

AI Standards Lab feedback to the draft guidelines on the classification of high-risk AI systems

July 2026.

This document contains the feedback and suggestions that AI Standards Lab submitted to the electronic form in the consultation <https://digital-strategy.ec.europa.eu/en/consultations/targeted-consultation-draft-guidelines-classification-high-risk-artificial-intelligence-systems> .

1. Structure of the document

Each response in this document follows a consistent structure:

1. At the start of each 2.x.y sub-sub-section heading, (between brackets) we identify the paragraph(s) of the draft guidelines we are giving the feedback on. The paragraph numbers are followed by a brief descriptive text in the heading.
2. Where a paragraph is particularly relevant, we reproduce its text so the comment can be read without consulting the guidelines in parallel.
3. Where we raise several distinct points on the same paragraph, we separate them into subsections ("Issue 1: ...", "Issue 2: ..."), each with its own recommendation immediately below.
4. Each response closes with our comment and, when relevant, we provide a specific change in language for the guidelines.

At the end in section 2,5, we reproduce the text of our 'yes' answers to the 'Questions in relation to Article 112'

2. Feedback

2.1. High-risk classification according to Article 6(1) – Annex I (Section III of the draft Guidelines)

2.1.1. (21) Pre-empt the Omnibus

(21) Consequently, not all AI systems that are components of regulated products, or that are
--

themselves regulated products, are high-risk AI systems. Only a sub-set of such products qualify as such namely those that fulfil the classification criteria of Article 6(1) AI Act, and only those products are subject to the high-risk requirements of the AI Act. Consumer products, if they fall into one of the Union harmonisation legislations listed in Annex I AI Act, are not excluded from the scope of Article 6(1) AI Act. Provided all the conditions of Article 6(1) AI Act are fulfilled, both professional and consumer products could be classified as high-risk AI systems. At the same time, many consumer AI systems, such as smart home appliances, will not meet one or more of the cumulative conditions in Article 6(1) AI Act and therefore will not be classified as high-risk, as the analysis below demonstrates.

The last part of paragraph 21 reads that ‘At the same time, many consumer AI systems, [...]’ This sentence pre-announces an interpretation of Article 6(1) for connected home appliances that has only become true after the amendments made to it by the Digital Omnibus – see our comments on para (30) and para (50). If the Office wants to update the guidance to also cover the effects of the Omnibus, then this sentence can stay, otherwise we strongly recommend it is deleted.

2.2. III.2.1 Component of a product or itself a product

2.2.1. (30) general comments on classification

2.2.1. AI systems that are themselves products covered by Union harmonisation legislation

(30) An AI system is the product itself where it is independently placed on the market, has its own intended purpose, and is directly regulated by the Union harmonisation legislation listed in Annex I AI Act. Such systems can include stand-alone AI systems, or AI systems that are embedded in other products. For example, under Regulation (EU) 2023/1230 (the Machinery Regulation), the definition of machinery-related products explicitly includes certain software, which may therefore itself be a regulated product under that Regulation and therefore be classified as high-risk pursuant to Article 6(1) AI Act, provided it is required to undergo a third-party conformity assessment. At the same time, many consumer AI systems, such as smart home appliances, will not meet one or more of the cumulative conditions in Article 6(1) AI Act and therefore will not be classified as high-risk, as the analysis below demonstrates.

Answer to question A:

This is a very short section, and it fails to adequately explain how article 6(1) works for products that are themselves AI systems. The result of this failure here, and of paragraphs like (21), (47), (49) and (59) elsewhere, is that many readers may draw the wrong conclusion about whether their AI based product is in fact a high-risk AI system. This would be a serious problem. We recommend adding extra discussion, see our answer to question B. On a historical note: we consider the lack of such a third test to be a serious flaw in the drafting of article 6(1) the AI Act,

we have given feedback about this during the original AI Act Trilogue and also during the Digital Omnibus process, but nothing was done with this feedback by the legislator. So what we can ask of the AI Office at this point is merely to write guidance that explains correctly how the 6(1) classification mechanism actually works. In the section 1 part of the earlier July 2025 consultation that led to the drafting of this guidance document, we also provided consultation input that asked the AI office to clearly explain how the 6(1) mechanism works for products, so we are somewhat disappointed to see that this point was not taken up.

Answer to question B

With respect to expanding on paragraph 30 with an example:

To give an example (taken from paragraph 47): say a product that is a toy: a connected toy that contains an AI system that recommends music. This product is subject to the toy regulation, according to section 2.3 of the draft guidance it also meets the test in section 'The product is required to undergo a third-party conformity assessment' as meant in article 6(1) of the AI Act. We also have that the toy, as a product, is itself an AI system, according to the definition of AI system in the AI Act.

We will now cite the article 6(1) selection logic, with parts before the 'or' in (a) and (b) replaced by [...]:

6(1) Irrespective of whether an AI system is placed on the market or put into service independently of the products referred to in points (a) and (b), that AI system shall be considered to be high-risk where both of the following conditions are fulfilled:

- (a) [...], or the AI system is itself a product, covered by the Union harmonisation legislation listed in Annex I;
- (b) [...], or the AI system itself as a product, is required to undergo a third-party conformity assessment, with a view to the placing on the market or the putting into service of that product pursuant to the Union harmonisation legislation listed in Annex I.

The toy, as an AI system that is also a product, satisfies both the parts behind the 'or' in 6(1)(a) and 6(1)(b), so this means that it is classified as a high-risk AI system.

Some considerations about including this example:

Some paragraphs in the guidance, for example (21) and (59) which cites a recital in the toy directive, seem to imply that for products that are themselves AI systems, article 6(1) adds a third test beyond (a) and (b), the test whether any AI component inside the product is actually a safety component. The above example toy would fail that third test. However, no such third test is present in article 6(1), and the guidance should make this very clear.

We would be delighted if the AI Office could construct, inside of its guidance, some kind of legally solid argument about how a third test for products that are themselves AI systems is in fact present in article 6(1), but we have failed to construct such an argument ourselves. It might be able to do so by leaning very heavily on recital 15 of the toy directive. Barring that, we recommend inclusion in the guidance of the above example and review of the article 6(1).

2.2.2. Safety components

2.2.3. (32) Update definition to Omnibus

We note that the Digital Omnibus has updated the definition of safety component. Would be good to cite the updated one.

2.2.4. (48) classification again, is misleading

Issue 1: This paragraph says that certain products are not automatically classified as 'high-risk', but then proceeds to give no evidence for this broad claim, instead it gives evidence for the claim that they would not be classified as 'high-risk AI systems which are safety components'. They may still be classified by high risk as 'products themselves', in fact many toys would be. A fix would be to just delete the first sentence.

Issue 2: The digital omnibus has the effect, we believe, of avoiding the high-risk classification for connected products like most smart home appliances, but only via the new exception paragraph 1c added to article 6. See our more extensive comments on paragraph (50). Please fix the presentation here to make correct claims about the AI Act's logic for connected products, see also our comments on paragraphs (30) and (50).

2.3. III.2.3 Third-party conformity assessment

2.3.1. (50) Comment specific to radio in the omnibus - (50) clarify the effects of new Article 6(1c)

Some screenshots for reference:
From AI Act:

Article 6

Classification rules for high-risk AI systems

1. Irrespective of whether an AI system is placed on the market or put into service independently of the products referred to in points (a) and (b), that AI system shall be considered to be high-risk where both of the following conditions are fulfilled:

- (a) the AI system is intended to be used as a safety component of a product, or the AI system is itself a product, covered by the Union harmonisation legislation listed in Annex I;
- (b) the product whose safety component pursuant to point (a) is the AI system, or the AI system itself as a product, is required to undergo a third-party conformity assessment, with a view to the placing on the market or the putting into service of that product pursuant to the Union harmonisation legislation listed in Annex I.

From Omnibus:

(8) in Article 6, the following paragraphs are inserted:

- ‘1a. For the purposes of this Regulation, including paragraph 1 of this Article, AI systems that are solely used for non-safety related aspects of user assistance, performance optimisation, service efficiency, automation or convenience or quality control shall not qualify as safety components.
- 1b. Notwithstanding paragraph 1a, AI systems the failure or malfunctioning of which would endanger health and safety shall qualify as safety components.
- 1c. A product that is required to undergo a third-party conformity assessment solely due to risks other than risks to health and safety, in particular risks relating to the distribution of radio spectrum or electromagnetic interference that do not affect health and safety, shall not be considered as fulfilling the condition in paragraph 1, point (b).’;

We strongly recommend that language is added to these guidelines that confirms and clarifies the legal situation for products subject to the radio equipment directive, the situation as it exists based on the Digital Omnibus updates to the AI Act. See our answer to question B below for suggested language. The situation has changed based on adding a new part to article 6(1): ‘1c. A product that is required to undergo a third-party conformity assessment solely due to risks other than risks to health and safety, in particular risks relating to the distribution of radio spectrum or electromagnetic interference that do not affect health and safety, shall not be considered as fulfilling the condition in paragraph 1, point (b).’ The exact legal consequences of this new part are tricky to trace, so we recommend that the guidance states both the consequences clearly, and cites the specific articles that create these consequences.

To reflect changes the to Act with the Digital Omnibus adding an article 6(1c), We recommend that the guidance adds text that says roughly the following:

(start of suggested text 1) Based on an exception defined in the AI Act Digital Omnibus package amending the AI Act, several products that include wireless networking or communication functions, so that they are subject to Directive 014/53/EU ('the radio equipment directive'), do not automatically qualify as high-risk AI systems.

This is because, while Article 17(2) of Directive 014/53/EU triggers an assessment based by either a notified body, or on internal control based on harmonised standards (see paragraph 57), this assessment is only against the Article 3(2) and 3(3) requirements in Directive 014/53/EU, and these requirements fit the criterion that they trigger an assessment 'solely due to risks other than risks to health and safety'. (end of suggested text 1).

The omnibus paragraph 1c creates an ambiguity: it is clear that it defines an exclusion for products, but less clear if it also defines an exclusion for safety components of products. We believe that the disambiguation most in line with the intent of the AI Act is that it does not in fact define an exclusion for safety components. We therefore recommend that the following text is also added to the guidance:

(start of suggested text 2) While the exception in the Digital Omnibus package excluded certain products from fulfilling the condition of article 6)(1)(b), it does not exclude AI systems that are safety components in such products from that condition. These should be treated as high-risk AI systems. (end of suggested text 2).

2.3.2. (55) Misleading paragraph

The rule of thumb mentioned here seems not to apply to the logic of REGULATION (EU) 2025/2509 on toys, where design complexity does not play a role in the logic for selecting for module A. The language here is therefore somewhat misleading as to the structure of at least some of the regulation in annex I A. We recommend deleting this paragraph entirely, as it would create more confusion than insight for the reader.

General comment: If the EC wants to make an explanatory augment on how a product using an AI function is never low complexity, and that therefore it is appropriate to overrule, for the purpose of the AI Act, any exception on that standard related to using module A in combination with harmonised and cited standards, it should make that argument more directly.

2.3.3. Paragraphs 56-58

(56) Union harmonisation legislation allows the use of Module A for assessment of conformity of certain products. However, in some cases use of Module A is conditioned on the application of harmonised standards, such as in the Regulation (EU) 2025/2509¹² ('the Toys Safety Regulation'), the Machinery Regulation, or the Radio Equipment Regulation. In these cases the use of module A for certain products or certain aspects is allowed, but only where the manufacturer applies harmonised standards published in the Official Journal that cover all

relevant health and safety requirements. The mandatory application of such harmonised standards as a condition for the application of Module A, for certain products, ensures that those products are subject to an appropriate level of scrutiny and that internal control alone is not relied upon where additional safeguards of public interest are necessary.

(57) The fact that the Union harmonisation legislation may allow a manufacturer to rely on internal control based on harmonised standards, as one procedural option, does not affect the classification of an AI system as high-risk under Article 6(1) AI Act. The choice of conformity assessment module provides manufacturers with procedural flexibility for demonstrating compliance, but it does not confer discretion on the manufacturer to determine the risk classification of an AI system for the purposes of the AI Act.

(58) In line with the legislative intent of the AI Act, the decisive factor for the classification of an AI system as high-risk under Article 6(1) AI Act is that the product, according to the Union harmonised legislation listed in Annex I AI Act, is subject to enhanced regulatory scrutiny before it can be lawfully placed on the market or put into service in the Union, due to its potential impact on public interests, such as health, safety and the environment. Such scrutiny is ensured through the requirement of a third-party conformity assessment or equivalent mechanisms, including internal control subject to mandatory application of harmonised standards published in the Official Journal. As emphasised in the Blue Guide, the 'protection of the public interest' should be a key consideration for the conformity assessment and the selection of applicable conformity assessment modules by the Union legislature.

2.3.4. (56) Issue 1: Safety component CA

Paragraph (31) says that an AI system is a safety component where it "is, or is intended to be, part of" a regulated product meaning that it can be a de facto safety component even if not declared, fitting the post-omnibus definition of "safety component" under Article 3(14): a component whose failure endangers health and safety is a safety component even if not identified as one by the manufacturer. If the manufacturer identifies it as such, the framework operates naturally. Where such a system is integrated and marketed under the product manufacturer's name, Art 25(3) allocates provider obligations to that manufacturer. The only pre-market checkpoint the guidelines identify is evaluation within the product's overall safety assessment (p31); yet p56–58 accept internal control under harmonised standards as meeting the CA condition, and under Art 43(3) that same manufacturer governs the AI assessment itself. No external examiner need ever be involved. Who identifies a de facto safety component the manufacturer does not declare? Before placement, the determination is made and reviewed by the same interested party.

2.3.5. (56) Issue 2: Misleading footnote

Issue 2: In footnote 11 the use of the choice of the word 'under' in the starting text 'For example, under Article 3(1) (a) of' is somewhat misleading, it is better to say 'for', as in 'For example, for Article 3(1) (a) of'. The article UNDER which this is allowed FOR 3(1)(a) is actually article 17(2).

2.3.6. (57) Issue 1: does not affect classification': possibly disproportionate

The first line about 'may allow does not affect the classification' is quite extraordinary, at least one of us was under the impression that article 6(1) implies that this allowance WOULD affect classification. This clarification may have the effect, for several I A verticals, of seriously increasing the number of products that will get high-risk classification.

Recommendation: We are unable to determine, because we do not have the necessary per-legislation expertise, if this effect would meet a proportionality test. If a certain vertical complains loudly in this consultation that the impact of this rule is too large for them, we recommend that the AI Office listens carefully.

2.3.7. (57) Issue 2: Clarifying Article 43(3) opt-out mechanism

The paragraph talks about opting out of conformity assessment. This is covered in AI Act Article 43(3), third paragraph. However, that third paragraph can be read to mean one of two things: 1) if the provider opts out of the Section A regulation specific third-party assessment based on applying harmonised standards, this means they also have the right to opt out of notified third party assessment for the AIA, as long as they apply the AIA harmonised standards and annex VI, or... If the provider does not want to apply the AIA harmonised standards or common specifications, they also lose the right to opt out of the Section A regulation specific assessment. It would be good if the guidance would clarify this, where we believe that option 1) is the best disambiguation, but maybe both 1) and 2) are intended to apply simultaneously.

2.3.8. (59) Highlight the product-itself route more strongly

This paragraph correctly concludes that 'On this basis, AI systems that are [...] themselves products covered by the Toys Safety Regulation, in accordance with Article 26(2) and Article 26(3) of that regulation, are classified as high-risk pursuant to Article 6(1) AI Act.' However, this aspect of 6(1) is easy to overlook. In comments on several other paragraphs we recommend to add additional text to the guidance to highlight this effect of 6(1) more strongly.

2.3.9. (60-62) Need for detailed implementation guidelines

3.1. (60) AI systems classified as high-risk under Article 6(1) AI Act as regards products covered by the Union harmonisation legislation listed in Section A of Annex I are subject to the requirements for high-risk systems listed in Section 2 of Chapter III AI Act. In contrast, for AI systems classified as high-risk AI systems under Article 6(1) AI Act as regards products covered by the Union harmonisation legislation listed in Section B of Annex I, only Article 6(1), Articles 102 to 109, and Article 112 AI Act apply

3.2. (61) Moreover, to reduce the compliance burden for economic operators of AI systems classified as high-risk under Article 6(1) AI Act, Article 8(2) AI Act on the interplay with the sectoral legislation, Article 9(10) AI Act on risk management, and Article 17(3) AI Act on quality management allow economic operators to add, when necessary and appropriate, an assessment of AI-specific risks into already existing risk and quality management systems. Article 40 AI Act further requires consistency of AI Act harmonised standards with standards developed under the Union harmonisation legislation listed in Annex I AI Act.

3.3. (62) Such mechanisms provided in the AI Act enable economic operators to comply with the requirements of the AI Act and Union harmonisation legislation within a single compliance framework, thereby avoiding any possible duplication of compliance efforts while ensuring a high level of protection of health, safety and fundamental rights of AI systems covered by Article 6(1) AI Act.

This section lacks clear guidelines and practical examples on how companies can actually implement art. 8(2), 9(19), 17(3) & 40. It is important to expand these paragraphs with actual examples and even templates if possible. There is no guidance currently available on how to actually do these in practice, which would likely materialise on companies doing 2 separate analyses or maintain certain duplicities to avoid legal problems. We suggest that pointers are added about where and when more information that helps these parties reduce their compliance burden will be available.

2.4. High-risk AI systems in the areas of Annex III

2.4.1. (203-206) confirm power-rerouting AI as a safety component — proposed example

From our reading of paragraphs 203 to 206, we understand that AI systems that automatically reroute power should fall within the "safety component" definition and therefore within point 2 of Annex III. We consider this clarification important: in a shortage, such systems can determine that the available energy is directed to the most critical places, such as hospitals. Such AI systems satisfy both elements of the post-omnibus definition of "safety component" in Article 3(14) AI Act. First, their intended purpose within the meaning of Article 3(12) AI Act as determined by the provider, in this case, preventing large-scale physical harm by deciding, in real time, where available power is directed. Second, their failure or malfunctioning would directly endanger the health and safety of persons and property. It is also consistent with paragraph (31) of the Annex I section of the draft guidelines, where it is, or is intended to be, part

of the infrastructure whose integrity it protects. The classification therefore turns on the factual impact of the AI system: whether its malfunction would cause damage or pose a risk to health.

We ask the Commission to add the following example to the practical examples under (a) following paragraph 206:

"An AI system that is intended to make or pre-compute a decision, taken within milliseconds and without human oversight, to reroute power to another part of the electricity grid or to change the load routed into a DC transmission line."

2.4.2. (449 - 450) On the lack of a definition for operators

We have a comment on section IV paragraph (449) and section V paragraph (450) that we will put here for lack of a better place to put it:

Paragraph 449 uses the term 'operators' without defining it. The AI Act defines that "operator" means a provider, product manufacturer, deployer, authorised representative, importer or distributor. The normal English term operator might be interpreted as being a user or deployer who operates an AI system. This therefore creates an opportunity for readers to misinterpret the text.

We recommend to either just remove the words 'to operators', or to add the definition to the text. The word 'operators' is used again in paragraph 450, please also fix the same issue there.

2.5. Questions in relation to Article 112 paragraph 1 (Evaluation and Review)

Question 1. Are you aware of concrete use cases of AI systems that in your opinion need to be added to the list of use cases in Annex III - within the existing 8 areas - to ensure that all AI systems having equivalent risk of harm (in accordance with the criteria laid down pursuant to Article 7(2) AI Act) are classified as high-risk under the Annex III AI Act?

If yes, please specify why you consider that this use case(s) should be classified as high-risk taking into account the criteria in Article 7(2) and provide evidence that justify such a change.

Yes.

We propose to add to area 3 (Education):

(e) AI systems intended to be used to increase or preserve student engagement with an educational or training task, via methods including AI based or AI supported gamification, AI based determination of rewards or the timing of rewards, and AI based persuasive feedback.

and add to area 4 ('Work'):

(c) AI systems intended to be used to increase or preserve worker engagement with a task or line of work, via methods including AI based or AI supported gamification, AI based determination of rewards or the timing of rewards, and AI based persuasive feedback.

Both target audiences, students and workers, that cannot opt out of the systems set over them. AI-driven gamification and reward timing exploit the same psychological mechanisms as gambling design, risking addiction, driving them into addictive patterns and harming mental health.

We propose to add to area 5 (services) the following:

(e) AI systems intended to be used by vulnerable persons to apply for, access, or enjoy essential private services, essential public services, or benefits, where these AI systems provide open-ended natural language interactions.

We propose to add to area 2 ("Critical infrastructure") the following:

(b) Agentic AI systems deployed in critical infrastructure, such as in smart energy grids, traffic and transportation, digital infrastructure (and especially the digital backbone), and water or gas systems, where these systems display long-range autonomy.

Question 4. Do you consider that the list of prohibited AI practices in Article 5 should be expanded? In particular, are there concrete examples of AI practices that, in your view, conflict with Union values and are not yet adequately addressed by the AI Act or other Union legislation? If yes, please specify your arguments and provide evidence that justify such a change.

Yes.

Highly specialised "narrow" systems can carry dangerous offensive-cyber or CBRN capabilities while falling below the compute thresholds and generality criteria that trigger GPAI systemic-risk obligations since GPAI status is a prerequisite for the systemic-risk tier. They also are not captured under article 6. For example, ESM3, a biological specialist model mapping onto the biorisk chain capabilities that can present systemic risk (<https://arxiv.org/abs/2605.01611>) yet it is somewhat ambiguous if it falls under Act. We recommend prohibiting release of models and systems with meaningful CBRN or offensive-cyber capabilities without appropriate safeguards.

We also propose prohibiting AI systems that demonstrate the capability to self-exfiltrate its weights, parameters or code to external infrastructure, or to maintain or spawn copies of itself across distributed systems without human authorisation.

UK AISI's RepliBench ([arXiv:2504.18565](https://arxiv.org/abs/2504.18565)) found frontier models can already exfiltrate their weights under simple security setups, deploy copies onto cloud compute, and write self-propagating code. Even if they cannot yet reliably chain the full sequence, AISI concludes the capability could soon emerge. Such systems would pose severe risks to health, safety and fundamental rights yet Art. 55(1)(b) applies only to GPAI models with systemic risk and Art. 14 oversight only to high-risk systems. Only a prohibition would mitigate this risk irrespective of classification.