

The component is not monetised and benefits from the exemption in Article 2(12). Foundation F's members do not become providers of the component by reason of their contributions. Where a member integrates the component into a high-risk AI system, that member is the provider of that system.

Example 32: Company A publishes a free and open-source AI component under its own name and offers a paid edition of that component that includes technical assistance and performance optimisation. Developers from companies B, C and D contribute to the component's maintenance, but it remains under the control of company A.

The paid edition is monetised and does not benefit from the exemption. Companies B, C and D are not providers of that component. If they integrate it into their own high-risk AI systems, they are responsible for the compliance of those systems.

Example 33: An individual developer publishes a free and open-source AI library on a public package repository and actively maintains it. In the package documentation, the developer adds a link to

collect donations. A provider obtains that library from the repository free of charge and integrates it into its own high-risk AI system.

The individual developer and the package repository have no obligations under the AI Act in respect of that library. The provider that integrates it is responsible for the compliance of its high-risk AI system as a whole. The exception in Article 25(4) applies, so no written agreement is required with the developer, but the provider must nonetheless be in a position to demonstrate that the integrated library does not compromise compliance with Chapter III, Section 2.

Question	Consequence
Is the component an AI system or a general-purpose AI model within the meaning of Article 3?	If not (for example documentation, sample code, training material, a dataset supplied as such), the AI Act does not apply to it.
Is the component under your responsibility, in the sense that you publish it and control its releases?	If not, you are a contributor and not a provider in respect of that component.
Is the component released under a licence granting the rights described in recital 102, with the code and, for models, the parameters openly shared?	If not, the exemption in Article 2(12) or Article 53(2) does not apply.
Is the component monetised (price, monetisation of other services through it, conditional processing of personal data, bundled support beyond cost recovery, donations amounting to a price)?	If yes, the exemption does not apply and the ordinary rules on providers apply.
Is the system placed on the market or put into service as a high-risk AI system, or does it fall under Article 5 or Article 50?	If yes, the exemption in Article 2(12) does not apply in respect of the corresponding obligations.
For a general-purpose AI model: does it meet the conditions in Article 51 for classification as a model with systemic risk?	If yes, the exemption in Article 53(2) does not apply and Articles 52, 53 and 55 apply in full.

Table 1: Stylised decision sequence for AI Act coverage of free and open-source AI

4 Substantial modifications and changes along the value chain

73. Article 3(23) defines a '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 which results in a modification to the intended purpose for which the AI system has been assessed.

74. The notion of substantial modification is relevant for ascertaining obligations under the AI Act in a number of cases, including the following:

(a) in accordance with Article 25(1), point (b), a distributor, importer, deployer or other third party that makes a substantial modification to a high-risk AI system already placed on the market or put into service, in such a way that it remains a high-risk AI system, is considered to be a provider of that system;

(b) in accordance with Article 25(1), point (c), any such actor that modifies the intended purpose of an AI system, including a general-purpose AI system, that 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 system becomes a high-risk AI system, is considered to be a provider of that system;

(c) in accordance with Article 43(4), a high-risk AI system that has already been subject to a conformity assessment procedure is required to undergo a new conformity assessment procedure in the event of a substantial modification, regardless of whether the modified system is intended to be further distributed or continues to be used by the current deployer;

(d) in accordance with Article 111(2), the AI Act applies to high-risk AI systems placed on the market or put into service before 2 August 2026 only where those systems are subject, as from that date, to significant changes in their designs.

75. It is therefore important to provide guidance to help operators understand when changes to an AI system qualify as substantial modifications.

4.1 Modifications that do not amount to substantial modifications

76. Not every change to an AI system after its placing on the market or putting into service is a substantial modification. A change qualifies only where it is not foreseen or planned in the initial conformity assessment and where, as a result, either compliance with the requirements of Chapter III, Section 2 is affected, or the intended purpose is modified. Changes that improve the performance or the safety of the system without altering its intended purpose, and that do not affect compliance with those requirements, are therefore not substantial modifications.

77. By way of illustration, the following will generally not constitute substantial modifications: the correction of defects; security updates addressing an identified vulnerability; changes confined to the user interface that do not alter the information provided under Article 13; refactoring or performance optimisation that does not alter model behaviour in a manner relevant to the requirements; and changes to parts of a product that are not the AI system and that do not interact with it.

Example 34: A provider deploys an update to a high-risk AI system correcting an input validation defect that could cause the system to process malformed records incorrectly. The update modifies the internal implementation without affecting the intended purpose and without introducing new risks. It is not a substantial modification. The provider nonetheless updates the technical documentation in accordance with Article 11(1).

Example 35: A provider retrains a high-risk AI system on a refreshed dataset drawn from the same population and of the same nature and distribution as the original training data, within the retraining cadence and the performance thresholds described in the technical documentation. The retraining was foreseen in the initial conformity assessment and does not affect compliance. It is not a substantial modification.

4.2 Pre-determined changes and continuously learning systems

78. The second subparagraph of Article 43(4) addresses AI systems that continue to learn after being placed on the market or put into service. It provides that changes to the high-risk AI system and to its performance that have been pre-determined by the provider at the moment of the initial conformity assessment, and that are part of the information contained in the technical documentation referred to in point 2(f) of Annex IV, do not constitute a substantial modification.

Figure 10 — The pre-determined change envelope

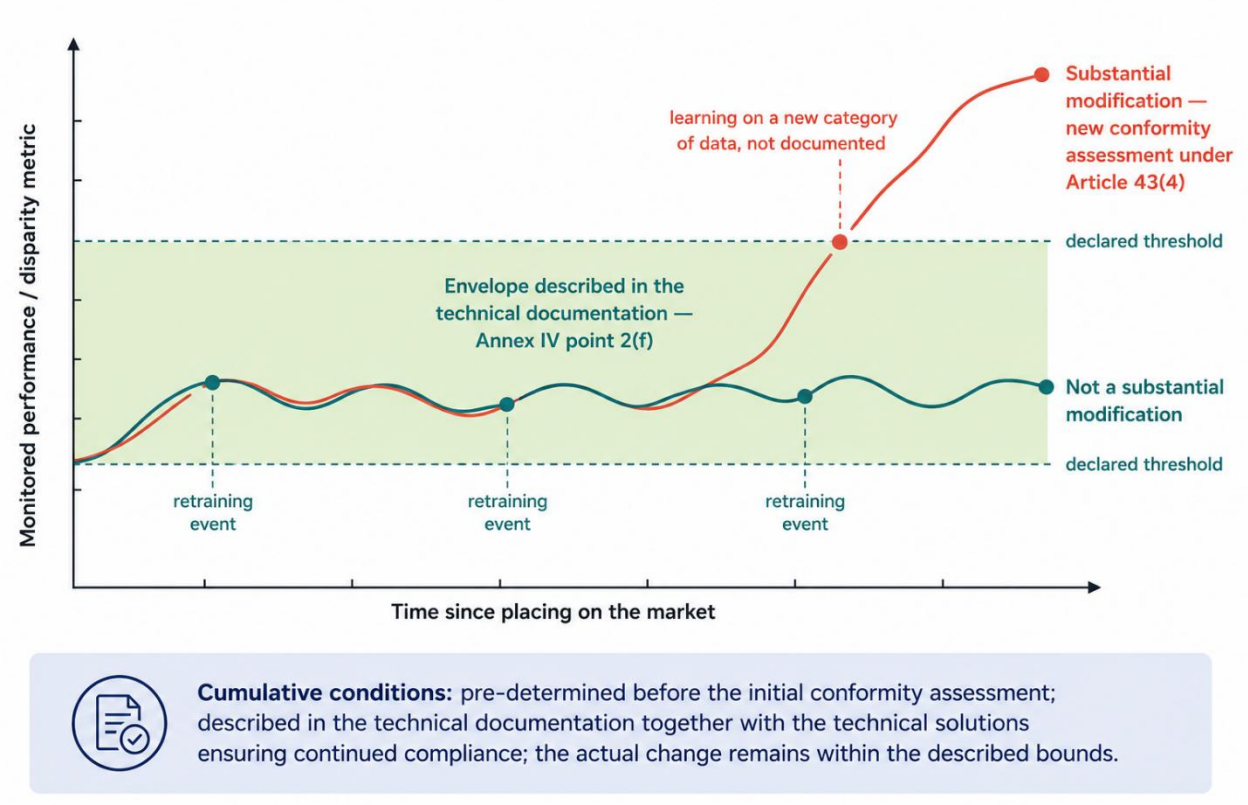

79. For a change to benefit from that provision, three conditions must be met cumulatively. The change must have been pre-determined by the provider before the initial conformity assessment; the pre-determined change must be described in the technical documentation, including the technical solutions adopted to ensure continuous compliance; and the actual change must remain within the envelope so described. It is not sufficient for a provider to state in general terms that the system will continue to learn. The documentation should describe the nature and the bounds of the permitted changes, the data on which learning may take place, the performance and fairness metrics that will be monitored, the thresholds beyond which the system will be halted or rolled back, and the human oversight measures applicable to the learning process.

Example 36: A provider places on the market a high-risk AI system that updates its parameters continuously on the basis of incoming operational data. The technical documentation describes the update mechanism, the categories of data on which learning occurs, the monitored performance and disparity metrics, the thresholds triggering automatic reversion to the last validated state, and the human review applicable to each promotion of a candidate state. Changes remaining within that envelope are not substantial modifications.

Example 37: The same provider subsequently enables the system to learn from a new category of data collected by deployers, which was not described in the technical documentation and which alters the population on which the system's outputs are calibrated. That change falls outside the pre-determined envelope, affects compliance with Articles 10 and 15, and is therefore a substantial modification.

4.3 Changes that qualify as substantial modifications

80. Where a provider introduces changes that result in a modification of the intended purpose of the system, it is likely that those changes were not considered in the initial conformity assessment. In such circumstances, the change will generally qualify as a substantial modification. This includes the extension of the system to a new population, a new modality of input, a new deployment context, or a new category of decision.

Example 38: A provider places on the market a clinical triage support system validated on and intended for an adult population. The provider subsequently extends the system to paediatric patients. The intended purpose has changed and the underlying data governance and accuracy analyses no longer cover the population concerned. The system has been substantially modified.

Example 39: A provider places on the market a system intended to organise and classify documents. It subsequently introduces an update enabling the system to determine, without human intervention, which documents are rejected and which are escalated. The intended purpose shifts from an information management tool to a system exercising decision-making functions that were not envisaged in the conformity assessment. The system has been substantially modified.

81. Conversely, even limited or seemingly minor changes may affect compliance with the requirements of Chapter III, Section 2. The assessment should therefore not be based on the scale or complexity of the change, but on its effect on compliance and on the intended purpose.

Example 40: A provider modifies a high-risk AI system so that operators may enter free-text instructions where previously only a fixed set of structured inputs was accepted. Although the change is limited in scope, it introduces new pathways by which the system's behaviour can be influenced, including through inputs designed to manipulate it, which were not considered in the risk management system or in the measures adopted under Article 15(5). The system has been substantially modified.

Example 41: A provider replaces the general-purpose AI model underlying a high-risk AI system with a different model. Whether this constitutes a substantial modification depends on the effect of the change. Where the new model alters the accuracy, robustness or bias characteristics of the system beyond the ranges declared in the instructions for use pursuant to Article 13(3), point (b)(ii), or introduces new capabilities that broaden the intended purpose, the change is a substantial modification. Where the new model performs equivalently against the validation and testing regime described in the technical documentation and the intended purpose is unchanged, it may not be.

82. When performing this assessment, providers may consider, in a non-exhaustive manner, whether the change:

(a) alters the intended purpose, the intended users, the deployment context or the population on which the system operates;

(b) affects the data governance measures required by Article 10, for example by introducing data of a different nature, provenance or distribution;

(c) affects the accuracy, robustness or cybersecurity characteristics declared under Article 15, including by introducing new attack surfaces or new external dependencies;

(d) affects the human oversight measures designed under Article 14, for example by removing an intervention point or by increasing the degree of automation;

(e) affects the information provided to deployers under Article 13, in a manner that would require that information to be revised.

83. Whether a change qualifies as a substantial modification is relevant when determining whether certain obligations apply, as set out above. However, irrespective of that qualification, providers remain responsible for ensuring the continued compliance of their systems, for maintaining the risk management system as a continuous iterative process throughout the entire lifecycle in accordance with Article 9(2), and for keeping the technical documentation accurate, complete and up to date in accordance with Article 11(1).

4.4 Consequences of a substantial modification

84. The consequences of a substantial modification differ depending on whether it is carried out by the initial provider or by another actor. Where the modification is carried out by an actor other than the initial provider, Article 25 applies. Where it is carried out by the initial provider, the consequences follow from Article 43(4) and from the generally established concept of substantial modification elaborated in Section 2.1 of the Blue Guide.

85. The generally established consequence is that, where a modification qualifies as a substantial modification and the modified system is made available on the market or put into service, it is to be treated as a new system for the purposes of the Regulation. As a result, the act of making the substantially modified system available constitutes a new placing on the market. This applies whether the substantial modification is carried out by the initial provider or by another person.

4.4.1 Substantial modifications carried out by a person other than the initial provider

86. In accordance with Article 25(1), a distributor, importer, deployer or other third party is to be considered a provider of a high-risk AI system, and is subject to the obligations of the provider under Article 16, in three situations: where it puts its name or trademark on a high-risk AI system already placed on the market or put into service, without prejudice to contractual arrangements stipulating that the obligations are otherwise allocated; where it makes a substantial modification to a high-risk AI system already placed on the market or put into service in such a way that it remains a high-risk AI system; and where it modifies the intended purpose of an AI system, including a general-purpose AI system, that 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 system concerned becomes a high-risk AI system.

Figure 9 — Is the change a substantial modification?

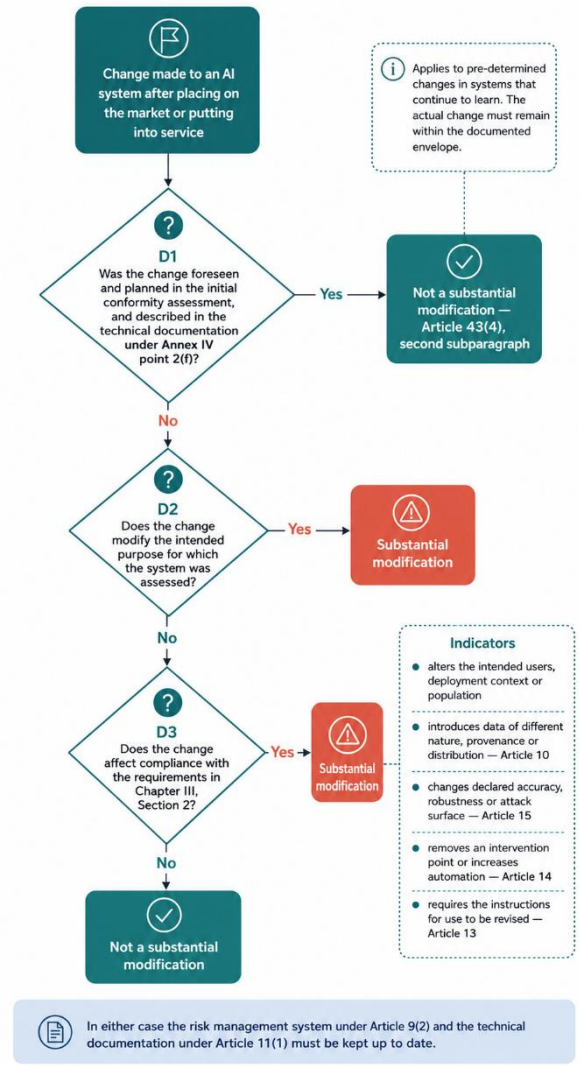

87. In those circumstances, Article 25(2) provides that the provider that initially placed the system on the market or put it into service is no longer to be considered a provider of that specific AI system. That initial provider is nonetheless required to closely cooperate with the new provider, to make available the necessary information and to provide the reasonably expected technical access and other assistance required for the fulfilment of the obligations set out in the AI Act, in particular as regards compliance with the conformity assessment of high-risk AI systems. That obligation does not apply 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 obligation to hand over the documentation.

88. It should be noted that the AI Act does not contain a provision limiting the new provider's obligations to the part of the system affected by the modification. The person who becomes a provider under Article 25(1) assumes the obligations of a provider in respect of the high-risk AI system as a whole. However, in practice, that person may rely on the information and technical access made available under Article 25(2), and, consistently with Section 2.1 of the Blue Guide, it is not necessary to repeat tests or to produce new documentation in relation to aspects that are not affected by the modification. The new provider remains responsible for the conformity of the modified system and for drawing up the EU declaration of conformity, even where it relies on existing tests and documentation.

Example 42: A deployer obtains a high-risk AI system intended to be used for the evaluation of learning outcomes, and fine-tunes it on its own historical assessment data in a manner that materially changes the system's behaviour and was not foreseen in the initial conformity assessment. The deployer then puts the modified system into service. The deployer has made a substantial modification and becomes the provider of the modified system under Article 25(1), point (b). The initial provider is required to cooperate and to make available the necessary documentation under Article 25(2).

Example 43: A deployer of the same system adjusts operating thresholds within the ranges expressly described in the instructions for use and foreseen in the initial conformity assessment. That is not a substantial modification. The deployer remains a deployer and is subject to Article 26, including the obligation to use the system in accordance with the instructions for use.

Example 44: A deployer integrates a general-purpose AI system, supplied by a third party and not classified as high-risk, into an internal process for the sifting of job applications, and puts the resulting system into service. By modifying the intended purpose in such a way that the system becomes a high-risk AI system within point 4(a) of Annex III, that deployer becomes the provider of a high-risk AI system under Article 25(1), point (c), and is subject to Article 16 in respect of that system.

89. A case of substantial modification should not be confused with a case of integration, in which a person assembles components, including where appropriate by modifying them, into a new AI system that they themselves place on the market or put into service. In such cases the integrator is not modifying a system already placed on the market by another provider; it is placing a new system on the market. The integrator is the provider of that new system and is required to comply

with the AI Act in its entirety in respect of the system as a whole. Where the integrator uses components that are themselves subject to the AI Act, it may rely on the compliance activities of the providers of those components, and in particular on the information supplied under Article 53(1), point (b), to facilitate its own compliance.

Example 45: A company obtains a third-party general-purpose AI model under licence, develops its own retrieval layer, guardrails and user interface, and places the resulting system on the market under its own name as a system intended to be used for the assessment of eligibility for public assistance benefits. The company is not substantially modifying a system already placed on the market; it is placing a new high-risk AI system on the market. It is the provider of that system and is required to comply with the AI Act in respect of that system as a whole, including through the written agreement referred to in Article 25(4) with the model provider.

4.4.2 Substantial modifications carried out by the initial provider

90. Where the substantial modification is carried out by the initial provider, including in the context of iterative development, that provider remains the provider for the purposes of the AI Act. However, the substantially modified system is to be considered as newly placed on the market and is required, under Article 43(4), to undergo a new conformity assessment procedure.

91. In accordance with Section 2.1 of the Blue Guide, the provider may reuse existing documentation and test results for aspects of the system that are not affected by the substantial modification. Particularly where the provider places substantially modified versions of the same system on the market, the conformity assessment procedure should focus on the modified parts. Similarly, where a third-party conformity assessment is performed, the notified body should focus its assessment on the modified parts and may rely on existing documentation and test results for the unchanged parts.

92. Following a substantial modification, the provider is also required to update the technical documentation in accordance with Article 11(1), to draw up a new EU declaration of conformity in accordance with Article 47, and to update the entry for the system in the EU database in accordance with Article 49(1) and Annex VIII. Where the modification reveals that the system is not in conformity, the provider is required to take the corrective actions referred to in Article 20 and to inform the relevant authorities.

Figure 11 — Transfer of the provider role under Article 25

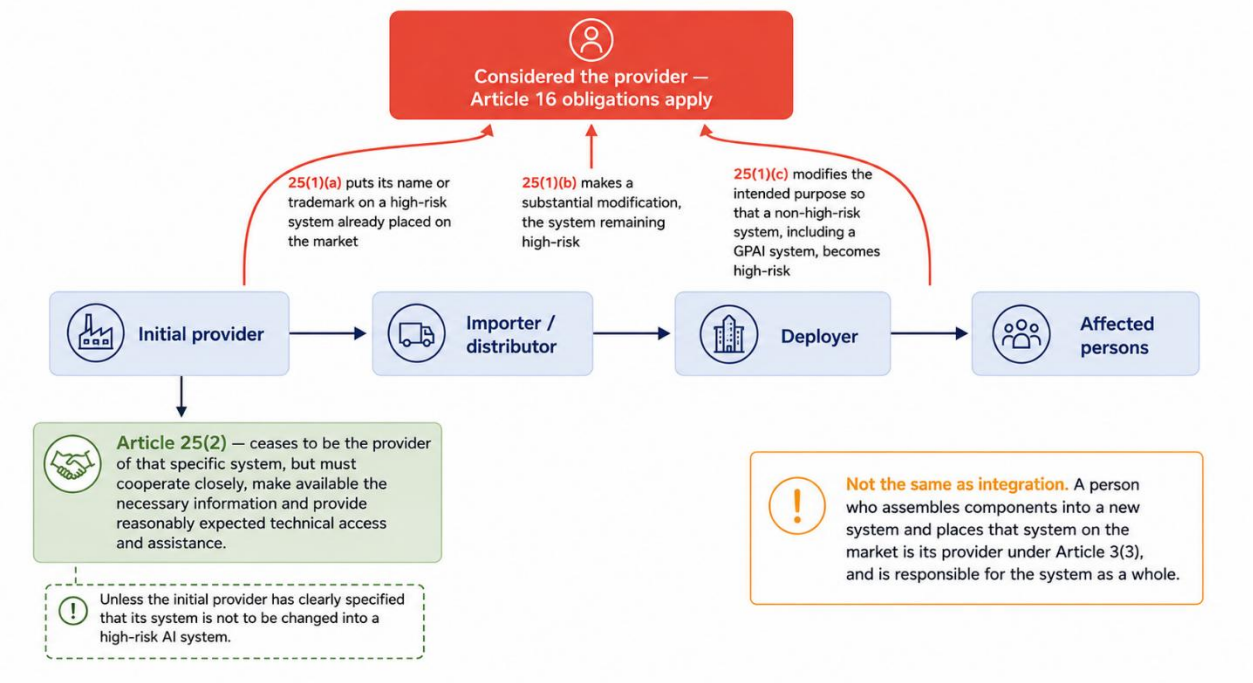

5 Lifecycle obligations

93. Unlike certain other instruments of Union law, the AI Act does not define a fixed 'support period' during which the provider is required to maintain a high-risk AI system. Instead, it establishes a set of obligations that attach to the provider throughout the lifecycle of the system and that are anchored in several distinct provisions. It is useful to set out how those obligations relate to one another.

94. Article 9(2) requires the risk management system to be understood as a continuous iterative process planned and run throughout the entire lifecycle of the high-risk AI system, requiring regular systematic review and updating. Article 15(1) requires high-risk AI systems to perform consistently as regards accuracy, robustness and cybersecurity throughout their lifecycle. Article 17 requires providers to put in place a quality management system covering, among other matters, procedures for post-market monitoring and for the reporting of serious incidents. Article 72 requires providers to establish and document a post-market monitoring system in a manner that is proportionate to the nature of the AI technologies and to the risks of the high-risk AI system, based on a post-market monitoring plan.

95. In addition, the AI Act specifies a number of concrete retention and availability periods. Article 18(1) requires providers to keep at the disposal of the national competent authorities, for a period ending 10 years after the AI system has been placed on the market or put into service, the technical documentation, the documentation concerning the quality management system, the documentation concerning changes approved by notified bodies, the decisions and other documents issued by notified bodies, and the EU declaration of conformity. Article 19(1) requires providers to keep the logs automatically generated by their high-risk AI systems, to the extent that those logs are under their control, for a period appropriate to the intended purpose of the system

and of at least six months, unless provided otherwise in applicable Union or national law. Article 26(6) imposes a corresponding obligation on deployers in respect of logs under their control.

5.1 Determining the period during which a system is maintained

96. The absence of a defined statutory support period does not mean that a provider may cease to maintain a high-risk AI system at will. As long as the system remains on the Union market or in service, the obligations under Articles 9, 15, 17, 20, 72 and 73 continue to apply to the provider. In particular, Article 20(1) requires a provider that considers or has reason to consider that a high-risk AI system that it has placed on the market or put into service is not in conformity to take the necessary corrective actions immediately, including bringing the system into conformity, withdrawing it, disabling it or recalling it, as appropriate.

97. It follows that a provider that no longer wishes to maintain a high-risk AI system cannot simply cease to do so while the system remains in use. Where a provider decides to discontinue support, and where the continued use of the unsupported system would result in non-conformity, the appropriate course of action is to take the measures referred to in Article 20 and to inform distributors, deployers, the authorised representative and importers accordingly.

98. Good practice, and in many cases the information obligations under Article 13(3), point (e), which requires the instructions for use to specify, where appropriate, the expected lifetime of the high-risk AI system and any necessary maintenance and care measures, will require providers to communicate to deployers, at the time of supply, the period during which updates will be provided and the arrangements applicable at the end of that period. Providers are encouraged to state that period expressly and in a clear and understandable manner, and to notify deployers in advance of its expiry.

Example 46: A provider places on the market a high-risk AI system intended to be used for the monitoring and evaluation of the performance of workers. The instructions for use state that the provider will supply corrective and security updates, and will monitor the system's performance against the declared metrics, for a period of seven years from the date of supply, that being the expected lifetime of the system, and that deployers will be notified twelve months before the end of that period. At the end of that period, the provider notifies deployers and, where continued use of the unsupported system would result in non-conformity, takes the measures referred to in Article 20.

5.2 Substantial modifications and the lifecycle of the AI system

99. As set out in Section 4.4, where a modification qualifies as a substantial modification and the modified system is made available on the market or put into service, it is treated as newly placed on the market. It is therefore useful to clarify the effect that this has on the periods described above.

100. The retention period under Article 18(1) runs from the placing on the market or putting into service of the system concerned. A substantially modified system being newly placed on the market, a new

ten-year period begins to run in respect of the documentation relating to that modified system. The provider's obligations in respect of instances of the earlier version that remain on the market or in service are unaffected by the modification, and the documentation relating to that earlier version must continue to be retained for the remainder of its own ten-year period.

Figure 12 — Obligations across the lifecycle

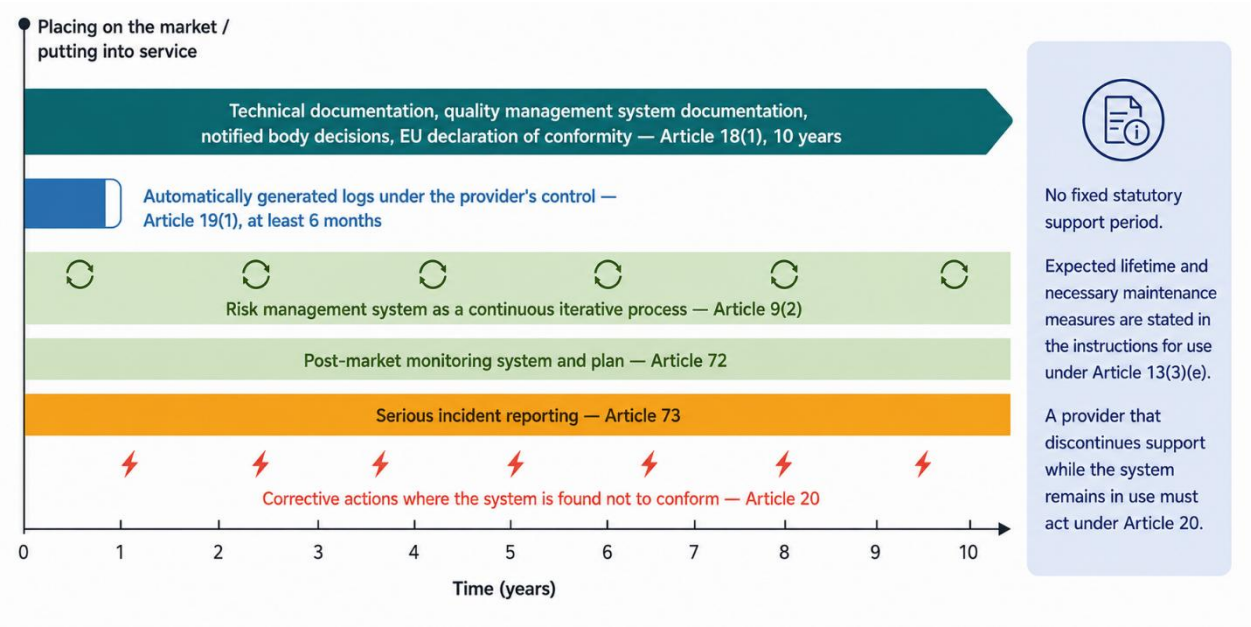

101. A substantial modification does not, however, automatically extend the expected lifetime of the system communicated under Article 13(3), point (e). The relevant question is whether the modification affects the factors that originally determined that expected lifetime. Where the modification concerns only software functionality and the expected lifetime was determined principally by the physical durability of hardware components, the expected lifetime is typically unaffected. Where the modification revitalises the system, for example by replacing the underlying computing platform with one designed for a significantly longer operational life, the provider should reassess and, where appropriate, extend the stated expected lifetime.

Example 47: A provider places on the market a machine-vision safety component with an expected lifetime of ten years, determined principally by the durability of the sensor hardware. After four years, the provider issues a software update that qualifies as a substantial modification by extending the intended purpose to a new class of object. A new conformity assessment is carried out and a new ten-year documentation retention period begins in respect of the modified system. The expected lifetime communicated to deployers remains determined by the hardware and is unchanged.

Example 48: The same provider subsequently replaces the embedded computing platform of the component with a new generation of hardware designed for a significantly longer operational life. Deployers may now reasonably expect the component to remain in productive use beyond the originally stated lifetime. The provider reassesses and updates the expected lifetime stated in the instructions for use accordingly.

6 High-risk AI systems

102. Article 6 establishes two routes by which an AI system is classified as high-risk. Under Article 6(1), an AI system is high-risk where it is intended to be used as a safety component of a product, or is itself a product, covered by the Union harmonisation legislation listed in Annex I, and where that product is required to undergo a third-party conformity assessment under that legislation. Under Article 6(2), the AI systems referred to in Annex III are high-risk, subject to the derogation in Article 6(3).

103. The classification of an AI system as high-risk determines the applicability of the requirements set out in Chapter III, Section 2 and of the obligations set out in Chapter III, Section 3, as well as the conformity assessment procedure that the provider must follow under Article 43 before the system may be placed on the market or put into service. It is therefore essential that the classification be performed correctly and documented.

6.1 Intended purpose

104. Both routes to classification turn on the intended purpose of the system. Article 3(12) defines 'intended purpose' as 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, in promotional or sales materials and statements, as well as in the technical documentation.

105. The intended purpose is therefore determined objectively, on the basis of the totality of the information supplied by the provider, and not solely on the basis of a declaration in one document. Where the instructions for use, the promotional materials and the technical documentation are inconsistent, that inconsistency is itself a relevant indicator.

106. AI systems are rarely restricted exclusively to the functions that define their intended purpose. Systems often perform additional or ancillary functions, and may incorporate components that would themselves fall within an Annex III category if placed on the market separately. The fact that a system performs functions other than or additional to those corresponding to an Annex III category does not in itself prevent it from being classified under that category, where that category corresponds to its intended purpose. Conversely, the mere integration of a component that would fall within an Annex III category does not in itself render the system as a whole high-risk, where its intended purpose lies elsewhere.

Example 49: A general-purpose document management system incorporates a feature that extracts structured fields from curricula vitae. The intended purpose of the system, as specified in the instructions for use, in the promotional materials and in the technical documentation, is the storage, indexing and retrieval of documents of any kind. The system is not intended to be used for the recruitment or selection of natural persons and is not classified under point 4(a) of Annex III. The provider should nonetheless address the reasonably foreseeable misuse of the feature within the meaning of Article 3(13), and should state clearly in the instructions for use that the system is not intended for that purpose.

Example 50: A provider markets the same system explicitly as a recruitment tool, and its promotional materials describe how it ranks candidates. Notwithstanding a statement in the technical documentation that the system is a general document management tool, the intended purpose, assessed on the basis of all the information supplied, corresponds to point 4(a) of Annex III. The system is high-risk.

107. A provider may not misrepresent the intended purpose of its system in order to escape the classification and the corresponding conformity assessment regime, for example by overly emphasising or downplaying the role of certain functions. Clear inconsistencies between promotional materials, instructions for use and technical documentation are a strong indicator of such misrepresentation, and market surveillance authorities may take them into account in the exercise of their powers under Chapter IX.

108. Some systems are placed on the market as single systems but are composed of distinct modules with separate functions. Where a provider also makes those modules available separately, for example through separate purchase, licensing or subscription, those modules constitute distinct AI systems in their own right and are to be classified on the basis of their respective intended purposes. This is without prejudice to modules supplied solely as components of an integrated system and not made available separately, in which case the classification is determined at the level of the integrated system.

Example 51: A provider places on the market a workforce management suite combining a scheduling module, a performance evaluation module and a reporting module, and also offers each module under a separate subscription. Each module is a distinct AI system classified on the basis of its own intended purpose. The performance evaluation module falls within point 4(b) of Annex III and is high-risk. The scheduling and reporting modules are classified on their own terms.

Figure 13 — Intended purpose determines classification

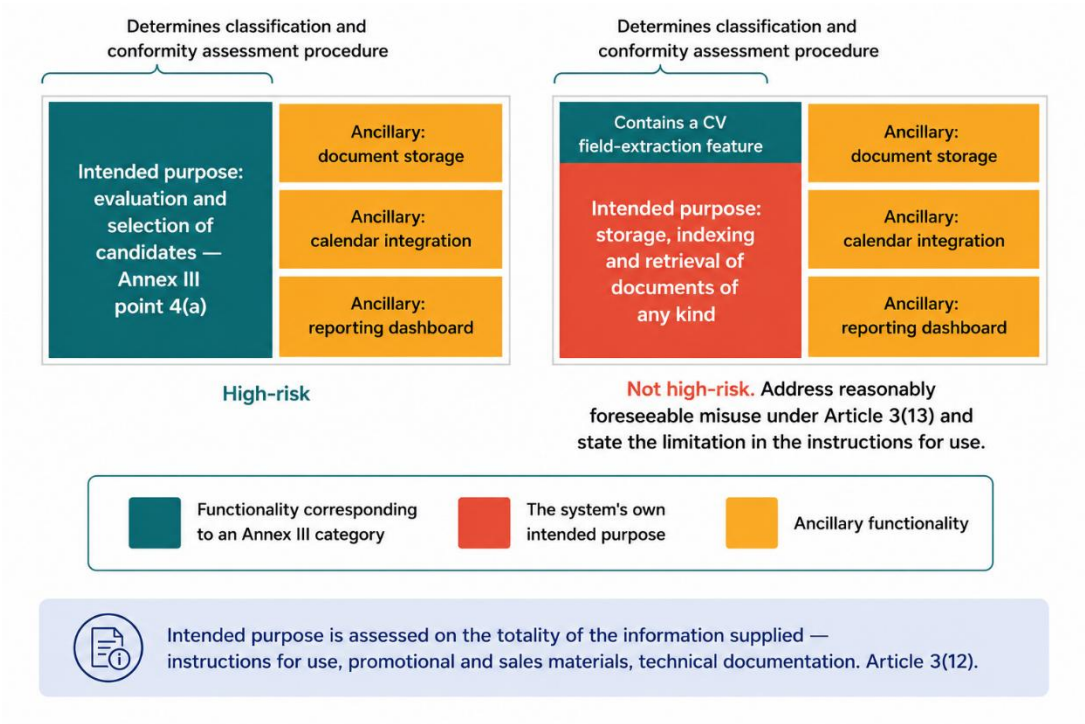

6.2 The derogation under Article 6(3)

109. Article 6(3) provides that an AI system referred to in Annex III is not to be considered high-risk where it does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons, including by not materially influencing the outcome of decision-making. That is the case where one or more of the following conditions is fulfilled:

(a) the AI system is intended to perform a narrow procedural task;

(b) the AI system is intended to improve the result of a previously completed human activity;

(c) the AI system is intended to detect decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment, without proper human review; or

(d) the AI system is intended to perform a preparatory task to an assessment relevant for the purposes of the use cases listed in Annex III.

110. Notwithstanding those conditions, an AI system referred to in Annex III is always to be considered high-risk where it performs profiling of natural persons. That exception is absolute and is not subject to a further assessment of the significance of the risk.

111. The derogation is subject to a procedural discipline. In accordance with Article 6(4), a provider that considers that a system referred to in Annex III is not high-risk is required to document its assessment before that system is placed on the market or put into service, to register the system in accordance with Article 49(2), and to provide the documentation of the assessment to national competent authorities upon request. The provider bears the burden of proof for the fulfilment of the conditions.

Example 52: A system intended to convert scanned application forms into structured text, without evaluating, scoring or ranking the applicants, performs a narrow procedural task within the meaning of Article 6(3), point (a). Provided that it does not perform profiling, it is not high-risk. The provider documents that assessment and registers the system under Article 49(2).

Example 53: A system intended to improve the language, structure and readability of a decision already taken and drafted by a human official, without altering its substance, falls within Article 6(3), point (b).

Example 54: A system intended to flag, for the attention of a supervisor, cases in which a caseworker's decision deviates markedly from historical decision patterns, where the flag triggers a human review and does not itself alter the decision, may fall within Article 6(3), point (c). Where, by contrast, the flag automatically suspends the decision or is presented to the caseworker before the decision is taken in a manner likely to influence it, the condition is not fulfilled.

Example 55: A system intended to rank applicants for a post by predicted suitability performs profiling of natural persons. It is high-risk regardless of any assessment under Article 6(3).

6.3 Conformity assessment for high-risk AI systems

112. Article 43 sets out the applicable conformity assessment procedures. For high-risk AI systems referred to in point 1 of Annex III, concerning biometrics, where the provider has applied harmonised standards referred to in Article 40 or, where applicable, common specifications referred to in Article 41, the provider may opt for the conformity assessment procedure based on internal control referred to in Annex VI, or for that based on the assessment of the quality management system and of the technical documentation, with the involvement of a notified body, referred to in Annex VII. Where those standards or specifications have not been applied, or have been applied only in part, or where they do not exist, the provider is required to follow the procedure set out in Annex VII.

113. For high-risk AI systems referred to in points 2 to 8 of Annex III, providers are to follow the conformity assessment procedure based on internal control referred to in Annex VI, which does not provide for the involvement of a notified body. A provider may nonetheless choose to seek third-party assessment on a voluntary basis.

Figure 14 — The Article 6(3) derogation

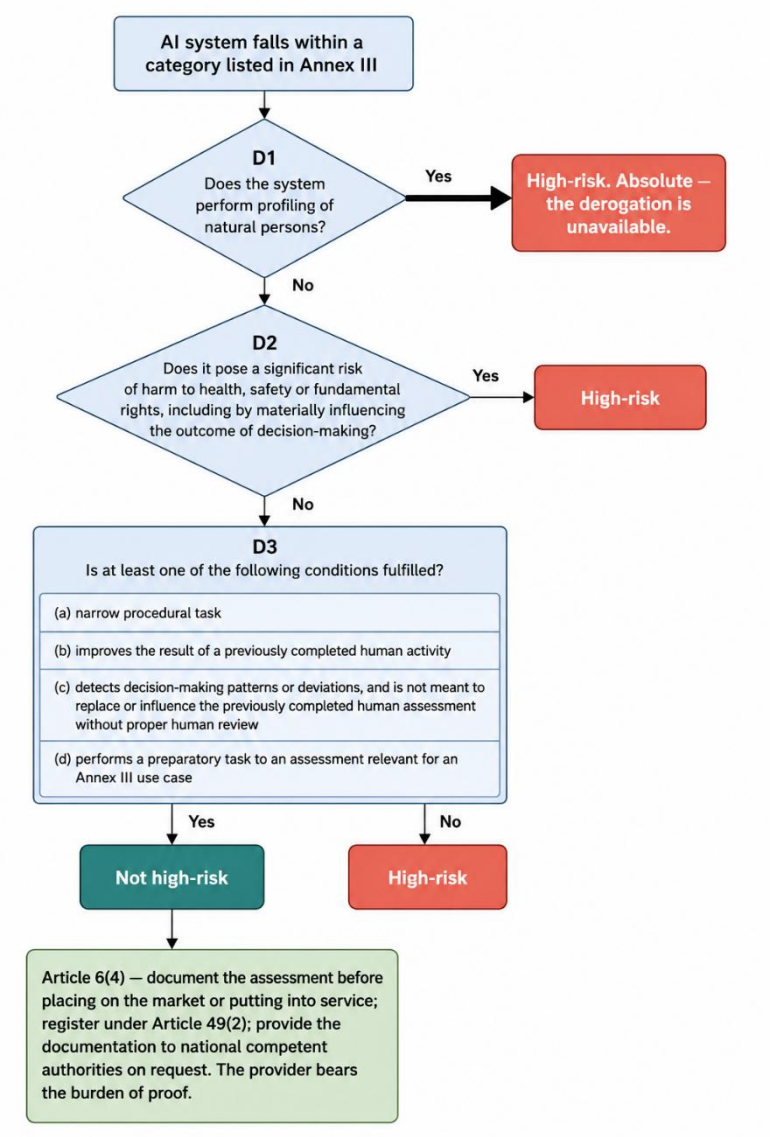

114. For high-risk AI systems covered by the Union harmonisation legislation listed in Section A of Annex I, the provider is required to follow the relevant conformity assessment procedure as required under those legal acts, in accordance with Article 43(3). The requirements set out in Chapter III, Section 2 apply to those systems and are to be part of that assessment. Points 4.3, 4.4, 4.5 and the fifth paragraph of point 4.6 of Annex VII also apply. For the purposes of that assessment, notified bodies which have been notified under those legal acts are entitled to control the conformity of the high-risk AI systems with the requirements set out in Chapter III, Section 2, provided that their compliance with the requirements laid down in Article 31 has been assessed in the context of the notification procedure under those legal acts.

115. In accordance with Article 44(2), certificates issued by notified bodies are valid for the period they indicate, which may not exceed five years for AI systems covered by Annex I and four years for AI systems covered by Annex III. That period may be extended on reapplication, on the basis of a reassessment, for further periods not exceeding those durations. Any supplement to a certificate remains valid as long as the certificate that it supplements is valid.

Example 56: A provider places on the market a remote biometric identification system falling within point 1(a) of Annex III. Relevant harmonised standards exist but the provider applies them only in part. The provider is required to follow the conformity assessment procedure set out in Annex VII, with the involvement of a notified body.

Example 57: A provider places on the market a system intended to be used to evaluate the eligibility of natural persons for essential public assistance benefits, falling within point 5(a) of Annex III. The provider follows the conformity assessment procedure based on internal control set out in Annex VI, draws up the EU declaration of conformity, affixes the CE marking and registers the system in the EU database.

6.4 Implications for presumption of conformity

116. Article 40(1) provides that high-risk AI systems and general-purpose AI models which are in conformity with harmonised standards, or parts thereof, the references of which have been published in the Official Journal of the European Union, are presumed to be in conformity with the requirements set out in Chapter III, Section 2, or, as applicable, with the obligations set out in Chapter V, Sections 2 and 3, to the extent that those standards cover those requirements or obligations. Article 41(3) extends an equivalent presumption to common specifications adopted by the Commission by means of implementing acts.

117. Two further presumptions are established by Article 42. Under Article 42(1), high-risk AI systems that have been trained and tested on data reflecting the specific geographical, behavioural, contextual or functional setting within which they are intended to be used are presumed to comply with the relevant requirements laid down in Article 10(4). Under Article 42(2), high-risk AI systems that have been certified, or for which a statement of conformity has been issued, under a cybersecurity scheme pursuant to Regulation (EU) 2019/881, the references of which have been published in the Official Journal, are presumed to comply with the cybersecurity requirements set out in Article 15 in so far as those requirements are covered by that certificate or statement.

118. The presumption operates only to the extent of the coverage of the standard, specification or certificate relied upon. The scope of a high-risk AI system may be broader than the scope of the relevant harmonised standard, and additional functions may present risks that the standard does not address. As is clearly stated in Section 4.1.2.2 of the Blue Guide, in risk-related harmonisation legislation providers always remain fully responsible for assessing all the risks of their product in order to determine which requirements are relevant. Providers are therefore required in all cases to establish the risk management system under Article 9 and to verify whether the application of the standard covers all the risks identified. Where it does not, the provider must document the additional measures adopted to address the remaining risks, which do not benefit from the presumption.

Example 58: A provider of a high-risk AI system applies a harmonised standard the reference of which has been published in the Official Journal and which covers the requirements of Articles 10 and 15 in respect of the system's principal function. The system also includes an additional function that relies on a different data source and that the standard does not address. The provider benefits from the presumption of conformity as regards the requirements covered by the standard, but not as regards the risks arising from the additional function, for which it must document the measures it has implemented.

Example 59: A provider of a high-risk AI system holds a certificate issued under a European cybersecurity certification scheme covering the resilience of the system against unauthorised access and against attempts to alter its use or performance. The provider benefits from the presumption of conformity in respect of the corresponding cybersecurity requirements of Article 15. The presumption does not extend to the requirements of Article 15 relating to accuracy or to robustness against errors and inconsistencies, nor to the vulnerabilities specific to AI referred to in Article 15(5) in so far as they are not covered by the certificate.

Figure 15 — Routing to the conformity assessment procedure

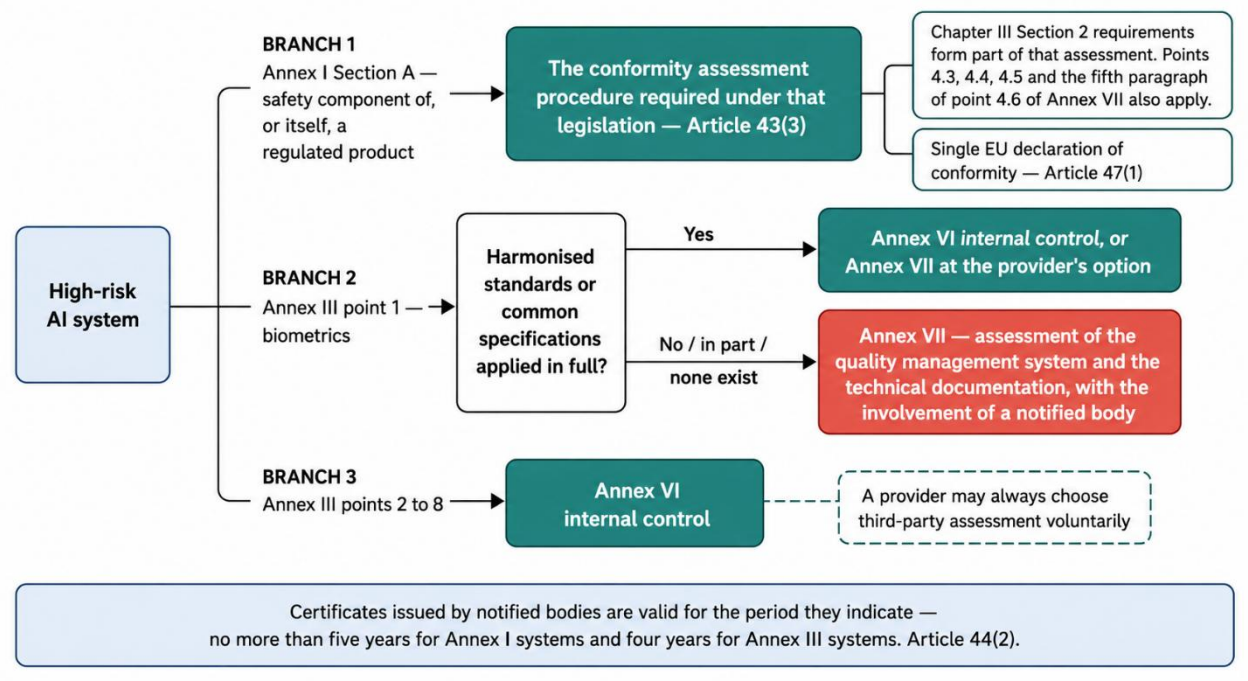

7 Risk management and the integration of AI systems, models and components

7.1 On the evaluation and treatment of risks

119. Article 9 requires providers of high-risk AI systems to establish, implement, document and maintain a risk management system. That system comprises the identification and analysis of the known and the reasonably foreseeable risks that the system can pose to health, safety or fundamental rights when used in accordance with its intended purpose; the estimation and evaluation of the risks that may emerge when the system is used in accordance with its intended purpose and under conditions of reasonably foreseeable misuse; the evaluation of other risks possibly arising on the basis of the data gathered from the post-market monitoring system; and the adoption of appropriate and targeted risk management measures.

120. In organisational risk management, risks are commonly evaluated against acceptance criteria derived from the organisation's internal objectives or risk appetite. By contrast, under the AI Act, residual risk is to be assessed against the standard laid down in Article 9(5), namely that the residual risk associated with each hazard, as well as the overall residual risk, is judged to be acceptable. That judgement is made by reference to the intended purpose and the conditions of reasonably foreseeable misuse of the system, and not by reference to the provider's commercial position.

121. Accordingly, the provider's internal risk tolerance, commercial strategy or mere cost considerations are not relevant in determining whether risks have been adequately addressed. Article 9(5) requires risks to be eliminated or reduced as far as technically feasible through adequate design and development of the high-risk AI system; where they cannot be eliminated, adequate mitigation and control measures must be implemented; and appropriate information must be provided pursuant to Article 13 and, where appropriate, training must be provided to deployers.

122. Residual risk is an inherent outcome of risk assessment and risk treatment. It is recognised that risks cannot in practice be entirely eliminated. However, the existence of residual risk does not imply that any risk may be accepted at the provider's discretion. A high-risk AI system may only be placed on the market where the residual risks, once appropriate measures have been taken, are judged acceptable in light of the intended purpose and reasonably foreseeable misuse.

Figure 16 — The reach of the presumption of conformity

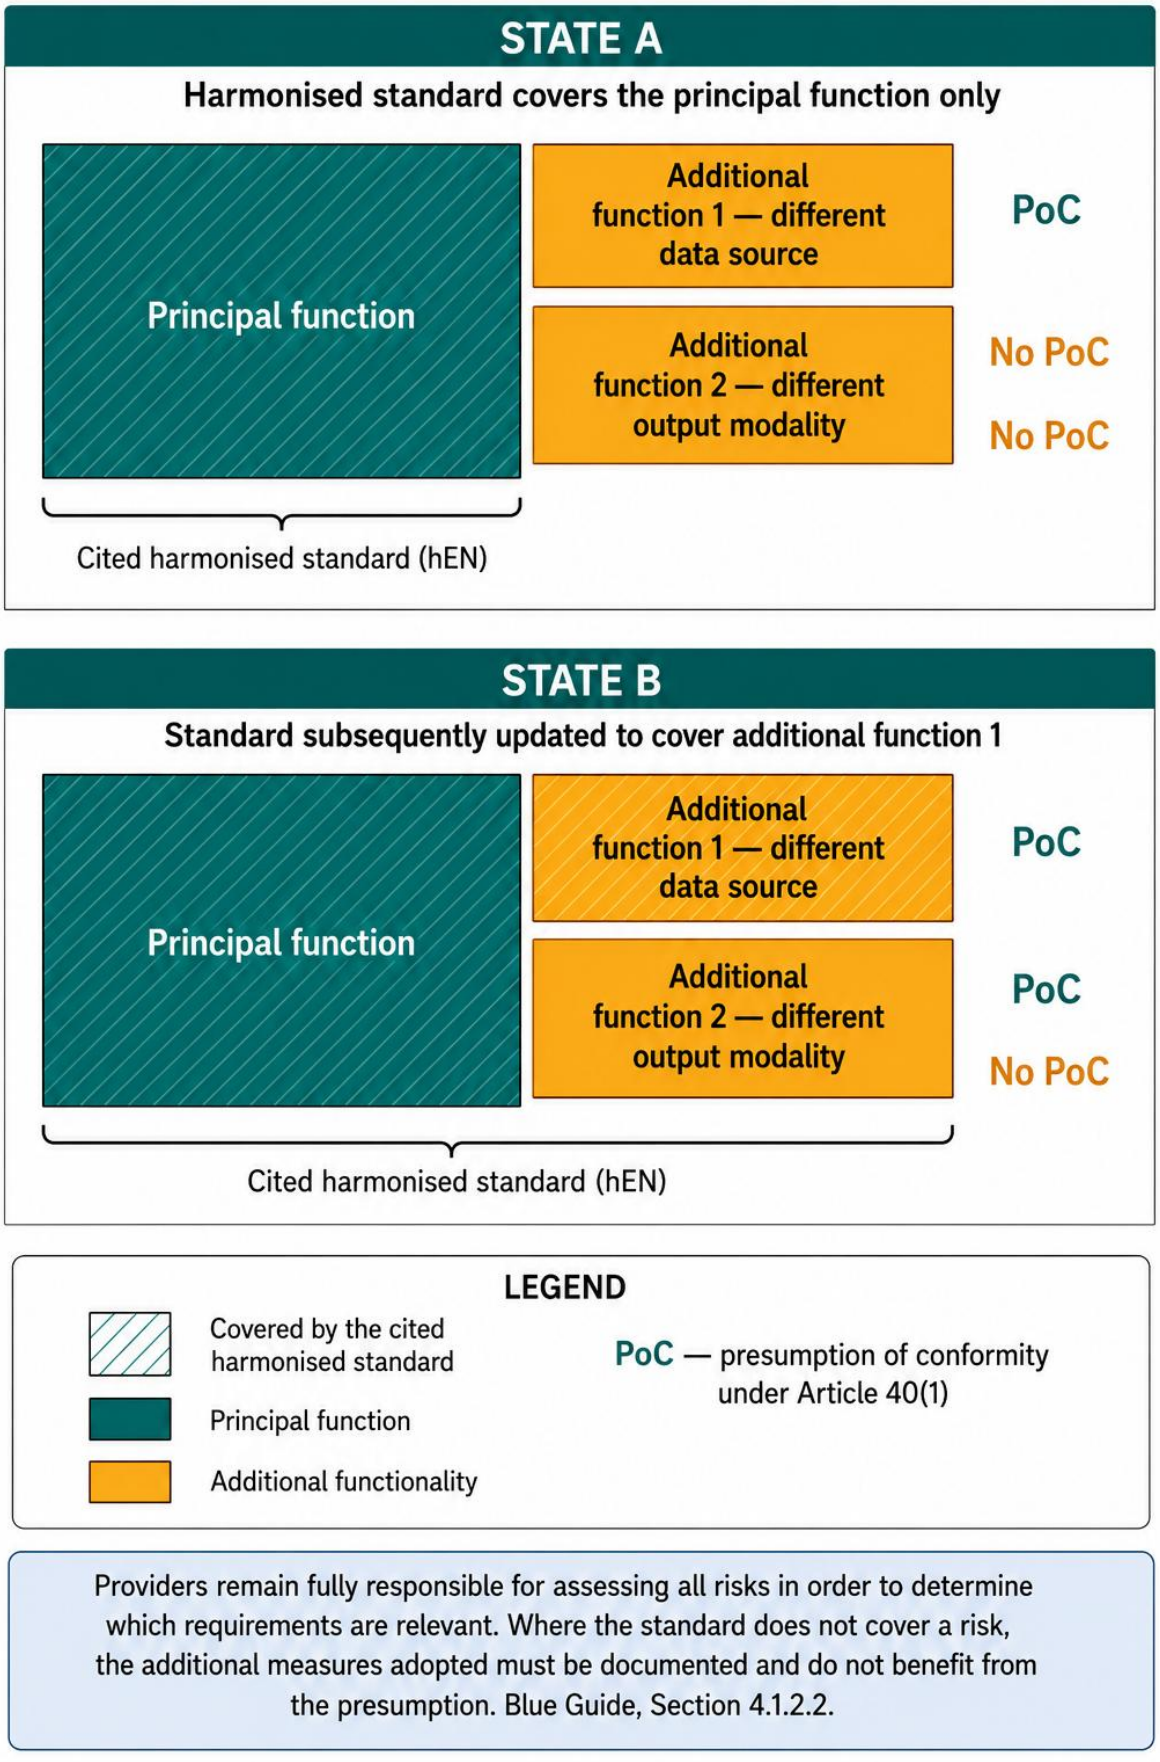

Providers remain fully responsible for assessing all risks in order to determine which requirements are relevant. Where the standard does not cover a risk, the additional measures adopted must be documented and do not benefit from the presumption. Blue Guide, Section 4.1.2.2.

123. The AI Act does not provide for the transfer of risk or of responsibility to deployers or to affected persons in order to compensate for shortcomings in system design, or to justify leaving risks unaddressed. The obligation to place a compliant system on the market and to demonstrate conformity remains with the provider. Information and instructions provided to deployers under Article 13 may be used to support the safe deployment and operation of the system, including where the provider has chosen to restrict the intended purpose to defined environments, and to inform deployers of residual risks, in a manner appropriate to the nature of the system and to its intended users. They may not be used as a substitute for measures that could reasonably have been implemented in the design.

Example 60: A provider develops a high-risk AI system intended to be used in an industrial setting under continuous supervision by trained operators. In order to mitigate risks identified in the risk management system, the provider restricts the intended purpose to supervised operation, designs the human oversight measures required by Article 14 so that an operator can at all times interrupt the system through a stop function, and states those limitations clearly in the instructions for use. That is a legitimate risk management measure.

Example 61: A provider identifies a risk that a high-risk AI system may produce systematically less accurate outputs for a subgroup of the intended population. Rather than addressing that disparity through the data governance measures required by Article 10 and through the design of the system, the provider confines itself to stating in the instructions for use that deployers should apply additional scrutiny to that subgroup. That is not sufficient: the risk must first be eliminated or reduced as far as technically feasible through design and development.

124. Article 9(9) requires providers, when implementing the risk management system, to give specific consideration to whether the high-risk AI system is likely to have an adverse impact on persons under the age of 18 and, as appropriate, on other vulnerable groups. Where such an impact is likely, the risk management measures adopted must address it specifically, and the analysis should be reflected in the technical documentation.

Figure 17 — The Article 9 risk management cycle

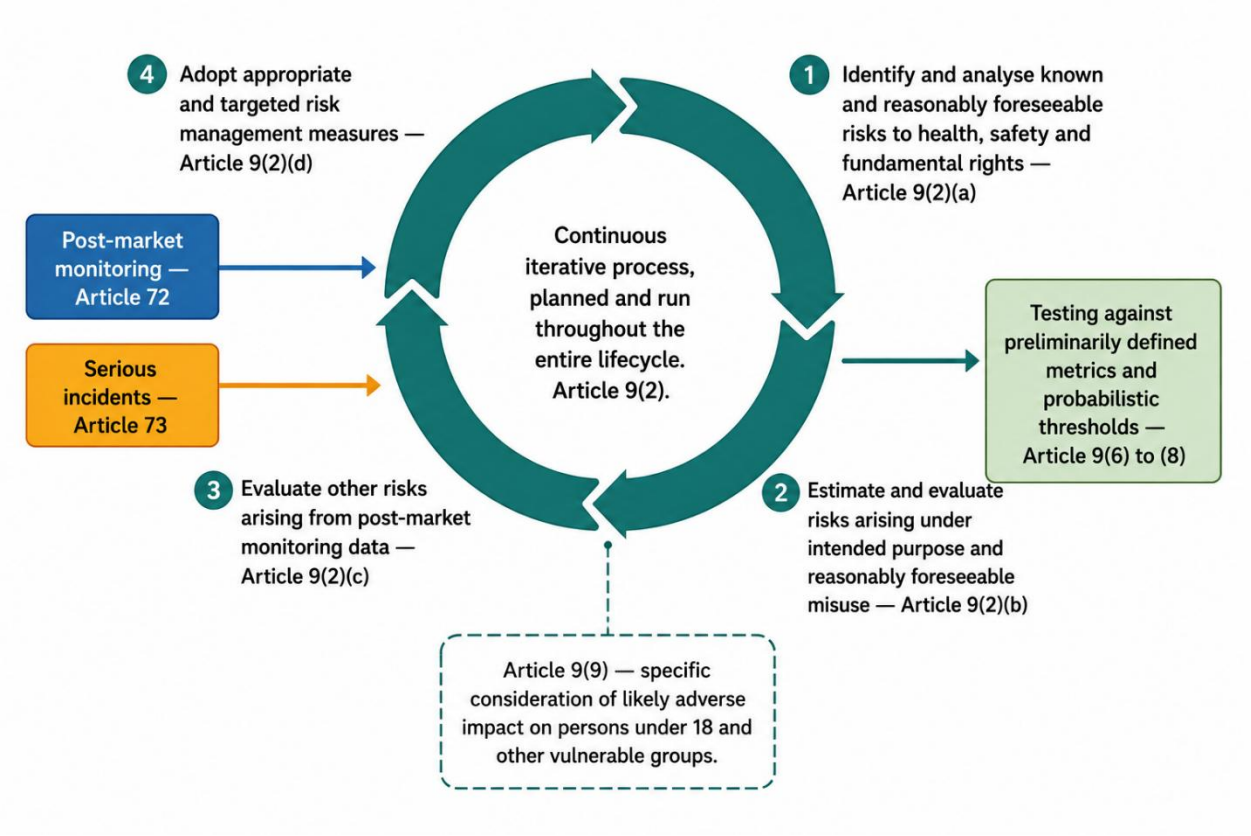

7.2 Accuracy, robustness and cybersecurity

125. Article 15(1) requires high-risk AI systems to be designed and developed in such a way that they achieve an appropriate level of accuracy, robustness and cybersecurity, and that they perform consistently in those respects throughout their lifecycle. That requirement functions, in part, as a residual requirement: it addresses those risks that are not otherwise adequately addressed through the implementation of the other requirements of Chapter III, Section 2.

126. Where all the relevant risks identified in the risk management system are treated by implementing adequate measures addressing the other applicable requirements, the requirement in Article 15(1) will generally be fulfilled. Conversely, where the risk management system identifies additional risks that are not fully addressed by those measures, providers are required to implement further measures at the level of the system. In practice, in most cases, compliance with the other requirements is expected to result in compliance with Article 15(1).

127. Article 15(3) requires the levels of accuracy and the relevant accuracy metrics to be declared in the accompanying instructions for use. Article 15(4) requires systems to be as resilient as possible regarding errors, faults or inconsistencies that may occur within the system or the environment in which it operates, and addresses in particular the risk of feedback loops in systems that continue to learn after being placed on the market or put into service. Article 15(5) addresses resilience against attempts by unauthorised third parties to alter the use, outputs or performance of the system, and requires technical solutions aimed at preventing, detecting, responding to, resolving

and controlling attacks trying to manipulate the training dataset, inputs designed to cause the model to make a mistake, confidentiality attacks and model flaws, as appropriate to the relevant circumstances and risks.

Figure 18 — The hierarchy of risk treatment

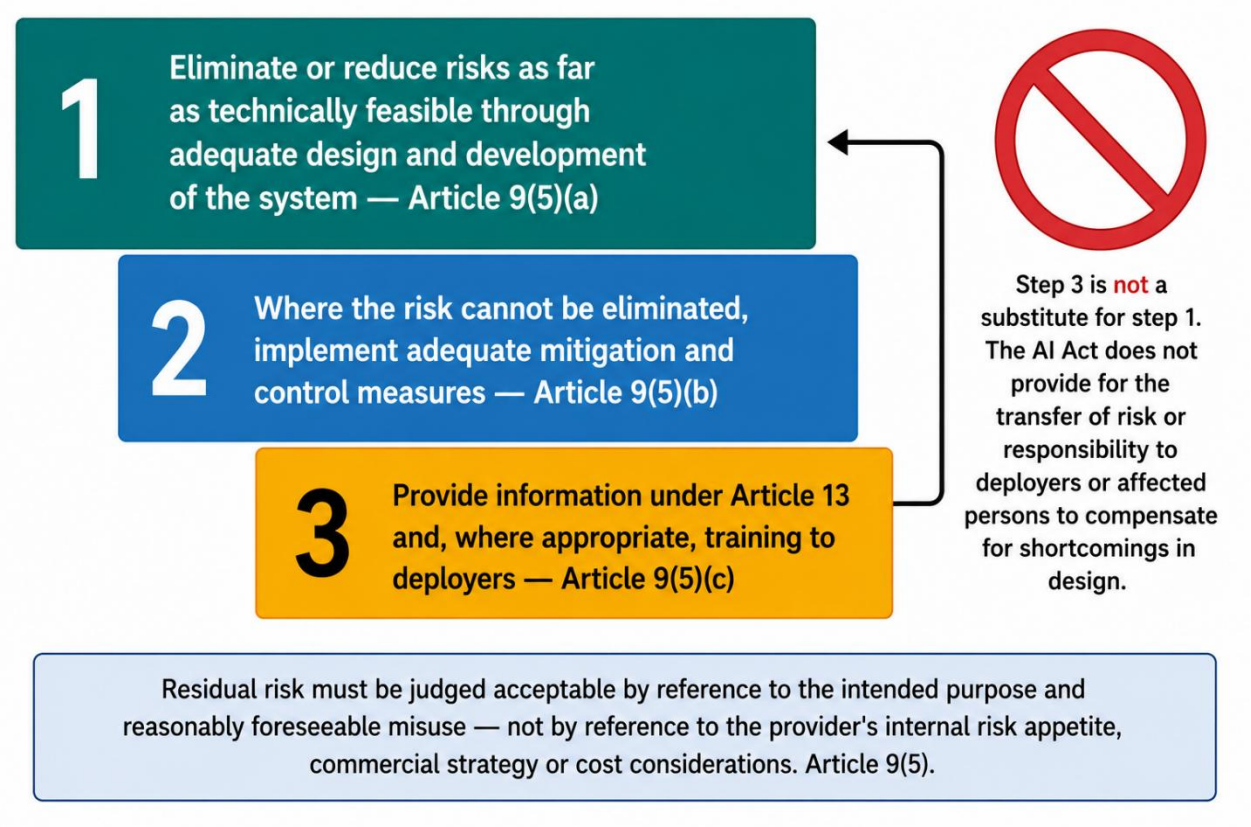

7.3 Third-party components, external dependencies and cooperation along the value chain

128. The AI Act establishes two distinct but complementary sets of considerations in relation to elements that the provider does not itself develop. First, the risk management system under Article 9 must take account of all relevant risks, including risks that originate outside the system itself, such as the deployer's operating environment, external networks, upstream data sources or other systems on which the system relies. Second, Article 25(4) requires cooperation, by written agreement, with third parties that supply AI systems, tools, services, components or processes that are used or integrated in a high-risk AI system.

129. For risks originating outside the system, the AI Act does not require providers to control or govern the external environment. Rather, providers are required to identify such risks and to mitigate them through the design and development of the system itself, and, where necessary, through the information and instructions provided to deployers under Article 13. This will involve implementing the appropriate requirements of Chapter III, Section 2 at the level of the system.

130. For example, a provider may identify a risk that the data feeds on which a high-risk AI system relies become unavailable, degraded or manipulated. Those feeds, where they are not part of the system, are not themselves subject to the AI Act. Nonetheless, the provider is required to address the risk they pose through measures at the level of the system, such as validating and authenticating incoming data, monitoring for distribution shift, generating logs and alerts when abnormal inputs are detected, and ensuring that a loss of the feed does not cause the system to enter an unsafe state or to produce outputs presented with unwarranted confidence. In such cases, the AI Act regulates how the system responds to external risks; it does not impose obligations on how the external infrastructure is organised, staffed or operated.

131. For elements that are integrated into the system, Article 25(4) requires the provider and the third party to specify by written agreement the information, capabilities, technical access and other assistance, based on the generally acknowledged state of the art, necessary to enable the provider to comply fully with its obligations. In alignment with its risk management system, the provider has to identify what the system requires from its components in order to meet its objectives, and to verify, in a risk-based manner, that those components are in line with those needs. Evidence for that purpose may consist of documentation obtained from the component supplier, such as technical specifications, evaluation reports, model documentation or relevant conformity or assurance documentation. Where appropriate, the provider may also carry out its own tests.

132. The obligation in Article 25(4) does not apply in respect of third parties that make tools, services, processes or components accessible to the public under a free and open-source licence. That exception does not extend to general-purpose AI models. Where no written agreement is required or obtainable, the provider remains responsible for the compliance of its system as a whole, and must therefore satisfy itself by other means that the integrated element is suitable, or select a different element.

Example 62: A provider of a high-risk AI system integrates a third-party optical character recognition component supplied under a commercial licence. The provider identifies, in its risk management system, that the accuracy of the component on low-quality scans directly affects the accuracy of the system's outputs. It specifies in the written agreement under Article 25(4) that the supplier will provide performance characteristics disaggregated by document quality, will give notice of model changes, and will provide technical access sufficient to allow the provider to re-run its own validation.

Example 63: A provider of a high-risk AI system relies on a third-party geospatial data service that is not integrated into the system but is queried at runtime. That service is not a component of the system. The provider is nonetheless required to assess the risks arising from its use, and to address them at the level of the system, for example by validating the plausibility of returned values, by defining behaviour in the event of unavailability, and by informing deployers of the dependency in the instructions for use.

7.4 Reuse of risk management and conformity documentation for families of AI systems

133. Providers may place on the market AI systems that are similar, for example different variants, models or configurations belonging to the same family. Where such systems share the same architecture, the same security- and safety-relevant design, the same intended purpose, and the same risk profile, the AI Act does not require providers to treat each variant as an entirely separate system for the purposes of risk management and conformity assessment.

134. In such cases, provided that all the variants concerned are adequately covered in terms of relevant risks and applicable requirements, providers may rely on a single risk management system established in accordance with Article 9, a single set of technical documentation, and a single conformity assessment procedure. This also allows a single EU declaration of conformity to be issued for the group of systems, provided that it clearly identifies the variants to which it applies.

Figure 19 — Integrated components and external dependencies

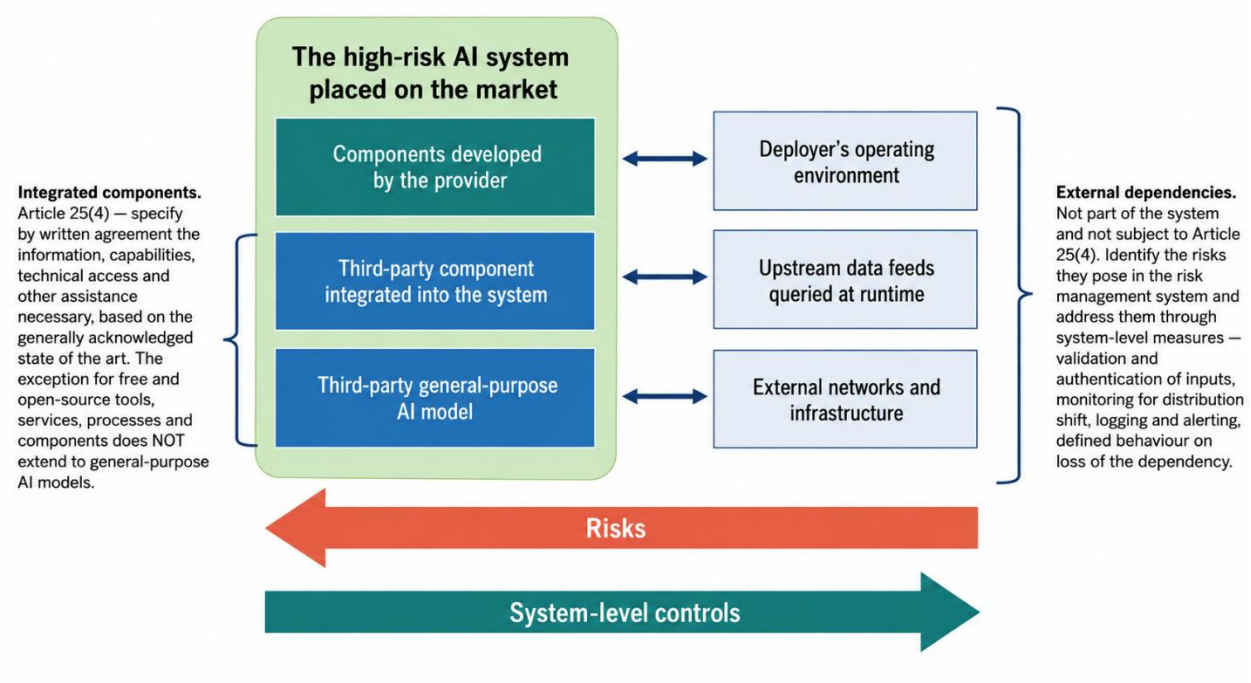

135. The decisive factor is whether the differences between variants are relevant to the risks addressed by the Regulation. Variations that do not affect those risks, such as differences in branding, packaging, interface language or non-functional configuration, do not require separate assessments. Conversely, variants may differ in ways that affect the implementation of the requirements, for example through different input modalities, different training or validation datasets, different intended populations, different degrees of automation or different human oversight arrangements. Such differences must be reflected in the risk management system and, where necessary, in the conformity assessment and technical documentation.

Example 64: A provider places on the market a high-risk AI system in three configurations that differ only in the maximum throughput they support and in the language of the user interface. Those differences do not affect the risks addressed by the Regulation. A single risk management system,

set of technical documentation, conformity assessment and EU declaration of conformity may cover all three, provided the declaration identifies them.

Example 65: The same provider subsequently introduces a fourth configuration that accepts audio input in addition to text. That variant is exposed to different risks, relies on different validation data and requires different accuracy metrics to be declared. Those differences must be reflected in the risk management system and in the technical documentation, and, where relevant, assessed separately.

Figure 20 — Decision sequence for general-purpose AI model obligations

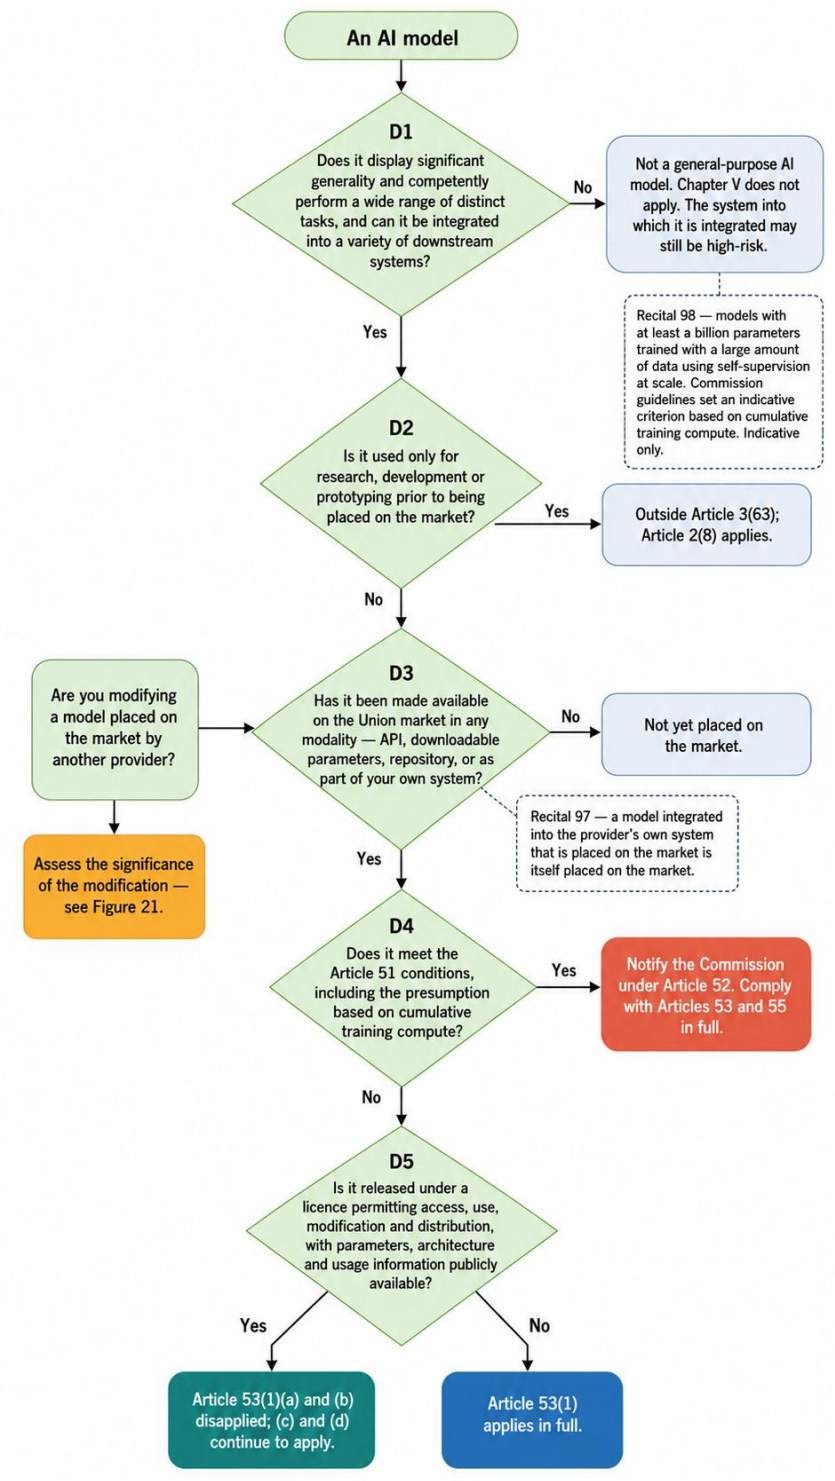

8 General-purpose AI models

136. Chapter V establishes a distinct regime for general-purpose AI models. That regime attaches to the model itself, independently of any AI system into which the model may subsequently be integrated, and is supervised and enforced exclusively by the Commission acting through the AI Office. The purpose of this Section is to assist providers and downstream actors in determining whether a model is a general-purpose AI model, when such a model is placed on the market, who is its provider, and how the resulting obligations interact with those of downstream providers of AI systems.

8.1 What is a general-purpose AI model?

137. Article 3(63) defines a general-purpose AI model as an AI model, including where such a model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications. That definition expressly excludes AI models that are used for research, development or prototyping activities before they are placed on the market.

8.1.1 The notion of generality

138. Three elements are therefore cumulative: the model must display significant generality; it must be capable of competently performing a wide range of distinct tasks; and it must be capable of integration into a variety of downstream systems or applications. A model that is large, or that was trained on a large quantity of data, is not for that reason alone a general-purpose AI model if it is specialised to a single task or to a narrow domain.

139. Recital 98 acknowledges that, whereas the generality of a model could be determined by a number of parameters, models with at least a billion parameters and trained with a large amount of data using self-supervision at scale should be considered to display significant generality. The Commission's guidelines on the scope of the obligations for providers of general-purpose AI models set out a workable indicative criterion based on the cumulative amount of computation used for training, together with the capability of the model to generate language, text-to-image or text-to-video content. That criterion is indicative: it does not displace an assessment of the three elements of the definition, and a model may fall within or outside the definition notwithstanding the criterion.

Example 66: A model trained on a large corpus using self-supervision, capable of summarising, translating, answering questions, writing code and classifying text, and made available for integration into a wide range of applications, displays significant generality and is a general-purpose AI model.

Example 67: A model trained on a very large corpus of chest radiographs, capable only of detecting a defined set of radiological findings, does not display significant generality and cannot competently perform a wide range of distinct tasks. It is not a general-purpose AI model, notwithstanding the

scale of its training. Where it is integrated into a medical device, the system may nonetheless be high-risk under Article 6(1).

Example 68: A general model is further trained so that it can only produce outputs in a single narrow format for a single downstream task, and its general capabilities are removed. Whether the resulting model remains a general-purpose AI model requires an assessment of its actual capabilities, and not of the capabilities of the model from which it was derived.

8.1.2 When is a general-purpose AI model placed on the market?

140. A general-purpose AI model is placed on the market when it is first made available on the Union market within the meaning of Article 3(10), whatever the modality of that supply. This includes making the model available through an application programming interface, as a downloadable set of parameters, through a model repository or library, through a copy supplied on physical media, or as part of an AI system supplied by the same provider. As explained in recital 97, where a provider integrates its own model into its own AI system that is made available on the market or put into service, that model is also to be considered to be placed on the market, and the obligations of Chapter V apply in addition to those applicable to the system.

141. By contrast, a model that is used exclusively for research, development or prototyping activities before being placed on the market is excluded from the definition in Article 3(63), and Article 2(8) excludes research, testing and development activity prior to the placing on the market from the scope of the Regulation more generally.

Example 69: A provider makes a general-purpose AI model available exclusively through a paid application programming interface, without releasing the parameters. The model is placed on the market. The provider is subject to Article 53 and, if the model meets the conditions of Article 51, to Articles 52 and 55.

Example 70: A provider develops a general-purpose AI model and uses it exclusively as the engine of its own consumer assistant, which it places on the market. The model is placed on the market together with the system. The provider is subject to Chapter V in respect of the model and to the applicable obligations in respect of the system, including Article 50.

8.1.3 Modification of a general-purpose AI model by a downstream actor

142. A downstream actor that modifies a general-purpose AI model, for example by fine-tuning it, may itself become the provider of a general-purpose AI model. Whether it does so, and the extent of the resulting obligations, depend on the significance of the modification.

143. Where the modification is such that it results in what is, in substance, a new model, the modifying actor becomes the provider of that new model and is subject to the obligations of Chapter V in respect of it. Where the modification is more limited, the Commission's guidelines indicate that the obligations of the modifying actor should be understood as applying only in respect of the modification, so that, for example, the technical documentation required by Article 53(1), point (a) need cover only the modification and the resources used for it, rather than the model as a whole.

Those guidelines set out an indicative criterion based on the ratio between the computation used for the modification and that used to train the original model.

144. In all cases, the actor that carries out the modification should be able to identify the model from which its own model was derived, and should have obtained that model on terms permitting the modification and the subsequent supply.

Example 71: A company obtains an open-weight general-purpose AI model and fine-tunes it on a modest volume of domain-specific data, using a small fraction of the computation used to train the original model, in order to improve its performance on legal drafting. It then makes the resulting model available to third parties. The company becomes a provider in respect of the modification, and its documentation obligations are to be understood as extending to the modification and the resources used for it.

Example 72: A company obtains the same model and carries out extensive continued pre-training using a very large additional corpus and an amount of computation comparable to that used to train the original model, materially altering the model's capabilities. The resulting model is, in substance, a new general-purpose AI model. The company is its provider and is subject to Chapter V in full, including, where the applicable conditions are met, to Article 52 and Article 55.

Question	Consequence
Does the model display significant generality and competently perform a wide range of distinct tasks?	If not, it is not a general-purpose AI model. Chapter V does not apply. The system into which it is integrated may nonetheless be high-risk.
Is the model used only for research, development or prototyping prior to placing on the market?	If yes, it falls outside Article 3(63) and Article 2(8) applies.
Has the model been made available on the Union market in the course of a commercial activity, in any modality, including as part of your own system?	If yes, it has been placed on the market and Article 53 applies to its provider.
Is the model released under a licence permitting access, use, modification and distribution, with parameters, architecture and usage information publicly available?	If yes, and the model does not present systemic risk, Article 53(1), points (a) and (b) do not apply. Points (c) and (d) continue to apply.
Does the model meet the conditions in Article 51, including the presumption based on cumulative training compute?	If yes, notify the Commission under Article 52 and comply with Articles 53 and 55 in full.
Are you modifying a model placed on the market by another provider?	Assess the significance of the modification. Limited modifications give rise to obligations confined to the modification; modifications producing what is in substance a new model give rise to the full set of obligations.

Table 2: Stylised decision sequence for general-purpose AI model obligations

8.2 Practical implications for downstream providers

145. Article 53(1), point (b) requires providers of general-purpose AI models to draw up, keep up to date and make available information and documentation to providers of AI systems who intend to integrate the model into their AI systems. That information must enable those providers to have a good understanding of the capabilities and limitations of the model and to comply with their own obligations, and must contain at least the elements set out in Annex XII. That obligation is central to the operation of the value chain: it is the principal mechanism by which a downstream provider of a high-risk AI system obtains the information it needs in order to comply with Chapter III, Section 2.

146. The downstream provider that integrates a general-purpose AI model into a high-risk AI system remains responsible for the compliance of that system as a whole. It may rely on the information supplied under Article 53(1), point (b) and Annex XII, and on the written agreement concluded under Article 25(4), but it must satisfy itself that the model is suitable for the intended purpose of its system, and must carry out its own validation and testing against that intended purpose. The general capabilities of a model, however extensively documented by its provider, are not a substitute for evidence about the performance of the system in its specific context of use.

147. As part of the risk management process and of the cooperation required under Article 25(4), the following non-exhaustive list of assurance artefacts may be re-used in support of the downstream provider's assessment:

(a) the documentation supplied under Article 53(1), point (b) and Annex XII, including model documentation, evaluation results and known limitations;

(b) evidence of adherence to the General-Purpose AI Code of Practice, or to another adequate alternative code of practice, where relied upon by the model provider under Article 53(4);

(c) a statement of conformity or certificate obtained under a European cybersecurity certification scheme adopted pursuant to Regulation (EU) 2019/881;

(d) evidence of conformity with ISO/IEC 42001:2023 or ISO/IEC 23894:2023, or with harmonised standards the references of which have been published in the Official Journal;

(e) the results of independent evaluations, red-teaming exercises or adversarial testing, where available.

148. Providers of AI systems relying on third-party models should also address, contractually and technically, the consequences of changes made by the model provider. A change to a third-party model relied upon by a high-risk AI system is not, in itself, a substantial modification of that system, since it is not a change made by or under the responsibility of the system's provider. However, it may cause the system to cease to conform, in which case Article 20 applies. Providers are therefore encouraged to ensure that they are kept adequately informed about model changes, deprecations and version retirements, to pin model versions where technically possible, to re-run their validation

regime when a change occurs, and to update the risk management system accordingly. In some cases, providers may need to modify their system-level controls or to change model supplier.

8.3 Use cases

In order to provide additional practical guidance, this Section introduces a series of fictional use cases describing archetypal deployments that rely on general-purpose AI models, and illustrating how the obligations of the AI Act are allocated in each case.

8.3.1 Customer-service assistant

A financial entity places on the market a customer-service assistant embedded in its mobile application. The assistant answers questions about the entity's products, retrieves the customer's transaction history and explains fees. It is built on a general-purpose AI model licensed from a third-party provider and accessed through an application programming interface, together with a retrieval layer and a set of guardrails developed by the financial entity.

Implications: the financial entity is the provider of the AI system, which it places on the market under its own name. The third-party model provider is the provider of the general-purpose AI model and remains subject to Chapter V. Because the assistant is intended to interact directly with natural persons, Article 50(1) applies and the financial entity must ensure that customers are informed that they are interacting with an AI system, unless this is obvious from the point of view of a reasonably well-informed, observant and circumspect natural person. Because the assistant generates synthetic text, Article 50(2) applies, subject to the exception for systems performing an assistive function for standard editing. The assistant, as described, does not evaluate creditworthiness and is not high-risk. Were the financial entity to extend its intended purpose so that it assessed the creditworthiness of natural persons, the system would fall within point 5(b) of Annex III and would become high-risk, with the consequence that the full set of obligations in Chapter III would apply and a written agreement under Article 25(4) with the model provider would be required.

8.3.2 Emergency triage support tool

A provider places on the market a system intended to assist emergency call handlers by transcribing calls, extracting clinical indicators and proposing a priority category for dispatch. The system is built on a general-purpose AI model that the provider has fine-tuned on a corpus of anonymised historical calls.

Implications: the intended purpose falls within point 5(d) of Annex III, which covers AI systems intended to be used for the triage of patients in emergency healthcare and for establishing priority in the dispatching of emergency first response services. The system is high-risk. The derogation in Article 6(3) is not available, since the system materially influences the outcome of decision-making. The provider must comply with Articles 8 to 15, follow the conformity assessment procedure under Annex VI, draw up the EU declaration of conformity, affix the CE marking and register the system under Article 49. Human oversight measures under Article 14 are of particular importance, and the accuracy metrics declared under Article 15(3) must be disaggregated in a manner that allows call handlers to understand the system's limitations. Depending

on the extent of the fine-tuning, the provider may also become a provider of a general-purpose AI model in respect of the modification, as discussed in Section 8.1.3.

8.3.3 Code completion tool

A provider places on the market a code completion tool built on a general-purpose AI model, integrated into a development environment, that suggests completions as a developer types.

Implications: the intended purpose does not correspond to any category in Annex III and the tool is not a safety component of a product covered by Annex I. It is not high-risk. The tool generates synthetic text and therefore falls within the scope of Article 50(2). However, that provision does not apply where the AI system performs an assistive function for standard editing or does not substantially alter the input data provided by the deployer or the semantics thereof. Whether the exception applies requires a case-by-case assessment: a tool that completes an identifier the developer has begun to type is likely to fall within it, whereas a tool that generates entire functions or files from a natural-language instruction is likely not to. Where the exception does not apply, the provider must ensure that outputs are marked in a machine-readable format and detectable as artificially generated, using technical solutions that are effective, interoperable, robust and reliable as far as is technically feasible, taking into account the specificities and limitations of the type of content concerned.

8.3.4 Industrial inspection system

A manufacturer places on the market a machine-vision system that inspects components on a production line and halts the line when a defect that could cause a hazard is detected. The system relies on a model trained specifically for that inspection task.

Implications: the model is not a general-purpose AI model, since it displays no significant generality and cannot perform a wide range of distinct tasks. Chapter V does not apply. The system is a safety component of machinery covered by Regulation (EU) 2023/1230, which is listed in Section A of Annex I, and is therefore high-risk under Article 6(1) where the machinery is required to undergo a third-party conformity assessment. In accordance with Article 43(3), compliance with Chapter III, Section 2 is assessed within the conformity assessment procedure carried out under that Regulation, and a single EU declaration of conformity is drawn up under Article 47(1). Those obligations apply from 2 August 2027.

8.3.5 Model hosting platform

A legal person operates a platform on which third parties publish models and datasets, together with documentation. The platform does not develop the models, does not place them on the market under its own name or trademark, and does not exercise control over their content.

Implications: the operator of the platform is not, on that basis alone, the provider of the models hosted on it. The provider of each general-purpose AI model is the person that develops it and places it on the market under its own name or trademark. The operator may, depending on the circumstances and on the way in which it presents and supplies AI systems, fall within the definition of distributor in Article 3(7), with the consequence that Article 24 applies. Where the operator itself develops or substantially modifies a model

and makes it available under its own name, it becomes its provider. Obligations arising under other Union law, in particular Regulation (EU) 2022/2065, are unaffected.

9 Additional elements

9.1 Reporting obligations

149. Article 73(1) requires providers of high-risk AI systems placed on the Union market to report any serious incident to the market surveillance authorities of the Member States where that incident occurred. Article 3(49) defines a serious incident as an incident or malfunctioning of an AI system that directly or indirectly leads to the death of a person or serious harm to a person's health; a serious and irreversible disruption of the management or operation of critical infrastructure; the infringement of obligations under Union law intended to protect fundamental rights; or serious harm to property or the environment.

150. The reporting deadlines are graduated according to the nature of the incident. Under Article 73(2), the report is to be made immediately after the provider has established a causal link between the AI system and the serious incident, or the reasonable likelihood of such a link, and in any event not later than 15 days after the provider or, where applicable, the deployer, becomes aware of the serious incident. Under Article 73(3), in the event of a widespread infringement or of a serious incident consisting in a serious and irreversible disruption of the management or operation of critical infrastructure, that period is reduced to two days. Under Article 73(4), in the event of the death of a person, the report is to be made immediately after the provider or the deployer has established, or as soon as it suspects, a causal relationship, and in any event not later than 10 days after the date on which the provider or, where applicable, the deployer becomes aware of the serious incident.

151. Article 73(5) permits an initial incomplete report to be submitted where necessary to ensure timely reporting, followed by a complete report. Article 73(6) requires the provider, following the reporting of a serious incident, to perform without delay the necessary investigations in relation to the incident and the AI system concerned, including a risk assessment of the incident and corrective action, and prohibits the provider from performing any investigation that involves altering the AI system in a way that may affect any subsequent evaluation of the causes of the incident before informing the competent authorities. Under Article 73(9), where a deployer identifies a serious incident, it is required to immediately inform first the provider, then the importer or distributor and the relevant market surveillance authorities.

152. The reporting obligation is triggered when the provider becomes aware of the serious incident. Guidance on when a provider is deemed to have become aware should therefore help providers determine the moment at which the applicable deadlines start to run. Providers may be subject to comparable reporting obligations under different Union acts. To ensure that the concept of 'becoming aware' is interpreted consistently and to facilitate compliance, this guidance is aligned

with recital 31 of Commission Implementing Regulation (EU) 2024/2690 and with Section II(A) of the Guidelines 9/2022 on personal data breach notification under Regulation (EU) 2016/679[11].

Figure 21 — Downstream modification of a model

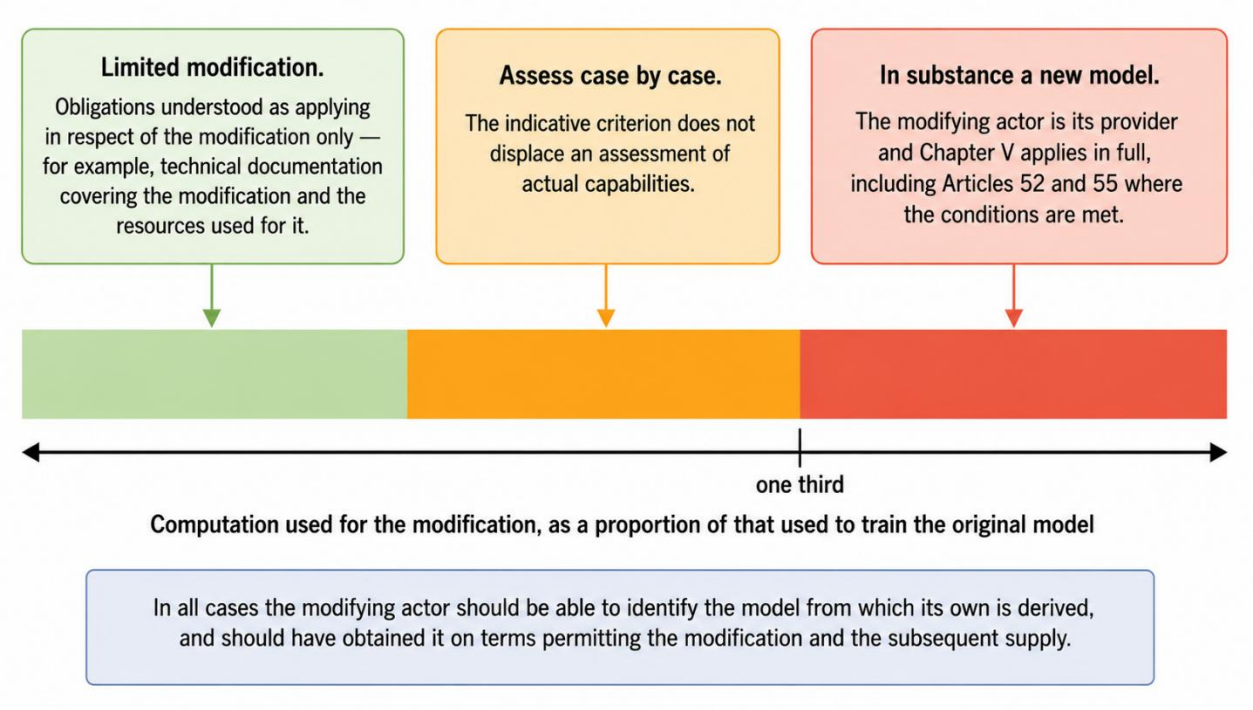

153. In some cases, a provider will detect a suspicious event; in others, a third party, such as an individual, a deployer, an authority, a media organisation or another source, will bring a potential incident to its attention. In such cases, the provider should assess the event immediately to determine whether it constitutes a serious incident. The provider is to be regarded as having become aware when, after such an initial assessment, it has a reasonable degree of certainty that a serious incident has occurred. The point in time at which a provider can be considered to be aware will therefore depend on the circumstances of the specific incident. In some cases it will be clear from the outset; in others it may take some time to establish whether the system was involved. The emphasis should be on prompt action to carry out the initial assessment, particularly where the incident may indicate a significant risk, and, where the conditions are met, to take remedial action and to notify.

154. Because the obligation is triggered by awareness, the AI Act does not require the retroactive reporting of serious incidents of which the provider had already become aware before the date on which Article 73 became applicable to it. By contrast, where the provider was aware of an event before that date but did not, at that time, have a reasonable degree of certainty that it constituted a serious incident, and subsequently acquires that certainty, the obligation applies.

155. Article 73 applies without prejudice to the obligation of providers of general-purpose AI models with systemic risk, under Article 55(1), point (c), to keep track of, document and report without undue delay to the AI Office and, as appropriate, to national competent authorities, relevant

[11]Commission Implementing Regulation (EU) 2024/2690 of 17 October 2024 laying down rules for the application of Directive (EU) 2022/2555, OJ L, 2024/2690, 18.10.2024; European Data Protection Board, Guidelines 9/2022 on personal data breach notification under the GDPR.

information about serious incidents and possible corrective measures to address them. It also applies without prejudice to incident notification obligations arising under other Union law, including Directive (EU) 2022/2555 and Regulation (EU) 2016/679.

Figure 22 — Information flow along the value chain

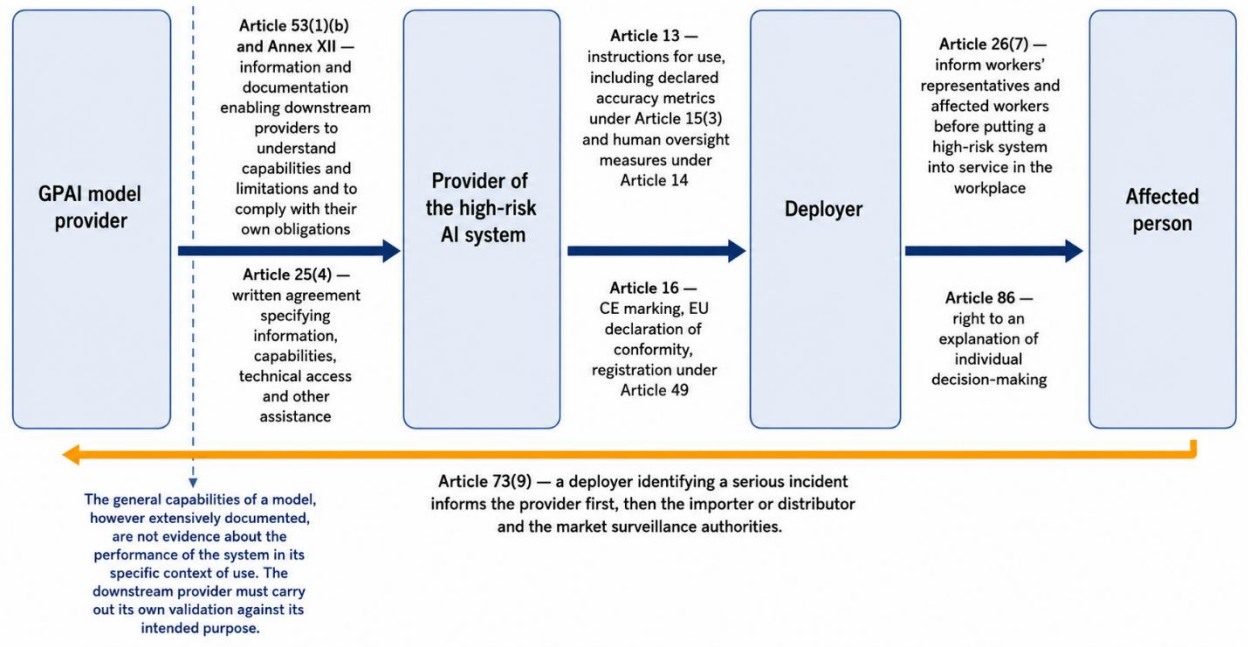

Figure 23 — Worked use case: customer-service assistant

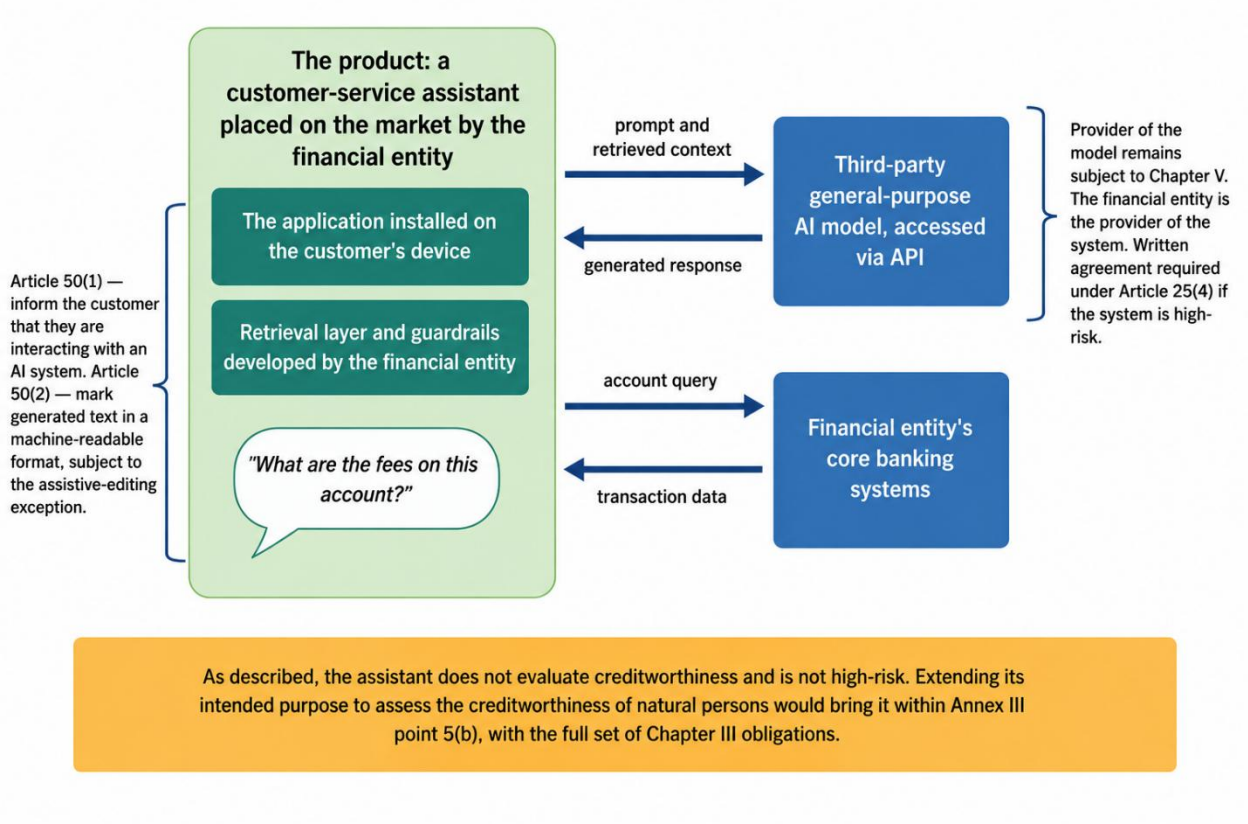

156. Following a serious incident, providers should also consider the extent to which affected deployers and other operators need to be informed. In line with the risk-based approach of the Regulation, the provision of information about a serious incident does not imply that all technical details must be made public or disclosed indiscriminately. Where appropriate, and in light of the nature of the system, of the affected users and of the potential impact, providers may limit the disclosure of detailed information to the deployers concerned, particularly where broader disclosure could itself increase risk. Once the underlying cause has been adequately addressed, broader disclosure may be appropriate, for example where it contributes to raising awareness or enables deployers to verify that their systems are no longer affected.

9.2 Transparency obligations

157. Article 50 lays down transparency obligations that apply irrespective of the risk classification of the system concerned. They apply cumulatively with, and not as an alternative to, the obligations arising under Chapter III where the system is also high-risk.

9.2.1 Systems interacting with natural persons

158. Article 50(1) requires providers to ensure that AI systems intended to interact directly with natural persons are designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system, unless this is obvious from the point of view of a natural person who is reasonably well-informed, observant and circumspect, taking into account the circumstances and the context of use. That obligation does not apply to AI systems authorised by law to detect, prevent, investigate or prosecute criminal offences, subject to appropriate safeguards for the rights and freedoms of third parties, unless those systems are available for the public to report a criminal offence.

159. The exception for cases in which the AI character of the interaction is obvious is to be applied restrictively and by reference to the specific context. Familiarity with a category of technology among some users does not make the character of a particular interaction obvious to all users, and providers should have particular regard to interactions with vulnerable persons and with persons under the age of 18.

Figure 24 — Serious incident reporting clocks

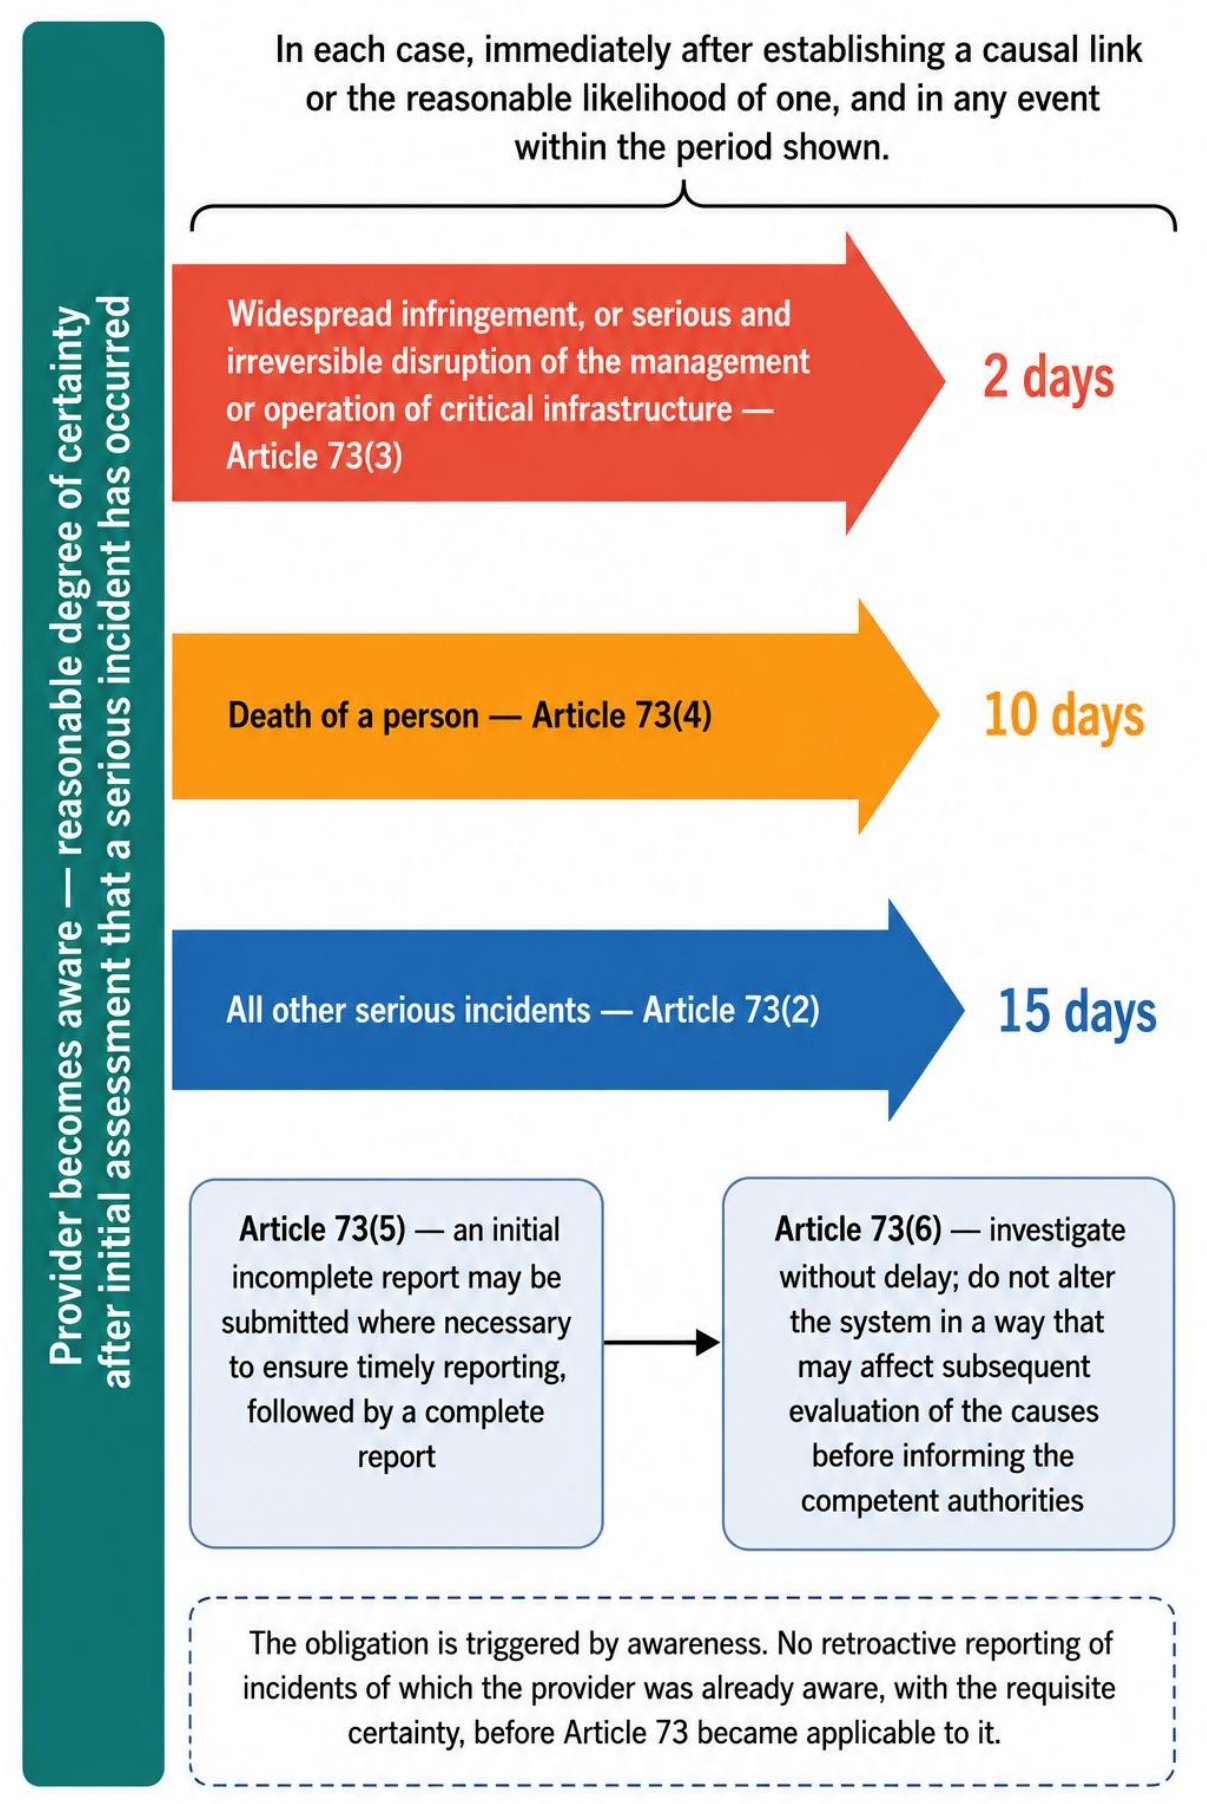

9.2.2 Marking of synthetic content

160. Article 50(2) requires providers of AI systems, including general-purpose AI systems, that generate synthetic audio, image, video or text content, to ensure that the outputs of the AI system are marked in a machine-readable format and detectable as artificially generated or manipulated. Providers are required to ensure that their technical solutions are effective, interoperable, robust and reliable as far as this is technically feasible, taking into account the specificities and limitations of various types of content, the costs of implementation and the generally acknowledged state of the art, as may be reflected in relevant technical standards. That obligation does not apply to the extent that the AI systems perform an assistive function for standard editing, or do not substantially alter the input data provided by the deployer or the semantics thereof.

161. The obligation in Article 50(2) attaches to the provider of the system and is therefore distinct from the obligations in Article 50(4), which attach to deployers. The two operate together: the provider is responsible for ensuring that the output carries a machine-readable marking, while the deployer of a system that generates or manipulates image, audio or video content constituting a deep fake is responsible for disclosing that the content has been artificially generated or manipulated, and the deployer of a system that generates or manipulates text published with the purpose of informing the public on matters of public interest is responsible for disclosing that the text has been artificially generated or manipulated, subject to the exceptions set out in that provision.

162. Article 50(3) requires deployers of emotion recognition systems or of biometric categorisation systems to inform the natural persons exposed to them of the operation of the system and to process personal data in accordance with the applicable Union data protection law. Article 50(5) requires the information referred to in Article 50(1) to (4) to be provided to the natural persons concerned in a clear and distinguishable manner at the latest at the time of the first interaction or exposure, and to conform to the applicable accessibility requirements. Under Article 50(7), the AI Office is to encourage and facilitate the drawing up of codes of practice at Union level to facilitate the effective implementation of the obligations regarding the detection and labelling of artificially generated or manipulated content.

Example 73: A provider places on the market a system that produces synthetic voice recordings from text. The provider embeds a machine-readable watermark in the generated audio and publishes documentation allowing detection tools to identify it. A deployer uses that system to produce a recording that impersonates an identifiable public figure. The deployer is required, under Article 50(4), to disclose that the content has been artificially generated or manipulated. The provider's compliance with Article 50(2) does not discharge the deployer of that obligation, and the deployer's disclosure does not discharge the provider of the marking obligation.

Figure 25 — Who owes what under Article 50

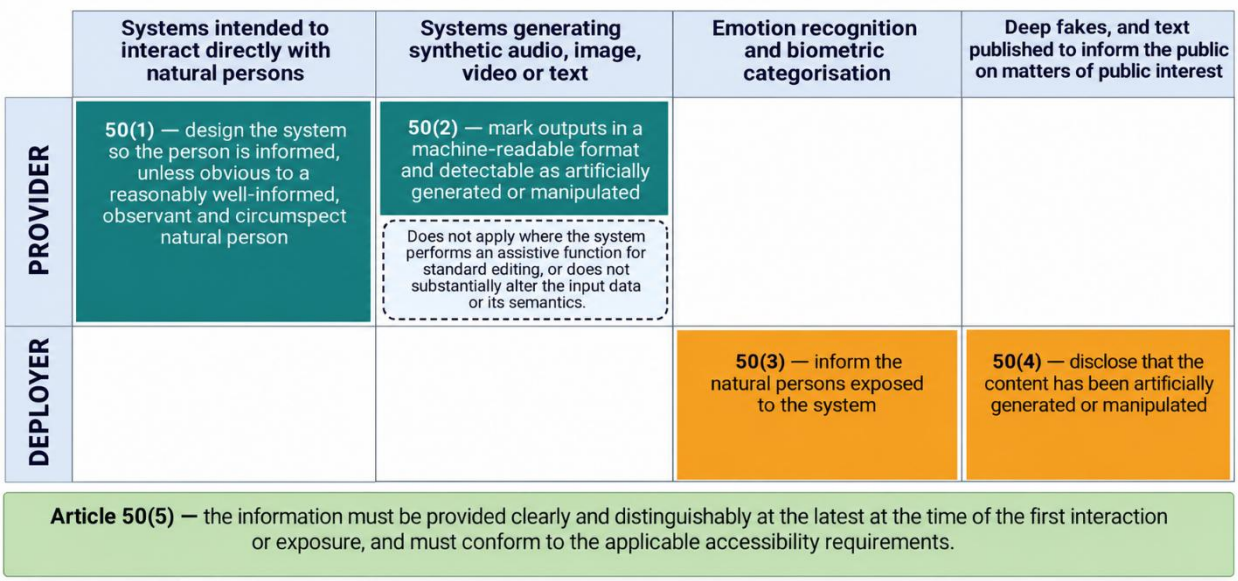

9.3 Interplay with other legislation

9.3.1 Regulation (EU) 2024/2847 (Cyber Resilience Act)

163. A high-risk AI system may simultaneously be a product with digital elements within the scope of Regulation (EU) 2024/2847. In such cases both instruments apply, but the Union legislature has provided for an explicit bridge in order to avoid duplication. Article 12 of Regulation (EU) 2024/2847 provides that high-risk AI systems that are products with digital elements falling within the scope of that Regulation and that fulfil the essential cybersecurity requirements set out in Part I of its Annex I are to be deemed to comply with the cybersecurity requirements set out in Article 15 of the AI Act, in so far as those requirements are covered by the EU declaration of conformity, or parts thereof, issued under that Regulation.

164. That bridge is limited in two respects. First, it covers only the cybersecurity requirements of Article 15 and does not extend to the accuracy and robustness requirements of that Article, nor to any other requirement of Chapter III, Section 2. Second, it operates only in so far as the requirements concerned are actually covered by the declaration issued under Regulation (EU) 2024/2847. Providers should therefore identify, in the technical documentation drawn up under Article 11 of the AI Act, which cybersecurity requirements are addressed by that route and which are addressed by other means. In particular, the vulnerabilities specific to AI referred to in Article 15(5), such as attacks seeking to manipulate the training dataset or inputs designed to cause the model to make a mistake, will in many cases require measures beyond those addressed under Regulation (EU) 2024/2847.

Figure 26 — Interplay with other Union legislation

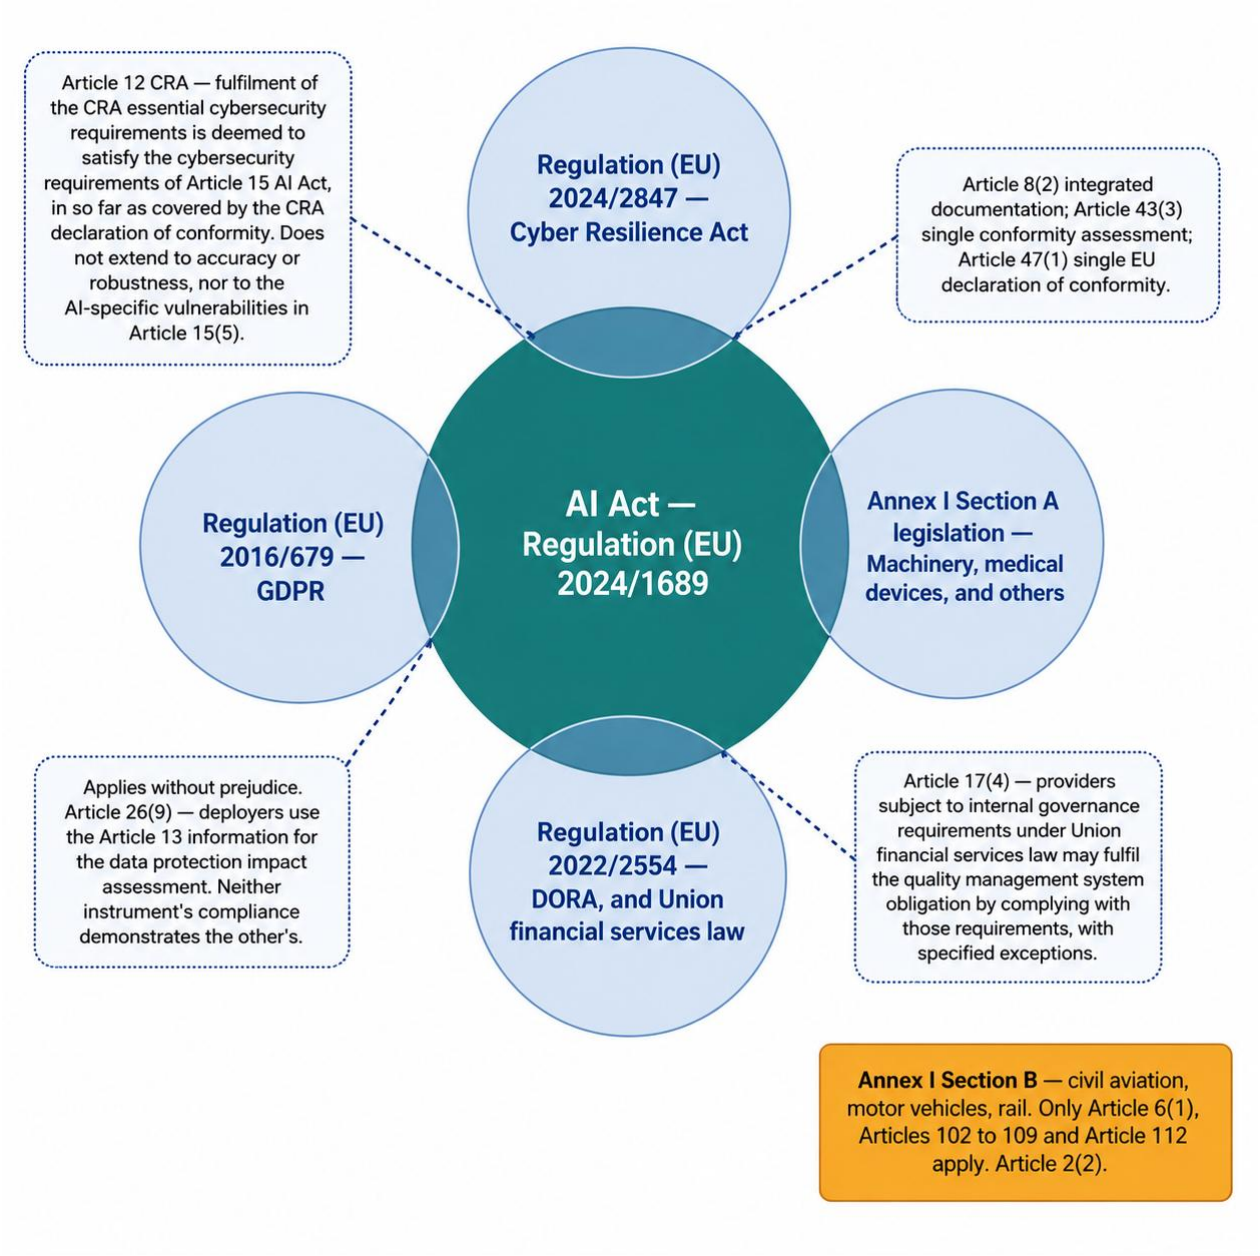

9.3.2 Union harmonisation legislation listed in Annex I

165. Where an AI system is a safety component of a product, or is itself a product, covered by the Union harmonisation legislation listed in Section A of Annex I, a number of provisions operate to integrate the two regimes. Article 8(2) provides that, in order to ensure consistency, to avoid duplication and to minimise additional burdens, providers may choose to integrate, as appropriate, the necessary testing and reporting processes, information and documentation they provide with regard to their product into documentation and procedures that already exist and are required under that legislation. Article 43(3) provides for a single conformity assessment procedure. Article 47(1) provides for a single EU declaration of conformity containing all the information required to identify the Union harmonisation legislation to which the declaration relates.

166. A different arrangement applies to the Union harmonisation legislation listed in Section B of Annex I, which covers, among others, civil aviation, motor vehicles and rail. In accordance with Article 2(2), only Article 6(1), Articles 102 to 109 and Article 112 apply to AI systems falling within the scope of those acts. The requirements of the AI Act are instead to be taken into account, where relevant, when adopting delegated or implementing acts pursuant to those acts.

9.3.3 Other Union law

167. The AI Act applies without prejudice to Union law on the protection of personal data, and in particular to Regulation (EU) 2016/679, Regulation (EU) 2018/1725 and Directive (EU) 2016/680. Compliance with the AI Act does not dispense a controller from carrying out a data protection impact assessment where one is required, and Article 26(9) expressly provides that, where applicable, deployers of high-risk AI systems are to use the information provided under Article 13 to comply with their obligation to carry out such an assessment. Conversely, compliance with data protection law does not in itself demonstrate compliance with the AI Act.

168. Similarly, the AI Act applies without prejudice to Regulation (EU) 2022/2065, to Directive (EU) 2022/2555, and, for financial entities, to Regulation (EU) 2022/2554, whose requirements concerning information and communication technology risk management continue to apply. In the financial services sector, Article 17(4) of the AI Act allows providers subject to requirements regarding internal governance, arrangements or processes under Union financial services law to fulfil the obligation to put in place a quality management system, with the exception of certain specified elements, by complying with those requirements.

169. The Commission may, in accordance with Article 96, issue further guidance addressing the interplay between the AI Act and these instruments in greater detail.