European Law Blog, 2026 · ISSN 2950-1970
Published: 19 October 2026
DATA PROTECTION AND DIGITAL GOVERNANCE

A Tempting Shortcut: How the Digital Omnibus Would Re-engineer the Legitimate Interest Test for AI

Pablo Trigo Kramcsák and Lorenzo Dalla Corte
Pablo Trigo Kramcsák — Research Group on Law, Science, Technology & Society (LSTS), Vrije Universiteit Brussel (VUB) · Lorenzo Dalla Corte — Tilburg University

To be lawful, processing of personal data must rest on one of the six grounds in Article 6 of the General Data Protection Regulation (GDPR). Among them, Article 6(1)(f), legitimate interest (LI), requires a controller, case by case, to show that it pursues a defined interest, that the processing is necessary to serve it, and that the data subject's rights do not override it. The Court of Justice (CJEU) treats these conditions as cumulative (see *Rīgas satiksme*, para 28; *M.I.C.M.*, para 106; and *Meta Platforms v Bundeskartellamt*, paras 106–112). This test, and the balancing exercise it relies on, has remained largely unaltered from the Data Protection Directive (DPD) to the GDPR. It may now change.

Following the Draghi and Letta reports, and aiming to foster AI development, the European Commission proposed the 'Digital Omnibus', a package of amendments meant to simplify a number of EU 'digital' laws, including the GDPR. Among its many changes, the package would insert a new Article 88c into the GDPR. Under this provision, where processing is necessary for the controller's interests in the context of developing and operating an AI system or model, such processing 'may be pursued for legitimate interests' within the meaning of Article 6(1)(f). The explanatory memorandum of the Digital Omnibus presents this as 'clarifying certain aspects as to the processing of data for AI training and development'.

The provision has had an eventful year in the Council of the European Union. The Cypriot Presidency compromise of February 2026kept it. In May the Presidency demoted it to a recital, and on 3 September the Irish Presidency, which took over on 1 July, brought the provision back

into the enacting terms, renumbered as ‘Article 88 bis’. At the time of writing, the latest version is the text of 2 October, the Presidency’s proposed mandate for negotiations with the European Parliament.

On one reading, the proposed AI clause simply codifies the position the European Data Protection Board (EDPB) reached in Opinion 28/2024, which accepted that LI can in principle support the development and deployment of AI models, subject to a case-by-case assessment and without treating that basis as a default (paras 60 and 66–69). If LI is already available for AI, and the provision creates no new rights or obligations, writing it into the GDPR would work no material change in the law. The provision may nonetheless do something more consequential than its drafters acknowledge. It would leave the wording of Article 6(1)(f) untouched, and the three-part test with it, while altering the conditions under which that test operates. The Council’s leaner Article 88 bis softens the problem without solving it, and its most useful clarification sits in a recital that will not appear in the consolidated GDPR.

Our concern is that the legislature would settle in advance part of what Article 6(1)(f) requires controllers to establish case by case. We begin with the tension between that legislative pre-assessment and a test the Court has always treated as contextual. We then consider how far Article 88c reaches and trace its trajectory through the Council up to the 2 October compromise, before asking whether our critique survives the provision in its latest form.

Legitimate interest and proportionality

Article 6(1)(f) gives controllers no shortcut to the balancing stage. They must first say what interest they pursue and where its limits lie, and they carry the burden of proof at every step. That interest must satisfy three requirements: it must be lawful under the wider legal order, real and present rather than speculative, and precisely articulated. A controller cannot weigh an interest it has not defined, and an interest pleaded in broad or shifting terms cannot be tested for necessity or balanced against the position of the data subject. This first assessment drives the necessity and balancing stages that follow.

The LI test is steeped in proportionality reasoning. Its conditions echo the tripartite structure of the *lato sensu* proportionality test (suitability, necessity, and proportionality *stricto sensu*)

that, while originating in German constitutionalism, is now used to resolve conflicts between constitutional and quasi-constitutional rights across many jurisdictions. In his Opinion in the *Rīgas satiksme* case, Advocate General Bobek made much the same point, describing the necessity requirement in Article 7(f) DPD as 'a type of proportionality' (para 61). If one frames the processing of personal data as an interference with the right to data protection, as the CJEU seems inclined to do, the parallel becomes clearer. Assessing whether the controller's interest is legitimate then amounts to asking whether the interference is *suitable*, i.e. whether it pursues a 'proper' purpose through means rationally connected to it. The processing must also be *necessary*, meaning that no suitable alternative interferes less with the rights of the persons involved (see, among others, AG Emiliou's Opinion in *RRC Sports*, para 175). Whether the data subject's rights override the controller's interest is, in turn, a question of proportionality *stricto sensu* (ibid, paras 180 ff).

Providing a (qualified) presumption of legitimacy for the interests underlying the development and operation of AI systems and models, as the Commission's Article 88c would have done, is analogous to presuming, under Article 52(1) of the Charter of Fundamental Rights, that a whole category of interferences pursues an objective the legal order recognises as worthy of protection, by means rationally suited to it. Formally, the presumption would be confined to the assessment of whether the interests pursued by the controller are legitimate, leaving the necessity and proportionality analysis unaffected. However, a default answer to that first question could colour the analysis that comes after it.

Article 88c would have settled the question of legitimacy for a macro-category, delimited only by a generic reference to a family of technologies. The controller would no longer need to articulate and defend a real, present and lawful interest, because the legislature would have done that work in advance for everyone developing or operating an AI system or model. AI covers an open-ended range of activities, which makes it the wrong unit of analysis for a test the CJEU and the EDPB insist must run case by case.

What the provision covers

The Commission's Article 88c spoke only of the interests of the controller, although Article 6(1)(f) also covers interests pursued 'by a third party'. Read literally, it would not reach a platform that licenses its users' posts to a model developer: the platform is the controller of that disclosure, but the interest in building the model belongs to the developer. The Council has closed the gap. Article 88 bis(1) refers to 'a legitimate interest of the controller or a third party'. Recital 33a, however, still speaks of 'a legitimate interest of the controller concerned' in the 2 October text.

Which activities the provision covers is less clear. Article 88 bis keeps the Commission's pairing of 'development' with 'operation', a term that neither the AI Act nor the Digital Omnibus defines. EDPB Opinion 28/2024 works instead with 'development' and 'deployment', treating the second as covering 'all stages relating to the use of an AI model', which 'may include any operations conducted after the development phase' (para 18). The Board treats a model's 'subsequent processing or operation' as following development (paras 7 and 13). Operation thus appears to sit within deployment. The Commission's own Recital 30, since deleted by the Council, spoke of the 'development or use (deployment)' of AI. In its September texts, the Council itself wavered: Recital 33a spoke of 'deployment' where Article 88 bis spoke of 'operation'. Only the 2 October text brings Recital 33a into line, and even then only partially. The text speaks of 'development and operation' (Article 88 bis and Recital 33a) and of 'development and *technical*operation' (Article 9(2)(k)), while Recital 33 also uses 'deployment'. Three formulations survive in a single instrument.

That choice matters because Opinion 28/2024 asks supervisory authorities, when verifying a controller's balancing test, to distinguish the risks facing data subjects whose personal data was used to develop a model from those facing data subjects whose personal data is processed once the model is in use (para 89). 'Operation' reaches the second group directly: the prompts, inputs and interactions of those who use the system, together with the personal data of third parties that those inputs may carry. Recital 33 limits the Council's Article 9(2)(k) derogation for residual special categories of data: it excludes data 'directly acquired by or provided to the AI system or AI model in the course of the deployment'. Article 88 bis contains no comparable limit. It gathers into a single formula two processing activities with different data subjects, different reasonable expectations and different risks, which a case-by-case balancing exercise

would assess separately.

The definitional link has also become less direct. The Commission's Article 88c cross-referred to the definition of an AI system in Article 3(1) of the AI Act. Article 88 bis drops that reference, although the Council keeps it in its Article 9(2)(k). What remains is the last sentence of Recital 33, which asks that the notions of AI system and AI model be 'understood in the same manner as in Regulation (EU) 2024/1689'. That works for AI systems. It helps little with AI models, since the AI Act defines only general-purpose AI models (Article 3(63)).

A rule written around a single technology also sits awkwardly in a regime designed to avoid exactly that. Recital 15 GDPR states that the protection of natural persons 'should be technologically neutral and should not depend on the techniques used'. Rules tied to specific technologies age badly. The proposed AI clause departs from that design, singling out AI through a dedicated statement about its lawful basis. Cloud computing also transformed the processing of personal data across borders. The Article 29 Working Party, and later the GDPR, dealt with it through the general framework rather than a cloud-specific rule. A dedicated AI clause would weaken the technological neutrality that has enabled Article 6 to govern forms of processing its drafters could not have foreseen, and nothing would then stop the legislature from accommodating the next heavily capitalised technology or politically favoured purpose. Article 6(1)(f) could drift towards a list of endorsed categories, with the balancing test retained as a formality.

A provision in search of its place: the Council's redrafting

So far we have addressed the provision largely as the Commission drafted it, drawing on the Council's texts only where they changed its reach. Our objection to legislative pre-assessment, however, has to hold against the version the Council is likely to carry into trilogue, and the Council has rewritten the provision several times since February. Along the way, its operative verb changed from 'may be pursued for' to 'may be regarded as' and then to 'may be carried out for'. Each verb claims a different degree of legislative endorsement.

The Council Presidency compromise text of 20 February 2026 left Article 88c intact. The compromise text of 21 May 2026 removed it from the enacting terms and moved a softened

version into a new Recital 33a, which said that AI-related processing ‘may be regarded as’ carried out for a legitimate interest of the controller. The GDPR uses that formula for direct marketing in Recital 47, conceding that the interest can qualify without asserting that it always does. For fraud prevention (Recital 47) and network and information security (Recital 49), by contrast, it says that the processing ‘constitutes’ a legitimate interest. The revised compromise text of 10 June, which the leaked text of 18 June left largely intact, then corrected two drafting slips. It replaced the conjunctive ‘and’ between the consent carve-out and the rights override with ‘or’, and the reference to ‘the most appropriate lawful ground’ with ‘an appropriate lawful ground’.

On 3 September the Irish Presidency changed course. Its revised compromise text reinstated the provision in the enacting terms as Article 88 bis, in three paragraphs. Paragraph 1 states that processing ‘in the context of the development and operation of an AI system or of an AI model may be carried out for a legitimate interest of the controller or a third party in accordance with Article 6(1)(f)’. Paragraph 2 adds that this ‘does not affect the obligation of the controller to choose an appropriate lawful ground of processing set out in Article 6’. Paragraph 3 subjects processing in the context of the development and operation of AI, ‘[r]egardless of the lawful ground of processing pursuant to Article 6’, to appropriate technical and organisational measures and safeguards. The new article omits the Commission’s qualifiers (‘where appropriate’, the consent carve-out and the rights override), which survive only in Recital 33a, and the Commission’s list of safeguards, including the unconditional right to object. The Commission’s Recital 31, which guided the balancing exercise for AI and also mentioned that right, had already gone. Recital 33a changed as well. It no longer says that the processing ‘may be regarded as’ carried out for a legitimate interest, but that it ‘may be carried out for’ one.

The text of 21 September, ST 13112/26 (not officially available), left the article untouched. It restored to the recital’s examples a safeguard that had previously been dropped: respect for data minimisation ‘during the stage of selection of sources’.

The 2 October text makes four changes relevant here. Article 88 bis(3) now ends with the words ‘including the right to object under Article 21’. In Recital 33a, ‘operation’ replaces ‘deployment’. A new sentence recalls that Article 6(1)(e) may be regarded as a lawful ground

where a public authority develops, ‘including training’, or operates an AI system or model used ‘in the performance of a task carried out in a public interest or in the exercise of official authority’. The recital then closes with a clarification: ‘the development and operation of an AI system or AI model does not in itself constitute a legitimate interest, and the applicability of Article 6(1)(f) to processing in the context of AI is to be assessed on a case-by-case basis, taking into account the underlying purpose of that processing’. Some slips, inherited from the Commission’s text, survive all this redrafting. The most visible is that Recital 33a still requires protection ‘against non-disclosure’ of residually retained data, when it plainly means disclosure.

What is left of the critique

The closing sentence of the 2 October recital comes close to conceding the core of our argument: where Recitals 47 and 49 GDPR state that processing strictly necessary for fraud prevention or network security ‘constitutes’ a legitimate interest, the Council now says that developing or operating AI ‘does not in itself constitute’ one. If the controller must still identify the purpose and assess it case by case, Article 88 bis(1) arguably does no legal work at all. The EDPB and the EDPS said much the same of the Commission’s version. In their Joint Opinion 2/2026 of 10 February 2026, they observed that Article 88c ‘does not bring any legal clarification following Opinion 28/2024’ (para 39).

Nonetheless, a provision in the enacting terms is unlikely to be read as empty. Controllers and courts are likely to look for its point, and the obvious candidate is that the legislature meant AI-related processing to be treated favourably.

That reading erodes case-by-case discipline. LI is a conclusion a controller must reach about a specific operation; it is not a status that attaches to a whole category of processing. Since *ASNEF and FECEMD*, the Court has consistently held that Member States may not, for particular categories of data, definitively prescribe the result of the balancing test Article 7(f) DPD required. That holding binds Member States rather than the Union legislature, which remains free to amend Article 6(1)(f), so the analogy must be drawn with care. The principle is nonetheless more than jurisdictional. A provision drawn at category level, in the enacting

terms, invites controllers to start from the category and work down to their case, whereas Article 6(1)(f) asks them to start from the case. The GDPR does identify categories of its own, fraud prevention and network security among them, though it does so in recitals and by reference to the purpose pursued rather than the technology deployed.

The burden that the LI test entails would also shift. Uncertainty currently falls on the party that benefits from the processing, which must show that the interests of data subjects do not outweigh its own. A legislative signal, even a soft one, moves part of that uncertainty onto data subjects and, in due course, onto courts and supervisory authorities. Anyone whose interests collide with those of a controller developing or operating an AI system would have to argue against a starting point carrying the legislature's endorsement, rather than against a controller's framing.

The right to object does not redress the balance. The reference to Article 21 added on 2 October changes nothing, because Article 21(1) already applies to processing based on Article 6(1)(f), as the EDPB itself recalled in Opinion 28/2024 (para 65). The right is also conditional: the controller may carry on if it demonstrates compelling legitimate grounds that override the interests, rights and freedoms of the data subject. And since Article 88 bis(3) applies '[r]egardless of the lawful ground', while Article 21(1) covers only processing under points (e) and (f) of Article 6(1), the cross-reference may even mislead. The unconditional right the Commission proposed has not returned, though it would not have sufficed either. Objection works after the fact and one person at a time. It presupposes that people know their data was used, which is unlikely where training data was scraped and the controller invokes the disproportionate effort exemption in Article 14(5)(b) (see EDPB Opinion 28/2024, para 63). It may also achieve little once a model has been trained. Recital 33 of the Council text accepts that the erasure of special categories of data memorised in a model 'would require re-engineering the AI system or AI model, or would be technically impossible'.

The clarification added on 2 October sits in the wrong place. Recitals can help us read the provisions they accompany, but they cannot lay down rules or override what those provisions say. Recital 33a also forms part of the preamble of the amending regulation, not of the GDPR. It may inform the construction of Article 88 bis and Article 6(1)(f), but from outside the GDPR, since it will not appear in the consolidated text that controllers, data subjects and their lawyers

actually consult. Someone who opens the GDPR after the reform will find an article announcing that AI-related processing may rest on LI, and nothing to tell them that developing or operating AI is not, in itself, a legitimate interest.

Conclusion

The Council's text goes some way towards addressing the concerns raised by the Commission's proposed AI clause. Article 88 bis now reaches interests pursued by third parties, and the 2 October recital confirms that developing or operating AI is not, in itself, a legitimate interest. But the article and its recital still pull in different directions. The recital speaks only of the controller's interest and keeps qualifiers the article has dropped. It alone contains the 2 October clarification. 'Operation' remains undefined.

This leaves controllers and supervisory authorities to work out what the provision adds to Article 6(1)(f). It may add nothing. If Article 6(1)(f) already accommodates AI-related processing, a special provision for AI has little to do, and one may fairly ask why the GDPR needs it.

If it does more, it answers in advance a question that Article 6(1)(f) leaves to each case. Endorsing a whole category of processing would tilt the test and shift part of the burden away from controllers. A rule written for one technology also strains the GDPR's neutrality, and invites the same request for the next. The shortcut remains tempting, and the Council has not resisted it.

AVAILABLE AT

https://europeanlawblog.eu/a-tempting-shortcut-how-the-digital-omnibus-would-re/

CC BY-SA 4.0