The AI Omnibus: a rollback of AI safeguards before they even apply

June 2026

This paper represents the joint position of European Digital Rights (EDRi), ARTICLE19, Access Now, AlgorithmWatch, Amnesty International, Danes je nov dan, the European Center for Not-for-profit Law (ECNL), Lafede - justícia global and Politiscope.

Thank you to the EDRi network and its AI working group for their long-running contributions to this analysis, which has been coordinated by the EDRi office.
In particular, we would like to highlight the significant contributions to this paper from our members Access Now, Amnesty International and the European Center for Not-for-profit Law (ECNL) and our partner AlgorithmWatch.

Overall assessment

The Digital Omnibus on AI Regulation Proposal, 2025/0359(COD), (the so-called 'AI Omnibus') should never have been put forward. The final text agreed on 7 May 2026 weakens the AI Act. It emboldens industry lobbying. It undermines the EU's credibility as a serious digital regulator. And it normalises the use of the Omnibus process to reopen hard-won fundamental rights protections and to introduce changes that go far beyond what this legislative instrument should be used for. This abuse of scope should concern all policymakers, regardless of where they stand on this file: once this procedural shortcut is normalised, it can be used by any political majority to reopen settled safeguards, weaken rights, or push through changes that would not survive proper scrutiny.

It is also evident that the Omnibus does not achieve the Commission's stated aim of simplification. In several areas, it does the opposite: creating inconsistent application across sectors, increasing reliance on future standards and guidance, weakening accountability mechanisms, and making the AI Act harder to scrutinise and enforce.

This matters beyond the AI Act. It is directly relevant to the Data Omnibus, the Digital Fitness Check and future attempts to reopen the EU digital rule-book. Industry actors now have a clear signal that implementation moments can be used to attack obligations they dislike. The pattern is clear: first frame safeguards as burdens, then call for 'targeted' simplification, then push for delayed application, reduced transparency, sectoral carve-outs and weaker oversight.

Beyond its borders, while the EU promotes itself as a pioneer in regulation for 'trustworthy AI', the AI Omnibus not only fails to resolve the shortcomings of the Act that stand in the way of it being truly rights-enhancing, but goes further in downgrading limited achievements in the original law even before they start to apply. The AI Omnibus reinforces a worrying trend globally towards deregulation, rather than strengthening rights-based AI governance, and the EU stands as a champion in a race to the bottom.

The AI Omnibus, like the other Omnibus proposals that make up the Commission's deregulatory agenda, should be rejected. We urge Members of the European Parliament to vote against the AI Omnibus. Despite some small mitigations of the worst parts of the Commission's original proposal, it remains a dangerous piece of legislation that should never have been proposed in the first place. Policymakers who care about fundamental rights must resist the normalisation of these deregulatory instruments by fully opposing them, rather than presenting limited damage control as victories.

Summary of main outcomes

- Co-legislators should not have agreed to the AI Omnibus. It reopens key parts of the AI Act before they had even started to apply, through a process presented as technical simplification but involving highly dangerous substantive deregulatory moves.

- One major rollback has been stopped: the final agreement does not fully delete the EU database registration obligation for providers who rely on Article 6(3) to classify systems as 'not high-risk'. However, the transparency obligation has still been weakened, because the final text reduces the information that providers must upload to the public database. This will make it hard to track AI technologies across Member States, particularly by human rights advocates and civil society organisations, and thus undermines public oversight and accountability.

• Annex I has not been fully deleted or relocated, which is positive, but machinery has still been moved from Section A to Section B, creating a dangerous carve-out for industrial AI and pushing those systems towards sectoral legislation. This outcome followed intense pressure from industry, Germany and supportive political actors in Parliament and Council.

• Key high-risk obligations have been delayed, extending the period before important safeguards apply. For many Annex III systems, obligations are delayed until 2 December 2027. For Annex I systems, the timeline is pushed even further, to 2 August 2028.

• Fundamental Rights Impact Assessments (FRIAs) have been preserved, despite pressure to weaken or remove them. However, the final text still allows cross-referencing to Data Protection Impact Assessments (DPIAs), so we will need to monitor whether this becomes a shortcut that narrows FRIAs to data protection risks.

• The Omnibus also weakens or complicates safeguards on AI literacy, special category data processing, Article 77 fundamental rights oversight, sandboxes and real-world testing.

• The new Article 5 prohibitions on non-consensual intimate material and child sexual abuse material are purported to address real harms, but were rushed, inserted through the wrong process (when a legitimate process already existed under Article 112). They also point to other issues that still await meaningful attention under the AI Act, such as the continued export of AI systems that are prohibited within the EU.

• Overall, the outcome weakens the AI Act, emboldens industry lobbying, and sets a dangerous precedent for the Data Omnibus, the Digital Fitness Check and future attempts to reopen the EU digital rulebook.

Process-related concerns

Before going into the substance, it is worth underlining that this file should not have existed in this form. The AI Act was adopted after years of negotiation and compromise, requiring immense human and financial resources and collective effort of national and EU level legislators, policymakers, civil society, academia, industry and other stakeholders. Some of its most important provisions, especially the high-risk obligations, had not yet started to apply. Instead of supporting implementation through guidance, resources and enforcement capacity, the Commission chose to reopen the law, as they are doing with many others (including the GDPR and ePrivacy as part of the same Digital Omnibus), under the banner of 'simplification'.

The Omnibus proposals are presented as a technical exercise with a clearly-limited scope, but have become vehicles for substantive political deregulatory changes. In the case of the AI Omnibus, this includes delaying obligations, weakening transparency, changing the relationship between the AI Act and sectoral legislation, softening obligations, and adding new prohibited practices in a very rushed way.

What's more, the Commission relied heavily on 'reality check' discussions with industry, did not make proper use of the Advisory Forum foreseen by the AI Act, and did not provide a robust impact assessment or public-interest consultation despite the substantive nature of the proposed changes. It is particularly striking that the Commission went as far as to claim that the proposed Omnibus would '[lighten the regulatory burdens on people]' (p.1) or '[reduce implementation challenges for citizens]' (Recital 23), whereas the law never intended to regulate individual use cases. The real beneficiaries of these delays and carve-outs are providers, deployers and industry actors seeking more flexibility and less public accountability.

The political dynamics have also been highly concerning. The pressure around Annex I and industrial AI was intense, including from industry, the government of Germany in the Council, and MEPs pushing similar deregulatory lines in Parliament. The machinery carve-out became the central political battleground of the final negotiations, with pressure from Berlin helping shift the debate towards industry's preferred framing and outcome.

This is a worrying precedent: if hard-won safeguards can be reopened before they apply whenever powerful actors complain, implementation becomes a de facto second legislative arena. This is the core political problem. The AI Omnibus not only weakens parts of the AI Act: it normalises the idea that the EU digital rulebook can be reopened through fast-moving Omnibus files before the rules have had a chance to be tested.

Main outcomes

1. Article 6(3) and EU database registration

The final agreement avoided the worst case scenario: the Commission's proposal to delete the EU database registration obligation for providers relying on Article 6(3) did not survive the trilogue negotiations.

This was one of the most dangerous and incoherent parts of the Commission's proposal. It would have removed the public registration obligation for systems used in high-risk areas where the provider unilaterally decides that the system should nevertheless not be classified as high-risk. In practice, this would have left only internal documentation that authorities could in theory request, with no meaningful public trace. Even enforcement authorities would not necessarily have known which providers had exempted themselves, rendering their ability to request documentation far less meaningful.

Civil society pushed hard against this, together with many partners. It is positive that this proposal did not survive the trilogue. However, it is also concerning that the Commission proposed and defended such a serious accountability loophole.

However, the final AI Omnibus text still weakens the AI Act as adopted. Providers will still have to register these systems, but the information uploaded to the public database will be reduced. In particular, information that would have been important for scrutiny - such as the Member States where the system is placed on the market and the reasons for the provider's self-exemption - is removed from the mandatory public entry.

This matters because Article 6(3) was already a huge victory for industry lobbyists during the original AI Act negotiations in terms of bypassing stringent public accountability. This is because the law already gives providers significant discretion to decide that certain systems used in high-risk areas are not high-risk. If that discretion is not matched by meaningful transparency, it becomes harder for regulators, researchers, civil society and affected people to scrutinise whether companies are classifying their systems correctly. Without public registration, this discretion would have been exercised with almost no public trace. Ultimately, even this 'victory' of the AI Omnibus negotiations to mitigate the worst possible outcome is still not to be celebrated. At most, it means that co-legislators avoided making the AI Act even harder to enforce, while still reducing the usefulness of one of its main public accountability tools.

2. Annex I, machinery and industrial AI

The push to fully delete or relocate Annex I was not successful, and this is important from a public interest and human rights perspective. However, the final agreement still moves the Machinery Regulation from Section A to Section B. In practice, this means that AI systems covered through machinery will be dealt with through the sectoral machinery framework rather than through the AI Act's horizontal high-risk approach.

This is a significant structural change. The AI Act was designed as a horizontal framework because AI-related risks do not neatly fit into sectoral product safety boxes. Machinery legislation may address important product safety issues, but it is not designed to cover the full range of AI-specific and fundamental rights risks, including opacity, bias, surveillance, worker monitoring, human oversight, data governance, robustness, or effective contestation.

This matters especially in workplace and industrial contexts. AI systems embedded in machinery can affect workers' safety, pace of work, task allocation, monitoring, autonomy and ability to challenge decisions, and can encode discriminatory treatment. If these systems are framed as optimisation, automation, efficiency, quality control or user assistance, the risk is that their rights impact becomes harder to capture under the AI Act's horizontal high-risk logic.

The final outcome is narrower than the most extreme proposals, which would have pushed many more regulated products out of the AI Act's horizontal framework. Machinery is one sector out of a broader set of regulated products that includes toys, medical devices and connected products. However, this does not make the carve-out harmless. It still introduces inconsistent treatment across sectors and creates a precedent for future attempts to move other areas away from the AI Act's horizontal safeguards.

The text also relies on future bridging standards and further work by the Commission to integrate AI-related requirements into the machinery framework. This means that, even where some connection with the AI Act is preserved, the protection of people affected by industrial AI becomes more dependent on future standardisation, sectoral implementation and administrative capacity. This is not a clear simplification but shifts complexity elsewhere.

3. Safety components[1]

The agreement also changes the definition and treatment of 'safety components', in particular through the amendments to Article 3(14) and the new Article 6(1a) to 6(1d). These provisions matter because they help determine when an AI system embedded in a product is treated as high-risk under Article 6(1). The text excludes AI systems used solely for purportedly non-safety aspects of user assistance, performance optimisation, service efficiency, automation, convenience or quality control from being treated as safety components. It does however state that AI systems whose failure or malfunctioning would endanger health and safety should still qualify as safety components.

That latter safeguard is important. However, the new wording still creates interpretative space around what is framed as 'non-safety'. In practice, systems that affect workers, users or other people could be presented as optimisation or efficiency tools, even where their operation, failure or malfunctioning has serious real-world consequences. The concern is not only a formal one: the more the AI Act relies on provider framing, sectoral classification and case-by-case interpretation, the less predictable the scope of high-risk obligations becomes.

---

1. This section should be read together with the Annex I carve-out. The combined effect of moving machinery to Section B and narrowing what counts as a safety component risks making it easier for industrial AI systems to fall outside the AI Act's high-risk route, especially where providers describe them as tools for optimisation, quality control, automation or convenience.

4. Delays to high-risk obligations

The final agreement also maintains a postponement of key high-risk obligations by 15 months for Annex III systems, until 2 December 2027. For Annex I systems, the timeline is pushed further, until 2 August 2028. This means that the period in which potentially harmful AI systems can be placed on the market, put into service or deployed without the full set of AI Act high-risk safeguards is lengthened – delaying guarantees of rights protections and justice in the context of AI.

The delay also creates a counter-productive incentive for providers to place systems on the market before the new deadline, since systems already placed on the market before 2 December 2027 may avoid the full high-risk obligations unless they undergo substantial modifications.

Public authorities and systems used on their behalf are subject to a different transitional timeline, with relevant obligations applying by 2 August 2030. Even with that clarification, the broader concern remains: the delay extends the period in which high-risk AI systems can affect people without the full accountability framework that the AI Act was supposed to provide.

This is especially concerning given the increasing deployment of AI systems in sensitive, high-risk areas such as workplaces, public services, health, education, policing, migration and the justice system. Delaying safeguards is not a neutral administrative choice. It postpones accountability and extends the period in which affected people lack the protections the AI Act was supposed to provide.

It also does not solve the legal certainty problem. Companies and public institutions will still need to prepare for compliance. Legal certainty would have been better served by timely guidance, proper resourcing of authorities, usable compliance tools and a clear implementation timeline.

The delay was the central goal of the AI Omnibus negotiations. It was justified through concerns about standards, guidance and implementation readiness. Yet these concerns should have been addressed through proper implementation support, not by reopening and weakening the law before it applies.

5. Small mid-caps and proportionality

The Omnibus extends several SME-style privileges to small mid-cap companies (fewer than 750 employees and annual turnover under EUR 150 million, with thresholds potentially increasing to 1,000 employees and EUR 200 million), including simplified technical documentation and other proportionality measures. This is not a marginal category: small mid-caps include companies with hundreds of employees and substantial turnover and they can still develop technologies used for risky purposes like surveillance, leading to discrimination and exclusion. Treating company size as a proxy for AI risk is therefore particularly problematic.

Proportionality matters, and smaller providers may need practical support to comply with the AI Act. However, exemptions or simplified obligations must remain proportionate to the actual risks of the AI system and not come at the cost of human rights. The current approach goes too far because it extends exemptions to a broad category of companies that can still develop and deploy systems with serious impacts on fundamental rights, including systems used for surveillance, exclusion or discriminatory treatment.

All providers of high-risk AI systems must remain subject to safeguards that are strong enough to protect people's rights. Compliance support should help smaller actors meet those obligations, not reduce the level of protection attached to high-risk AI systems.

6. Article 4a and special categories of personal data

The final agreement still introduces a worrying derogation for the processing of special categories of personal data for bias detection and correction.

Some safeguards are included, including necessity, limits on re-use, security measures, deletion and documentation. However, the provision still allows sensitive data to enter the AI value chain for debiasing purposes, including by deployers, providers of non high-risk AI systems, and providers or deployers of AI models, where the conditions of the new derogation are met.

Pseudonymisation also remains listed among the safeguards, despite long-standing concerns that it should not be treated as sufficient protection for high-dimensional sensitive data as well as recent research showing that advanced AI tools can easily re-identify pseudoynmised data and make supposedly protected datasets more vulnerable. This remains a serious data protection concern, especially because the provision sits at the intersection of the AI Act and General Data Protection Regulation Article 9.

This derogation creates a new route for processing highly sensitive data in AI development and deployment contexts. Even with safeguards, it risks making the use of special categories of personal data in AI pipelines more normal, rather than exceptional. This further creates an oversight gap for AI systems that are not designated as high-risk. Such systems do not have to meet the majority of requirements under the AI Act, making it harder for regulators and authorities to ensure that they do not misuse sensitive data.

The connection with the so-called 'Data Omnibus' is also important here. The AI Omnibus already expands the possibility to process special categories of personal data for bias detection and correction in AI systems. The Data Omnibus risks going further by creating additional derogations for the use of special categories of personal data in AI development, including where removing such data from datasets is considered impossible or disproportionately difficult. Taken together, these provisions risk normalising the presence of highly sensitive data in AI pipelines, rather than treating it as exceptional and subject to strict necessity, minimisation and deletion requirements.

This is part of a broader move to relax the data protection safeguards that currently limit how AI systems are trained, tested and deployed. If the Data Omnibus narrows the practical scope of personal data or expands the routes through which personal data can be reused for AI development, this will not only weaken data protection. It will also weaken the practical effectiveness of AI Act safeguards that depend on GDPR concepts such as personal data, profiling, biometric data and special categories of data.

7. Fundamental rights impact assessments (FRIAs)

We welcome the preservation of FRIAs for deployers of high-risk AI systems in the final agreement, despite industry pressure to remove or replace them. The final text allows deployers to cross-reference or quote relevant parts of a data protection impact assessment when conducting a FRIA, provided that the obligations laid down in Article 27 are met through the Data Protection Impact Assessment (DPIA).

This is not necessarily problematic if it avoids duplication where the DPIA genuinely covers the same issue and embeds the necessary elements of the FRIA listed in Article 27(1). However, the risk is that, in practice, deployers could assert that the DPIA covers all FRIA requirements and cross-reference the DPIA without conducting the FRIA at all, or not to the full extent.

It's important to emphasise that DPIAs and FRIAs have different scopes. A DPIA focuses on impacts to rights and freedoms linked to the processing of personal data. While a good DPIA should look into all fundamental rights, in practice a vast majority of data controllers tend to prioritise aspects related to data security. In addition, some fundamental rights impacts can stem not from processing personal data but from the automation, AI system design or the use of non-personal datasets. DPIA methodologies are also not necessarily aligned with the requirements of Article 27 of the AI Act, e.g. when it comes to identifying affected groups, assessing deployment context, or considering fundamental rights impacts that are not limited to personal data processing.

A FRIA should instead assess wider impacts on all fundamental rights and should consider the institutional context and deployment conditions, system design, access to services, relevant power imbalances, working conditions and the ability of affected people to challenge outcomes. The final wording of the AI Omnibus does not fully address this concern, therefore the AI Office template and further guidance will play a key role in clarifying the links between the two assessments.

The preservation of FRIAs is one of the few important safeguards that survived the negotiations. However, this should not obscure the fact that the final wording still creates risks of procedural shortcuts. The next phase of implementation must ensure that FRIAs remain meaningful tools for fundamental rights accountability, not a box to be ticked through existing DPIA paperwork.

8. Article 77 and fundamental rights bodies

The final text remains disappointing on Article 77. Fundamental rights bodies will be able to access information and documentation through market surveillance authorities, and the recital clarifies that existing powers under other Union or national law are not limited. This is supposedly better than a reading that would block direct access where such powers already exist.

However, the Omnibus still moves the AI Act further towards an indirect cooperation model, rather than preserving strong, direct access powers under the AI Act itself. This may complicate oversight, especially for bodies monitoring sensitive high-risk systems used by law enforcement, migration authorities or other public bodies. This complicates oversight rather than simplifying it.

These bodies play a crucial role in monitoring AI systems that can affect fundamental rights in highly sensitive contexts. Requiring them to rely more heavily on market surveillance authorities risks adding procedural friction, delaying access to relevant information and weakening independent oversight in practice.

9. AI literacy

The final agreement also weakens the AI literacy provision. The original AI Act imposed an obligation on providers and deployers to ensure a sufficient level of AI literacy among staff and others dealing with AI systems on their behalf. The Omnibus shifts the centre of gravity towards the Commission and Member States supporting and facilitating AI literacy, with providers and deployers taking measures to support its development.

This may reduce the directness and enforceability of the obligation. AI literacy should not be treated as a marginal administrative burden. It is one of the basic conditions for responsible deployment, especially in high-risk contexts. Weakening AI literacy can undermine the ability of people operating AI systems to challenge AI outputs in contexts such as welfare allocation or predictive policing, and can weaken the quality of FRIAs.

We would not accept safety-critical professionals being asked to operate complex systems without proper training. The same logic should apply to AI systems used in contexts where decisions can affect people's rights, access to services, working conditions or exposure to state power.

10. a) Article 5 prohibition on non-consensual intimate material

The final agreement adds two new Article 5 prohibitions on AI systems generating or manipulating (1) non-consensual intimate material and (2) child sexual abuse material. The harms addressed here are real and severe. Non-consensual intimate image generation is a form of gender-based abuse, harassment and intimidation. AI-generated or manipulated child sexual abuse material is also an extremely serious criminal issue.

However, in relation to the prohibition on non-consensual intimate material generation and manipulation - in addition to questions of effectiveness and feasibility - a ban that is drafted too broadly can create serious side effects for people it should not target. The proposed provision relies on problematic and potentially overbroad concepts of intimacy and sexual content, and risks affecting consensual adult sexual expression, sex workers, content moderation, harm-reduction practices, sexual and reproductive health information, research, security-testing, dual-use tools and general-purpose systems. Furthermore, the emphasis on identifiability of a person and body parts does not address the real issue for adults who are victims of these practices, which is rather an issue of consent.

It can also encourage over-filtering and over-removal, while failing to address the broader accountability failures that allow abuse to spread.

Substantively, the final wording attempts to distinguish between systems intended to generate such material, systems where such generation is a reasonably foreseeable and reproducible outcome without adequate safeguards, and deployers who intentionally use systems for the prohibited purpose. This is theoretically important, especially for open-source, research, security-testing, dual-use and general-purpose systems. However, we remain cautious. Liability should not attach simply because a general-purpose system can technically be jailbroken, adapted or misused by third parties.

Any prohibition in this area must be narrowly targeted at non-consensual abuse and child sexual abuse material. It should not create liability simply because a general-purpose system can technically be jailbroken, adapted or misused by third parties, nor should it enable over-removal, over-filtering or restrictions on lawful consensual adult content.

10.b) Undermining established procedural mechanisms to introduce prohibitions

From a process perspective, the AI Omnibus also sets a problematic precedent of sidestepping the established mechanism under Article 112 for introducing new prohibitions on harmful practices needing urgent attention. Whereas the AI Act allows the European Commission to annually propose amendments to prohibited practices as soon as the AI Act came into force, new additions were inserted into an Omnibus file that was supposedly limited to technical simplification and implementation adjustments. Instead, political objectives have been pursued without the proper procedures to assess the best regulatory mechanisms to address harmful use cases through adequate public consultation, particularly with impacted communities.
There is also a wider point. The same deregulatory agenda that weakens the AI Act, the GDPR or ePrivacy can remove safeguards that help prevent, detect and challenge these harms. Strong data protection, transparency, accountability, oversight and redress are part of the framework needed to protect affected people. Weakening those safeguards while celebrating a new ban is politically incoherent and legally short-sighted.

10.c) Unresolved gaps in human rights protections that are being ignored by the EU and Member States

What's more, the AI Act still has serious gaps in protections against prohibited systems, such as emotion recognition tools which are still allowed to be used in the migration and policing contexts, or for opaque 'health and safety' reasons, which should also be urgently addressed by lawmakers looking to protect people from discriminatory and harmful AI. And EU-based companies can still develop AI systems which they are not allowed to sell within the EU but can instead export elsewhere, profiting from human rights violations abroad. What this reluctance to address these other issues reveals is less a hierarchy of issues than a selective recognition of harms, whose experiences attract political and media attention, and whose are quietly overlooked.

This does not mean that non-consensual intimate image generation or AI-generated child sexual abuse material should be treated as secondary issues. It means that the legislative route and political context matter. A rushed Omnibus file should not become the substitute for a proper review of prohibited practices, nor should a headline ban be used to make a broader deregulatory package appear rights-protective.

Conclusion

The AI Omnibus weakens the AI Act and sets a dangerous precedent for EU digital rights regulation. It delays safeguards, reduces transparency, fragments the horizontal logic of the AI Act, weakens accountability and gives industry lobbyists a clear signal that implementation moments can be used to reopen rules they dislike. The problem is not limited to one file. The same logic is already visible in the Data Omnibus and the Digital Fitness Check. If policymakers accept this model, the EU digital rulebook will become permanently unstable, with fundamental rights protections repeatedly reopened before they have had the chance to work.
Members of the European Parliament and Members States in the Council of the EU should reject the AI Omnibus. Policymakers who are committed to upholding fundamental rights must resist the normalisation of deregulation through Omnibus procedures and defend the safeguards that people rely on in practice.