Meta Submission in Response to the EDPB Guidelines 3/2025 on the interplay between the DSA and the GDPR

Executive Summary

While Meta Platforms Ireland Limited ("**Meta**") recognises the EDPB's role in providing data protection guidance under GDPR, we have serious concerns that these guidelines misinterpret key GDPR provisions and inappropriately extend data protection concepts into content regulation domains outside of EDPB's remit. The guidelines treat content regulation as subordinate to data protection rather than recognising the distinct competencies and objectives of each. Additionally, operational realities are dismissed in favour of theoretical privacy maximalism that creates more complexity and ultimately undermines both safety and user experience. This approach risks creating fundamental conflicts between safety and privacy objectives rather than enabling their harmonious application.

The draft guidelines' expansive interpretation of Article 22 GDPR, which extends automated decision-making restrictions far beyond the letter of the law and established jurisprudence is highly problematic. Article 22 GDPR is misapplied to routine platform operations essential for user safety that lack the "legal effects" or "similarly significant impacts" the provision was designed to address. This misinterpretation reflects a broader pattern in the draft guidelines, where privacy considerations are treated as automatically overriding other legitimate objectives, including safety, contrary to the GDPR's own balancing principles and the proportionality principle in the Charter of Fundamental Rights of the EU. The operational implications are severe given that large platforms process billions of content moderation decisions. The guidelines' approach would require human review of automated safety functions at an operationally impossible scale and runs directly counter to the EU Digital Services Act's (DSA) requirements for swift responses to illegal content.

Equally problematic is the draft guidelines' introduction of novel and unsupported expectations for recommender systems, particularly by prohibiting any collection or processing of personal data for profiling purposes when a non-profiling option is active. The guidelines further overreach by mandating equal presentation of profiling and non-profiling options and misapplying Article 22 GDPR to ordinary recommender system operations. These suggestions are not found in either the DSA or the GDPR, create legal uncertainty, and threaten the design of intuitive, user-friendly interfaces, ultimately frustrating the DSA's goal of fostering a more open and user-centric digital ecosystem.

These draft guidelines risk transforming the GDPR from an intended balanced and proportionate data protection framework into a general restriction on any algorithmic processing, thereby compromising safety, stifling innovation and violating the right to conduct a business. Meta strongly urges substantial revisions to ensure the guidelines serve their intended purpose without creating unworkable regulatory conflicts and operational impossibilities.

1. Shortcomings of the Regulatory Approach

1.1 Regulatory Overreach: Exceeding Mandate and Process Shortcomings

The EDPB has exceeded its mandate by issuing guidelines that stray far beyond the scope of personal data protection into content regulation. The DSA and GDPR regulate different objectives; the DSA aims to create a safe, predictable, and trustworthy online environment, while the GDPR ensures robust protection of personal data. Many service providers, including Meta, are subject to both frameworks. Unfortunately, the draft guidelines conflate the distinct regulatory objectives of the DSA and GDPR. Blurring the line between content regulation and data protection undermines regulatory clarity, may risk violating the rule of law, and creates unnecessary complexity for platforms and users alike.

Absence of Meaningful Collaboration or Industry Pre-Consultation

Disregarding questions regarding the EDPB's competence to issue these guidelines, the EDPB appears to have taken this action without sufficient collaboration with other relevant EU authorities, including the Digital Services Coordinators. This stands in marked contrast to the approach taken in recent joint guidelines relating to the Digital Markets Act (DMA).

In addition, the EDPB should have ensured meaningful industry pre-consultation during the development phase. The absence of such engagement resulted in a failure to address industry and stakeholder requests for clarification on the GDPR implications of key DSA provisions, where the interplay between the two frameworks remains genuinely unclear and in need of further guidance. For instance, Article 21 DSA (out-of-court dispute settlement) and Article 40 DSA (researchers' access to data) both involve complex data processing considerations, raising fundamental questions about data sharing, data minimisation, and the implementation of appropriate safeguards, as well as significant operational measures for relevant online platforms. Platforms have specifically sought guidance on these issues, but the lack of meaningful pre-consultation has left unresolved critical questions regarding how the DSA obligations can be implemented while complying with GDPR.

Instead of addressing these increasingly important implementation challenges, the EDPB chose to focus on areas where existing legal frameworks and established practices already provide reasonable clarity, such as the scope of article 22 GDPR, basic transparency obligations and well-established data protection principles. This approach demonstrates a disconnect between the EDPB's priorities and the genuine compliance challenges that platforms encounter when seeking to implement GDPR requirements in the context of DSA compliance.

Finally, the guidelines fail to establish the necessary binding mechanisms or clear, structured processes for consultation, coordination, or dispute resolution. This procedural gap risks fragmented enforcement, inconsistent outcomes, regulatory overreach, and conflicting actions against platforms. More specifically, the guidelines omit practical detail necessary for global platforms to manage overlapping DSA and GDPR mandates, particularly regarding cooperation protocols between the EDPB and EBDS.

Dark Patterns: Overly Broad and Subjective Standards

The draft guidelines adopt an expansive and subjective definition of "deceptive design patterns", not based on any scientific research, that threatens legitimate user experience practices with legal uncertainty. By suggesting case-by-case assessment based on broad criteria such as "whether personal data is being processed," the guidelines risk capturing virtually all modern digital design practices within GDPR scope, creating regulatory overlap, legal uncertainty and violating the right to conduct a business.

The absence of objective criteria exposes legitimate UX design practices to arbitrary enforcement and wrongly equating legitimate digital commercial practices with coerced nefarious manipulation. For example, standard industry practices that enhance usability and engagement – such as gamification in educational applications, autoplay, or countdown timers for live events – are at risk of misclassification as "deceptive patterns" despite providing genuine user benefits. The draft guidelines conflate any behavioural influence with manipulation, failing to distinguish between harmful deception and persuasive design that preserves user autonomy and choice.

This approach demonstrates the draft guidelines' broader flaw: extending data protection regulation into design domains where safety harms, not privacy, are the primary concern. Effective regulation requires demonstrable evidence of privacy harm and objective criteria distinguishing harmful manipulation and deception from legitimate engagement. The draft guidelines' reliance on theoretical assumptions without empirical validation creates compliance uncertainty that will chill innovation and reduce service quality.

The EDPB's attempt to regulate "addictive" design features through the GDPR – when these are properly addressed by DSA provisions – exemplifies the inappropriate conflation of content regulation with data protection, adding regulatory complexity without addressing underlying issues. Similar concerns arise from the EDPB’s 2020 Guidelines on the targeting of social media users, which also adopted an expansive approach by treating persuasive design and engagement as data protection issues, with limited distinction between harmful manipulation and legitimate user experience design. Therefore the present regulatory conflation and overreach constitutes a pattern.

Recommendation

- Restore proper regulatory boundaries and acknowledge distinct purposes that content regulation and data protection serve. The DSA governs platform safety obligations; the GDPR governs personal data protection. Neither framework should be interpreted to undermine the other's legitimate objectives.
- Establish formal, binding processes for consultation, coordination, and dispute resolution between DSA and GDPR enforcement authorities, including inter-board protocols (EDPB/EBDS), to prevent fragmented enforcement, regulatory overreach, and conflicting actions.

- Undertake genuine consultation with industry to understand their operational realities and actual interpretive challenges with the GDPR and DSA.

1.2 Data Protection Maximalism and Failure to Consider Operational Realities

The draft guidelines treat data protection as paramount, thereby disregarding Recital 4 GDPR (which affirms the need to balance the right to protection of personal data with other fundamental rights) and the objective of the DSA to set out "harmonised rules for a safe, predictable and trusted online environment" (Article 1 DSA). Data protection maximalism cannot be the single standard against which all user and societal needs are measured – such an approach leads to disproportionate outcomes where user safety, product reliability, and democratic discourse are subordinated to theoretical privacy interests, rather than balanced against actual risks in a proportionate manner.

Data Protection Maximalism at the Expense of Operational Realities

The draft guidelines fail to acknowledge operational realities and to provide for practical guidance. Instead, the draft guidelines merely restate the GDPR and DSA texts, conflating obligations and missing the opportunity to provide practical guidance on how the GDPR should be interpreted in the context of the DSA. Importantly, the EDPB should focus on offering sensible GDPR guidance without over-emphasizing privacy at the expense of other considerations, and should refrain from interpreting the DSA itself.

For example, the expectation that providers should "avoid processing of personal data insofar as possible," misinterprets the GDPR principles of data minimization and privacy by design. These principles are grounded in proportionality and context, not in the absolute minimisation of data processing. In practice, most content moderation technologies – particularly those based on machine learning and deployed in environments where vast amounts of user-generated content are created every second – require access to significant amounts of data for training, operation, monitoring, and improvement. The draft guidelines' expectation that providers should "avoid processing unless strictly necessary," and always consider less intrusive means, fails to account for the DSA's mandate to, for example, inform issuing authorities without undue delay of the effect given to orders to act against illegal content or to enforce the services' terms and conditions.

In addition, data retention is particularly crucial for DSA-mandated systemic risk assessment and mitigation, where platforms must analyze patterns across time to identify emerging threats and coordinate responses. Strict processing limitations would prevent the longitudinal analysis necessary for effective systemic risk management and would ultimately undermine the effectiveness of integrity and safety tools, limiting platforms' ability to protect users from harm.

Inappropriate Expansion of Risk Management: DPIAs for Non-High Risk Activities

The draft guidelines inappropriately expand the requirement for Data Protection Impact Assessments (DPIAs) by making them a mandatory step for any identified systemic risk under the DSA - even when those risks do not constitute a high data protection risk. This goes beyond the GDPR's established

thresholds for DPIA necessity and risks making DPIAs a routine requirement for all systemic risk assessments, regardless of the actual privacy risk involved. Privacy is already considered within Meta's systemic risk landscape, as required by Article 34(1) DSA, and mitigation measures must be 'reasonable and proportionate' under Article 35 DSA. Imposing DPIAs for every systemic risk creates undue and significant compliance and operational challenges for global platforms, especially when managing a wide range of risks and regulatory requirements under the DSA rather than the GDPR. The absence of clear, GDPR-consistent criteria for when DPIAs are truly necessary increases legal uncertainty and may result in unnecessary and misdirected assessments.

Recommendation

- Data protection measures must be proportionate to actual risks and compatible with legitimate safety functions. Therefore, the draft guidelines should ensure that interpretations of GDPR principles - including of data minimisation and privacy by design - do not undermine the goals of the DSA by suggesting that the processing of personal data should be limited insofar as possible.
- The EDPB should not expand the role of DPIAs to non-high risk activities or make them a blanket requirement for DSA systemic risk management.

2. Three Critical Examples of the Guidelines' Shortcomings

2.1 Misapplication of Article 22 GDPR undermines Safety and Integrity

The guidelines adopt an impermissibly expansive interpretation of Article 22 GDPR that conflicts with the letter of the law, established jurisprudence, regulatory precedent, and the provision's legislative intent, threatening to capture automated processing explicitly excluded from its scope.

Article 22 GDPR establishes a narrow, sui generis prohibition on specific categories of automated decision-making ("ADM"). The provision's scope of application is limited by strict *cumulative* criteria: (1) a determinative decision affecting the data subject (C-634/21, paras. 44-46); (2) processing based solely on automated means, excluding meaningful human intervention (WP251rev.01, p. 21); and (3) legal effects or similarly significant effects meeting the threshold established in Recital 71 and refined through regulatory practice.

Legislative intent and judicial interpretation confirm Article 22's narrow scope. Recital 71 GDPR emphasises that the provision addresses "evaluation of personal aspects" with potential for discrimination, not routine processing operations. The CJEU's SCHUFA judgment (C-634/21, paras. 58-73) clarified that preparatory data analysis without determinative third-party reliance falls outside Article 22's prohibition.

Regulatory guidance and enforcement practice likewise adopt a restrictive reading, reinforcing that Article 22 is not a catch-all for automated analytics but a targeted safeguard for decisions with significant impact. The Article 29 Working Party and EDPB's consolidated guidance (WP251rev.01) confirms that all three Article 22 elements must be satisfied simultaneously, that most content personalization does not

constitute qualifying ADM, and that meaningful human involvement satisfies the intervention requirement. National DPA practice – including CNIL deliberation n° 2022-103 and Irish DPC enforcement practice – also emphasizes high thresholds for both "solely automated" processing and "similarly significant effects."

Overreach in the Draft Guidelines

The draft guidelines' interpretation exceeds these well-established regulatory boundaries. By suggesting that content moderation decisions, recommender system outputs, and platform integrity measures presumptively fall under Article 22, the guidelines: (1) conflate routine processing with qualifying ADM; (2) ignore meaningful human oversight through appeals processes and algorithmic governance frameworks; and (3) misapply the "significant effect" threshold to decisions that do not reach the level of "*decisions that affect someone's financial circumstances, such as their eligibility to credit; decisions that affect someone's access to health services; decisions that deny someone an employment opportunity or put them at a serious disadvantage; or, decisions that affect someone's access to education, for example university admissions.*"

This approach risks diminishing the established roles of the current DSA enforcement entities, including the European Commission and Digital Services Coordinators, potentially leading to reduced regulatory consistency and legal certainty by shifting authority away from these key regulators.

Conflation of Routine Processing with Qualifying ADM

The draft guidelines classify a wide range of core content moderation activities – such as automated detection and enforcement of illegal and harmful content, spam and fraud prevention, and other integrity tasks – as "solely automated" decision-making with legal or similarly significant effects under Article 22 GDPR. This approach disregards both legal precedent and operational necessity, treating established industry-wide trust and safety best practices as legally dubious and threatening the practical viability of core platform operations. Equally, the draft guidelines fail to recognise that many large-scale content moderation activities do not involve the processing of personal data as defined by the GDPR. For example, AI technologies may analyse content in an isolated and anonymized manner to identify harmful content and violations of applicable terms and conditions, without processing identifiable personal data. As a result, the guidelines undermine the balance between existing compliance frameworks, risk chilling innovation in content moderation technologies and methods, and run counter to the objectives of the DSA.

Automation is Key to DSA Compliance

The operational reality is clear: large-scale online platforms make billions of content moderation decisions annually – removing violating content, suspending malicious actors, and preventing spam or fraud. While the underlying process as such is subject to privacy review, only a small fraction of these decisions can reasonably be reviewed by human moderators or escalated for privacy assessment. The speed and scale required by the DSA are only achievable through automation; human review alone is insufficient. If the centrality of automated enforcement is not recognised, delays in enforcement will allow harmful and violating content to spread rapidly, reduce the volume of violations that can be

actioned, reduce the effectiveness of abuse and fraud prevention, and compromise user safety by prioritising a maximalist privacy interpretation over practical risk mitigation. Furthermore, the DSA already addresses transparency and accountability by requiring explicit disclosure of automated moderation decisions (Article 17) and ensuring meaningful human involvement through robust appeals and complaint mechanisms.[1]

Misapplication of the "Significant Effect" Threshold

The draft guidelines trivialise the threshold of impact required to trigger Article 22 GDPR, which is intended for decisions that affect a person's legal rights or have a substantial impact on their life – such as access to pensions, credit, or essential medicines. In the context of the DSA, safety or integrity decisions (e.g., content downranking, removal, or even user account suspension for relevant Terms violations) do not rise to the level of impact contemplated by Article 22 and have corresponding mitigations in the DSA for user redress. This expansive interpretation, which treats any decision that might upset or disappoint an individual as significant, risks undermining the effectiveness of content moderation – a core objective of the DSA. For example, excessive disclosure of automated decision-making logic could enable bad actors to game the system, undermining the very integrity and security objectives that both the DSA and platform contractual obligations seek to protect. Additionally, the draft guidelines' broad Article 22 interpretation inappropriately captures routine advertising functions. By suggesting that ad targeting constitutes automated decision-making with "significant effects" based on criteria including "intrusiveness of profiling, cross-device tracking, and use of known vulnerabilities," the draft guidelines ignore operational complexity and the established high thresholds for qualifying automated decision-making.

Recommendation

- The draft guidelines' interpretation of automated decision-making provisions contradicts the letter of the law, established jurisprudence and regulatory practice. Therefore, the guidelines should clarify that automated decision-making restrictions under Article 22 GDPR apply only to decisions that are both solely automated AND produce significant impacts on individuals' legal rights or fundamental life circumstances. More specifically, the guidelines should acknowledge that safety and integrity functions performed under DSA obligations do not constitute Article 22 qualifying decisions and recognise that content analysis using anonymised AI technologies falls outside GDPR scope entirely.

2.2 Recommender Systems: User Choice vs. Regulatory Overreach

Prohibition on Data Collection for Profiling

The draft guidelines suggest that any collection or processing of personal data for profiling purposes, while a non-profiling option is active, should be prohibited. That restriction is not found in either the DSA or the GDPR. On the contrary, the draft guidelines are improperly used as a vehicle to create entirely

[1] Please refer to Meta's Transparency Centre for more information about Meta's transparency reports mandated by the DSA or information about content enforcement and related user notifications.

novel expectations outside the legislative framework. Adding further requirements creates significant technical and product challenges, and negatively impacts user experience and expectations without delivering meaningful privacy benefits.

Among the DSA's core objectives are to enhance content-related control, transparency, and safety in the online environment, empowering individuals to make informed decisions about their digital experience. However, the prohibition on data collection for profiling purposes would actually limit users' control and experiences. In practice, users are accustomed to switching between different types of feeds on their own volition – algorithmic, chronological, or close friends – without expecting that using one option will degrade their experience with another. And yet, the approach in the draft guidelines would mean that users who choose non-profiling options will receive a degraded experience each time they switch back to the standard recommender system.

In addition, immediately ending data collection each time a user toggles between profiling and non-profiling options presents substantial technical hurdles for systems designed to provide seamless experiences, particularly when users switch between these modes multiple times within a single session.

Mandatory Equal Presentation and Opt-in Requirements

The draft guidelines further suggest that (i) profiling-based and non-profiling-based recommender options shall be presented equally to users on first use, and (ii) a profiling-based recommender system may only be used after the user has actively chosen it. These provisions fall outside the GDPR's remit and go beyond the DSA's text, which only mandates that at least one non-profiling option be provided and accessible, without specifying when options must be presented or requiring prior opt-in. This represents an overreach, creates legal uncertainty and may hinder the design of intuitive, user-friendly interfaces, affecting a key aspect of the service Meta aims to provide.

Moreover, the guidelines appear to assume that a service would, by default, become a generic information platform devoid of social context, rather than a personalized social network. The 'equal presentation' expectation would force fundamental redesigns of user interfaces that are less aligned with user expectations and market practices, effectively transforming personalized social networks into generic information services. This would force platforms to redesign user flows in ways that are misaligned with user expectations, reduce service quality, and restrict Meta's ability to offer a legitimate, personalized experience in accordance with its Terms of Service and unjustly interfere with Meta's fundamental right to conduct a business.

Misapplication of Article 22 GDPR to Recommender Systems

Finally, the draft guidelines repeatedly imply that any recommender system could fall under Article 22 GDPR, which is incorrect. This interpretation wrongly expands Article 22's scope to include ordinary profiling activities. Such a broad reading could capture a wide range of content decisions not intended to have legal or similarly significant effects. The absence of clear criteria for what constitutes a "significant effect" further increases legal uncertainty and operational risk.

Recommendation

- The draft guidelines introduce new compliance expectations that are not grounded in the text of either DSA or the GDPR – namely: (1) prohibiting data collection for profiling while non-profiling options are active; and (2) mandating "equal presentation" of recommender options. These points should be removed to avoid implying new legal obligations through soft-law guidance, as their inclusion represents regulatory overreach that transforms interpretive guidelines into unauthorized lawmaking and undermines both GDPR legal certainty and DSA compliance.

2.3 Protection of Minors: Minors' Safety Undermined by Privacy Absolutism

The draft guidelines do not adequately recognise the need for robust, ecosystem-wide age assurance to protect minors online. In order to comply with the requirements of the DSA and other relevant laws to ensure the safety and protection of minors on their platform, companies must have the flexibility to choose and implement age assurance methods that are necessary and proportionate, in absence of uniform and specific rules and standards.

Age Data Restrictions Impact Minor Safety

By dissuading the permanent storage of age or age range data, the draft guidelines make it impossible for platforms to comply with the DSA and other legal obligations that require evidence of users' ages. This restriction undermines any business' ability to provide a safe and age-appropriate experience for minors, including the application of differentiated settings, content controls, interaction controls, parental supervision tools, etc. For example, age data storage enables critical differentiated protections between early and late teens, such as automated content controls that downrank diet tips or exercise content inappropriate for younger teens, while allowing age-appropriate wellness content for older adolescents, as well as between minors and adults, which is essential to understand, prevent or avoid risky interactions where minors' vulnerability might be exploited by bad (especially among adult) actors. These core protections require persistent age data to ensure consistent application across sessions and to manage transitions as users mature. Without age data retention, platforms cannot provide the age-appropriate safeguards that effective minor protection requires in accordance with the DSA.

The inability to retain age data also increases the risk of circumvention by minors and limits the effectiveness of parental controls and transitions as users reach adulthood. Furthermore, the draft guidelines' suggested prohibition makes simultaneous compliance with divergent regulatory requirements impossible, creating legal uncertainty and fragmented user protection standards across jurisdictions.

While the draft guidelines require that all processing for age assurance be strictly necessary and proportionate, they do not acknowledge the respective trade-offs between privacy, safety, and regulatory obligations. On the contrary, a consistent, EU-wide, privacy-preserving approach – such as age assurance at the app store or operating system level, with age attributes shared securely with relevant services – could help balance these competing interests.

Prohibition of Automated Moderation for Minors

The draft guidelines also state that moderation decisions based solely on automated processing, including profiling, are generally prohibited for minors. This is not rooted in the GDPR and disregards the fact that automated tools are critical for ensuring safety and age-appropriate experiences at scale. Therefore, it contradicts the DSA's objective. For instance, automated content controls that downrank material unsuitable for teens, such as certain extreme diet contents, would not be feasible without automation. Requiring human review in all cases is not practical and could reduce the effectiveness of content controls. It would create additional compliance challenges and, most importantly, deprive minors of the protection they deserve – protection that only automated tools can provide at scale.

Lack of Clarity on "Necessary and Proportionate" Processing

Although the draft guidelines recognize that Articles 28(1) and (2) DSA can serve as a legal basis for processing under Article 6(1)(c) GDPR, they fail to provide legal certainty on what constitutes "necessary and proportionate" processing for moderation and age assurance activities. This discourages the implementation of robust assurance methods and further increases uncertainty for platforms seeking to comply with both the DSA and other child safety laws.

Need for Unified Age Verification Regulation

Robust legislation is needed to specify how age should be verified, ideally as part of a digital age of majority framework for accessing online services, with parental approval required for teens below this threshold. Effective age assurance requires the introduction of unified regulation to clearly define criteria and methods, enabling platforms to select proportionate solutions that comply with legal requirements and ensure child safety. The most privacy-protective and effective solution would be to verify age once—at the operating system or app store level—and allow relevant services to rely on a secure age attribute or group.

Recommendation

- The draft guidelines should support age assurance approaches that balance fundamental rights, proportionality and effective safety protection. They should also recognise that age data retention is indispensable for minor protection and compliance with multiple regulatory frameworks, and acknowledge that retaining verified age attributes is necessary and proportionate to implement consistent differentiated content controls and manage user transitions (e.g., early vs. late teens), ensuring persistent safeguards.

4. Conclusion

The current draft guidelines contain fundamental interpretive errors that require substantial correction to avoid creating unworkable regulatory conflicts with both the GDPR and DSA, compromising regulatory clarity and user safety.

Meta remains committed to industry-leading privacy protection and user safety, demonstrating through our practices that these objectives are complementary when approached with appropriate balance. However, interpretations that create false choices between privacy and safety or that extend data protection restrictions beyond their proper scope cannot be accepted.

While the EDPB has a role in clarifying genuine areas of regulatory uncertainty within the scope determined by the GDPR, these draft guidelines exceed that role by misinterpreting established provisions and creating new restrictions that lack legislative basis. Substantial revisions are necessary to ensure the guidelines contribute to, rather than undermine, the coherent application of EU digital regulation.

We believe that, in order to achieve workable, proportionate, and effective guidelines and to meet the objectives of both the GDPR and DSA, a genuine and thorough process of engagement with industry, the European Commission, and relevant stakeholders is required. Such engagement is vital to ensure that guidance reflects the practical realities of content moderation and strikes a better balance between the need to protect personal data and promote user safety.