AI Standards Lab Recommendations for the EU AI Act Digital Omnibus Trilogue

Contact person for this document: Marcel Mir, marcel@aistandardslab.org
7 April 2026.

1. Executive summary

The Council and European Parliament (EP) have now published their respective initial positions for the EU AI Act Digital Omnibus trilogue:

A. The Council position is at https://www.consilium.europa.eu/en/press/press-releases/2026/03/13/council-agrees-position-to-streamline-rules-on-artificial-intelligence/
B. The EP position is at https://www.europarl.europa.eu/doceo/document/A-10-2026-0073_EN.pdf
C. The original Commission proposal is at https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-ai-regulation-proposal

This document provides the AI Standards Lab (AISL) recommendations for the trilogue negotiations. It builds on our previous analyses of the Commission proposal and EP proposed amendments.[1] Our main concern in making these recommendations is that, while it is positive if the Omnibus can remove some administrative burdens, it should not lower the level of protection of health, safety, and fundamental rights compared to what is now offered by the EU AI Act.

Our main recommendations for the trilogue are the following. We have sorted them based on a combination of our priority, and the remaining divergence between the co-legislators.

1. We strongly welcome that both the Council and EP reinstate the registration requirement under Article 6(4). This relatively minor burden for companies is essential for proper enforcement of the regulation and oversight of AI systems in the EU.

2. We strongly welcome the Parliament's proposed Article 64(2a), which will legally require an adequate resourcing of the AI Office. If the AI Office takes on expanded enforcement responsibilities, ensuring sufficient capacity is essential.

[1] Our previous analysis of proposals is at https://aistandardslab.org/recommendations-on-the-digital-omnibus-amendments-to-the-eu-ai-act/ and https://aistandardslab.org/recommendations-on-the-european-parliament-amendments-to-the-eu-ai-act-in-the-digital-omnibus/. Our initial input to the consultation done before the Commission wrote its proposal is at https://aistandardslab.org/our-input-to-the-european-commission-on-the-digital-simplification-package-and-omnibus/.

3. On Article 75 which defines AI Office oversight powers, we welcome both the Council and EP proposals as strong improvements over the Commission's initial centralisation proposal. If centralisation with exclusivity is chosen in the trilogue, our strong recommendation to ensure that the AI Office has adequate resources through the EP's proposed Article 64(2a) becomes all the more necessary.

4. On the EP's proposal to exclude Annex I.A high-risk AI systems from many obligations, we have a weak preference for the Council's approach that maintains the current obligations. To inform the trilogue negotiations, our writeup below contains some detailed analysis of the pros and cons of the EP proposal. To avoid causing delays in standards writing for the AI Act, it is crucial that whichever position prevails, the final text must avoid ambiguities. We make specific recommendations on this below.

5. We generally support the EP's extension of value chain obligations to providers of GPAI models under Article 25, as long as these extensions are limited to exclude providers who have clearly specified that their models or systems are not intended for integration or conversion into high-risk AI systems. We also make recommendations for the sub-case of open source.

6. On Article 4a, related to bias mitigation and privacy, we stress that this addition to the AI Act is only positive under the premise that the GDPR safeguards it references are not diluted or eliminated through the parallel GDPR Omnibus procedure. We support both co-legislators in reinstating the "strictly necessary" threshold consistent with Article 10(5) of the current AI Act.

7. For the new Article 5 prohibitions, NCII and CSAM, we recommend the EP's cleaner formulation: a direct prohibition paired with a safe harbour. We specifically recommend against the Council's inclusion of "reproducing" among prohibited functions, which risks capturing legitimate operations such as cloud backups and content moderation, and against its "use of an AI system capable of" wording, which could make any user of a general-purpose system liable regardless of their actual use. We note that the Council's structural separation of NCII and CSAM merits consideration in the final text.

8. On Article 72(3), post-market monitoring, we support the shift from implementing act to guidance but back the EP's inclusion of a template and its earlier deadline of February 2027 to ensure providers have clear reference in time.

Our full analysis and recommendations cover additional points not included in the above executive summary. See the table of contents below for all proposed changes to the AI Act we cover, sorted by Article number.

1.1. Table of contents

2. Discussion of proposed amendments to the AI Act

2.1. On Articles 2(2), 110(a)-(l), and Annex I.A: Removal of Annex I.A obligations

On the issue of changing Annex I.A, we see a significant difference between the Council and the EP proposals going into the trilogue.

1. The Council (and Commission) propose to leave the current situation for Annex I.A high-risk AI systems as is, except for requesting via a new Article 96(1)(g) that that the Commission writes some additional guidelines
2. The EP proposes that Annex I.A high-risk AI systems will no longer be subject to the high-risk AI requirements in the AI Act (articles Articles 8-15, or Section III Chapter 2), while also creating the possibility for the Commission to add some of these requirements, at a later time, to Annex I.A sectoral safety legislation.

We have a weak preference for the Commission stance, but only a weak preference. Our main recommendation to the parties entering the trilogue is that **whatever they decide, they should make sure to avoid ambiguities in the final text.**

To support a high quality discussion between participants in the trilogue who have to make choices, we include several subsections with analysis below. We also include further subsections with specific proposals for a middle ground, and with recommendations for avoiding ambiguities in the final language. A specific concern is that the trilogue result should avoid creating ambiguities that will negatively affect the speed of the work on the AI Act standards that are now being written in CEN-CENELEC JTC21.

2.1.1. Analysis: impact of article 6(1) overclassification on administrative burdens

According to our analysis in the AI Standards Lab (AISL), the current Article 6(1) of the AI Act is very badly written. We have often communicated that Article 6(1) over-classifies certain products (e.g. smart phones, connected smart laundry machines) as high-risk AI systems, and thereby also creates unnecessary administrative burdens.

We have communicated this finding and recommendations to fix the problem at various earlier stages of the digital Omnibus process[2]. Neither the Commission, the Council, or the EP has used our information to propose amendments that would elegantly fix the overreach problems in Article 6(1) that we identified. However, the EP has made a proposal for a more radical fix, one that removes the entire Annex I.A.

[2] We communicated to the Commission in October 2025 (see section 2.6 of the PDF contribution we link to in https://aistandardslab.org/our-input-to-the-european-commission-on-the-digital-simplification-package-and-omnibus/ , and to the EP and the Council in January 2026, see section 2.4 of the PDF at https://aistandardslab.org/recommendations-on-the-digital-omnibus-amendments-to-the-eu-ai-act/.

Given that it is too late for more elegant solutions, we recommend that the trilogue negotiations focus only on considering the choice between the EP proposal and the Council proposal.

2.1.2. Analysis: benefits and drawbacks of the horizontal approach chosen for the AI Act and required of for its standards

The AI Act was drafted to choose a 'horizontal' approach to regulating the safety of AI technology when used in Annex I.A products classified as high-risk AI systems. The EP amendments, in practice, reject this choice, restoring a 'vertical' approach where the safety regulation and related standards for Annex I.A products are each individually updated instead.

To discuss the impact of the horizontal approach, it is useful to consider a specific case where the creation of an AI-based product is regulated by the AI Act and several other Annex I.A regulations. We consider the (somewhat extreme) case of a CT scanner with AI components in the picture below. The scanner is subject to the Medical Device Regulation (MDR), as it is a medical device, the Machinery Directive (MD), as it has motor-driven moving parts, and the Radio Equipment Directive (RED), because it emits radio waves. The AI components, combined with conformity assessment triggered under the MDR and the RED, make it subject to the EU AI Act.

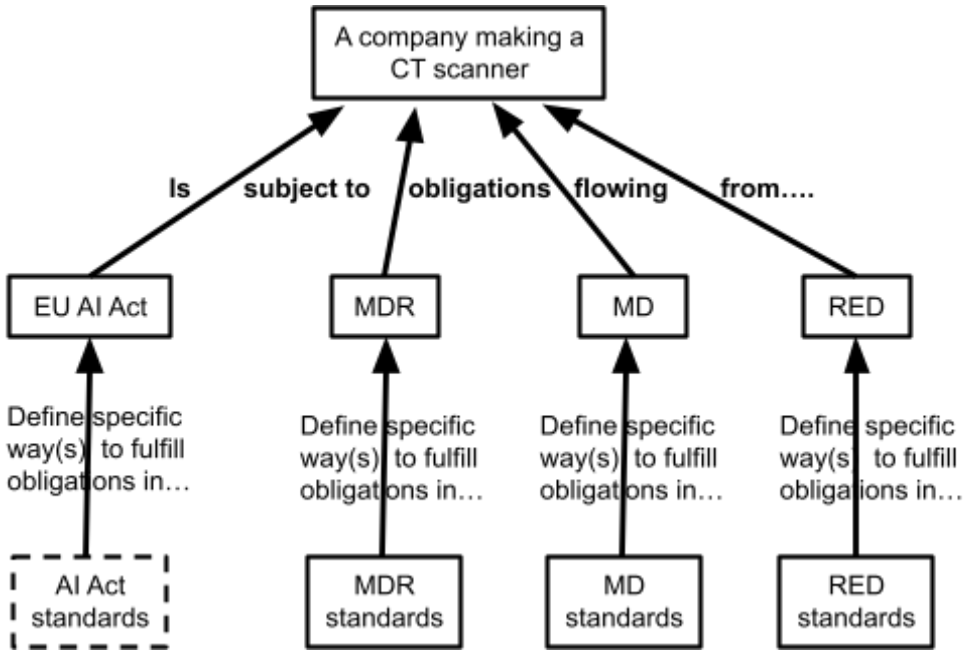

The above picture also shows how each of these regulations is expected to come with a set of standards explaining how the obligations in them can be fulfilled. The CEN-CENELEC JTC21 committee is now writing these standards for the AI Act.

A common industry concern, if multiple regulations apply to the same product, is that the regulations will each contradict each other, or be written in such a way that two separate sets of quality management process documents will have to be produced to satisfy each.

The same concern also applies to the standards coupled to these regulations: if they are not well written these might also end up having the effect of requiring two separate sets of documentations.

The AI Act acknowledges these concerns and tries to prevent such outcomes, e.g. article 9(10) states that

> *For providers of high-risk AI systems that are subject to requirements regarding internal risk management processes under other relevant provisions of Union law, the aspects provided in paragraphs 1 to 9 may be part of, or combined with, the risk management procedures established pursuant to that law.*'.

Still, the Act gives little guidance about how this combining could be done in practice. Article 40(2) tries to arrange that the above JTC21 standards make sure that they are consistent and compatible with other applicable standards :

> *When issuing a standardisation request to European standardisation organisations, the Commission shall specify that standards have to be clear, consistent, including with the standards developed in the various sectors for products covered by the existing Union harmonisation legislation listed in Annex I, [....]*

As participants in JTC21, we observe that JTC21 has been having trouble writing such 'vertical' standards that ensure consistency and avoid requirements that trigger the need for two separate sets of documentation.

The vertical nature of the AI Act, and the need to consider both Annex I.A and annex III type high-risk AI systems simultaneously, has made the work in JTC21 more difficult. This vertical nature is one of the causes why JTC21 has faced delays. We expect that the standards that will eventually be published by JTC21 will not provide a lot of detail that will help Annex I.A providers figure out how they might produce a single set of documentation to fulfill all obligations. JTC21 relies for the most part on volunteer manpower, and there has been a lack of volunteers with sufficient time and expertise.

From the above picture, it should be clear that if the EP proposal removes the flowing from the AI act for Annex I.A product providers, then this will lead to a lowering of their administrative burdens.

2.1.3. Analysis: EP proposal's impact in lowering fundamental rights protections

A downside of the EP position is that it risks creating a small negative impact on the protections to health, safety, and fundamental rights provided by the AI Act.

If we look at the main features that the AI Act adds to the regulatory landscape, they are:

1. Bans on certain AI practices (Article 5)
2. More legal clarity on how to build and deploy AI systems and models in a way that appropriately manages risks to health, safety, and fundamental rights, clarity because

the Act encodes triggers that require certain organisational and technical steps to be taken
3. A new oversight regime that defines rights and duties for new oversight bodies

Based on our knowledge of technical and market conditions, we consider that the need to introduce points 2 and 3 exists primarily towards Annex III high-risk AI systems, and towards the providers and modifiers of General Purpose AI models and systems. For these parties, nothing was defined before. For Annex I.A high-risk AI systems, existing sectoral regulation and standards already provide a great deal of clarity on point 2, and an existing oversight regime as under point 3. Therefore, the EP proposal to omit points 2 and 3 for Annex I.A high-risk AI systems, while maintaining point 1, has a relatively low impact.

As an example of the impact that the EP proposal does have, we consider the case of medical devices classified as high-risk AI systems. The current medical device regulation and its harmonised and cited standards provide a level of clarity that is equivalent to the content of, for example, Chapter III Articles 8 and 9 of the AI Act, but they do not provide legal quality equivalent to the content of Articles 10-13. Excluding medical devices from being subject to Articles 10-13 therefore creates a greater legal uncertainty in the market about what practices would be necessary and sufficient for AI-based medical devices. This uncertainty affects good-faith manufacturers, increasing their compliance costs. It may also encourage bad-faith manufacturers to cut corners, leading to a lower level of protection to health, safety, and fundamental rights. **We therefore welcome the EP proposal that compensates for these effects, if Annex I.A obligations are deleted, by empowering the Commission to integrate certain Articles taken from the AI Act into sectoral legislation.**

2.1.4. First potential middle ground in the tradeoff: stronger obligation to amend sectoral legislation

Given our above analysis of the upsides and the downsides of each position, **a potential middle ground** for reducing administrative burdens while protecting health, safety, and fundamental rights, would be for the trilogue to **build in stronger obligations, compared to the EP proposal, about when sectoral law would need to be updated**. The language in the EP proposal for the new Articles 110(a)-(l) now reads:

> The Commission is empowered to adopt delegated acts [...].

This could be strengthened to

> The Commission **shall adopt, by [date],** delegated acts […].

2.1.5. Second potential middle ground in the tradeoff: radio directive only

Another **potential middle ground would be to move only Directive 2014/53/EU (on radio equipment) to annex I.B, and leave the rest in Annex I.A.** Our analysis shows that the Radio Directive combined with Article 6(1) triggers the majority of the unnecessary burdens via mis-classification.

2.1.6. Analysis: impact of the EP proposal on the speed of standardisation

We have seen publications by several parties that speculate about what would happen to the speed of the JTC21 committee finishing the AI Act standards, in case the EP proposal is accepted. As participants in JTC21, we are bound by confidentiality, so we cannot share a detailed analysis on this question. We can however give some high level remarks.

For context, we should state that, based on our experience, we consider the standards process run by CEN-CENELEC to be not fit for purpose, when it comes to the timely, efficient, and inclusive delivery of harmonized standards for new digital or green legislation. The current system creates incentives for industry stakeholders to abstain from making volunteer contributions to the standards process, hoping that standards will be late, so that they then can argue for an extension of entry into force dates.[3] This has happened in JTC21. In our view, the announcement by the Commission of the Digital Omnibus process for the AI Act has been a contributing factor to further delaying the standards process in JTC21. That being said, none of these problems can be fixed at this late stage in the trilogue discussions.

If the EP position is accepted, we expect the following effects on the speed of the JTC21 process:

1. Negative effect (short term): JTC21 will be somewhat delayed while its participants figure out the impact of the changes to the AI Act on their work. This negative effect could be more severe if the legal text coming out of the trilogue leaves ambiguities on the status of the standards. We make specific proposals further below about how the legal text could avoid ambiguities.
2. Negative effect (longer term) : JTC21 is composed of volunteer experts, including volunteer experts who are paid by Annex I.A stakeholder organisations to participate. If the EP changes are accepted, we would expect most of these Annex I.A experts to leave, or to or seriously reduce their level of participation, which would slow down the work.
3. Positive effect (short and long term): At the same time, if these changes are accepted, there is no longer a need for the standards writing to carefully avoid inconsistencies with existing Annex I.A legislation or its standards. This will accelerate the work, as problems in this area, which are present in the current drafts, will no longer have to be found, commented on, and resolved while finalising the work.

Overall, it is hard to predict whether the negative effects will outweigh the positive effects on speed. We definitely recommend that the co-legislators seek to limit the negative effects from point 1 above.

2.1.7. Specific recommendations to avoid ambiguities: Article 2(2)

We first consider the EP amendment 27 below:

[3] We have discussed these points in more detail in our consultation input available at https://aistandardslab.org/our-input-to-the-european-commission-on-the-future-of-european-standardisation/.

Amendment 27

Proposal for a regulation
Article 1 – paragraph 1 – point 2
Regulation (EU) 2024/1689
Article 2 – paragraph 2

Text proposed by the Commission	*Amendment*
2. For AI systems classified as high-risk AI systems in accordance with Article 6(1) related to products covered by the Union harmonisation legislation listed in ***Section B of*** Annex I, only Article 6(1), Article 60a, Articles 102 to 109 and Articles 111 and 112 shall apply. Article 57 shall apply only in so far as the requirements for high-risk AI systems under this Regulation have been integrated in that Union harmonisation legislation.;	2. For AI systems classified as high-risk AI systems in accordance with Article 6(1) related to products covered by the Union harmonisation legislation listed in Annex I, only Article 6(1), Article 60a, Articles 102 to 109**, *Articles 110a-110l*** and Articles 111 and 112 shall apply. Article 57 shall apply only in so far as the requirements for high-risk AI systems under this Regulation have been integrated in that Union harmonisation legislation.;

If the trilogue adopts this change, we prefer that two clarifying parts are added to this paragraph 2, as follows:

> 2. For AI systems classified as high-risk AI systems in accordance with Article 6(1) related to products covered by the Union harmonisation legislation listed in Annex I, only **Article 5,** Article 6(1), Article 60a, Articles 102 to 109, **Articles 110a-110l** and Articles 111 and 112 shall apply. **Any presumption of conformity with the requirements of Chapter III Section 2 of this legislation, that is created by Article 40(1) via standards requested under Article 40(2), shall not apply to these high-risk AI systems.** Article 57 shall apply only in so far as the requirements for high-risk AI systems under this Regulation have been integrated in that Union harmonisation legislation.;

We believe both these clarifications would reflect the EP's intent. The proposed clarification on presumption of conformity would disambiguate matters for JTC21.

2.1.8. Specific recommendations to avoid ambiguities: Articles 110a-110l

The EP's proposed new Articles 110a-l empower the Commission to amend sectoral legislation which all contain similar language. We propose the following disambiguating clarifications (bold, underlined) in that language:

> 2. The Commission is empowered to

adopt delegated acts in accordance with
Article [number] to amend the essential health
and safety requirements set out in [location]
in order to adapt them to scientific or
technical progress or to international
developments or to add requirements in
relation to emerging risks or technologies.
When it is making such amendments for high-risk AI systems referred to in
Article 6(1) of Regulation (EU)2024/1689
the relevant requirements set out in
Chapter III, Section 2 of (EU) Regulation
2024/1689 shall be deemed **by the Commission** to constitute
essential health and safety requirements
for the purpose of this Regulation.
[...]

This extra language clarifies who shall deem these requirements essential. Without this language a legal interpretation would be possible where *all readers of the amended legislation in question* shall deem that all the requirements in Chapter III, Section 2 of (EU) Regulation 2024/1689 shall already constitute essential safety requirements for the purpose of the respective Regulation when the Omnibus legislation goes into force, so long before the Commission gets around to amending the legislation with more specific language.

2.2. On Article 4a: Processing of sensitive personal data

We support both the EP and the Council in using the more restrictive "*strictly necessary*" wording consistent with Article 10 of the current AI Act. For a detailed analysis of why we consider this change important, we refer to our previous assessments of the Commission Omnibus proposal and the EP amendments. However, it remains worth emphasising that this improvement only holds under the premise that the GDPR safeguards referenced in the article are not diluted or eliminated through the parallel GDPR Omnibus procedure currently ongoing.

2.3. On article 5a: Addition of the NCII and CSAM bans

AISL understands the rationale for the inclusion of prohibitions on AI-enabled generation of non-consensual intimate imagery (NCII) and child sexual abuse material (CSAM) in Article 5.

We recommend the EP's formulation as the starting point for drafting the final text. Its approach, a clear prohibition paired with a safe harbour for providers who implement effective safety measures, is legally robust and achieves the intended objective. The Council's formulation is more convoluted. As we detail below, certain wording choices risk being misinterpreted or producing effects the Council did not intend.

We recommend against including "reproducing" among the prohibited functions. The Council uses "generating, manipulating or reproducing" while the EP omits reproduction. Including "reproducing" risks forbidding AI systems that merely copy files without any awareness of the nature of the content of these files, as part of legitimate operations. For example, AI enhanced cloud services that generate thumbnails, compress images, or create backups technically "reproduce" user-uploaded content without any awareness of its nature.

We recommend against the Council's "use of an AI system capable of" wording. This formulation would risk making any professional user liable for simply using a system that has the capability to generate prohibited content, regardless of what they actually use it for. For instance, earlier versions of Grok were capable of generating intimate imagery of identifiable persons. Under the Council's wording, any user of such a version would have theoretically been in breach of Article 5 independently of what they used it for.

The Council's structural separation of NCII and CSAM would be a good approach to drafting if CSAM is included. The Council distinguishes between NCII (ba) and CSAM (bb), with a CSAM subparagraph linked to Directive 2011/93/EU. This separation reflects that the two categories involve different legal frameworks, different victim profiles, and different policy considerations, notably the absence of any consent-based exception for CSAM. We have not made any independent analysis on the question of whether adding CSAM provisions to the AI act would actually create an improved level of protection, given the state of current CSAM related legislation.

2.4. On Article 6(1)(a): Safety Components

We support the proposed EP amendment to Article 6(1)(a) that adds more detail on the classification of safety components.

We however caution that this amendment does not resolve the overclassification problems we discussed above in section 2.1.1 (Analysis: impact of article 6(1) overclassification on administrative burdens), even though we have seen signs that the EP believes it does so. Ahead of the IMCO/LIBE vote in the EP, a verbal statement was made by Kokalari, stating that the proposed amendments would ensure that a connected laundry machine would not automatically be a high-risk AI product. We support the sentiment that a connected laundry machine should not be a high-risk AI product. But we also believe that the IMCO/LIBE amendments to Article 6(1) do not achieve this outcome**.**

The Annex III mechanism in combination with Article 6(2), where AI systems are classified as high-risk based on the intended function, is a much more refined classification tool than Article 6(1) and Annex I are. This mechanism could in future be used by the Commission, via delegated acts, to create outcomes where certain types of Annex I.A products become subject to the obligations for Annex III high-risk AI systems. These obligations would include registration requirements (according to the EP and Council proposals), a requirement to self-assess compliance to several Chapter III articles, and requirements to provide information if requested by competent authorities overseeing the market.

2.5. On Article 6(4): Registration requirement for non-high risk systems

We strongly welcome that both the Council and EP intend to reinstate in identical form this obligation. This relatively minor burden for Companies is key for proper enforcement of the AI Act. The removal of this obligation would also send the wrong signal to the market. In our initial Commission analysis we provide more details on the reasons why.

We don't oppose the simplification to the registry proposed by the EP and Council as it keeps the main point of maintaining visibility, but reduces further the small burden for companies, keeping only the main objective of registration.

2.6. On Article 25: EP proposal to extend some value chain obligations to providers of GPAI models

The EP proposes to extend article 25 as follows, by adding the text in bold.

Proposal for a regulation
Article 1 – paragraph 1 – point 9 a (new)
Regulation (EU) 2024/1689
Article 25 – paragraph 2

Amendment

(9a) Article 25(2) is replaced by the following:

"2. Where the circumstances referred to in paragraph 1 occur, the provider that initially placed the AI system on the market or put it into service shall no longer be considered to be a provider of that specific AI system for the purposes of this Regulation.

That initial provider, as well as providers of general-purpose AI models whose models are integrated into high-risk AI systems, shall closely cooperate with new providers and shall make available the necessary information and provide the reasonably expected technical access and other assistance that are required for the fulfilment of the obligations set out in this Regulation, in particular regarding the compliance with the conformity assessment of high-risk AI systems.

This obligation shall include:

(a) the provision of technical documentation sufficient to assess compliance with Article 16 requirements;

(b) the disclosure of known limitations and failure modes that could affect high-risk applications;

(c) the provision of reasonable technical access for testing and validation purposes.

This paragraph shall not apply in cases 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.'

We conditionally support this extension of obligations, under the condition that some necessary exceptions as discussed below are added. We believe this extension would create a safer value chain. We also believe that the added language 'This obligation shall include (a)-(c)' is a useful clarification, in particular (a) is useful.

However, we recommend that the co-legislators add some clarifying language to these obligations. While it is clear that the provider of an AI system can clearly specify against having it changed into a high-risk AI system, this level of clarity is not present for providers of general-purpose AI models.

The proposed text as written above implies instead that, if a downstream actor takes a general-purpose AI model and integrates it into a high-risk AI system against the express wishes or recommendations of the model provider, the model provider nevertheless has an obligation to 'closely cooperate' with that downstream actor. This burden is too much for the providers of GPAI models which are specifically not intended for high-risk applications, and also too much for open source model providers.

We therefore also propose that the last part the new article 25(2) is extended to read:

> This paragraph shall not apply in cases where the initial provider has clearly specified that its AI system is not to be changed into a high-risk AI system, **or its general-purpose AI model is not to be integrated into a high-risk AI system,** and therefore does not fall under the obligation to hand over the documentation. **In the case that the initial provider is providing an free and open source AI system or free and open source general-purpose AI model, this paragraph shall also not apply in cases where this initial provider has clearly specified in the documentation that its AI system is not designed to be changed into a high-risk AI system, or its general-purpose AI model is not designed to be integrated into a high-risk AI system.**

We recommend the last sentence on open source above specifically because many open source projects will generally avoid 'clearly specifying' limitations on downstream use, as many open source theorists would say that if such limitations are included in an open source license, it is no longer a valid free and open source license. The workaround we propose is that providers can specify their design intent instead, outside of the license.

We also recommend that the entire block of text for the end of 52(2) above is also added at the end of the first paragraph of 53(4), as it seems like the legislator originally forgot to repeat that language there.

On article 27(4) and 96(1)(fa): EP proposal to integrate the DPIA with the FRIA

We recommend adopting the EP's approach to 27(4), provided that it also includes their proposed Article 96(1)(fa), mandating the commission to specifically clarify how the data protection impact assessment (DPIA) can be reused for the fundamental rights impact assessment (FRIA).

In addition, we stress that the guidelines should acknowledge the fundamental difference in the object of analysis between both assessments, to prevent the FRIA from being reduced to a cross-referencing exercise. We propose that the following change is made to article 96(fa):

> "*the application of the obligations referred to in Article 27, including the possibility to reference or include relevant sections or parts of the data protection impact assessment into the fundamental rights impact assessment,* ***as well as clarifying their respective objects****, pursuant to Article 27(4) of this Regulation, using, where relevant, standardised templates.*"

We acknowledge the practical demand behind the EP's amendment to Article 27(4), which would allow deployers to incorporate relevant parts of their DPIA into the FRIA and the addition of article 96.1 (g) mandating the Commission to develop guidelines on the application of Article 27 requirements. Data protection officers have been vocal about the burden of overlapping assessments and the possibility of reusing relevant DPIA sections for

the FRIA, while the EDPB is already working with the Commission on guidelines for the interplay between the GDPR and the AI Act.[4]

However, the object of analysis is not the same: a DPIA under Article 35 GDPR examines risks arising from a specific data processing operation for a specific use case. A FRIA under Article 27 of the AI Act requires a broader assessment of the AI system as a whole and its deployment context including, but not limited to, data protection risks. Even where both assessments address data protection, the FRIA can capture risks that a narrower, operation-specific DPIA does not.

The risk, therefore, is that deployers treat the DPIA as a substitute rather than a complement effectively ticking the FRIA box for data protection by referencing an assessment that does not cover the full scope of deployment-level risks. This risk is compounded by the absence of a single harmonised DPIA methodology: different DPOs use different approaches, meaning the scope and depth of any reused DPIA can vary significantly.

2.7. Article 43(3): On the 1.5 years notified body application deadline

We recommend either accepting the Commission's revised Article 43(3) without the 18-month application deadline, or, as the EP proposes, rejecting it entirely and reinstating the original provision. We have a slight preference for the former, as it preserves the Commission's useful procedural improvements while keeping the application process open-ended.

Our concern is with the fixed deadline to become a notified body:

> "*[...] at the latest [18 months from the entry into application of this Regulation].*"

This fixed deadline may prevent organisations that develop expertise later from entering the system, reducing the long-term pool of notified bodies. This could limit competition, slow conformity assessments as demand grows and reduce the long term scalability of the European AI oversight ecosystem. Ideally, the network of notified bodies should grow organically over time to meet demand, as the AI adoption increases across industries. Therefore, a more flexible, ongoing application process would better support the development of a robust auditing infrastructure capable of meeting regulatory and market demand.

2.8. On Article 50(7): The Board's Role in the adequacy of the Transparency Codes of Practice

We welcome that the Council in its Article 50(7) text reinstates the power of the Board in assessing the adequacy of the Codes of Practice on Transparency of AI-Generated

[4] https://www.edpb.europa.eu/system/files/2026-01/edpb_edps_jointopinion_202601_proposal_ai-omnibus_en.pdf

Content together with the Commission. This is something we had previously asked for in our first analysis of the Commission Omnibus proposal as we believed that giving the Commission sole responsibility for assessing their adequacy could tip the power balance too much, especially as other changes proposed also give more power to the Commission and the AI Office.

2.9. On Article 56(6): The Board's Role in the adequacy of the Codes of Practice for GPAI models

Both the EP and the Council maintain the changes to article 56(6) which gives the Commission ultimate authority to unilaterally assess and public its assessment on the adequacy of the codes of practice while reducing the Board's role to issuing a non-binding opinion.

We consider that this concentration of power has its risks , as in our Commission Omnibus analysis, **recommend to reject this change to Article 56(6)** in order to better balance the power between the Commission and the Council.

2.10. On Article 60a: Real world testing for regulated products

We recommend basing the final Article 60a text on the Council's version of Article 60a, as it adds significant procedural safeguards that were largely absent from the Commission's proposal. We flagged these as necessary to protect health, safety and fundamental rights in our Commission's Omnibus analysis.

The Council replaces the Commission's voluntary bilateral agreement between Member States and the Commission with a structured national framework mechanism. Member States choosing to permit real-world testing must now lay down the detailed conditions under which testing may take place, as well as the requirements, governance, and accountability arrangements necessary for implementation (2a). The Council adds an explicit harm minimisation requirement: both the framework design and individual testing plans must ensure that any risk of harm to health, safety, or fundamental rights of natural persons is minimised (4(c)). Furthermore, individual testing plans must be agreed between the provider and the national competent authority (4(b)), and applicable Union and national law applies to the testing (4a).

These additions address several concerns we raised in our Commission Omnibus analysis, particularly regarding the need for minimum procedural safeguards conditioning the removal of obstacles, the involvement of competent authorities in testing oversight, and the protection of third parties exposed to testing. However, the Council text still does not explicitly address the degree of public disclosure of testing plans, which we continue to recommend clarifying.

2.11. On Article 63(1): QMS simplification requirements

We welcome that the Council proposes keeping the expansion of the simplification opportunities for quality management system (QMS) to SMEs from microenterprises but reinstating the safeguard of "provided that they do not have partner enterprises or linked enterprises within the meaning of Recommendation 2003/361/EC".

We previously analysed that eliminating this safeguard in the first place was a mistake. This safeguard made sure that larger actors, which are not meant to benefit from such a simplification mechanism, could not do so through linked enterprises or corporate law engineering.

2.12. On Article 64(2a)): EP proposal to legally require adequate resources for the AI Office

We strongly welcome the EP proposal to strengthen the AI Office through additional resources through Amendment 57 which adds Article 64(2a) as an obligation in the AI Act:

> "*Without prejudice to the budgetary procedure and through existing financial instruments, the AI Office shall be allocated with adequate human, financial and technical resources, and with infrastructure to fulfil their tasks, to effectively perform its duties and exercise its powers in respect of the enforcement of Regulation (EU) 2024/1689. In particular, the AI Office shall have a sufficient number of personnel permanently available with in-depth competences and technical expertise. The AI Board shall assess competence and resource requirements*."

This language mirrors similar obligations in Article 70(3) applying to member states, where member states are required to ensure adequate resources for their national competent authorities. Adding this language to safeguard effective enforcement is made somewhat urgent in light of the Article 75 proposal by the Commission to create exclusive competences for the AI Office in a way where the national competent authorities can no longer step in if the Office fails to act – see our discussion in section 2.14 on Article 77 below.

2.13. On article 72(3): Post-market monitoring plan

All three texts replace the original implementing act mechanism with non-binding guidance from the Commission. We have no objection to this change of instrument, as what matters is that the outcome helps providers obtain clarity on how to comply with their monitoring obligations.

We support the EP reinstating the template as part of the guidance. A standardised template, while non-binding, can provide significant practical value for providers by reducing uncertainty about what a compliant post-market monitoring plan should contain and lowering the cost of compliance. The template can also support more consistent supervisory practice

across Member States by establishing a common baseline for what market surveillance authorities can expect to see.

We also support the EP's deadline of February 2027 over the Council's September 2027 and the Commission's lack of deadline. The original AI Act already set a February 2026 deadline for the implementing act, which passed without such an implementing act being written. We believe it is important for the guidelines and the template to be ready for providers and deployers ahead of the expected new date for the entering into force of high risk systems obligations, August 2027.

2.14. On Article 75: New Powers for the AI Office and their exclusivity

In earlier reports[5], we considered the Commission's proposal to centralise oversight over many GPAI models based systems in the AI Office. Our main concern was that, while centralisation can have some benefits, the Commission was likely underestimating the resources and competences needed to exercise these new powers properly, while the exclusivity written in the Commission proposal made it impossible for national competent authorities to step in if the Commission failed to investigate a case.

We welcome both the Council and the EP proposals as they are both strong improvements over the Commission proposal. We also welcome that the Council text adds procedural detail on how the AI Office is to exercise its supervisory and enforcement powers.

To ensure adequate resources for enforcement with respect to fundamental rights risks from Annex III high-risk AI systems, **we strongly recommend that in the trilogue, at least one of the two safeguards against a lack of resources in the AI Office is adopted:**

1. **The safeguard which is a combination of**
    a. **Of using the language the EP amendment 60 proposal for Article 75, the language that deletes the word 'exclusive' so that the AI Office gets enforcement powers but not exclusive enforcement powers; and**
    b. **the trilogue refraining from inserting the following EP amendment 60 paragraph into the text:** *'Where the Commission has not initiated proceedings for the same infringement, the competent authority of a Member State in which the main establishment of the provider of very large online platform or of very large online search engine is located, or where their legal representative is established, may have the powers to supervise and enforce the obligations under this Regulation.'* This paragraph could be read to constrain the general powers now conveyed by the AI Act, where every national competent authority has the power to supervise and enforce the

[5] https://aistandardslab.org/recommendations-on-the-digital-omnibus-amendments-to-the-eu-ai-act/ section header '(25) Article 75', and https://aistandardslab.org/recommendations-on-the-european-parliament-amendments-to-the-eu-ai-act-in-the-digital-omnibus/ section 4.7.

regulation, if it independently deems this to be necessary. We believe that the insertion of the above paragraph would encourage unwanted regulatory shopping that could prevent national authorities from intervening if they see a development that will specifically affect local residents.This would create the same bottleneck the GDPR suffers from where the Irish DPA remains the primary authority with powers to supervise most large tech companies. This would cause a negative impact on how much human rights will be enforced in practice.

2. **The safeguard which is a combination of**
    a. **the safeguards in the Council proposal for Article 75, where the exclusive competences of the AI Office only apply to the oversight of providers, not third party deployers of the systems concerned, and**
    b. **the safeguard in the EP amendment 57, which creates an Article 64(2)(a), which forces the allocation of sufficient resources to the AI Office to fulfill its obligatory enforcement tasks.**

Our preference is for the first option above (this preference is independent from us supporting EP amendment 57 also when the first option is used). The first option avoids a single point of failure as several authorities could act and, in the long run, would likely foster cooperation between the AI office and National authorities.

We also welcome the proposals from the EP and the Council to create specific exceptions to the Commission text, exceptions where several types of Annex III high-risk systems would remain within the competence of national competent authorities.

2.14.1. Analysis: examples of high-risk AI systems that would be in scope of exclusive authority of the AI Office

According to the original Commission proposal, an Annex III high-risk AI system made by a certain party would fall within the exclusive surveillance and enforcement scope of the AI Office whenever that system contains a GPAI model made by the same party. As mentioned, we believe the Commission has under-estimated the enforcement implications and manpower needed. We now give some examples of oversight actions that would fall inside of the exclusive scope of the AI Office if the Commission proposal were to be accepted without modification.

First example, how some deployers of ChatGPT would fall under exclusive AI Office oversight. Say that an educational institution has a policy encouraging or allowing its teachers to use ChatGPT via its conversational interface, to grade written work submitted by students. Teachers would manually enter a prompt that makes ChatGPT to compute and assign a grade to each document. This would make the educational institution a deployer of ChatGPT, one which deploys it as an Annex III high-risk AI system.

The use of ChatGPT to grade work could lead to violations of the fundamental rights of the students, especially after a group of students catches on and starts using AI tools themselves to optimise their essays, so that they work as adversarial inputs creating higher than average grades on ChatGPT. We assume that the teachers have not been warned

about, and are skilled enough, to analyse such attacks and to compensate for them. This is definitely an area where oversight is needed.

Under the Commission proposal, the AI Office would have exclusive authority to do surveillance and enforcement to ensure that the above situation does not lead to fundamental rights violations. We believe the implied workload would be very high. It would be more efficient if the national competent authorities, who still need to surveil these institutions for other uses of AI tools anyway, would have the authority to do the surveillance. **So we welcome the solutions offered by a) the Council proposal, that would limit the surveillance powers and exclusivity to providers only, not deployers, and b) the EP proposal, that removes exclusivity entirely.**

Second example, how many AI systems made by small companies would fall under exclusive AI Office oversight. Say that an educational publisher starts creating and selling an essay grading tool, tailored to a specific language used within an EU member state.

In this case, a potential design approach for the tool is that the publisher will create from scratch a language-specific LLM, which is then further customised to do fair grading, with some robustness against adversarial attacks. According to the current guidelines published by the Commission[6], that LLM will be called a GPAI model if the training compute exceeds 10^23 FLOP and if the customisation does not remove its general capabilities.

According to Epoch data[7], a 10^23 FLOP GPAI model currently costs about $100K in 2023 dollars, of compute power to train. If the publisher just re-tunes the parameters of an existing GPAI model, the Commission guidance says that the publisher will definitely be considered the provider of that re-tuned model if the re-tuning takes more than ⅓ of the original training cost of that model, so the re-tuning cost might be as low as $33K in 2023 dollars. These amounts are well within the budget range of medium and even small size companies.

From this example, we conclude that many Annex III high-risk AI systems, hundreds to thousands of them, might end up being systems that would have to be exclusively overseen by the AI Office under the Commission proposal. This means a potentially large expansion of the workload of the Office. In an earlier analysis, we proposed that exclusivity is limited to systems incorporating GPAI models with Systemic Risk, as these have much greater associated training costs. But this proposal was not adopted by either the EP or the Council. As an alternative, **we welcome the EP proposal to remove exclusivity in the oversight duties towards the above providers.**

[6] Available at https://digital-strategy.ec.europa.eu/en/news/learn-more-about-guidelines-providers-general-purpose-ai-models

[7] See the data collected by Epoch and available at https://epoch.ai/data/ai-models/, which includes training costs and FLOP estimates for many models.

2.14.2. Analysis: concerns if the AI office has the competence to judge local expectation with respect to fundamental rights

If it were to get exclusive competence for some Annex III high-risk AI systems, we also have a concern beyond capacity concerns. For several types of country-specific AI systems, we worry that the AI Office will not have the expertise and local knowledge to deal with all fundamental rights complaints.

We foresee that many Annex III high-risk AI systems will be specifically designed to support government bodies or small and medium companies operating inside of a single Member State. To investigate claims that those systems create unfair (biased) outcomes which violate fundamental rights, it will often be necessary to look specifically at what is considered fair or unfair according to the customs and laws of that Member State. National competent authorities have the knowledge and network needed to carry out such local assessments. We expect that the AI Office would struggle in developing sufficient member state specific expertise. While we recognise that fundamental rights apply to all of the EU, we must also stress a key insight from the AI fairness field, that there are often differences between countries about what is considered 'fair' in a certain specific use case.

We also believe that, if the EU wants to promote trust in AI, it may simply be a bad idea to assign the exclusive competence to judge local human rights complaints to a centralised regulator. **We therefore welcome the EP proposal to remove exclusivity in the oversight duties, so that local competent authority expertise can be used when necessary, and so that there are also stronger guarantees that local expertise will not be overruled.**

2.15. On Article 77, powers of fundamental rights authorities

We recommend rejecting the Commission proposal entirely and maintaining the current version of article 77 in the AIA. Alternatively, if the co-legislators want the fundamental rights authorities to have to go through the relevant authority (either the national competent authority or the AI Office) we recommend that the version proposed by the EP with the key addition of "without undue delay" is chosen. This addition would, in theory, help in lowering the risk of that authority being slow to respond or act on the FRA requirement, one of the negative aspects of this change to the AIA that we highlight bellow:

We understand that the rationale behind this change is streamlining the number of points of contact with the AI providers. However, the current Article 77 already addresses coordination by requiring the fundamental rights authority to inform the relevant market surveillance authority of any request. The change from a notification duty to full intermediation introduces risks that, in our view, outweigh the procedural benefits.

The core problem is that the fundamental rights authority loses its autonomous access to information and becomes dependent on the authority's cooperation, timing, and interpretation. The authority could slow investigations by taking time to process requests, hence the EP's addition of "without undue delay," which itself signals that this risk is

recognised. The authority could theoretically also narrow the scope of what it requests from providers on the fundamental rights authority's behalf, by applying the "necessary for effectively fulfilling their mandates" threshold in Article 77(1) or the confidentiality and trade secret protections under Article 78. Under the current text, the fundamental rights authority applies these standards itself; under the proposed change, this judgment shifts to an institution that may not share the same mandate or priorities.

The change also seems to create concerning national situations where a fundamental rights authority investigating the use of an AI system by a government body, such as a police force deploying a surveillance system, would no longer be able to request documentation directly. It must instead go through a third public institution within the same Member State. This adds an unnecessary layer between two domestic institutions and risks reducing the independence of fundamental rights oversight, which is precisely the function Article 77 was designed to protect.

2.16. On Article 113: Entry into force

We welcome that both the EP and Council retain the revised entry into force dates proposed by the Commission while removing the discretionary mechanism allowing the Commission to advance such entry into force which would create legal uncertainty and unease.