EU Digital Omnibus

-

General views

Main messages

- **On GDPR**
  - **AI & Legitimate Interest:** The proposal creates a restrictive "dual regime" for AI with undefined conditions that diverge from standard GDPR rules. To ensure legal certainty and prevent Single Market fragmentation, Article 88c must be harmonized with established case law by removing these additional conditions and the provision allowing national consent requirements.
  - **Special Category Data & AI:** The requirement to strictly "avoid" collecting sensitive data from the open web is technically unworkable and paradoxically necessitates invasive monitoring. This obligation must be replaced with a practical "risk mitigation" standard of technical and organisational measures, while the mandatory output filters that undermine the functional utility of AI tools should be removed.
- **On ePrivacy directive**
  - **Browser-Level Consent:** Article 88b must be deleted, as mandating centralized browser signals fails to practically solve user fatigue and instead creates "consent confusion" while functionally undermining the ad-supported web, including privacy-safe contextual ads models.
  - **Unified GDPR Enforcement:** To achieve genuine simplification and prevent enforcement chaos, the "split regime" between personal and non-personal data must be resolved by moving all terminal equipment rules into the GDPR under a single, consistent mechanism.
  - **Incentivizing Contextual Ads:** To incentivize privacy-by-design, Article 88a must be expanded to explicitly exempt low-risk, non-profiling activities—specifically contextual advertising and its essential safety features—from consent requirements.
- **On Cyber**
  - **Level up the ambition:** by standardising reporting timelines around the 72-hour reporting deadline, harmonising the trigger point for these timelines, aligning thresholds for reportable incidents and harmonising reporting templates.

    - While a single-entry point is a **good first step,** it should be integrated with the CRA reporting platform, and entities submitting notifications should be consulted on the design and functionality of the tool, as well as having a formal feedback mechanism to flag any potential issues with the tool.
    - **Further simplification efforts** need to be taken beyond the incident reporting, including simplifying NIS2 and CRA compliance.
- **On AI Act**
    - **Extend transparency deadlines (Art. 50):** Increase the compliance period to one year for all systems (new and existing) covering both Art. 50(2) and 50(4) to align with the Code of Practice.
    - **Set firm application dates for high-risk systems:** Fix the deadlines to December 2027 (Annex III) and August 2028 (Annex I), removing dependencies on administrative decisions.
    - **Allow choice of conformity assessors:** Enable providers to select between EU and Member State bodies to avoid bottlenecks and prevent lock-in to the Commission's entity.

The Digital Omnibus represents a strategic pivot toward a pragmatic digital rulebook. It aims to balance the highest standards of protection for European citizens with the legal predictability essential for innovation. However, significant gaps remain between this political ambition and operational reality.

GDPR

1. On the use of Legitimate Interest for AI training

The new GDPR Art.88c is a critical step forward for Europe's AI ambitions. It seeks to codify recent EDPB guidance (Opinion 28/2024), confirming the appropriateness of the Legitimate Interests legal basis for developing AI models. This welcome amendment by the European Commission is a first step toward the harmonization and legal certainty developers need to build made-in-Europe AI models.

However, the current drafting creates a "special regime" for AI that introduces new, untested conditions to the long-standing GDPR legitimate interest test. This could create unnecessary confusion and new complexity.

- The proposal effectively creates two separate regimes—one for AI, one for other technologies—risking competing requirements for controllers relying on legitimate interest. This bifurcated regime affects many existing technologies widely deployed in Europe. SMEs would be disproportionately impacted, needing to invest in new LIA procedures specifically for AI models. This added complexity undermines the intended simplification of administrative overhead.

- While “AI System” is defined under the AI Act, “AI model” is not defined in EU law, nor is there a centralized understanding of what it comprises. Imposing a new LI test for training or deploying undefined “AI models” creates significant confusion. It could inadvertently capture decades-old technologies that have been available in the European market, such as chatbots or fraud detection systems.
- The current wording refers only to “the interests of the controller,” which is unduly narrow. This does not align with Article 6(1)(f), or with extensive guidance and case law that permit the interests of a third party to be pursued.
- Finally, allowing national laws to override Article 88c and mandate consent opens the door to 27 different national regimes for AI training. This undermines the EU Single Market, forcing businesses to navigate divergent rules for the same products. It runs directly contrary to the CJEU ruling in *ASNEF*, which ruled against national laws that mandate consent and preclude reliance on legitimate interest.

Recommendations:

- Co-legislators should ensure the conditions for legitimate interest are harmonised between AI development and other processing activities. We recommend deleting the additional conditions in Article 88c to ensure the test remains consistent for all technologies.
- Regulation should focus on the activity (data processing) rather than undefined labels like "AI models." Alternatively, co-legislators should reference the definition of a “general-purpose AI model” in Article 3, point (63) Regulation (EU) 2024/1689.
- The proposal should explicitly recognise that the interests pursued may include those of a third party, in addition to the controller.
- To preserve the Digital Single Market and ensure harmonization, the provision allowing national laws to override the GDPR and mandate consent must be removed."

2. Ensuring neutrality in AI storage

The GDPR should remain neutral on evolving technical questions to avoid freezing a disputed scientific position into law.

However, the current drafting of Recital 33, Article 9(5), and Article 88c appears to codify a contested factual premise: that AI models “store” or “retain” personal data. This position is technically and legally disputed. Leading authorities, such as the Hamburg DPA, have recognized that AI models function by learning patterns rather than storing copies of training data. Similarly, in the recent *Getty Images vs. Stability AI* decision, the UK High Court accepted technical evidence that the AI model did not store copies of the training data.

Legislating based on a disputed technical assumption creates significant risks. In particular, codifying the view that AI models “store” data could trigger unintended ripple effects across other legal regimes,

particularly copyright law, where "'storing" has an important meaning. Since Member States retain significant autonomy in copyright enforcement, inadvertently defining AI model weights as "storage" of personal data could prejudge complex IP questions currently being debated in separate forums. This creates potential legal conflict and risks "backdoor" regulation of intellectual property issues through data protection law.

Recommendation: References to personal data being "retained," "stored," or "memorized" within an AI model should be removed from the text (specifically in Recital 33, Articles 9(5), and 88c).

3. On the use of special category of data with AI

We welcome the proposal's recognition that processing Special Category Data (SCD) is essential for Europe to reap the societal benefits of AI, such as powering cancer research or ensuring models are trained on representative data to mitigate bias. The introduction of a specific derogation for the *incidental* processing of SCD is a positive step.

However, the current conditions attached to this permission effectively negate its utility by imposing technically contradictory requirements:

- Mandating that developers "avoid" the incidental collection of SCD from the open web implies an impossible obligation to pre-scrub the entire internet before training begins. This effectively prohibits the use of large-scale public datasets.
- This requirement to identify and remove SCD implies a proactive monitoring duty. This creates a "privacy paradox" where developers are forced to scan and process *more* sensitive data to identify and remove it, increasing privacy risks rather than mitigating them.
- Finally, mandatory filters to prevent SCD in outputs would undermine the functional utility of AI tools. For example, a strict output filter would prevent a chatbot from answering factual questions about a public figure's political opinions or health history, even when such information is public.

Recommendations:

- The requirement to strictly "avoid" collecting SCD should be replaced with a standard of "mitigating the risk" of processing via technical and organisational measures. This is a workable standard that ensures high protection without demanding the impossible.
- The obligation to remove data should be triggered when a controller is explicitly "notified" of its existence, rather than implying a general, proactive monitoring obligation.
- We recommend deleting the requirement to strictly filter sensitive data from outputs, as this undermines the utility of AI tools for answering legitimate user queries.

4. On the explicit mention of PETs

The proposal's move toward a risk-based approach—acknowledging that data exists on a spectrum of risk rather than a binary state—is a welcome alignment with recent CJEU case law. The introduction of Article 41a, empowering the Commission to set criteria for when pseudonymised data is no longer considered "personal," provides a pathway for legal certainty.

However, to make the most of this new definition, the regulation must explicitly incentivize the tools that make this safety possible. Privacy-Enhancing Technologies (PETs)—such as homomorphic encryption and synthetic data—are state-of-the-art solutions that allow useful insights to be extracted while maintaining deep privacy protection. Currently, the drafting assesses the "state of the art" broadly. Recognizing that pseudonymised or secured data carries different risks than direct identifiers creates a powerful incentive for safety. If the law treats safe data exactly the same as risky data, companies have no motivation to invest in these advanced, resource-intensive technologies. Explicit regulatory recognition is needed to drive the market to mature these technologies, making them scalable solutions for all companies, including SMEs.

Recommendations: Article 41a(2)(a) should be amended to explicitly reference "Privacy-Enhancing Technologies" as examples of the state-of-the-art techniques the Commission must assess.

5. On Data Subject Access Rights

We welcome the Commission's proposal to address the growing issue of abusive Data Subject Access Requests (DSARs), particularly those used for purposes unrelated to data protection (e.g., litigation discovery). Strengthening the mechanism to refuse "manifestly unfounded" or "excessive" requests is a pragmatic step that protects the integrity of the GDPR.

However, the current drafting inadvertently creates a "clarity gap" and inconsistency. The proposal applies the "reasonable grounds" standard only to excessive requests, while leaving manifestly unfounded requests subject to an undefined evidentiary threshold. This inconsistency introduces significant legal uncertainty. It effectively forces controllers to fulfill abusive requests due to the difficulty of proving "absolute" unfoundedness, thereby neutralizing the simplification goal.

Recommendation: We recommend amending Article 12(5) to explicitly state that the controller bears the burden of demonstrating "reasonable grounds to believe" that a request is *either* manifestly unfounded *or* excessive. This ensures the reasonableness threshold applies consistently to both scenarios.

6. On Scientific Research

We welcome the introduction of a clear definition for "scientific research" in Article 4(38), which now includes research that may further a commercial interest. This is a vital step for the long-term

competitiveness of the EU's R&D ecosystem. Recognizing the unique needs of the research community is essential for ensuring Europe remains a global hub for discovery.

However, the current text leaves ambiguity regarding whether this new definition fully encompasses privately funded scientific research. Privately funded research drives the vast majority of innovation investment in Europe, from green technology to pharmaceutical breakthroughs. For example, AI weather forecasting tools or mental health research are often developed through private funding but serve clear public interests. If interpreted narrowly to exclude corporate R&D, the rule would fail to support critical sectors and place European industry at a disadvantage compared to other jurisdictions.

Recommendation: We recommend to co-legislators to make clear that the new definition does not exclude that the research may be privately funded. This would align the European Union with other forward-looking jurisdictions like the United Kingdom, who has adopted a revamped definition of scientific research in its Data Use and Access Act.

7. On Data protection impact assessment (DPIAs)

The proposal to harmonize the list of processing operations requiring a DPIA via a central EU mechanism is a positive step. This standardization has the potential to significantly reduce compliance costs and complexity for pan-European companies.

However, there is a risk that the new EDPB list could diverge significantly from the existing lists established by national supervisory authorities. If the EDPB amalgamates all national requirements into a "super-list," it would create a far more onerous burden than currently exists. Furthermore, the proposal fails to clarify how new lists apply to existing products. It is logically impossible to conduct a DPIA "prior to processing" (as required by Art. 35) for operations that have been running safely for years. Without a clear "grandfather clause," businesses face legal limbo regarding established services, risking retroactive non-compliance.

Recommendations:

- Co-legislators should mandate that the EDPB's harmonized list reflects the existing consensus of national supervisory authorities rather than creating new, expansive obligations.
- We recommend that the text explicitly states that new DPIA requirements apply only to new processing operations. Existing processing, which is already compliant under current national guidance, should not be subject to retroactive assessments.

ePrivacy directive

1. On the integration of the “cookie rule” in GDPR

Integrating the ePrivacy "cookie rules" (Article 5(3)) into the GDPR is the correct ambition. It aims to align the rules for tracking technologies with general data protection standards and centralize enforcement under Data Protection Authorities.

However, the current drafting creates a problematic "split regime." It moves the regulation of personal data on terminal equipment to the GDPR, but leaves non-personal data under the legacy ePrivacy Directive. In practice, a single cookie often contains both types of data, leaving businesses facing parallel enforcement from its supervisory GDPR authority and 27 national telecom regulators for the exact same file. This would lead to overlapping and chaotic enforcement, complexifying compliance by requiring extensive guidance to determine which specific tracker constitutes personal or non-personal data.

Recommendation: to achieve genuine simplification, all rules regarding terminal equipment must move to the GDPR. This ensures a single, consistent enforcement mechanism without compromising security, as unauthorized access to devices remains strictly prohibited under existing national computer misuse laws.

2. On the exemptions to the new Article 88a

We agree with the Commission’s assessment that the cookie consent system is broken. Under current legal framework both low-impact like contextual advertising and high-impact profiling operations trigger the same user consent requirement. This creates a systemic disincentive for businesses to invest in and develop more privacy-oriented products, while simultaneously fatiguing users with pervasive consent banners.

However, the Digital Omnibus' current level of ambition is not sufficient. The two additional exemptions — security and audience measurement — are far too narrow to meaningfully alter the status quo. They will not lead to a meaningful reduction of consent banners, nor will they incentivize privacy-centric business models such as contextual advertising. This is because the audience measurement exemption is limited to first-party data (ignoring SME reliance on third parties) and the security exemption merely codifies existing practice. Crucial 'plumbing' for contextual ads (like frequency capping) remains non-exempt, leaving banners in place even for more privacy-centric models.

The Digital Omnibus presents a crucial opportunity to address both systemic issues simultaneously. It must introduce a broader range of exemptions to enable businesses to conduct important, low-impact processing activities and at the same time reduce the volume of consent banners. This represents a critical opportunity to incentivize the shift toward contextual advertising. By allowing businesses to adopt these privacy-safe models without the need for consent banners, the regulation creates a tangible

market reward for privacy-by-design—simultaneously supporting business viability and delivering the banner-free experience users expect.

Recommendation:

- We encourage co-legislators to expand Article 88a to exempt low-risk, non-profiling activities that are essential for the digital ecosystem. Specifically, this should cover:
- **Contextual Advertising**: A privacy-centric model where advertisements are not behavioural and mostly selected based solely on the specific content the user is currently viewing (e.g., displaying running shoes on a marathon training article).
- **Associated Processing**: These are the essential "plumbing" functions required to deliver *any* digital advertising—even contextual ones—safely and efficiently including (but not limited to):
    - *Frequency Capping:* Strictly limiting the number of times a user sees the same advertisement to prevent saturation and user fatigue.
    - *Ad Fraud & Brand Safety:* Detecting and blocking invalid traffic (bots) to prevent financial theft, while ensuring ads do not appear alongside illegal or harmful content.
    - *Measurement*: Verifying that an advertisement was actually loaded and seen by a human (viewability), providing the essential "proof of delivery" that allows advertisers to validate their spend.
    - *Third-party Analytics*: Gathering aggregate statistics on campaign performance and audience reach.

3. Explicitly recognize cookie banners are not required where only essential cookies are used

Current regulatory ambiguity creates a transparency trap. Even when cookies are strictly necessary and legally exempt from consent, organizations often display banners anyway just to meet unclear "notice" requirements. This defensive compliance adds a completely unnecessary layer of fatigue for users, cluttering the web with pop-ups for technical processes that do not require user action.

Recommendation: Clarifying that transparency obligations for non-consent cookies can be satisfied via the website's privacy policy, rather than a dedicated banner. This simple clarification provides immediate legal certainty for SMEs, reduces visual clutter for users, and ensures all data processing information is centralized in one easy-to-find location without lowering privacy standards.

4. On the “Automated and machine-readable indications of data subject’s choices” (i.e., browser-level consent)

The European Commission has correctly identified the problem: Europe needs a simpler cookie rulebook that reduces "consent fatigue" for users. However, by resurrecting the idea of a browser-level consent, the European Commission fails to learn lessons from prior reform efforts, where this approach failed

twice (i.e. 2017 ePrivacy Regulation Proposal; 2023 Cookie Pledge). The lack of legal and technical workability of Art.88b introduces profound uncertainty that risks undermining existing ads-funded business models and the open internet.

- **A browser-level consent fails to meet fundamental GDPR consent requirements and will do little to reduce cookie banners.** From a user perspective, a binary browser signal cannot meet the GDPR's requirement for "specific" consent without presenting an interface so complex that it becomes unusable. A simple "Accept All" setting renewed every six months fails to provide the granular detail required for "informed" consent across millions of unique websites and purposes. Consequently, websites will still be legally entitled to display banners to request specific permissions or overrides, meaning the number of banners is unlikely to decrease. Finally, browsers will never be in a position to monitor websites' individual practices, opening the door to fraud and misuse of data.
- **A poorly designed centralised consent mechanism will have consequences for the economic viability of the ad-supported open web**. It will drastically impair the ability of most websites to monetize content and drive client acquisition. The resulting low consent rates and severely restricted data access—driven by limits on personalization—will precipitate a sharp decline in advertising value and sales conversion.
- **This will also inadvertently jeopardize contextual advertising models.** Essential "plumbing" tools for contextual ads—like frequency capping, fraud prevention, and measurement—rely on cookies and would be blocked from operating. The resulting hit to revenue will inevitably force publishers towards paywall models, creating a two-tiers internet and undermining the free, open web.

Recommendation: We strongly urge the co-legislators to delete Art.88b. They should focus their efforts on expanding Art.88a to better incentivize contextual advertising and exempt low-risk activities from consent banner requirements.

5. On the media exemption to the new browser-level consent (Art.88b)

We note that the European Commission has proposed a "media exemption" in an attempt to protect press publishers from the clear economic impact of its browser-level consent mechanism. This would allow media websites to ignore users' browser signals and initiate their own consent request.

However, this carve-out fails on its own terms because the digital ecosystem is interconnected. Media publishers rely on third-party advertisers and ad-tech vendors (who are not exempt from Art.88b) to place cookies for essential functions like retargeting, ads measurement or frequency capping. Furthermore, publishers rely on cross-site data from non-exempt sectors (like retailers) to validate and value their ad inventory. If the retail sector is blocked from using measurement tools, the media sector loses the data flows necessary to monetize their content effectively. This creates a market distortion that

risks concentrating ad spend on large, privileged players while crushing independent publishers and SMEs.

Recommendation: We strongly caution co-legislators against this approach, which risks creating a two-tier internet, arbitrarily picking winners and losers without achieving long-term positive change. We believe the Digital Omnibus should focus on creating a level playing field that enables safe data use for all actors. This means focusing on the root cause of consent fatigue by exempting low-risk activities—such as contextual advertising—from consent requirements entirely.

Data Act

While we welcome the European Commission's ambition to streamline the Data Act, the proposed revision unfortunately misses the mark on several critical points, leaving significant legal conflicts unaddressed and maintaining barriers to innovation. To truly unlock the value of data in the EU, the co-legislators must address these structural flaws.

1. Resolving the Data Portability Conflict Between the Data Act, DMA, and GDPR

The Digital Omnibus fails to address a significant legal contradiction that paralyzes the GDPR right of data portability. Article 5(2) of the Data Act bars companies designated as gatekeepers under the DMA from acting as data recipients. This prohibition is in direct conflict with a user's fundamental right to port their data under Article 20 of the GDPR and with the continuous data portability obligations enshrined in Article 6(9) of the DMA.

This legislative clash places data holders in an impossible position. When a user requests to transfer their data to a service operated by a designated gatekeeper, the data holder is faced with a legal trilemma: honoring the user's request could breach the Data Act, while refusing it would violate their clear obligations under the GDPR and the DMA. This stalemate negates the user's control over their own data, creates friction in the market, and directly undermines the EU's goals of fostering competition and preventing user lock-in.

Recommendation: We encourage the co-legislators to resolve this legal conflict via the Digital Omnibus. The text should establish a clear hierarchy of rules, affirming that the restrictions in the Data Act cannot invalidate the pre-existing, fundamental rights of users under the GDPR or the explicit obligations for gatekeepers under the DMA. It is essential to clarify that companies should facilitate all user-directed portability requests to ensure that the right to data portability remains effective and that the pro-competitive objectives of the digital rulebook are upheld.

2. Address Duplicative Rules for Consumer Data

The Data Act's obligations for consumer products (Chapter II) overlap almost entirely with consumer rights already established under the GDPR, such as data access and portability. This creates conflicting legal regimes that are difficult for businesses to navigate and confusing for consumers, without providing meaningful additional protections. The current layering of rules adds compliance costs without tangible benefits for the user, effectively penalizing EU businesses with double regulation for the same data.

Recommendation: We encourage the co-legislators to simplify the framework by carving consumer data out from the scope of the Data Act's Chapter II portability and information requirements. The GDPR already provides a robust, well-understood framework for consumer data rights. Relying on the GDPR as the single standard would eliminate legal fragmentation and reduce unnecessary compliance overhead.

3. Clarify and Narrow the Scope of "Connected Product"

The Act's definitions of "connected product" and "related service" are currently so broad that they could apply to nearly any item that generates data, including general-purpose devices like smartphones and PCs. This ambiguity expands the Act's scope beyond its intended focus on industrial data, creating legal uncertainty and significant compliance burdens for a vast range of consumer technologies. Applying industrial data sharing rules to personal consumer devices ignores the fundamental differences in how these products are used and regulated.

Recommendation: We encourage the co-legislators to amend the definition of "connected product" in Article 2(5) to explicitly exclude devices whose primary function is not the storing and processing of data on behalf of others, such as personal computers, tablets, and smartphones. This provides immediate legal certainty and focuses the Act's obligations where they are most relevant.

4. On the limitations for re-use of open government data

The European Commission proposal allows public sector bodies to charge higher fees for the re-use of open government data to very large enterprises and gatekeepers. We believe this approach is counterproductive to the EU's goals of fostering a data-driven economy.

Discriminatory pricing creates artificial barriers that prevent some of the most capable technical actors from building valuable services on top of public data, ultimately reducing the societal return on that data. This issue extends far beyond the developers of the models themselves. Frontier AI models built by the leading AI Labs form the foundational layer for a vast ecosystem of European developers, startups, and established companies. If these foundational models are prevented from learning from European public data due to prohibitive costs, their utility within the EU is severely diminished. This trickle-down effect undermines the entire European ecosystem that depends on these base models, hurting local innovation and placing domestic companies using these models at a competitive disadvantage.

Recommendation: We encourage the co-legislators to remove discriminatory fee provisions for companies designated as gatekeepers under the DMA. Ensuring a level playing field where all market actors can access open government data on equal terms is essential to support the entire AI value chain and ensure that models accurately reflect European society.

Cybersecurity

1. On the Single-entry point

We support the introduced single-entry point for incident reporting unifying NIS2, GDPR, DORA, eIDAS, and CER into one EU-level reporting system established and managed by ENISA. This is a good **first step** to streamline the reporting process and is essential for reducing compliance costs and fragmentation.

It is also encouraging that the proposed Article 23a(3)(f) explicitly affirms the 'report once' principle, stating that a single notification submitted via the single-entry point must be capable of satisfying the reporting requirements for that same incident under any other relevant Union legal acts. This ensures that information flow to competent authorities is maintained without duplicating the workload for regulated entities.

While the proposal acknowledges the CRA, the current drafting of Article 23a(1) states only that ENISA 'may ensure' that the single-entry point builds on the single reporting platform established under the CRA. To prevent technical fragmentation and support a truly unified reporting landscape, this provision should be strengthened. The regulation should mandate that the single-entry point be fully integrated with, or built directly upon, the CRA platform. Ensuring the underlying platforms are technically identical or fully integrated is the logical next step to streamline incident reporting.

The proposal mandates that ENISA develop the technical, operational, and organizational specifications for the single-entry point in cooperation with 'the Commission, the CSIRTs network and the competent authorities'. Currently, the text does not mention consulting entities that will be required to use the system to submit incident notifications. The entities submitting notifications must be involved in designing and assessing the functioning of the single-entry point. This is particularly critical given the requirement in Article 23a(3)(d) that the system must be interoperable with European Business Wallets for identification and authentication. Excluding the primary users (reporting entities) from the design phase risks creating a system that is theoretically sound for authorities but operationally burdensome or technically incompatible with industry workflows. We call for an amendment to Article 23a(3) to explicitly include industry stakeholders in the consultation process regarding technical specifications and API standards.

The proposal establishes a mechanism for the Commission to assess the 'proper functioning, reliability, integrity and confidentiality' of the single-entry point. If the system is found lacking, ENISA is required to take corrective measures. While top-down assessment is necessary, it is insufficient for real-time operational stability. There should be an official mechanism in place that allows stakeholders, specifically entities submitting notifications, to flag issues directly to ENISA or the European Commission regarding

the functioning, reliability, integrity, or confidentiality of the single-entry point. Operational bugs or security flaws in the reporting pipeline may be detected first by reporting entities rather than the Commission. A formal feedback loop is essential to ensure that if the single-entry point fails, entities can swiftly transition to 'alternative means' of reporting as referenced in Recital 57, without facing legal liability for missed deadlines.

Recommendation: To ensure the system is technically robust and operationally efficient, we recommend strengthening the single-entry point provisions regarding **integration with the CRA single reporting platform**, mandated **stakeholder consultation** during the design phase, and the establishment of formal **feedback mechanism** on the functioning, reliability, integrity, or confidentiality of the single-entry point for reporting entities.

2. On the scope of the proposal

While the proposed Digital Omnibus Regulation introduces a single-entry point for incident reporting, which we strongly support, this measure is only a good first step and primarily a technical solution rather than a substantive harmonisation of rules. True simplification requires addressing the underlying legal fragmentation. The current proposal is **disappointing in its ambition** because it leaves the distinct reporting logic of various regulations intact. For example, it amends Article 33 of the GDPR to introduce a 96-hour notification window for personal data breaches likely to result in high risk, while NIS2 generally operates on a 24-hour early warning and 72-hour incident notification timeline. This creates a regulatory landscape where an entity suffering a cyberattack involving personal data must track two different countdown clocks for the same event.

The proposal also stops short of addressing diverging severity thresholds: determining if an incident is 'significant' under NIS2, 'severe' under CRA, 'major' under DORA, or poses a 'high risk' under the amended GDPR. This lack of alignment diverts critical resources away from actual incident response toward complex administrative assessments of which legal threshold has been met.

The proposal acknowledges the success of the financial sector's harmonisation, noting in Recital 54 that the Commission should 'take due account' of the common templates developed under DORA when developing specifications for the single-entry point. However, this falls short of a binding mandate to harmonise the reporting templates. It is important that entities are able to submit one unified report using a common template rather than merely a single submission flow that still requires completing different templates for different legal acts.

To make simplification meaningful, the regulation should expand its ambition by standardising the reporting clock and ensuring reporting templates are harmonised. We advocate for a universal 72-hour reporting deadline: to promote consistency while allowing scarce incident response personnel to focus on responding to an incident. Furthermore, the 'reporting clock' should only start once an incident is confirmed, preventing precautionary over-reporting that floods authorities with noise and takes away from incident response itself.

Despite our concerns regarding the ambition of the current text, we are hopeful for further simplification efforts. We specifically welcome the prospective review of the **Cybersecurity Act (CSA)** mentioned in Recital 33, and look forward to a more ambitious approach to simplifying the cybersecurity framework looking at risk aligning management measures, documentation, audits and other compliance obligations beyond reporting, including in the **implementation of CRA** (e.g., aligning the application with the availability of standards and removing dangerous obligation to report unpatched vulnerabilities)**.** These future steps are vital, including to reduce the burden for regulatory authorities themselves, however, they should not serve as a justification for delaying the necessary harmonisation of timelines and thresholds in the current Digital Omnibus text.

Recommendation:

- We urge the co-legislators to standardise reporting timelines around a single **72-hour reporting deadline** for all regulations covered by the single-entry point. Furthermore, the regulation should **harmonise the trigger point for these timelines** ensuring the clock starts only when an incident is confirmed, rather than merely suspected. The Digital Omnibus proposal should also **align thresholds** for reportable incidents across the digital acquis and mandate the creation of a **harmonised reporting template**. Importantly, further simplification efforts need to be taken beyond the incident reporting, including **simplifying CRA compliance**.

AI Act

The Commission's focus on creating a clear and innovation-friendly AI framework is a positive step that will help Europe lead in responsible AI. We welcome the goal of refining the AI Act through targeted simplifications, such as streamlined high-risk assessments and broader data processing for bias mitigation. We also welcome the pragmatic amendment to Article 4 (AI Literacy), which shifts the burden from a direct obligation on providers to a state-led encouragement model, reducing administrative liability for businesses. These are crucial steps toward providing the legal certainty necessary for the development and deployment of world-class AI in the Union.

Three additional improvements to the AI Omnibus can further simplify the AI Act and provide greater legal clarity without fundamentally changing the law's objectives:

To ensure AI systems providers and deployers can effectively implement complex provisions that are not yet finalized, the Omnibus must address unrealistic compliance deadlines in a principled, streamlined way across the Act. Specifically, for Article 50 regarding generative AI provenance and transparency, the current proposal only grants six months for compliance, applies only to a small fraction of requirements (Art. 50.2, i.e. the marking obligation), and covers only AI systems already on the market. Given that the Code of Practice - which is intended to enable the implementation of both of these requirements - is likely to be finalized just weeks before August 2026, this timeline is insufficient. In particular, the current August 2026 timeline for Article 50(4) - i.e. the deep fake labelling obligation - is impossible from a technical

development and implementation point of view. Both providers and deployers will need to understand the specific requirements of what constitutes a “deep fake” first, and the limited time between the conclusion of the Code and compliance date will be insufficient to develop and integrate the right technical tools.

We support the proposal's mention of December 2027 as a compliance milestone for high-risk AI systems under Annex III. However, Article 113 currently structures this as a conditional timeline: compliance is triggered 6 months after a Commission decision confirming that 'adequate measures' (standards) are available. If no decision is made, a 'backstop' applies (Dec 2, 2027 for Annex III systems; Aug 2, 2028 for Annex I systems). This dual-trigger mechanism creates a 'moving target' that lacks the firm predictability essential for compliance planning. Leaving this date dependent on administrative decisions lacks clear guardrails and hinders industry readiness (and will have a severe impact on public authorities as well).

Under the proposed amendments to Article 75, the Commission is set to become the sole organizer of pre-market conformity assessments for certain high-risk AI systems (namely where based on 1P GPAI models or themselves constituted or integrated into VLOPs and VLOSEs under the DSA). While the Commission may entrust notified bodies to assist, this is purely at the discretion of the Commission. We are also concerned about the lack of independence (including from the Commission and AI Office themselves) and quality safeguards in the current drafting. Under product safety rules, such bodies are traditionally subject to stringent requirements, as also evident from the rules in Section 4 of Chapter 3 of the AIA.

The current approach creates a single point of failure and potential bottlenecks for high-risk systems. We recommend amending this to ensure providers retain the flexibility to choose accredited third-party conformity assessors directly, avoiding a lock-in to the European Commission that is inconsistent with standard EU product safety legislation.

Recommendations:

1. **Extend deadlines for transparency provisions (Art. 50 of the AI Act) to align with the full scope of the currently negotiated Code of Practice**: To provide legal clarity, the Omnibus should extend the compliance period for Article 50 to one year (rather than six months). This extension should apply to new systems as well as those already on the market and must cover all relevant transparency obligations (Article 50(4) in addition to Article 50(2)), and require changes to Article 113 of AI Omnibus.
2. **Establish a firm application date for high-risk AI systems:** We recommend that the Omnibus establishes December 2027 (for Annex III) and August 2028 (for Annex I) as the firm application date for high-risk systems, rather than leaving it as a flexible target dependent on administrative decisions.

3. **Ensure flexibility for choosing the high-risk conformity assessors to prevent bottleneck and delays in Article 75**: In addition to applying the traditional rigorous quality and independence standards to the Commission's own conformity assessment body, the law should ensure flexibility for providers to choose from among the different conformity assessors (both at the EU level and in Member States), avoiding a lock-in to the European Commission that could lead to bottlenecks and inconsistencies with EU product safety legislation.