A LAIBRARY WHITEPAPER

The 2027 Deadline Trap

What the Digital Omnibus Actually Changed — and the Four Obligation Sets That Never Moved

By Tanner Beetge

Published September 2026

Research current to 14 September 2026

Author bio to be supplied.

Executive summary

On 24 July 2026, Regulation (EU) 2026/1744 was published in the Official Journal. It entered into force three days later, six days before the EU AI Act's original high-risk compliance deadline. Known as the Digital Omnibus on AI, it amended **42 articles and 3 annexes** of Regulation (EU) 2024/1689.

The headline that travelled was simple: high-risk AI obligations delayed to December 2027.

That headline is accurate and, taken alone, dangerous. Our analysis of the amended text against the original produces five findings that materially change how an organisation should be planning.

FINDING 1 **The deferral is narrow.** Annex III standalone high-risk obligations moved from 2 August 2026 to 2 December 2027. Annex I embedded high-risk obligations moved to 2 August 2028. Nothing else in the Regulation was postponed.

FINDING 2 **Four obligation sets are already enforceable.** Article 5 prohibitions have applied since February 2025. AI literacy duties since February 2025. General-purpose AI model obligations since August 2025. Article 50 transparency obligations since 2 August 2026. None of these were touched by the Omnibus.

FINDING 3 **The Omnibus added obligations as well as deferring them.** Two new prohibited practices were inserted into Article 5, covering AI systems that generate non-consensual intimate imagery and child sexual abuse material. These phase in on 2 December 2026.

FINDING 4 **Enforcement capability arrived on schedule.** The Commission's powers to investigate and fine general-purpose AI model providers became exercisable on 2 August 2026, ending an informal grace period that had run since August 2025. Penalties reach €15 million or 3% of worldwide turnover.

FINDING 5 **Sixteen months is less than it appears.** As of publication, no CEN-CENELEC harmonised standard has been cited in the Official Journal. No provider currently has access to the Article 40 presumption of conformity. The standards that would make conformity assessment routine are targeted for availability in Q4 2026 at the earliest, with Official Journal citation a separate and later step. Organisations planning to build against finished standards have materially less runway than the calendar suggests.

The practical conclusion: the Omnibus reset a clock. It did not reduce a single requirement.

Any organisation that stood down its AI Act programme in July 2026 is now behind on duties that carry penalties today, and has consumed part of a critical path that was already tight.

1. How the misreading happened

There was a genuine crisis in the first half of 2026, and it was not manufactured.

Article 40 of the AI Act grants a presumption of conformity to high-risk systems built in accordance with harmonised standards whose references have been published in the Official Journal of the European Union. This is the mechanism by which EU product regulation normally works: the legislature sets essential requirements, the European standardisation organisations write the technical specifications, and a manufacturer who follows the specifications is presumed to meet the requirements.

By early 2026 it was clear that the technical specifications would not exist in time. CEN-CENELEC Joint Technical Committee 21, established in June 2021 and now carrying over 300 experts across five working groups, had not delivered a citable standard. The original standardisation request, M/593, had been repealed and replaced by M/613, which extended the timeline to February 2027. In October 2025 the CEN and CENELEC Boards adopted exceptional acceleration measures, permitting direct publication after a positive Enquiry vote without a separate Formal Vote, and convening a small drafting group to complete the six most delayed drafts.

The Commission therefore faced an uncomfortable position: a compliance deadline of 2 August 2026 for a regime whose conformity infrastructure did not yet exist. The Digital Omnibus was the answer.

It was proposed in November 2025, reached provisional political agreement on 7 May 2026, was approved by the European Parliament on 16 June and the Council on 29 June, published on 24 July and in force from 27 July — six days before the deadline it was designed to move.

That compressed sequence explains the misreading. Most organisations encountered the story as a single line in a trade publication in late July: *high-risk delayed to 2027*. Very few read the amending regulation. Fewer still compared it against the original text to establish what had not changed.

This whitepaper is that comparison.

2. What the Omnibus changed

2.1 The deferrals

OBLIGATION SET	ORIGINAL DATE	AMENDED DATE	CHANGE
Annex III standalone high-risk (Arts 6–49)	2 Aug 2026	**2 Dec 2027**	**+16 months**
Annex I embedded high-risk (Arts 6–49)	2 Aug 2027	**2 Aug 2028**	**+12 months**
Legacy high-risk systems used by public authorities	2 Aug 2030	**2 Aug 2028 aligned**	Realigned

Annex III covers the eight categories of standalone high-risk use: biometrics; safety components in critical infrastructure; education and vocational training; employment and worker management; access to essential private and public services including creditworthiness and insurance pricing; law enforcement; migration, asylum and border control; and administration of justice and democratic processes.

Annex I covers AI that is a safety component of, or is itself, a product already regulated under EU harmonisation legislation — machinery, medical devices, in-vitro diagnostics, toys, lifts, radio equipment, personal protective equipment, and the transport sector instruments in Section B.

2.2 The additions

Two new prohibited practices were inserted into Article 5, addressing AI systems designed or used to generate:

- non-consensual sexually explicit imagery of real persons, commonly described as "nudification" applications; and
- child sexual abuse material.

These phase in on **2 December 2026**. They sit alongside the existing Article 5 prohibitions, which carry the Regulation's highest penalty band.

2.3 The simplifications

Several changes genuinely reduce administrative burden and are worth understanding because they may alter your compliance route:

Small mid-cap category	Simplified obligations previously available only to SMEs were extended to a new "small mid-cap" (SMC) classification, capturing organisations that have outgrown SME thresholds but lack large-enterprise compliance capacity.
Registration simplified	Providers claiming the Article 6(3) exemption — the provision allowing a system falling within an Annex III category to be treated as not high-risk where it does not pose a significant risk of harm — face a lighter registration process in the EU database.
AI literacy recast	The Article 4 obligation on organisations was simplified, with a stronger promotional and supporting role assigned to the Commission and Member States. The duty on organisations persists; its administrative shape changed.
Bias testing unlocked	Processing of special categories of personal data is now expressly permitted where strictly necessary for bias detection and correction. This resolves a genuine tension between the AI Act's fairness requirements and GDPR Article 9, and it is the most operationally useful amendment in the package.

2.4 The expansions

The Omnibus also extended the AI Office's oversight powers, including over AI systems built on general-purpose models and systems embedded in large online platforms. Read together with the enforcement powers commencing in August 2026, the direction is toward a more capable central supervisor, not a lighter one.

3. The four obligation sets that never moved

This is the core of the analysis. Each of the following has applied since the date stated, was untouched by the Omnibus, and carries penalties today.

3.1 Article 5 prohibitions — applicable since 2 February 2025

The prohibited practices have been enforceable for over eighteen months. They cover subliminal or manipulative techniques causing significant harm; exploitation of vulnerabilities based on age, disability or social or economic situation; social scoring leading to detrimental treatment in unrelated contexts or disproportionate to behaviour; predictive policing based solely on profiling or personality traits; untargeted scraping of facial images to build recognition databases; emotion inference in the workplace and in education, outside medical and safety uses; biometric categorisation to infer sensitive attributes; and real-time remote biometric identification in publicly accessible spaces for law enforcement, subject to the narrow Annex II exceptions.

> **Penalty band: up to €35 million or 7% of total worldwide annual turnover, whichever is higher.**
>
> This is the highest exposure in the Regulation, and the one most likely to arise from a product decision made without legal review.

The workplace emotion-inference prohibition deserves particular attention because it catches ordinary commercial software. Sentiment analysis applied to employee communications, engagement-scoring tools, and candidate assessment products that infer emotional state all sit close to or across the line.

3.2 AI literacy — applicable since 2 February 2025

Article 4 requires providers and deployers to take measures ensuring a sufficient level of AI literacy among staff and others operating AI systems on their behalf, having regard to their technical knowledge, experience, education and training, and the context of use.

The Omnibus simplified how this obligation is discharged. It did not remove it. Evidence of training delivery, role-appropriate content and currency is the artefact an authority will ask for.

3.3 General-purpose AI model obligations — applicable since 2 August 2025

Chapter V has bound GPAI model providers for over a year. Core duties under Article 53: maintain technical documentation per Annex XI; provide downstream integrators with the information set in Annex XII; implement a copyright policy respecting the text-and-data-mining reservation under the DSM Directive; and publish a sufficiently detailed summary of training content.

Providers of models designated as posing systemic risk carry additional Article 55 duties: model evaluation including adversarial testing, systemic risk assessment and mitigation, serious incident tracking and reporting, and adequate cybersecurity protection for the model and its physical infrastructure.

Designation runs through Annex XIII criteria, with a presumption attaching at cumulative training compute exceeding 10^{25} floating point operations. Article 52(1) requires notification to the AI Office within two weeks of meeting the criteria, and that duty applies irrespective of whether a provider has signed the Code of Practice.

Models placed on the market before 2 August 2025 must be brought into compliance by 2 August 2027.

3.4 Article 50 transparency — applicable since 2 August 2026

This is the obligation most commonly assumed to have been deferred. It was not.

Article 50 imposes four transparency duties:

50(1)	Providers must design systems that interact directly with natural persons so that those persons are informed they are interacting with an AI system, unless obvious from context.
50(2)	Providers of systems generating synthetic audio, image, video or text must mark outputs in a machine-readable format detectable as artificially generated or manipulated, with technical solutions that are effective, interoperable, robust and reliable as far as technically feasible.
50(3)	Deployers of emotion recognition or biometric categorisation systems must inform exposed persons.
50(4)	Deployers generating or manipulating deepfake content must disclose it; those publishing AI-generated text on matters of public interest must disclose unless human editorial review and responsibility apply.

The short grace period for machine-readable marking of output from systems already on the market before 2 August 2026 **expires on 2 December 2026**.

Penalty band: up to €15 million or 3% of worldwide turnover.

4. December 2026: the month nobody diarised

Three things land within eight days of each other, and none of them featured in the summer's coverage.

2 Dec 2026	2 Dec 2026	9 Dec 2026
Article 50(2) marking grace expires Any generative system placed on the EU market before 2 August 2026 must by this date produce output carrying machine-readable provenance marking.	**New Article 5 prohibitions phase in** The nudification and CSAM prohibitions take effect at the top penalty band.	**Product Liability Directive transposition** Directive (EU) 2024/2853 must be transposed into national law by this date.

For organisations running generative features inside established products, the marking deadline is an engineering deliverable, not a policy one. Provenance marking touches the output pipeline, any downstream editing or export path, and the metadata handling of every integration that consumes the output.

For the new prohibitions, most organisations face a screening question rather than a programme: does anything in our portfolio, or anything a customer could readily configure our product to do, generate this class of content? The answer needs to be documented rather than assumed.

The Product Liability Directive repeals the 1985 directive and expressly extends the definition of "product" to software, AI systems and digital manufacturing files. It introduces disclosure obligations and rebuttable presumptions of defectiveness where a claimant faces excessive difficulty proving their case.

The significance of the third item is routinely underestimated in AI governance discussions. The **AI Liability Directive**, designed specifically to handle AI harm claims, was withdrawn by the Commission in February 2025 and never replaced. Claims will therefore run through the revised Product Liability Directive, national tort law, the AI Act itself, contract, data protection and anti-discrimination law.

The consequence for documentation practice is direct. Technical files, risk assessments, test records and incident logs assembled for regulatory conformity become discoverable evidence in civil litigation, under a regime that includes presumptions operating against the defendant where disclosure is inadequate. Documentation quality stops being a compliance question and becomes a liability question.

5. Why sixteen months is not sixteen months

The deferral reads as generous. Examined against what has to happen inside it, it is not.

5.1 The standards are not finished

No CEN-CENELEC JTC 21 deliverable has been cited in the Official Journal as of publication. Until citation occurs, **no provider has a presumption of conformity for any high-risk AI system**, and conformity must be demonstrated directly against the Articles.

The current position:

- Standardisation request **M/613** governs the work and expires 28 February 2027.
- The October 2025 acceleration measures target availability of prioritised deliverables in **Q4 2026**.
- **prEN 18286**, covering quality management systems for EU AI Act regulatory purposes and supporting Article 17, was at Approval and Formal Vote stage in mid-2026 and is the most likely first citation.
- Availability is not citation. Publication by CEN-CENELEC, assessment by the Commission, and citation in the Official Journal are sequential steps, each taking months.

A realistic planning assumption is that citable standards covering the full requirement set are available during 2027, leaving a compressed window before December.

Aligning to a draft prEN remains worthwhile. It produces audit-ready evidence and reduces rework. It confers no legal presumption.

5.2 Notified body capacity is untested

Systems requiring third-party conformity assessment under Annex VII — including certain biometric systems under Annex III point 1, and any system where harmonised standards have not been applied — depend on notified bodies designated by Member State notifying authorities.

That designation infrastructure was required to be in place by 2 August 2025. Implementation has been uneven, with several Member States having missed authority-designation deadlines. The number of designated bodies with AI scope remains small relative to the expected volume of assessments arriving in late 2027.

This matters because the standards gap and the capacity constraint compound. Where harmonised standards have not been applied, the Annex VII route becomes more likely, which increases demand on a constrained resource at precisely the moment it is least able to absorb it.

5.3 The work itself takes longer than the calendar

Back-planning from 2 December 2027 against realistic organisational lead times:

ACTIVITY	REALISTIC DURATION	MUST START BY
AI system inventory and discovery	2–4 months	Q1 2027 at the latest
Risk classification against Annex III	1–2 months	Q2 2027
Data governance remediation (Art. 10)	4–8 months	Q1 2027
Risk management system design and operation (Art. 9)	3–6 months, then continuous	Q1 2027
Technical documentation assembly (Annex IV)	3–6 months	Q2 2027
Quality management system (Art. 17)	**6–12 months**	**Q4 2026**
Human oversight design and operator training (Art. 14)	2–4 months	Q2 2027
Conformity assessment	2–6 months	Q3 2027
EU database registration (Art. 49)	2–4 weeks	Q4 2027

The quality management system is the constraint. Article 17 requires a documented QMS covering regulatory compliance strategy, design control, verification and validation procedures, technical specifications, data management, risk management, post-market monitoring, incident reporting, record-keeping and accountability. Organisations without an existing ISO 9001 or ISO/IEC 27001 management system typically need nine to twelve months to stand one up and demonstrate it operating.

Counting back from December 2027, that means **starting in Q4 2026**. Which is now.

5.4 Data governance is the hidden long pole

Article 10 requires training, validation and testing datasets to be relevant, sufficiently representative, and to the best extent possible free of errors and complete in view of the intended purpose. It requires examination of possible biases likely to affect health and safety or lead to prohibited discrimination, and appropriate measures to detect, prevent and mitigate them.

Where the necessary remediation is discovered late, the remedy is frequently retraining. Retraining has a cost and a lead time that no amount of documentation effort can compress. Organisations that defer data governance assessment until 2027 may find that the only compliant path within the remaining window is withdrawal of the system from the EU market.

This is the argument for running data governance assessment first rather than last, against the sequence most programmes naturally adopt.

6. What to do in the next ninety days

A defensible programme opening, ordered by dependency rather than by difficulty.

1 **Establish the inventory.** Every AI system, model and agent in use or in development, with owner, purpose, data inputs, deployment geography and user population. Nothing downstream is possible without it, and it is the only activity that does not become easier with delay.

2 **Screen for Article 5 exposure immediately.** Prohibited practices are enforceable now at the highest penalty band. Screen specifically for workplace and educational emotion inference, biometric categorisation inferring sensitive attributes, social scoring patterns, and — ahead of December — any capability to generate the newly prohibited content categories.

3 **Audit Article 50 compliance.** Every system interacting with natural persons, every generative output pipeline, every deepfake or synthetic media capability, every emotion recognition or biometric categorisation deployment. Confirm machine-readable marking is implemented or scheduled ahead of 2 December 2026.

4 **Confirm your role, per system.** Provider, deployer, importer, distributor or authorised representative. The obligation set differs materially by role, and a single organisation commonly occupies different roles for different systems. Fine-tuning or substantially modifying a third-party model can convert a deployer into a provider.

5 **Classify against Annex III.** Provisional classification is sufficient at this stage. The purpose is to size the population requiring the full high-risk programme, so that the QMS work can be scoped and started in Q4 2026.

6 **Start the quality management system.** This is the long pole. If an ISO 9001 or ISO/IEC 27001 system exists, scope the extension. If not, begin now, and consider ISO/IEC 42001 as the vehicle since it shares the harmonised management-system structure and produces evidence usable across several frameworks simultaneously.

7 **Assess data governance early.** Specifically because the remedy may be retraining, and retraining has a lead time that compliance effort cannot shorten.

8 **Diarise December 2026 and January 2027.** Both months carry obligations that arrive well ahead of the high-risk deadline, and both have been widely overlooked.

Appendix A — Consolidated timeline

DATE	EVENT	STATUS
1 Aug 2024	AI Act enters into force	Complete
2 Feb 2025	Art. 5 prohibitions; Art. 4 AI literacy	**Enforceable**
2 Aug 2025	Ch. V GPAI obligations; governance; penalties; notifying authorities	**Enforceable**
24 Jul 2026	Regulation (EU) 2026/1744 published in OJ	Complete
27 Jul 2026	Omnibus enters into force	Complete
2 Aug 2026	Art. 50 transparency; national sandboxes; Commission GPAI enforcement powers	**Enforceable**
2 Dec 2026	**Art. 50(2) marking grace expires; new Art. 5 prohibitions phase in**	**Imminent**
9 Dec 2026	**Product Liability Directive transposition deadline**	**Imminent**
Q4 2026	CEN-CENELEC target for prioritised deliverable availability	Target
20 Jan 2027	Machinery Regulation (EU) 2023/1230 applies	Scheduled
28 Feb 2027	Standardisation request M/613 expires	Scheduled
2 Aug 2027	GPAI models placed on market before 2 Aug 2025 must comply	Scheduled
2 Dec 2027	**Annex III standalone high-risk obligations apply**	Scheduled
11 Dec 2027	Cyber Resilience Act full application	Scheduled
2 Aug 2028	**Annex I embedded high-risk obligations apply**	Scheduled
31 Dec 2030	Annex X large-scale IT systems	Scheduled

Appendix B — Penalty schedule

BREACH	MAXIMUM
Article 5 prohibited practices	**€35,000,000 or 7% of total worldwide annual turnover**
Most provider and deployer obligations	€15,000,000 or 3%
Supplying incorrect, incomplete or misleading information to authorities	€7,500,000 or 1%
GPAI model provider obligations (Art. 101, Commission-imposed)	€15,000,000 or 3%

For SMEs and start-ups, the lower of the percentage or the fixed amount applies. Member States set penalties for certain provisions in national implementing law, so national variation exists below these ceilings.

Verification note

This analysis reflects research current to 14 September 2026 and was compiled against primary sources including the Official Journal, the European Commission's AI Act Service Desk, and CEN-CENELEC publications.

Two items warrant direct verification before operational reliance:

1. **Harmonised standards status.** CEN-CENELEC Enquiry and Approval windows continue to move, and Official Journal citation is a separate Commission decision. Confirm current citation status directly before assuming any presumption of conformity is available.
2. **National implementation.** Member States set penalty levels and designate competent authorities for several provisions. Confirm the position in each Member State where you place systems on the market.

In compiling this research we identified errors in several widely used commercial regulatory trackers, including sources describing repealed instruments as in force. Where sources conflicted we relied on official publications.

This whitepaper is research and analysis. It is not legal advice, and it does not establish a lawyer-client relationship. Organisations should obtain advice from qualified counsel in the relevant jurisdictions.

laibrary builds AI governance software that treats compliance as a workflow and evidence problem.
Golden, Colorado.