

Playbook

The EU AI Omnibus and the Living EU AI Act

A Playbook for Enterprises



Table of contents

Who this playbook is for, and how to read it

Executive summary

Part I. The EU AI Omnibus as a living instrument

1.	What the EU AI Omnibus is, and what it is not	19
2.	What the Omnibus does not change	23
3.	The EU AI Act as a living instrument, and how its five layers work	34
4.	The new master calendar	41

Part II. The amended EU AI Act by actor

5.	Providers of high-risk AI systems	47
6.	Providers of general-purpose AI models and GPAI-based systems	54
7.	Deployers	60
8.	Horizontal expectation applicable to both providers and deployers (regardless of risk-classification)	63
9.	SMEs, start-ups and small mid-caps	65
10.	Importers, distributors, third-party suppliers and the conformity infrastructure	69
11.	The new enforcement reality	73

Part III. The transparency obligations

12.	How Article 50 (Transparency obligations for providers and deployers of certain AI systems) works in operation under the final Commission Guidelines	78
13.	The Code of Practice on Transparency of AI-Generated Content	83

Part IV. Prohibitions for AI Systems with Unacceptable Risk

14. The Article 5 prohibitions, the guidelines and the safeguards you'll need	86
---	----

Part V. Risk Classification

15. High-risk classification rules under the draft guidelines on Article 6	95
--	----

Part VI. The rest of the living instrument

16. The General-Purpose AI (GPAI) Code of Practice	102
17. The Commission guidelines inventory	104
18. Delegated acts, implementing acts and templates	106
19. European harmonised standards	108
20. Institutional build-out to monitor	111

Part VII. Using the playbook

21. Action checklist by horizon	114
22. The judgments that matter	117
23. Where does Credo AI fit?	121

Annex: Detailed Calendar

Who this playbook is for, and how to read it

This playbook is written for the people who sit between the legal text and the rest of the organisation (compliance, legal, risk and AI governance leads) who must give a board a defensible account of where the company stands,

and tell the teams building and deploying AI what to do next.

It follows the order the work has to happen in.

Part 1

Part I sets out what the Omnibus changed, what it left untouched, and how the Act's five layers work together, because no obligation makes sense read in isolation from the instruments beneath it. It ends with the master calendar.

Part 2

Part II takes each actor in turn (high-risk providers, GPAI model providers, deployers, smaller enterprises, importers and distributors, and the conformity infrastructure) so a reader can go straight to their own position and the new enforcement reality attached to it.

Part 3 to 5

Parts III to V cover the three areas that are live or imminent: the Article 50 transparency duties, in force since 2 August 2026; the Article 5 prohibitions and the safeguards stack due on 2 December 2026; and the high-risk classification decision that governs everything arriving in 2027 and 2028.

Part 6

Part VI maps the pipeline still being written (codes of practice, guidelines, delegated acts, templates and harmonised standards) and what to watch in each.

Part 7

Part VII converts all of it into a checklist by horizon, six judgments to take to a board with the evidence attached, and where Credo AI fits.

Two things run throughout. Every position that will not survive the final texts is flagged where it appears. And every deliverable is expressed as something written down,

owned by a named person, and re-runnable when the basis for it changes because that is the form in which a regulator will eventually read it.



The EU AI Omnibus and the Living EU AI Act: Executive summary

The bottom line

On 27 July 2026 the AI Omnibus moved the AI Act's largest compliance deadline back by sixteen months. Six days later, on 2 August, the Act's enforcement architecture came into operation as originally scheduled, together with the Article 50 transparency duties and the Commission's power to fine providers of general-purpose AI models. Relief and exposure arrived in the same week. Coverage has attended almost exclusively to the first.

Four corrections follow, and each one changes what an organisation should be doing now.

The obligations did not change; some of their start dates did.

The Act still sorts AI into four risk tiers, with the same governance framework and the same core requirements adopted in June 2024.

Nothing in Articles 9 to 15 was softened. Resource the programme against the original rules, not the new dates.

This is an implementation correction, not a change of policy.

The Act depends on machinery that was not ready: harmonised standards, designated national authorities, notified bodies, official guidance. The extra time is for that ecosystem as much as for industry — which means the period to 2028 will be dense with delegated acts,

guidelines, codes of practice, templates and standards. Most of the detail that determines what compliance actually requires will be settled there, not in the amended articles. Track the pipeline, not only the text.

The dates are fixed, and this corrects a widely repeated error.

The Commission's November 2025 proposal tied the high-risk deadlines to an assessment that standards and support tools were ready, which would have made the calendar impossible to plan against. The adopted text drops it. Article 113(3)(c) sets calendar dates: 2 December 2027 for stand-alone high-

risk systems under Annex III, 2 August 2028 for AI embedded in regulated products under Annex I. No trigger, no assessment, no conditionality. Commentary still describing a conditional delay is describing the proposal, not the law.

The relief is uneven, and 2026 is not a quiet year.

Three sets of obligations arrived on schedule. Article 50 transparency and Article 101 (the Commission's power to fine model providers) took effect on 2 August 2026. On 2 December 2026 the two new Article 5 prohibitions apply, at the Act's highest penalty tier, and the extension for legacy generative systems to mark their outputs expires. Clear the 2026 obligations

before turning to 2027. So the immediate task is to establish which of your systems fall where. Inventory your AI systems, each with its intended purpose recorded, and classify against Annex III, Annex I and the general-purpose AI obligations. That tells you what moved to December 2027 or August 2028, and what is enforceable today.

Who this playbook is for

Credo AI's EU AI Omnibus Playbook

explains how the Act is built, because no obligation within it makes sense in isolation. This playbook is written for the managers who sit between those facts and the rest of the organisation. You must give a board a defensible account of where the company stands while

giving the teams building and deploying AI instruction on what to do next. Six judgments follow, each a position that can be defended in a meeting, with its evidence attached. Then the dates, and then the work, in sequence and with ownership assigned.

The six judgments

1. The deferral is preparation time, not a pause

The extra time exists because the ecosystem the Act depends on (harmonised standards, enforcement structure, notified bodies and official guidance) was not ready. Operators who use 2026 and 2027 to settle classification, contracts and architecture will absorb the standards and guidelines as refinements. Those who wait will rebuild under them, against dates that will not move again.

2. Much of the Act is already in force, and untouched

The Article 5 prohibitions have applied since 2 February 2025 and have carried penalties since 2 August 2025. The general-purpose AI model regime under Articles 51 to 55 has applied since 2 August 2025, with its documentation, copyright and training-content duties intact and its systemic-risk obligations unchanged. What the Omnibus altered around that regime is supervision and procedure, not substance. Any account of the reform as rollback is simply wrong where it reaches these parts of the Act, and the programmes built on them must keep running.

Two dates sit inside the supposed quiet period and are easily missed. The Commission's power to fine model providers under Article 101 took effect on 2 August 2026, and with it ended the AI Office's first-year posture of treating signatories still working toward their commitments as acting in good faith — that stance was expressly tied to the period before that date. And models placed on the market before 2 August 2025 must reach full compliance by 2 August 2027.

3. 2026 is a transparency and prohibitions year

Two hard dates fall inside what was widely described as a relief period, and the first has passed.

Since 2 August 2026, providers must ensure people know they are dealing with an AI system and must mark and make detectable the synthetic content their systems generate, while deployers must label deep fakes and AI-generated text published to inform the public on matters of public interest. Breach sits at 15 million euros or 3 percent of worldwide turnover. Three points are routinely missed. Marking without an available means of detection does not comply. The Article 11(2) grandfathering for legacy high-risk systems does not extend to Article 50, so an older system that is interactive owed disclosure on 2 August. And text generated before that date but published on or after it must be labelled.

The second date is 2 December 2026, when the new prohibitions on AI-generated non-consensual intimate imagery and child sexual abuse material apply, at 35 million euros or 7 percent. Providers are caught where such generation is the intended purpose, or where it is a reasonably foreseeable and reproducible outcome and the system lacks reasonable and adequate safeguards. Deployers are caught only where they use a system for that purpose. Because the safeguards are the exculpatory condition, the compliance artefact is a documented and tested safeguards stack, calibrated to how the system is released, open weights included. There are no Commission guidelines on these prohibitions to calibrate against; Recital 12 of the Omnibus offers a non-exhaustive list of qualifying measures and is currently the only anchor. The same date closes the window for legacy generative systems to meet the Article 50(2) marking duty.

4. Risk classification is the gateway decision, and the draft reading is broad

Risk classification decides whether the Chapter III apparatus applies to a system at all, which makes it the single most consequential judgment an operator makes. The Omnibus moved the perimeter in both directions at once.

The Act narrows it. Systems used solely for user assistance, optimisation, efficiency, automation, convenience or quality control fall outside safety-component status, and products needing third-party assessment only for non-safety reasons no longer qualify on that basis. But the carve-out has a floor: a system whose failure would endanger health or safety remains a safety component regardless.

The draft Article 6 guidelines of 19 May 2026 widen it again, on six fronts.

- Products under enhanced regulatory scrutiny are read into the third-party-assessment limb.
- Modular and agentic systems are assessed as one functional whole rather than component by component.
- Configuring or fine-tuning a general-purpose system to perform an Annex III task typically makes the configuring party the provider of a high-risk system. This is the most common route by which a deployer becomes a provider, and it is triggered by ordinary business activity.
- A human in the loop is an Article 14 oversight duty, not a classification escape.
- Intended purpose must be consistent across marketing, instructions and technical documentation, so disclaimers will not defeat realistic use.
- The Article 6(3) filter is read narrowly, carries its own documentation and registration duty, and never applies where a system profiles people.

The consultation closed on 23 July 2026 and the final text is due at year end. Reclassify now, in writing, and register every reliance on the filter.

5. Soft law is the real rulebook, and standards are its hardest edge

Two codes of practice with published adequacy assessments, six sets of Commission guidelines, a delegated-acts pipeline and the CEN-CENELEC catalogue will together carry more day-to-day compliance content than the amended articles.

Only one of those layers reverses the burden of proof. Under Article 40, a harmonised standard cited in the Official Journal gives a presumption of conformity: a challenging authority must then show the standard inadequate or misapplied, rather than the operator proving conformity from first principles. That is a materially better position to be in, and no AI Act standard has reached it yet.

EN 18286 on quality management was approved on 12 July 2026 and published on 22 July. It matters more than its subject suggests, because the quality management system is the spine through which every other Chapter III requirement is evidenced. Until the Commission cites it, it is the best available benchmark but carries no legal weight. Adopt it now, so the presumption attaches the day citation lands. Give each instrument a named owner, and watch the citation calendar as closely as the legislative one.

6. Enforcement has a new centre of gravity

For operators under its exclusive competence (systems built on their own general-purpose models, and systems inside designated very large platforms) the AI Office is now a regulator rather than a coordinator. Under Articles 75a to 75d it can enter and seal premises, order access to systems, make commitments binding, impose the Article 99 fines and levy daily penalties calculated on average daily turnover, and it must publish its decisions. Several categories stay national, including Annex I products, critical infrastructure, law enforcement, migration, justice and supervised financial institutions.

Resourcing is subject to the ordinary budgetary procedure, with no guaranteed allocation, which points toward a few high-visibility published cases rather than broad inspection. Prepare for the visible case, not the routine audit — and treat publication as a reputational exposure to be managed alongside the legal one.

For everyone else the national authority remains the body you deal with, but exposure travels further than before. Anyone affected can complain under Article 85, national authorities can refer matters to the AI Office with a guaranteed four-month response, and equality and fundamental-rights bodies now receive enforcement documentation promptly and in machine-readable form.

Five Layers Comprise the Living Instrument

The AI Act can no longer be read, or complied with, as a single text. It is the top layer of a five-layer structure, and each layer binds differently.

The Act itself is the only layer that creates obligations. Beneath it, delegated and implementing acts carry the same binding force and fill in the detail the Act leaves open. Codes of practice are voluntary, but they are how a provider demonstrates compliance in the absence of anything better, and the Commission now assesses and publishes their adequacy on a rolling

basis rather than approving them once. Commission guidelines are non-binding, but they are the benchmark national regulators will apply, so departing from one means being ready to defend the departure. Harmonised standards are the only layer that reverses the burden of proof: follow a standard cited in the Official Journal and you are presumed to conform, and it falls to the regulator to show otherwise.

Two consequences follow:

First, most of the detail that determines what compliance actually requires will be settled in those four layers, not in the amended articles (which is why the layers, not the legislation, are what need tracking).

Second, the layer that helps you most is the least developed: no standard has yet been cited for the AI Act, so until then, providers must prove conformity from scratch on every requirement. The deferrals to 2027 and 2028 exist largely to close that gap (alongside the gap of effective enforcement).

Importantly: What the Omnibus changed, and what it left untouched

	Changed	Remained unchanged
Timelines	Annex III stand-alone high-risk systems move to 2 December 2027, sixteen months; AI embedded in regulated products to 2 August 2028, twelve months. Article 113(3)(c) sets calendar dates — the readiness trigger in the November 2025 proposal did not survive. National sandboxes deferred a year to 2 August 2027.	The requirements waiting at those dates. Articles 9 to 15 are unchanged: lifecycle risk management, data quality and governance, technical documentation, record-keeping and logs, transparency and instructions to deployers, human oversight, accuracy, robustness and cybersecurity. So are Article 16, the conformity assessment logic, the EU declaration of conformity, CE marking and database registration under Articles 49 and 71.
Scope and classification	"Safety component" narrows in Article 3(14). Systems serving only non-safety functions (user assistance, optimisation, efficiency, automation, convenience, quality control) fall outside it under Article 6(1a), subject to a health-and-safety floor in 6(1b). Products requiring third-party assessment for non-safety reasons alone no longer qualify under 6(1c). SME and small mid-cap definitions added at Article 3(14a) and (14b).	Annex III. No use case was removed, and the Article 7 power to update the list still cannot be used to weaken protections. The Article 3(1) definition of an AI system, the extraterritorial reach, the provider, deployer, importer and distributor roles, and the Article 25 trigger by which a downstream actor becomes a provider all stand as adopted.

Prohibitions	Two added, effective 2 December 2026: AI systems generating or manipulating non-consensual intimate material and child sexual abuse material. New calibration paragraphs 5(1a) and (1b) apply different tests to providers and deployers.	The original eight, in force since 2 February 2025 and fineable since 2 August 2025. None removed, none narrowed. The Commission's prohibited-practices guidelines predate the Omnibus and address none of the new material.
Transparency (Article 50)	Two narrow changes only. Generative systems already on the market get four months to 2 December 2026 for the marking duty, under the new Article 111(4). Codes of practice under Article 50(7) are now assessed differently, with a new Commission power to impose common rules by implementing act if a code falls short.	Every duty, as originally written, applying from 2 August 2026: interaction disclosure, marking and detection of synthetic content, emotion recognition and biometric categorisation notices, deep fake and public-interest text labelling. Article 111(2) grandfathering does not extend to Article 50 at all.
Enforcement and supervision	The AI Office becomes the exclusive supervisor for systems built on the provider's own GPAI models and for AI in designated very large platforms. From 27 July 2026 it holds powers under Articles 75a to 75d: binding information requests, on-site inspection, sealing premises, binding commitments, fines and daily penalties, with publication of its decisions mandatory. Equality and fundamental-rights bodies gain stronger access under Article 77. Member States may add warnings and non-financial measures.	The enforcement calendar and the ceilings. Article 101 was always scheduled for 2 August 2026; Chapter IX market surveillance and Chapter XII penalties apply from that date. Maximum fines stay at 35 million euros or 7 percent, 15 million or 3 percent, and 7.5 million or 1 percent. National authorities remain the interlocutor for everyone outside the AI Office's exclusive competence.
Value chain and data	Written agreements with third-party suppliers into a high-risk system become mandatory under Article 25(4), including for open-source AI models; the hand-over duty in Article 25(2) becomes a specific list; both are separately fineable at 15 million euros or 3 percent. New Article 4a provides a legal basis for using special-category data in bias detection and correction, extending to deployers and to non-high-risk systems and models.	Data protection law. The GDPR, e-Privacy Directive and Law Enforcement Directive continue to apply in full under the amended Article 2(7). Article 4a is a permission rather than a duty, and answers only whether special-category data may be used for this purpose — lawful basis, purpose limitation, minimisation, security and DPIA obligations are all where they were.

Sectoral alignment and soft law

AI in machinery moves onto the Machinery Regulation's conformity route. Article 42(3) treats Cyber Resilience Act compliance as satisfying Article 15. Article 43(3) resolves overlapping conformity routes. Notified bodies gain combined applications, an expedited route and the Annex XIV taxonomy, with the bridge expiring 28 January 2028. Codes of practice lose their one-off approval: adequacy is now assessed continuously and published. The post-market monitoring template becomes guidance.

The status of the instruments. Guidelines remain non-binding, and only the Court of Justice can interpret the Act authoritatively. The AI Board, advisory forum and scientific panel continue. Member States must still designate notifying and market surveillance authorities, and the EDPS still supervises the EU institutions.

The reform is therefore neither relief nor rollback. It is a change of sequence. The obligations arriving in 2026 are live and enforceable, the obligations arriving in 2027 and 2028 are the ones adopted in 2024, and the time in between is preparation, not pause.

The architecture, and why it decides everything else

The Act sorts systems by risk, not by technology. Four tiers apply to AI systems, with a parallel track for general-purpose AI models.

1. Unacceptable risk is prohibited outright under Article 5, now covering ten practices after the Omnibus added two. Maximum fines are 35 million euros or 7 percent of worldwide turnover.
2. High risk covers both the Annex III use cases under Article 6(2) and AI acting as a safety component of regulated products under Article 6(1). It carries obligations in the Chapter III apparatus (risk management, data governance, technical documentation, logging, human oversight, accuracy and cybersecurity, conformity assessment, CE marking, registration and post-market monitoring) at 15 million euros or 3 percent.

3. Transparency risk is addressed in Article 50: interactive systems, synthetic content, emotion recognition and biometric categorisation, deep fakes and public-interest text. Same ceiling, 15 million euros or 3 percent.
4. Minimal risk is everything else, which owes no tier-specific obligations. The Article 4 AI literacy duty is horizontal and applies across all four tiers.

Running alongside this, general-purpose AI models are classified by capability rather than by use under Articles 51 to 55, with a systemic-risk sub-tier attracting evaluation, adversarial testing, mitigation and incident reporting.

Three consequences follow:

A single product can sit in more than one tier at once. A high-risk system that is also interactive owes Chapter III and Article 50 cumulatively, on their separate timetables; and the grandfathering relief for existing high-risk systems does not extend to Article 50.

The boundary of the high-risk tier is contested and moving. The Omnibus narrowed it, by carving non-safety functions out of the "safety component" definition. The Commission's draft Article 6 guidelines widen it, by reading the third-party assessment route broadly and by assessing modular and agentic systems as a functional whole rather than component by component. A classification settled before mid-2026 has been overtaken on both sides.

Because tiering drives everything downstream, an inventory that has not been risk-classified in writing is not a compliance programme.

Where Credo AI fits

Under the AI Act you comply by producing a record. A classification with its reasoning. A registered Article 6(3) reliance. A tested safeguards stack. An open-weights decision documented when it was taken. A regulator arriving in 2027 reads that file, and gives no credit for analysis done but never written down.

The difficulty is that the file has to keep moving. Final guidelines, revised codes and newly cited standards change what you must evidence without changing a single article of the Act. Classification alone has to be re-run when the Article 6 guidelines are finalised at year end — and again after that.

Credo AI provides the system of record across the portfolio: every position dated, owned and traceable, and re-runnable when the basis for it changes. To help you operationalize your EU AI Act compliance seamlessly, Credo AI has launched the EU AI Act Policy Pack Bundle, up to date with Omnibus amendments and translating each component: what the Act, its guidelines, its Codes of Practice, templates and its standards require of each AI system, expressed as controls, evidence and sign-offs. The packs are revised as the EU AI Act instruments are evolving so that you continue to stay ahead of your obligations, rather than playing catch up.

Start with the map. Which systems are exposed today, which move to December 2027, which to August 2028, and who owns each obligation. Credo AI's EU AI Omnibus Playbook sets out that analysis in full, and the platform helps you operationalize your organizational strategy to stay ahead of EU AI Act compliance, even as the rules underneath the living instrument change.

www.credo.ai

Talk to us

Part | •

The EU AI Omnibus as a living instrument

Part I • Section 1

What the EU AI Omnibus is, and what it is not

What the EU AI Omnibus is, and what it is not

The EU AI Omnibus (“Omnibus”) is an amending regulation.

It is not a recast, not a new EU AI Act, and not a deregulation effort. The EU AI Act, [Regulation \(EU\) 2024/1689](#), remains in force in its entirety, and the AI Omnibus, [Regulation \(EU\) 2026/1744](#), adjusts specific provisions, timelines and

The Omnibus also amends two other regulations: [Regulation \(EU\) 2018/1139](#) on civil aviation and [Regulation \(EU\) 2023/1230](#) on machinery. The machinery change is important, because before a machine is sold in the EU, someone has to confirm it meets the law. When that machine has AI built into it,

It was published in the Official Journal on 24 July 2026 and entered into force just three days later on 27 July 2026. The speed was deliberate. The general application date the Omnibus operates on, was the second of August 2026,

implementation mechanisms within it. One thing that is easy to miss is that no single document now states the law. You have to read the [consolidated version](#), and it gets rebuilt every time the law is amended again.

there are two possible rulebooks for that check: the AI Act’s, or the one for machinery. The machinery amendment shifts AI-in-machinery out of the AI Act’s own conformity process and onto the machinery rulebook, carrying with it the machinery regime’s detailed rules and the 2 August 2028 compliance date.

and the normal twenty-day wait between publication and entry into force would have pushed it past that date. It had to take effect first, so it wasn’t changing rules that were already live.

The practical consequence is narrow but real. Enterprises (referred to as operators in the AI Act, meaning providers, product manufacturers, deployers, authorised representatives, importers or distributors) who spent

June and July hedging against the possibility that the Omnibus would miss its window can stand that hedge down. The high-risk deferral is now set. Four points frame what substantively actually changed.

What did the Omnibus change?

First, the basic structure hasn't changed.

The AI Act still works the same way: It sorts AI by risk into four groups (prohibited practices, high-risk systems, transparency obligations and general-purpose AI rules) and keeps the same governance framework and core obligations it had when it was adopted in June 2024. Anyone reading the deferrals as an invitation to pause compliance work is misreading the instrument. The obligations have not changed. Only some of their start dates have moved. Resource your compliance programme around the original rules, not the delayed dates.

Second, the reform is largely an implementation correction, it doesn't reverse the policy.

The Act depends on supporting machinery to work in practice (harmonised standards, designated national market surveillance authorities who enforce the rules, notified bodies that certify products, and official guidance). Much of that was not ready for the original 2 August 2026 general applicability date. So the extra time is as much for the ecosystem as for businesses. Because of this, the period between now and 2028 will be full of follow-on rules, like delegated and implementing acts, guidelines, codes of practice, templates and standards. Most of the detail that actually affects compliance will be decided there, in that stream of documents, rather than the amended articles themselves. Parts III to VI of this playbook map it out. Track that pipeline, not only the amended text.

Third, the dates are fixed dates (not conditional), and this clears up a mistake that's been repeated widely.

The European Commission's November 2025 proposal had tied the high-risk deadlines to a Commission assessment that standards and support tools were ready. That would have made the calendar contingent and, in practice, impossible to plan around. The adopted text drops that approach. Article 113(3)(c) now sets fixed calendar dates: 2 December 2027 for stand-alone high-risk AI systems in areas like hiring, credit, and policing (listed in Annex III) and 2 August 2028 for systems built into products (Annex I). There is no trigger, no assessment, no conditionality. Any commentary still describing a conditional delay is describing the old proposal, not the law. Plan against these dates.

Fourth, the relief is uneven, and 2026 is not a quiet year.

The high-risk obligations did move back: sixteen months for Annex III stand-alone systems (now December 2027) and by a year for Annex I product-embedded ones (now August 2028). But three sets of obligations still arrive in 2026, and none of them was delayed. The Article 50 transparency duties begin 2 August 2026. On the same day, the enforcement architecture for general-purpose AI model providers switches on. This comes from Article 101, which was always set for 2 August 2026 and the Omnibus did not alter it. And the two new Article 5 prohibitions, along with the Article 50(2) rule requiring older (legacy) generative systems to make their systems, begin on 2 December 2026. The easiest mistake to make right now is to treat the Omnibus as across-the-board relief and miss the 2026 cluster. Clear the 2026 items first, before turning to 2027.

Part | • Section 2

What the Omnibus does not change

The political noise around
“simplification” has left
many people with the
impression that the rules
were rolled back.

The actual legal text does not support that. The parts of the AI Act listed in the next
pager are unchanged, and several of them are already in force today.

Banned practices

The original prohibitions are in force and unchanged.

The practices banned under Article 5, since 2 February 2025, remain exactly as adopted. They are:

1. manipulating or deceiving people in ways that seriously distort their behaviour
2. exploiting people's vulnerabilities tied to age, disability or their social or economic situation
3. social scoring
4. predictive policing based solely on profiling
5. scraping facial images from the internet in bulk to build recognition databases
6. recognising emotion in workplaces and educational institutions
7. using biometric categorisation to infer sensitive traits about people
8. real-time remote biometric identification (such as facial recognition) in public spaces for law enforcement (with a few narrow exceptions).

Fines for breaking these rules have applied since 2 August 2025. The Omnibus adds two new prohibitions (covered in Parts IV of this playbook). It removes none, narrows none, and leaves the maximum fine untouched: 35 million euros or 7 percent of worldwide annual turnover. Enforcement of the original prohibitions is live now; the two new prohibitions take effect on 2 December 2026 (Section 14.2).

GPAI models

The rules for general purpose (GPAI) models are in force and their substance hasn't changed.

The obligations for providers of these models (Articles 51 to 55) have applied since 2 August 2025 and the Omnibus doesn't change what they require. Each model provider still owes:

1. technical documentation (Annex XI);
2. information and documentation for the businesses that build on their model (called "downstream providers") (Annex XII);
3. a copyright policy; and
4. a public summary of what the model was trained on, using the European Commission's template.

Providers of the most powerful models (those flagged as posing "systemic risk") owe more on top: testing the model, probing it for weaknesses (adversarial testing), assessing and reducing systemic risks, reporting serious incidents, and protecting against cyberattacks. The thresholds for deciding which models count as systemic-risk, and the Commission's power to designate them, are also unchanged.

What the Omnibus does change is the supervision around these rules (the AI Office's powers and how the codes of practice are assessed), not the obligations themselves. It also leaves the enforcement calendar alone: the Commission's power to fine (Article 101) begins on 2 August 2026, and the deadline for models already on the market (Article 111(3)) stays at 2 August 2027.

Transparency rules

The substance of Article 50's transparency rules hasn't changed.

Every part applies exactly as originally written. In short:

1. Providers must make sure people know when they are dealing with an AI system
2. Providers of generative AI must mark what their system produces so a machine can recognise it as AI-generated
3. Deployers must disclose when emotion recognition or biometric sorting is used, and must label deep fakes and AI-written text that's published to inform the public on matters of public interest.

The Omnibus only makes two small changes here. First, generative systems already on the market before 2 August 2026 get four extra months (until 2 December 2026) to meet the marking duty. Second, it changes how codes of practice under Article 50(7) are assessed. Everything else under Article 50 applies from 2 August 2026, exactly as originally scheduled.

High-risk system requirements

The high-risk requirements are delayed, not weakened.

Articles 9 to 15 keep everything they required initially. For high-risk systems, that still means:

1. managing risk across the system's whole lifecycle;
2. meeting quality standards for data and how it's governed;
3. keeping technical documentation (Annex IV);
4. keeping records and logs;
5. being transparent and giving instructions to businesses that use the system (deployers);
6. real and effective human oversight;
7. accuracy, reliability, and cybersecurity.

The same is true for the provider's list of obligations (Article 16), the conformity assessment logic (this is the process for proving a system meets the law) and the EU declaration of conformity, the CE marking, and the duty to register the system in the EU database (Articles 49 and 71). The deployer duties (Article 26) and the fundamental rights impact assessment (Article 27) also survive in full. What the Omnibus changes is when these provisions apply, and it adds simpler and more proportionate formats for smaller businesses. It does not lower the bar.

Protections

Annex III is intact.

No high-risk use case was removed. All of them are still on the list: recruitment and managing workers, education, credit scoring, essential services, biometrics, critical infrastructure, law enforcement, migration and border management, and justice system. The Commission still has its existing powers to update the list (Article 7), and, as before, it can't use that power to weaken protections.

High-risk system requirements

The rules on scope, definitions and reach are clarified.

These all stay as adopted:

1. What counts as an AI system (Article 3(1));
2. the reach beyond the EU - the rules apply to providers and deployers based outside the EU when their systems, or the outputs of those systems, are used inside the Union;
3. the definitions of different roles - provider, deployer, importer and distributor and the points at which a company further down the chain becomes a "provider" itself (Article 25)

The Omnibus makes two small changes: it narrows the definition of "safety-component" and it adds definitions for "SME" (small and medium-sized enterprises) and "small mid-cap" companies.

Data protection (GDPR)**Data protection law is expressly preserved, with one narrow permission added.**

The GDPR ([Regulation \(EU\) 2016/679](#)), the e-Privacy Directive ([Directive 2002/58/EC](#)) and the Law Enforcement Directive ([Directive 2016/680](#)) all continue to apply in full. The amended Article 2(7) confirms this, and updates across-reference so it now points to Articles 4a and 59 instead of the deleted Article 10(5). None of this excuses an AI project from the normal data protection requirements. You still need: a lawful basis for the data processing involved, to use the data only for the stated purpose and no more of it than necessary, transparency towards the people whose data is used, proper security, and a data protection impact assessment where one is required.

The addition is Article 4a, and it is more significant than it looks. It looks like a simple relocation, but it isn't. The old Article 10(5) let providers of high-risk systems use special categories of personal data (data on things like race, health, or beliefs) where strictly necessary to find and fix bias. But it sat in a part of the Act that hadn't taken effect yet (Chapter III, Section 2), and it was deleted on 27 July 2026 without ever having produced legal effect. In practice, the Act had never actually provided this permission until now.

Article 4a moves it into Chapter I, which is already in force, and broadens who can use it:

1. The original conditions stay the same. Using special categories of personal data this way must be exceptional and strictly necessary; synthetic or anonymised data must not be enough to do the job; and the data must be pseudonymised where possible, kept under access controls, not shared with anyone else, and deleted once the bias is fixed.

2. More people can rely on it. The same permission now extends to deployers of high-risk systems, and to providers and deployers of other AI systems and models, on the same conditions plus one more: the processing must be strictly necessary to address biases that could affect health and safety, harm fundamental rights, or lead to discrimination that EU Law prohibits.

Two limits keep this in proportion:

- It's a permission, not a duty. Article 4a states that it does not require anyone to carry out bias detection or correction.
- It works inside data protection law, not alongside it. It answers one question (whether you can use special category data for this purpose) and leaves every other data protection requirement exactly where it was.

The penalty structure

The penalty structure stands.

The maximum fines are unchanged (the percentages are of a company's worldwide annual revenue):

1. 35 million euros or 7 percent – for prohibited practices;
2. 15 million euros or 3 percent – for breaching the main obligations including Article 50;
3. 7.5 million euros or 1 percent – for supplying incorrect information.

EU institutions, bodies and agencies still face their own, lower caps, including 750,000 euros for transparency breaches.

The Omnibus makes three narrow changes. It lets small mid-cap companies use the more favorable “lower of the two” cap (paying whichever of the euro figure or the percentage is smaller). It lets Member States include warnings and other non-financial measures in how they penalize. And it makes one new type of breach fineable, for value-chain duties. It does not reduce the exposure.

The foundations of governance and interpretation are unchanged.

The European AI Board, the advisory forum, and the scientific panel all continue, with some adjusted tasks. Member States must still name the notifying authorities and market surveillance authorities (the national authorities that oversee and enforce the rules). EU institutions remain supervised by the European Data Protection Supervisor (EDPS).

And all the guidelines, from how systems are classified to transparency, remain non-binding.

Only the EU's top court, the Court of Justice, can interpret the Act authoritatively. So every interpretation in the current stream of guidance is the Commission's considered view, not settled law.

The bottom line.


A business's 2026 obligations (under the original prohibitions, the GPAI regime and most of Article 50 on transparency) are the same with or without the Omnibus. Compliance programmes built on the 2024 text are still the right foundation. The Omnibus updates the calendar and the way the rules are put into practice around that foundation.

Part | • Section 3

The EU AI Act as a living instrument, and how its five layers bind

The EU AI Act doesn't work on its own. It's the top layer of a five-layer structure.

When the Omnibus was adopted, four of those layers were still being built. Reading the amended articles alone won't tell a business what to do. What tells a business what to do is how the layers work together and each layer is binding in a different way.



Layer one

The Regulation as amended

This is the EU AI Act (Regulation (EU) 2024/1689), as amended by the AI Omnibus (Regulation (EU) 2026/1744). It applies directly, it is binding and enforceable, and only the Court of Justice can give the final word on its interpretation. This is the only layer that creates obligations. Every other layer below does one of four things: specifies detail, presumes something, grants an exemption or explains.

Layer two

Delegated and implementing acts.

These are binding secondary law or rules the Commission is authorised to write itself using powers the Act hands it. They carry the same legal force as the Act itself.

The Omnibus adds two new powers of this kind and grants them for five years from 27 July 2026:

- one to set limits on when a sector's own rules count as equivalent (Article 2(13));
- one to amend the new list of code categories in Annex XIV (Article 30(2)).

It also makes two smaller adjustments: it turns one binding rule into non-binding guidance (the post-market monitoring template, Article 72(3)) and it creates a new backup power to set common transparency rules if a code of practice is found to fall short (amended Article 50(7)).

This is a layer to watch. Delegated acts are where real relief may still appear and where real new burden may still appear. They are binding.

Layer three

Codes of practice

These are voluntary. They're drafted through processes that bring many stakeholders together, run by the AI Office, and then assessed by the Commission and the AI Board. Two exist and both are live now:

- the GPAI Code of Practice (covering Articles 53 and 55) - published 10 July 2025 and endorsed 1 August 2025; and
- the Code of Practice on Transparency of AI-Generated Content (covering Articles 50(2), (4) and (5)) - published 10 June 2026, found adequate by Commission Opinion on 8 July 2026 and the AI Board on 9 July 2026.

Codes don't create obligations and following one isn't conclusive proof of compliance. What they do is change where you stand with regulators: a business that follows a code of practice has an easier time showing its acting properly and attracts less scrutiny, while one that doesn't has to prove compliance in another way. That's a different thing from an obligation - and in practice, a more important one. Parts III and VI look at this closely.

Layer four

Commission guidelines

These are non-binding documents that explain how to interpret the Act - most issued under Article 96. Six important ones now exist, covering: what counts as an AI system, prohibited practices, GPAI models scope, serious incident reporting, Article 50 transparency and high-risk classification (still in draft form). At least four more are on the way (discussed in Section 17 of this playbook). Guidelines can't make something legal that the Act prohibits, and they can't make something illegal that the Act allows. But they are the benchmarks that the national enforcers (national market surveillance authorities) will actually apply - so if a business departs from them, it should be ready to defend that choice.

Layer five

European harmonised standards.

These are technical standards developed by the European standards bodies (in this case, CEN-CENELEC) at the Commission's request and, crucially, officially listed in the Official Journal.

This is the only layer that flips the burden of proof. Normally a provider has to prove that its system meets the law. But under Article 40, a provider that follows a cited harmonised standard is assumed to meet the requirements that standard covers. If a regulator wants to challenge that, the regulator has to show that the standard is inadequate or wasn't applied properly - the provider no longer has to prove conformity from scratch.

One catch: It isn't enough for CEN-CENELEC to adopt a standard.

It only counts once it's officially listed in the Official Journal and for the AI Act, that hasn't happened for any standard yet. Section 19 of this playbook looks at this closely.

More on the next page →

Three consequences follow, and they should shape how a compliance function is organised.

1. The layers have different failure modes, so they need different monitoring.

- Delegated acts - through the Commission's comitology register and the Official Journal.
- Codes of practice - through the AI Office's published adequacy assessments and the work of the signatory-taskforces.
- Guidelines - through consultations, and consultations close. The Article 6 classification window closed on 23 July 2026 and is gone.
- Standards - through national standards bodies and, decisively, the Official Journal citation calendar.

A business that only watches the legislation will miss the moving pieces that make the AI Act actually work.

2. Soft law is doing hard work, and the Omnibus deliberately gave it more.

By removing the approval step from Articles 50(7) and 56(6), codes of practice are no longer approved once and left alone. The Commission now judges whether a code is adequate on an ongoing basis and publishes its assessment. So a code is a moving object with a public scorecard, and being a signatory is a standing governance position rather than a completed project. At the same time, the Omnibus turned the binding post-market monitoring template into guidance and expanded the Article 96 guidance role across Articles 8 to 15 and Articles 25 and 26. In short, the detail that matters for compliance has shifted downward - out of the articles and into the layers beneath them.

3. Only one layer reverses the burden of proof, and it is the slowest.

This is the central imbalance in the whole structure. The layer that would help businesses most, cited harmonised standards, is the least developed, and the delays to 2027 and 2028 exist largely to buy time. Until a standard is cited, a business has to prove compliance on every requirement in Articles 9 to 15 and 17. That is why the standards layer, and not the amended articles, deserves the most senior attention in the next twelve months.

Part | • Section 4

The new master calendar

Dates below reflect the Omnibus as published, with entry into force on 27 July 2026.

Date	What applies	Who is affected
Since 2 February 2025	Original Article 5 prohibitions; Chapters I (General Provisions) and II (Prohibitions) generally; AI literacy duty (Article 4), now recast as an obligation of means.	Everyone placing on the market, putting into service or using AI systems in the EU
Since 2 August 2025	GPAI model obligations (Articles 51 to 55); Chapter III Section 4 on notifying authorities and notified bodies; Chapter VII governance; Chapter XII penalties (excluding Article 101); confidentiality rules (Article 78).	GPAI model providers; Member States; notified bodies; enforcement structures
27 July 2026	Entry into force of the Omnibus. Articles 102 to 110, the amendments to other EU acts apply. The amended governance and enforcement structure (including the AI Office powers in Articles 75a to 75d) takes effect. The two new delegated powers run for five years from this date.	AI Office; national authorities; operators under central supervision; product sectors affected by the Annex I reshuffle
2 August 2026	General Applicability of the EU AI Act. But, also Article 50 transparency in full (other than the marking obligation for older generative systems). Interaction disclosure under 50(1), machine-readable marking and detection under 50(2) for systems newly placed on the market from this date, emotion recognition and biometric categorisation disclosure under 50(3), deep fake and public-interest text labelling under 50(4), and the horizontal requirements in 50(5). The Commission's power to fine GPAI model providers (Article 101). Chapter IX market surveillance and post-market monitoring (including Article 72 and the Article 73 serious incident regime). Chapter X codes of conduct. Article 4a as a legal basis for processing special category personal data.	Providers and deployers of interactive and generative AI; GPAI model providers; all operators now facing an operational market surveillance system. None of this was deferred

Date	What applies	Who is affected
2 December 2026	New Article 5(1)(ba) and (bb) prohibitions on non-consensual intimate material and AI-generated CSAM, (with the calibration paragraphs 5(1a) and (1b)). Marking and detection for generative systems placed on the market or put into service before 2 August 2026, under the new Article 111(4) and 50(4).	All providers and deployers of generative systems. The prohibitions carry the top fines tier
2 August 2027	National regulatory sandboxes operational, deferred by a year (amended Article 57(1)). Deadline for the sectoral-equivalence delegated acts (Article 2(13)). Commission guidelines on the practical implementation of Articles 8(2), 9(10) and 17(3) due by 1 August 2027 (amended Article 96). Deadline for GPAI models placed on the market before 2 August 2025 to reach compliance (Article 111(3)).	Member States; the Commission; Annex I providers watching the equivalence perimeter; providers of legacy GPAI models
2 September 2027	Commission post-market monitoring guidance and voluntary template due (amended Article 72(3)).	Providers of high-risk systems
2 December 2027	Annex III stand-alone high-risk systems under Article 6(2), covering classification, the Articles 9 to 15 requirements, and the Articles 16 to 27 provider and deployer obligations, covering recruitment, credit scoring, education, critical infrastructure, biometrics, law enforcement, migration and justice (Chapter III, Sections 1 to 3).	Providers, deployers, importers and distributors of Annex III systems
28 January 2028	Deadline for notified bodies designated under Annex I Section A sectoral legislation, and relying on the bridge to assess high-risk AI systems, to apply for AI Act designation (amended Article 43(3)).	Notified bodies; operators depending on their capacity
2 August 2028	Annex I product-embedded high-risk systems under Article 6(1) (Chapter III, Sections 1 to 3). Machinery Regulation delegated acts importing the AI health and safety requirements into Annex III of Regulation (EU) 2023/1230 apply.	Product manufacturers acting as providers; sectoral value chains

Date	What applies	Who is affected
2 August 2030	Compliance deadline for legacy high-risk systems intended for use by public authorities (Article 111(2)).	Public-sector suppliers and deployers
31 December 2030	AI systems that are components of the large-scale EU information systems listed in Annex X (Article 111(1)).	Operators of EU large-scale IT systems

Reading the 2 August 2026 line correctly.

There is an imbalance here worth naming, because it's easy to miss and it changes the enforcement picture. The enforcement machinery is not delayed. Chapter IX (which covers post-market monitoring, information sharing and market surveillance) applies from 2 August 2026. So does all of Chapter XII on penalties, including Article 101. But the high risk obligations that this machinery exists to enforce (in Chapter III) don't apply until December 2027 and August 2028, respectively, as illustrated above. It would be a mistake to read that as enforcement has nothing to do for sixteen months. The opposite is true. From 2 August 2026 there is a fully working market surveillance system (including complaint rights under Article 85) and it applies to the obligations that are already in force, namely the Article 5 prohibitions, the Article 4 literacy duty, the Article 50 transparency duties and the GPAI rules. In other words, 2 August 2026 is the date the enforcement machine starts running, even though its largest single workload arrives later.

Grandfathering under Article 111(2), as clarified.

"Grandfathering" means letting systems that are already on the market keep operating under the old rules. The Omnibus ties this grace period for existing high-risk systems to the new Chapter III application dates, and makes one thing explicit: it does not touch Article 5. Prohibitions apply regardless of when a system was placed on the market. Otherwise the logic is unchanged and it operates per type and model. A unit identical to one that was lawfully placed on the market before the applicable date may continue to be placed in the market or put in service with no new obligations – as long as the design is unchanged. A significant design change triggers full compliance, and that threshold remains undefined. Retraining, material parameter updates or scope extensions may cross that line sooner than providers expect. Treat every model or system update from now on as something that could potentially reclassify it and write down your assessment each time.

The Article 111(2) grace period does not cover Article 50.

This is the single most-missed connection in the whole calendar, and the final Article 50 guidelines state it explicitly. The extra time under Article 111(2) applies only to high-risk requirements. Therefore where a system is both high-risk and covered by Article 50, both sets of duties apply. The grandfathering covers the high risk side, not the transparency side. What this means in practice is that an existing high-risk system that happens to be interactive still owes Article 50(1) disclosure from 2 August 2026, even though its high risk (Chapter III) obligations are years away. The only transparency-side relief is a separate, narrow one under the new Article 111(4), which covers the Article 50(2) marking duty for pre-existing generative systems, and it expires on 2 December 2026.

Part II •

The amended EU AI Act by actor

Part II • Section 5

Providers of high-risk AI systems

Providers of high-risk AI systems

What you gained.

1. Time

The single most consequential relief in the package is time, with sixteen additional months for Annex III standalone high risk systems and a deadline of August 2028 for Annex I product-embedded high risk systems, on fixed dates rather than the conditional trigger the Commission originally proposed.

2. Fixes for being wrongly classified as high risk

Beyond time, the Omnibus delivers a practical test to avoid being wrongly classified as high risk. Three new provisions work together:

- Article 6(1a) - what doesn't count. An AI system used only for things unrelated to safety (user assistance, performance optimization, service efficiency, automation, convenience or quality control) does not qualify as a "safety component";
- Article 6(1b) - the floor. Despite that carve-out, a system whose failure or malfunctioning would endanger health and safety does qualify. This is a baseline that the carve-out can't undercut;
- Article 6(1c) - fixing an absurdity. Sometimes a product needs to undergo third-party conformity assessment, but only because of risks that have nothing to do with health and safety (radio spectrum use is the standard example). That kind of check no longer drags the product into high risk classification under Article 6(1)(b).

The amended Article 3(14) definition of "safety component" supports all of this by adding that a component counts as a safety function where its intended purpose is to prevent or mitigate risks to health and safety.

3. No more proving cybersecurity twice where the Cyber Resilience Act applies

The new Article 42(3) builds the bridge directly into the AI Act itself, so that where a high-risk system falls within the scope of the Cyber Resilience Act (Regulation (EU) 2024/2847) and meets the conditions in its Article 12(1), the system is automatically treated as meeting the AI Act's cybersecurity requirements in Article 15.

4. Conformity routes are clarified

The rewritten Article 43(3) states two things:

- Manufacturers of Annex I Section A products don't have to choose a conformity assessment procedure involving third-party assessment just because the product includes a high-risk AI system as a safety component (Annex I Section A addresses the EU's modern product-safety laws - machinery, medical devices, toys, radio equipment, and the like - where an outside body typically certifies a product before it goes on sale).
- When a system falls under both Annex I Section A and an Annex III category, the sectoral conformity process is the one that applies.

5. Real-world testing expands

Two changes widen where AI can be tested in real-world conditions (rather than only inside regulatory sandboxes) - a controlled, supervised test environment):

- Article 60 (amended) - Real word testing outside sandboxes used to be available only for Annex III standalone high risk systems. It now also covers high-risk systems covered by Annex I Section A legislation.
- Article 60a (brand new) - This goes further. It lets Member States set up frameworks for real-world testing of AI-enabled products under Annex I Section B sectoral law (aviation, vehicles and now machinery) against the Articles 8 to 15 requirements. The testing comes with safeguards: a required testing plan, notice to the Commission, and a high standard of protection.

Article 76 is adjusted so the right authorities are in charge of overseeing this.

6. Monitoring and registration burdens ease

- The rigid post-market monitoring template used to be rigid; it now becomes Commission guidance with a voluntary template due by 2 September 2027 (see Article 72(3)). That gives you room to tailor your monitoring to your own organisation.
- Registration is also trimmed a little. Points 7 and 9 of Annex VIII Section B are deleted, so providers of systems that were filtered out of high risk under Article 6(3) now have less to file in the EU database. But the underlying Article 6(4) assessment still has to be documented and handed over if a regulator asks.

What you did not gain.

No cut to the substance of Articles 9 to 15 and 17 to 25 (as Section 2 of this playbook explains).

In future, the equivalence mechanism in Article 2(13) (a future power to disapply specific AI Act requirements where a sector's own safety law already gives equal or better protection) may limit the application of specific requirements under those articles for Annex I Section A systems – but only where sectoral legislation provides equal or a higher level of protection and the limitation does not reduce overall protection. That relief, though, depends entirely on delegated acts due by 2 August 2027 that don't exist yet. When they arrive, they'll specify the systems, requirements, conditions and scope. So the prudent planning assumption is full compliance. Treat equivalence as potential upside for specific product families later, not as a general exemption you can count on now.

What you now owe that you did not owe before.

1. The value chain is a matter of contract,

and this is the most underestimated new burden in the package.

- Article 25(4) - written agreements with your suppliers. You now need a written agreement with every third party that supplies something built into your high risk system: systems, AI models, tools, services, components or processes. The Omnibus expressly adds AI models to that list. Each agreement has to specify the information, capabilities, technical access and assistance needed, measured against the current accepted best practice(s). Public free and open-source resources remain exempt, but GPAI models are not.
- Article 25(2) - what the original provider must hand over. This used to be a vague duty to “closely cooperate”. It’s now a concrete list. When a new provider takes over a system - by rebranding it, substantially modifying it, or repurposing it - the original provider must: make available enough technical documentation to assess check with Article 16, inform the new provider about known limitations and the ways the system can fail, and give specific technical access, including for testing and validation.


Both duties can now be fined. Under the new Article 99(4)(da), breaching either one can cost up to 15 million euros or 3 percent. That is new exposure attached to your contract paperwork (unusual in this Act) and means how well your supplier agreements are drafted is now directly a penalty question.

2. An option to switch the duty off.

There is also a useful protective option here, carried over from the original text. A provider can state up front that its system is not to be turned into a high-risk system - and doing so switches off the cooperation duty. So if you don’t intend for others to repurpose your system, build that exclusion into standard terms now. Make sure your marketing and solution-engineering teams know they can’t contradict it in practice, because as Section 15 of this playbook explains, the draft high-risk classification guidelines say a system’s intended purpose has to be consistent across its instructions for use, its marketing and its technical documentation.

How to use the runway.

The delay exists because standards, guidance and working enforcement structures weren't ready, and that gap is closing on a schedule you can plan around. Sequence the two years in a deliberate order.



Throughout, document your classification assessment, even for systems you conclude are not high-risk. Relying on the Article 6(3) filter comes with its own duties: you have to document the assessment (Article 6(4)) and register in the EU database, the draft guidelines read the filter more narrowly than many expected, and regulators can demand to examine the assessment.

One live trap on incident reporting.

The obligation to report serious incidents (Article 73) sits in Chapter IX, so it applies from 2 August 2026 - not the later Chapter III dates.

The Commission's draft guidance of 26 September 2025, and the reporting template, that came with it, tell us three things:

- 1. The deadlines are short** - reports are due within two to fifteen days, depending on how serious the incident is and what type it is.
- 2. The trigger is wide** - even an indirect causal link between the system and the harm is enough to trigger a report.
- 3. It overlaps with other laws** - The guidance maps how it interacts with DORA, the Critical Entities Resilience Directive and the Medical Devices Regulation.

The final guidance is a pipeline item to watch. In the meantime, build your incident-detection and escalation process on two assumptions: that the clock is short and that a “serious incident” is defined broadly. The definition under Article 3(49) includes breaches of fundamental rights obligations, so its broader than a technical idea of something “failing”.

If you sell to the public sector.

The 2030 deadline for legacy systems is a hard backstop, but don't plan to it. Public buyers will start requiring proof of compliance in their tenders well before then. So treat 2030 as the outer limit; your real timeline will be set by when your buyer starts asking.

Part II • Section 6

Providers of general-purpose AI models and GPAI-based systems

Providers of general-purpose AI models and GPAI-based systems

Your actual obligations haven't changed and they already apply. The Articles 51 to 55 rules have been in force since 2 August 2025 and the Omnibus leaves what they require untouched. What does change for you is who supervises you, how codes of practice work,

the density of hard dates in the next fourteen months, and three concrete new demands: documented safeguards against the new prohibited-content bans, written agreements down your value chain, and a fresh legal basis for bias-testing work.

The supervisory ground has shifted.

The amended Article 75(1) makes the AI Office the exclusive body responsible for supervising and enforcing two groups of AI systems:

- systems built on a general-purpose AI model where the provider (or another company in the same corporate group) made both the model and the system; and
- newly, systems that are (or are integrated into) a very large online platform (VLOP) or very large online search engine (VLOSE) designated under the Digital Services Act (DSA).

Some things are carved out and stay with national or sectoral regulators: Annex I product systems, critical infrastructure (Annex III point 2), law enforcement, border management, financial institutions supervised under Article 74(6), and the justice system.

What this means in practice:

- If you're a frontier developer offering consumer systems built on your own models, you exchange twenty-seven national authorities for a single EU-level supervisor, and that supervisor now carries the enforcement powers described in Section 11 of this playbook.
- For AI built into a platform, the DSA's own risk assessment, mitigation and audit obligations are the first thing you face, though the AI Act can still be enforced against you afterward.

The supporting machinery matters operationally. Under the new Article 75(1a) to (1e) and (2a):

- serious incidents for these systems are reported to the AI Office, not a national authority;
- national authorities must assist;
- the AI Office is responsible for pre-market conformity assessments and tests, carried out through notified bodies under Commission responsibility, with the provider paying the fees; and
- national market surveillance authorities can formally refer suspected breaches through designated contact points, and the AI Office has to respond within four months.

That last channel is the one to note: a complaint filed in any Member State can surface in Brussels on a guaranteed timetable.

Codes of practice change in legal character.

The implementing-act approval step has been removed from Article 56(6). Instead, the Commission now assesses on an ongoing basis whether a code covers the Articles 53 and 55 obligations (giving great weight to the Board's opinion), monitors whether the code is actually working, and publishes its assessment. Codes are still the practical way to show compliance, but “adequate” is now a living, published judgment rather than a one-off stamp of approval. For more information on the GPAI Code of Practice, see Section 16.

This means both your signatory status and the code’s ongoing changes need someone inside your organisation who owns them permanently. The Signatory Taskforce set up under the GPAI Code, and chaired by the AI Office, is where the real interpretation is being worked out, so taking part in it is now genuine compliance work, not a public-affairs exercise.

Your key dates.

2 August 2026

is your enforcement date, and it was never delayed.

Article 113 always excluded Article 101 from the August 2025 group, scheduling it for 2 August 2026 instead. From that date, the Commission can fine providers of GPAI models up to 3 percent of worldwide annual revenue or 15 million euros, whichever is higher. It can also start using its full toolkit of powers: demanding information, getting access to models, running evaluations, and ordering corrections. In its first year, the AI Office took a lenient stance – it treated signatories who hadn’t yet fully met their commitments as acting in good faith and worked with them toward compliance. But that stance was explicitly tied to the period up to 2 August 2026. If you’re still relying on it, assume it’s over.

2 August 2027

is your legacy-model date.

Under Article 111(3), if you placed a GPAI model on the market before 2 August 2025, you have to bring it into compliance by 2 August 2027. By then, documentation, the copyright policy, the training-content summary and, for systemic-risk models, the safety and security framework, all have to be there. The Omnibus didn’t change this.

Two hard deadlines on transparency

These are about marking AI-generated content (synthetic audio, image, video or text):

- Systems placed on the market or put into service **before 2 August 2026** must meet the Article 50(2) marking and detection duty by **2 December 2026** (the new Article 111(4)).
- Systems newly placed on the market from **2 August 2026** must comply straight away.

The final Article 50 guidelines are clear that there are two parts and you need both: a machine-readable mark and a working way to detect the mark. A mark with no detection route doesn't comply. Section 12 of this playbook has the details – including that you have to use publicly available, industry-standard detection tools, and that today's permission to use your own or shared alternatives is strictly temporary.

Watch how this interacts with the model layer. Article 50 applies to the system; Article 53 applies to the model – so a single product can trigger both. The final guidelines say you can put the marking solution at different points in the value chain: added afterward at the system level, built into the underlying model, or built into the inference process (how the model generates output). You can also rely on an upstream model provider's or a third party's solution, as long as that solution complies. But relying on someone else doesn't transfer the responsibility; as the system provider, you still have to show your product complies.

If an in-scope system runs on your own GPAI model, the guidelines say you could put the transparency measures at the model level. For models you supply to others, marking at the model-level is encouraged but not required, and for models with systemic risk, it can count as part of your Article 55(1)(b) risk-mitigation measures. One thing to watch: The Transparency Code may ask signatories to do more on upstream marking than the guidelines' minimum, and that gap is the real “how much do we need to build?” question for model providers.

Three new concrete demands.

On the prohibitions

1.

by 2 December 2026 you need a documented set of safeguards against two things: generating non-consensual intimate imagery (NCII) and generating generating child sexual abuse material (CSAM). Section 14 of this playbook sets out the full legal test; the short version is that the prohibition applies when you put a system on the market or into use and either:

- the system is intended to generate that material; or
- the system's design, training, architecture, capabilities or user-facing functionalities makes generating that material a reasonably foreseeable and reproducible outcome – and the system lacks reasonable and adequate technical safety measures to prevent it.

So the thing you produce to show compliance is the documented safeguards themselves, tailored to how you release the system (and this expressly includes open-source releases). The fines here are the highest in the Act.

Value chain

2.

Open-source is normally exempt from the Article 25(4) duty to have a written agreement, but that exemption does not cover GPAI models. The Omnibus added AI models to the provision precisely to close that gap. So if your model gets built into someone's high-risk system, expect to sign a contract even if the model is open-source. Prepare standard information and technical-access packages up front, so each integration isn't a one-off negotiation. And bear in mind that the other side will push hard for this: under Article 99(4)(da), they can be fined if the agreement isn't in place.

Bias work

2.

Article 4a(2) gives providers whose systems and models are not high-risk a lawful way to use special categories of personal data (data on things like race, health, or beliefs) for bias detection and correction. You have to meet all of the safeguards in Article 4a(1), together:

- use it only when synthetic or anonymised data won't do the job;
- pseudonymise the data;
- keep it under access controls;
- Don't share it with anyone else; and
- delete it once the bias is corrected.

Importantly, this is a permission, not a duty - it does not require you to carry out bias detection and correction. It is a legal basis and it works within the GDPR. Recital 9 makes it available ahead of the Chapter III deadlines, so you can rely on it now rather than waiting for December 2027 or August 2028. It does not reach backwards: the old Article 10(5) was deleted on 27 July 2026 without ever having applied, so bias work carried out before that date falls to be assessed under the GDPR alone.

Part II • Section 7

Deployers

Deployers

Prohibitions bite on intent and circumvention. From 2 December 2026, using an AI system to generate or manipulate non-consensual intimate material or CSAM is prohibited. The calibration paragraph 5(1a(b)) limits the deployer-side prohibition to purposeful generation, so accidental output and lawful uses aren't caught, but deliberate misuse is banned no matter what tool is used, including when someone gets around the safeguards on an otherwise lawful system.

Update acceptable-use policies, monitoring and disciplinary frameworks to match. And make sure your contracts pass the point through to contractors and agency staff who use systems on your behalf, since the Article 50 guidelines are clear that your organisation stays the “deployer” even when a third party operates the system for you and under your control.

Not having to duplicate the FRIA and the DPIA is a real relief. The Article 27 fundamental rights impact assessment (FRIA) duty itself is untouched, but the amended Article 27(4) flips how it relates to the data protection impact assessment (DPIA). Previously, the FRIA had to complement the data protection impact assessment. Now, wherever a DPIA already covers one of the Article 27 requirements – a DPIA under Article 35 GDPR or under Article 27 of the Law Enforcement Directive – you can just point to the relevant sections of that DPIA, or copy relevant parts of it into the FRIA.

So restructure your assessment process now, so the DPIA is built in a way that feeds straight into the FRIA. And when the AI Office publishes its questionnaire template – which is required to support this and to include an automated tool. Use it, rather than building your own custom format.

Supervision. If you're an ordinary business using frontier or platform AI, your national market surveillance authority stays the body you deal with. You come under the AI Office's exclusive control only if one of these is true: you're also the provider, you're part of the same corporate group as the provider, or the system is (or is integrated into) a designated VLOP or VLOSE.

The Article 50 guidelines add a useful detail on the transparency side:

- For AI systems built on GPAI models from the same provider, the AI Office is the market surveillance authority (which covers the provider-side duties in Article 50(1) and (2))
- It takes over the deployer-side duties in Article 50(3) and (4) only where that same model-and-system provider is also the deployer.

Be disciplined about repurposing. If you modify or repurpose a system in a way that makes you the new provider, all of the Chapter III provider obligations land on you. The draft Article 6 high risk classification guidelines confirm that configuring or fine-tuning a general-purpose system to perform an Annex III (high risk) task typically does exactly that. The Article 50 guidelines say the same thing for transparency: a company that takes an existing generative system, modifies it, and puts it out under its own name becomes the provider of the new system.

The other perspective is that you now get the benefit of the initial provider's cooperation duties under Article 25(2) – they must give you enough documentation to assess Article 16 compliance, tell you about known limitations and the ways the system can fail, and provide targeted technical access. So before any integration project, work out whether it trips the Article 25 trigger, and lock the documentation hand-over into the contract rather than assuming you'll get it.

Timing. The deployer duties for high-risk systems under Article 26 follow the delayed dates – December 2027 and August 2028. Those duties cover human oversight, controlling the input data, keeping logs, informing workers, monitoring, and (under Article 26(11)) telling people when they're subject to a high-risk decision December 2027 and August 2028. Deployers running legacy systems for public authorities have the 2030 backstop.

But the Article 50(3) and (4) duties are different: they're yours from 2 August 2026, and they were not delayed. So they apply to you now if you do any of these anywhere in your organisation – publish, advertise, broadcast, run a helpdesk, operate a newsroom, use synthetic presenters, or run emotion recognition or biometric categorisation.

Part II • Section 8

Horizontal expectation applicable to both providers and deployers (regardless of risk-classification)

Horizontal expectation applicable to both providers and deployers (regardless of risk-classification)

AI literacy softened, not removed. Article 4 used to require you to “ensure a sufficient level” of AI literacy. It now requires you to “support the development” of AI literacy among your staff and the people who run the systems on your behalf – and it states expressly that you are not required to guarantee that any particular person reaches any particular level. In practice: keep the training programme, but stop worrying about proving each individual’s proficiency.

Also worth knowing: the European Commission and Member States have to help with this, and the Board is due to issue recommendations. Practical material will appear on the [AI Act Single Information Platform](#), and using that is cheaper than building your own.

Part II • Section 9

SMEs, start-ups and small mid-caps

SMEs, start-ups and small mid-caps

The Omnibus makes small mid-cap enterprises (SMCs) an officially recognised group that gets special treatment across the Act, alongside small and medium-sized enterprises (SMEs). (SMCs are defined by reference to point (2) of the Annex to Recommendation (EU) 2025/1099; SMEs by Article 2 of the Annex to Recommendation 2003/361/EC. Both definitions are now written into Article 3, as points (14a) and (14b), and SMC status runs through the Act in several places: the innovation goal (Article 1(2)(g)), technical documentation (Article 11(1)), quality management (Article 17(2)), national guidance duties (Article 70(8)), codes of conduct (Article 95(4)), priority access to sandboxes (Article 57(3a)), and penalties (Article 99).

The offer on help is substantial:

- **Simpler technical documentation.** You can file it on a Commission form. If an eligible business chooses to use it, the notified body has to accept it for the purposes of conformity assessment, and it can't insist on the full Annex IV documentation set.
- **Right-sized quality management.** The quality management system has to be scaled to the provider's size - especially for a SME, a start-up, or an SMC - while still meeting the required rigour and level of protection. The simpler route under Article 63(1) now extends beyond micro-enterprises to all SMEs including start-ups (without partner or linked enterprises), with Commission guidelines to come.
- **Priority in the sandbox.** Smaller players get reserved priority access to the new EU-level AI Office sandbox (Article 57(3a)), and the sandbox's purpose has been rewritten around helping them to get to market faster.

- **Targeted national guidance.** National authorities get stronger powers to issue guidance aimed specifically at SMEs, start-ups and small mid-caps.
 - **Softer penalties.** Member States have to take account of smaller businesses' interests and whether they can survive a fine. Penalties can include warnings and non-financial measures. And for small mid-caps, each fine under Article 99(4) and (5) is capped at the lower of the percentage or the fixed amount (the treatment SMEs already had).
-

Two cautions

- **Lighter treatment isn't a free pass.** You will still have to prove you comply. The simpler form changes the format of documentation, not the duty itself.
 - **Check that you actually qualify.** Eligibility turns on precise ownership and linkage tests in the recommendations these definitions point to. So confirm your status before you build a compliance strategy around it. And check again after any funding round or acquisition, because a change in who owns the company can silently remove the benefits at the worst possible moment.
-

What to watch. Three things:

1. the Commission's simplified documentation form;
2. the Article 63 guidelines, which will define what the light-touch quality management route actually involves; and
3. the criteria for getting into the EU sandbox's.

For a smaller provider, priority access to a supervised channel with the AI Office does two jobs at once: it lowers your risk, and it's a mark of credibility with big enterprise customers. One more thing worth knowing: SMEs are actively helped to take part in setting the standards (through the standardisation request) and that's the cheapest way there is to influence the technical detail you'll later be audited against.

Part II • Section 10

Importers, distributors, third-party suppliers and the conformity infrastructure

Importers, distributors, third-party suppliers and the conformity infrastructure

Importers and distributors. The verification and gatekeeping duties in Articles 23 and 24 keep their content and follow the delayed dates, so they apply from 2 December 2027 for Annex III systems, and from 2 August 2028 for Annex I systems. The duties themselves haven't changed. Before an importer puts a system on the market, it has to confirm that:

- the conformity assessment was carried out;
- the technical documentation exists;
- the system carries the CE marking; and
- it comes with the required documentation and instructions for use.

And the importer has to refuse to place non-conforming systems on the market. Distributors owe the corresponding downstream checks.

The re-anchored grandfathering rule (the grace period that lets existing systems keep operating under the old rules) helps this part of the chain in a very practical way. Because the grace period works per type and model, further identical units of a system lawfully placed on the market before the applicable date can continue to be distributed without new obligations – for as long as the design is unchanged. That keeps supply flowing for legacy systems. But the moment a supplier ships a significantly changed design, the full set of rules kicks in.

So your import and distribution contracts should require suppliers to notify you about any design changes, and should allocate the resulting compliance risk explicitly. And note (as detailed in Section 6 of this playbook) that the grandfathering does not cover Article 50, so a legacy generative or interactive system moving through your channel carries transparency duties on the 2026 dates regardless.

Third-party suppliers into high-risk systems. If you supply anything that gets built into someone's high risk system – an AI system, an AI model, tools, services, components or processes – you now come under the Article 25(4) written-agreement duty. The agreement has to specify information, capabilities, technical access, and assistance you'll provide, measured against the generally acknowledged state of the art. Public free and open-source suppliers are exempt, unless what they supply is an AI model. And a breach can be fined on its own, up to 15 million euros or 3 percent.

Suppliers should put together standard information and technical-access packages now. Every high-risk provider in the EU will soon be asking for essentially the same thing, and a supplier who can hand over a ready-made package will win deals against one who negotiates each request from scratch.

And if you're a supplier who takes over a system and becomes the new provider (by rebranding it, substantially modifying it, or repurposing it) you now get the benefit of the original provider's cooperation duties under the fleshed-out Article 25(2). That significantly lowers the risk in acquisition and integration projects, so it's worth factoring into your deal terms.

Notified bodies and the assessment capacity question. ("Notified bodies" are the independent organisations that certify products by conducting conformity assessments.) Whether the 2027 and 2028 deadlines are realistic depends less on businesses than on whether there are enough notified bodies to assess everyone in time. The Omnibus attacks that bottleneck from four directions.

- 1. One application, one combined assessment.** Under the new Article 28(8) and (9), a conformity assessment body that wants to be approved under both the AI Act and Annex I Section A sectoral legislation can submit a single application and go through one combined assessment. A body already approved under several Section A laws only has to apply once.
- 2. A faster route.** The amended Article 29(4) lets a body use the documents and certificates from its existing sectoral approvals not just to support its AI Act application but to expedite it, with that single application routed to the notifying authority that handles the sectoral side.

3. A precise map of what each body can assess. The amended Article 30(2) requires approvals to be based on a new list of codes, categories and corresponding types of AI systems in a brand new Annex XIV. In plain terms, it's a catalogue of AI system types:

- AIP 0102 to 0112 - systems tied to Section A product legislation;
- AIB 0201 to 0203 - the biometric systems of Annex III point 1;
- AIH codes - cross-cutting technology types, including symbolic AI, machine learning (by data type), and generative and GPAI-based systems (AIH 0301);
- AIH 0401 - a catch-all for all other emerging AI technologies not covered above, expressly including agentic AI.

Annex XIV can be amended by delegated act. Its practical function is to make sure that a body is only approved for what it can actually assess – and, as a side effect, it gives everyone the EU's first formal legal taxonomy of AI technology types, which will get reused elsewhere.

4. The bridge. Under the rewritten Article 43(3), notified bodies already approved under Annex I Section A legislation can assess the conformity of high-risk AI systems now, but they must apply for full AI Act approval by 28 January 2028.

What this means for you: Treat the flow of new approvals as the leading indicator of how much capacity exists, and book your assessment slots early, rather than assuming availability in late 2027. Two more points:

- Notified bodies can't refuse the simplified technical documentation form from eligible SMEs, start-ups and small mid-caps.
- For systems under the AI Office's exclusive control, the pre-market third-party conformity assessment is run under the Commission's responsibility (through notified bodies), with the provider paying the fees. That creates a new centrally-run stream of assessment work (and a new cost) for centrally supervised providers.

Part II • Section 11

The new enforcement reality

The new enforcement reality

This is the most underestimated change in the package. For businesses under its exclusive control, the AI Office gains an

enforcement toolkit deliberately modelled on competition (antitrust) law, in force from 27 July 2026.

Under the new Articles 75a to 75d, the AI Office can:

- request information either by simple request or by a binding decision backed by the threat of fines;
- investigate on its own initiative or after a complaint;
- conduct remote and on-site inspections, including entering any business premises, land or property located in the EU; examining books, data and other material; taking copies and extracts; asking anyone for oral or written explanations; and sealing business premises, books or records;
- call on national assistance, including police support if its obstructed, (with court approval where national law requires it);
- order access to AI systems and explanations of how they work, and require data and documents to be retained;
- bring in external experts and auditors.

A business can offer commitments to resolve concerns. The AI Office can make those binding and then close the case – but it can reopen it if the facts change significantly, the commitments are broken, or the decision was based on misleading information. If it finds non-compliance, the AI Office can impose the Article 99 fines itself. It can also charge a daily penalty – up to 5 percent of average daily income or worldwide annual revenue from the previous year – for each day of continued non-compliance.

Other key points:

- a five-year time limit applies;
- fines go to the EU budget;
- decisions can be fully reviewed by the Court of Justice;
- businesses keep their rights of defence and negotiated access to the case file, protecting business secrets;
- under Article 75d, the AI Office must publish the decisions it adopts under Articles 75b and 75c, which is a reputational risk in its own right.

There is a structural criticism of this design worth taking seriously, because it will shape how enforcement actually feels in practice. The AI Office has been given the powers that resemble those of an antitrust authority – but the resourcing provision behind them (the new Article 64(3)) states that the AI Office shall be allocated adequate resources to do its job and even that is subject to the normal EU budget process. In other words, there is no guaranteed money, and a power like sealing premises is only as real as the staff and machinery available to use it.

So the likely result is not under-enforcement, but concentrated enforcement: a small number of high-profile cases against centrally supervised businesses, with published decisions doing the work that a broad inspection programme can't.

For businesses in-scope this translates into concrete steps.

- Build an inspection protocol that covers AI records – not just email. That means model documentation, training-data records, evaluation logs, red-team (adversarial testing) results, and versions of your safety framework.
- Set up record-keeping that can preserve data when litigation requires it (a “litigation hold”) – and square that with the duties to delete data in Article 4a and the GDPR, so the two don't collide.

- Decide your organisation's position on the commitments route ahead of time. It will come up as an option in the middle of an investigation, when there's little time to think it through – and a commitment decision is made public.
- Treat that publication as an ongoing reputational risk, to be managed alongside the legal one.

For everyone else, your national regulator stays the body you deal with, and from 2 August 2026 that system is fully up and running. The right to complain under Article 85 is open to anyone affected, or to any person or organisation with reason to believe a breach has occurred.

Two channels widen your exposure at a national level:

1. Referral up to the AI Office. National authorities can pass matters to the AI Office through designated contact points, and the AI Office must respond within four months. So a complaint filed in your own country can end up in Brussels.
2. Access for rights and equality bodies. Under the amended Article 77, fundamental-rights and equality bodies get stronger access to information and documentation created or kept under the Act. Its routed through the market surveillance authority, which has to hand it over promptly, in plain language, and in a machine-readable electronic format – and there's a new duty for these bodies to cooperate and exchange information with each other layered on top.

Expect non-discrimination authorities to become active users of AI Act enforcement material – especially around employment and essential-services, and expect them to ask for it in a form that's easy to analyze.

Part III •

The transparency obligations

Part III • Section 12

How Article 50 (Transparency of AI- Generated Content) works in operation under the European Commission Guidelines

The Commission approved the final Article 50 guidelines...

for transparency of AI-generated content on 20 July 2026 ([C\(2026\) 5054 final](#)). They are non-binding but they are the benchmark authorities will apply. Two framing points matter:

1. The four duties below apply cumulatively and a single product can trigger them for the provider and the deployer at the same time.
2. Meeting Article 50 doesn't make a system lawful everywhere else. A deep fake can be labelled perfectly and still be prohibited, high-risk or breach another rule.

The four duties under Article 50:

Article 50(1): disclosures for interactive systems. This applies where a system is intended to interact directly with people in a genuine back-and-forth conversation. That covers chatbots, voice assistants, helpdesks, robots, companions, avatars and coding agents, but not recommenders, spam filters, translation, search, autocomplete or behind-the-scenes decision support.

For AI agents, you have to disclose two things: that the agent is artificial and who it is acting for. Where you can't predict the interaction in advance, build that disclosure into the system itself; and disclose it to people directing the agent at the authorisation and validation steps.

Five techniques are not enough on their own:

- 1) burying the disclosure in the terms and conditions;
- 2) machine-readable marks that a person can't notice at the point of interaction;
- 3) vague signals like calling it an "assistant";
- 4) platform-wide notices such as "services on this website use AI"; and
- 5) technology descriptions such as "this system uses LLMs."

The obviousness exception is narrow. There is an obviousness exception, but it's read narrowly - limited to cases where there's essentially no doubt, such as developer tools, internal staff assistants and clinician-only systems. It expressly does not cover consumer helpdesk chatbots or immersive avatars.

Article 50(2): content marking and detection. Both are required, not just one: AI-generated content must carry a machine-readable mark and there must be a way to detect that mark and show a human-readable result available to the people who encounter it. A marking with no detection route doesn't comply.

It applies across content types - audio, image, video and text, including multimodal, 3-D and VR - and what you mark is the final asset.

Out of scope:

- short strings, captions and UI labels;
- source code and APIs;
- machine-to-machine traffic;
- an agent's reasoning and browser actions;
- sensor observations'; and
- internal closed-loop production workflows (where only the final output is marked).

Editing: Standard editing is exempt - grammar, formatting, compression, minor cropping and colour work, transcription and translation. But **summarisation, paraphrasing, object or face replacement, voice synthesis and composites all require marking.**

Detection has to use publicly available, industry-standard solutions. Using your own proprietary detection is expressly only a temporary bridge.

Three narrow carve-outs exist:

- less robust metadata for closed instructive products such as vehicle navigation;
- no marking for strictly technical industrial or B2B output; and
- Real-time ephemeral game or VR content, where disclosure is made once per session.

Article 50(3): specifics for emotion recognition or biometric categorisation. If you use either of these, you have to tell the people exposed – in real time or afterwards and including children. Two points: emotion recognition is always high-risk (unless it's outright prohibited, as it is in workplace or education); and biometric categorisation is covered **whether or not** it's high-risk.

Article 50(4): specifics for deepfakes and text outputs.

Deepfakes: A deepfake, for this purpose, is content that meets criteria: It noticeably resembles something existing (or plausibly existing), whether a person, object, place, entity or event, and that falsely appears authentic or truthful.

This is judged objectively – you don't need any intent to deceive. And unlike Article 50(1), the test is not whether an average person would be fooled; it's enough if a vulnerable part of the likely audience would be. Standard film backgrounds and effects aren't deep fakes; AI-generated actors, digital replicas, de-aging and fake documentary content are. **A deployer can't satisfy this duty with the provider's machine-readable marking** – the label must be visible or audible without any tools. Works that are clearly artistic, satirical or fictional get a lighter duty (not an exemption), and that's read strictly; where a work is both informative and creative, the informative side wins.

Text outputs: This applies to published text that informs the public on matters of public interest – from politics and administration to health, environment, consumer safety and financial news. It covers things like AI-generated news summaries, health-related lifestyle articles, corporate and investor reports, and public warnings. It does not cover fantasy novels, product copy with no health, safety or sustainability claims, single-user chatbot replies, or client advice.

There is an editorial exception, but it's demanding. It needs:

- real substantive review, with fact-checking as a minimum; and
- a named person who is publicly identified as taking editorial responsibility.

And it falls away if any substantive AI change is made after sign-off – which catches generating headlines, rewriting for SEO, and auto-translating a piece after its been approved.

Timing and adjacent regimes: The disclosure has to be clear, distinguishable and hard to miss – given the first time each person interacts with or sees the content, and per output. If people join partway through, the opening disclosure has to be topped up again later.

On dates:

- Article 50 applies from **2 August 2026, no matter when a system was placed on the market.**
- The only relief is the Article 111(4) window to 2 December 2026 for Article 50(2) marking on pre-existing generative systems. A system that's partly interactive and partly generative still owes Article 50(1) from August.
- Content generated before 2 August doesn't need marking after the fact, **but text generated before and published on or after that date has to be labelled.**
- Article 111(2) grandfathering doesn't cover Article 50 at all.

Separately, under the Consumer Rights Directive AI functionality may count as an “essential characteristic” you have to disclose before the contract – regardless of **whether the AI is obvious**. Fines here run to 15 million euros or 3 percent, whichever is higher.

Part III • Section 13

The Code of Practice on Transparency of AI- Generated Content

Published 10 June 2026;

found adequate by the Commission on 8 July 2026 and by the AI Board of 9 July 2026. It is the only EU-wide recognised way to show you comply with Article 50(2), (4) and (5) – wherever you’re based and whichever regulator oversees you. But following it is expressly not conclusive proof of compliance: signing it does not mean you’re automatically treated as complying). The Code has two parts:

Section 1 **for providers**

Marking and detection, based on secured metadata plus watermarking.

Section 2 **for deployers**

Labelling, and it comes with an official EU icon set.

As with codes generally now, “adequate” is an ongoing and published assessment – and if a code is deemed inadequate, the Commission can step in and impose common rules through an implementing act.

Enforcement works differently depending on whether you sign the Code of Practice.

- If you sign, supervision focuses on whether you’re actually following it, and your commitments can reduce a fine under Article 99(7)(e). You can opt out of a specific section, but you visibly lose the benefit for that section.
- If you don’t sign, you have to prove you comply by some other adequate way, check yourself against the Code to find any gaps, and expect more requests for information and access.

Either way, the Code is the benchmark you’ll be assessed against. So not signing still leaves you doing the analytical work – just without the supervisory benefit.

Part IV •

Prohibitions for AI Systems with Unacceptable Risk

Part IV • Section 14

The Article 5 prohibitions, the guidelines and the safeguards you'll need

The prohibitions are the part of the Act that have been in force the longest and have the best guidance behind it.

And, with the Omnibus additions, they're also the part with the soonest new deadline, carrying the highest tier of fines. They're worth understanding for another reason too: they show a built in quirk of a law that keeps evolving. The guidance can run a year and a half ahead of the legislative in one respect, and behind it in another.

14.1 The existing prohibitions and the Commission's guidelines

The eight original prohibitions have applied since 2 February 2025, and fines for breaching them since 2 August 2025.

They prohibit:

1. subliminal, purposefully manipulative or deceptive techniques that seriously distort people's behaviour and causing (or are reasonably likely to cause) significant harm;
2. exploiting people's vulnerabilities due to age, disability or a specific social or economic situation, with the same harmful effect;
3. social scoring that leads to unjustified or disproportionate treatment of people;
4. predicting how likely someone is to commit a crime based only on profiling or personality traits;
5. scraping facial images in bulk from the internet or CCTV to build or expand facial recognition databases;

6. inferring people's emotions in workplaces and education institutions, (subject to narrow medical and safety exceptions);
7. using biometrics to infer sensitive traits (such as race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation); and
8. real-time remote biometric identification in publicly accessible spaces for law enforcement (allowed only under the narrow conditions in Article 5(2) to (6)).

The Commission's Guidelines on prohibited AI practices, approved in February 2025, run well over a hundred pages – the most detailed piece of guidance in the whole structure. Four features are worth knowing, even if the substance is already familiar to you.

- 1. They set out the general concepts everything else rests on:** “placing on the market”, “putting into service”, and “use”, and point to the Commission's Blue Guide on how EU product rules work (the same source the Article 50 guidelines rely on). The point: it's product-safety concepts, not AI-specific ones, that set the boundaries of what's covered.
- 2. They assign responsibility by role:** Providers and deployers have different responsibilities depending on their role and how much control they have over a system's design, development and actual use. For each prohibition those roles are read proportionately – asking who along the value chain is best placed to take specific preventive and mitigating steps. That principle is the ancestor of the “calibration” paragraphs the Omnibus has now written into Article 5 itself, and it is the frame within which the new safeguards duty should be read.
- 3. They map what Article 2 leaves out:** national security, military and defence, judicial and law enforcement cooperation with non-EU countries, research and development, and purely personal non-professional use. The Article 50 guidelines cross-refer to this for this same purpose, which is a good example of the different guidance documents lining up with each other.

4. The gap to watch. These guidelines came before the Omnibus, so they say nothing about the two new prohibitions in Article 5(1)(ba) and (bb) (discussed below). Right now, there is no Commission guidance on these prohibitions, on the calibration paragraphs, or on what constitutes “reasonable and adequate technical safety measures” – and the prohibitions take effect on 2 December 2026.

Recital 12 of the Omnibus partly fills the only gap. It gives a non-exhaustive list of measures that count, and it touches on what makes measures adequate, on foreseeable ways people might get around them, and on how you could release a system. It’s the best available anchor until real guidance arrives. A revision or expanded version of the prohibited-practices guidelines is the single most valuable piece missing from the pipeline. Until it lands, you have to build your safeguards straight from the statutory text and reason by analogy from the approach the existing guidelines already use. So assign someone to watch for it, and design your safeguards so that they can absorb the guidance without being rebuilt. (See the below for more details about the safeguards stack).

14.2 The two new prohibitions

Article 5(1)(ba): non-consensual intimate material. From 2 December 2026, it is prohibited to place on the market, put into service, or use an AI system that generates or manipulates realistic images, videos, audio or similar material of an identifiable natural person's intimate parts, or showing an identifiable natural person engaged in sexually explicit activities, without that person's consent – and that consent has to be freely given, specific, informed, unambiguous, explicit, and given for that purpose. The clear target is “nudification” apps, but the prohibition is written around what a system can do, not around a particular product.

Four elements to note:

1. It has to be realistic. This pulls in reasoning close to the deepfake analysis in the Article 50 guidelines.
2. The person must be identifiable.

3. The consent standard is very high. It's the full GDPR-style consent standard, plus explicit and plus tied to this specific purpose - which in practice means running such a system on a consent basis is only workable in narrow, well-documented commercial settings.
4. "Similar material" keeps the list of content types open - unlike Article 50(2), which uses a fixed, closed list.

Article 5(1)(bb): child sexual abuse material. A matching prohibition covers AI systems that generate or manipulate the material or performances defined in Article 2, points (c) and (e) of Directive 2011/93/EU - except where a "without right" defence applies under national law.

That reference to national "without right" defences is what keeps lawful work lawful: law enforcement, judicial proceedings and - importantly for model developers - red-teaming and safety evaluation conducted under national legal authority. But if you rely on that route, you have to be able to point to the specific national legal basis for it. A general "this is research" rationale is not enough.

Article 5(1a) and (1b): the calibration paragraphs. These are the paragraphs that actually drive how you design compliance, and they apply a different test to providers and deployers.

For providers, placing on the market or putting into service is prohibited only if either:

1. generating or manipulating is the intended purpose of the AI system, or
2. the system's design, training, architecture, capabilities or user-facing functionalities makes generation or manipulation a reasonably foreseeable and reproducible outcome - and the system lacks reasonable and adequate technical safety measures.

Read the second limb carefully. It has three parts, each is something you can act on.

- **Reasonably foreseeable** is an objective standard, assessed against those listed features of the system. You can't argue you didn't anticipate what your architecture plainly allows.
- **Reproducible** filters out the trivial. An isolated jailbreak that can't be reliably repeated is not the same as a capability that a competent user can trigger at will.
- **Lacks reasonable and adequate technical safety measures** is the part that lets you off the hook. The prohibition does not attach to a general-purpose generative system just because misuse is possible; it applies where misuse is foreseeable and reproducible and the safeguards are inadequate.

So what you produce to show compliance is the safeguards themselves, and the evidence that they exist and work.

For deployers, use is prohibited only where the deployer purposefully generates or manipulates such material. Accidental generation is not caught. Deliberate misuse is caught no matter what tool is used, including where the deployer gets around the safeguards on an otherwise lawful system.

One more clarification in these paragraphs matters for image-editing and photography products: simply enhancing an image, without revealing more of the body is not “manipulation” for this purpose. That protects ordinary retouching – but the line is narrow, so document your work against the list of standard-editing operations in the Article 50 guidelines, which is the closest comparison available.

14.3 Building the safeguards stack

Because the prohibition depends on whether your technical safety measures are adequate, the thing a generative model or system provider has to produce is a “safeguards stack” – a set of safety measures that is written down, tested, version-controlled, and tailored to how you release the system.

There is no Commission guidance yet, so the defensible approach is a stack that covers the system's whole lifecycle, with evidence at each layer.

Data layer. Clean and filter your training data. That means: a written method for detecting the material; hash-matching against known CSAM databases through the proper national and international channels; and removing intimate-imagery datasets and known nudification sources. Your evidence here is the dataset documentation, the filter settings, and a record kept of the filtering (the audit trail).

Watch the overlap with Article 4a and the GDPR: the processing you do to detect and remove this material itself involves special-category data considerations, so it needs its own legal analysis.

Model layer. Train the model to refuse the prohibited categories, do adversarial testing and red-teaming (and keep the results), and check whether refusals actually hold up when someone realistic tries to get around them. For systemic-risk models, this work is already required under the Article 55 obligations and the GPAI Code's safety and security chapter - so document it once and reference it in both places, rather than doing it twice.

System layer. Add controls around the model:

- classify inputs and filter prompts;
- classify and block outputs;
- restrict the features that make the prohibited result reproducible - image-to-image generation, inpainting (editing on part of an image) on human subjects, and features that key off of a person's identity; and
- limit rate and watch for suspicious patterns.

Operational layer. Run it responsibly once it's live:

- detect and monitor fraud;
- enforce at the account-level;
- provide reporting channels with set response times;

- escalate incidents - through the Article 73 route when a serious incident arises, and to the AI Office under Article 75(1a) if the system is centrally supervised; and
- re-test regularly as the state of the art moves on - mirroring the ongoing-adaptation duty the Article 50 guidelines set for marking.

Tailoring to how you release, including open source. This is where the analysis is hardest and where the Act offers the least reassurance. Article 2(12) does not let open-source systems off Article 5, and the Article 50 guidelines take the same line in their own domain.

Once you release the weights openly, you lose control of the system-layer and operational-layer safeguards - anyone can strip them out. So you have to conduct the “foreseeability and reproducibility analysis” on the assumption that those downstream guardrails will be removed. That means, if you release openly, expect to have to show that your data-layer and model-layer safeguards are strong enough on their own. And document it at the time: the decision to release, the capability testing behind it, and the mitigations you chose.

No guidance does not mean no risk, and the fines here are top tier: 35 million euros or 7 percent.

One connection to take in. Marking and labelling don’t make something legal. The Article 50 guidelines say this plainly: if a labelled deep fake is flagged as possible CSAM or as infringing trademark or copyright, the label makes no difference to whether it’s illegal under those frameworks. And the duty to be transparent does not mean harmful and unlawful deep fakes (the guidelines name CSAM and non-consensual intimate images) are allowed to be generated and shared. Put simply: a fully Article 50-compliant nudification product is a prohibited product.

Part V •

Risk classification

Part V • Section 15

High-risk classification rules under the draft guidelines on Article 6

This is the most important draft document currently in play, and its direction of travel is to cover more systems, not fewer.

Status. The Commission published draft guidelines on the classification of high-risk AI systems under Article 6 on 19 May 2026 – later than the statutory deadline of 2 February 2026 in Article 6(5), to allow incorporation of earlier stakeholder input. They come in three documents, following the structure of Article 6: general principles for classification, the Article 6(1) product route, and the Article 6(2) Annex III route. A focused consultation ran until 23 July 2026 (extended by four weeks from the original 23 June date), and has now closed. Final guidelines are expected by the end of 2026. The same content is also on the AI Act Single Information Platform, with summaries, examples and a searchable explorer organised by area and use case, which is the practical way to work through it.

The guidelines are non-binding guidance; only the EU's top court, the CJEU, can interpret the Act authoritatively. But they're the benchmark regulators will actually apply, and the examples (expressly not a complete list and set to be updated over time) are how classification arguments will be won and lost.

Why does this matter more than the delay? Before anything else, you have to answer one question about each system: is it high risk or not? That single yes-or-no question switches the entire high-risk rulebook on or off – all the core requirements, the conformity assessment route, CE marking, registration, post-market monitoring, and the delayed 2027 and 2028 deadlines. (In Article terms, that's Articles 9 to 15, 16 to 25, 26 and 27.) If a system isn't high risk, none of it applies; if it is, all of it does.

Here's the catch: the Omnibus moved the line between "high-risk" and "not" in both directions at once.

- **The law itself narrowed the line.** Article 6(1a) to (1c) and the amended definition of "safety component" (Article 3(14)) keep AI that only does non-safety jobs out of the high-risk category. So some systems that looked high-risk before may not be now.
- **The draft guidelines widen it.** The Commission's draft Article 6 guidelines read the category broadly (see Section 15 of this playbook). So some systems you'd assume are safe might actually be caught.

Because the line moved both ways, you can't assume last year's classification still holds. Any business that hasn't re-checked its systems against both the narrowing law and the widening guidance doesn't know what it owes.

The substantive positions to plan against.

- 1. The third-party-assessment limb is read broadly.** Under Article 6(1)(b), the product route to high risk (or the AI system sold as a product) has to go through a third-party conformity assessment under the Annex I laws. The draft reads this as covering products subject to enhanced regulatory scrutiny, which captures more product-embedded AI than the many manufacturers who assumed they could self-assess (also called internal control) expected. Combined with the Omnibus's Article 6(1c) (which removes products requiring third-party assessment only for non-health-and-safety risks), the overall scope for product-embedded AI is narrower in one specific way, but wider in general. If you concluded you were outside Article 6(1) because you could self-assess, run the analysis again.
- 2. Modular and agentic systems are assessed as a whole.** Where separate components combine to materially shape a high-risk decision, they are treated as one system, and a component cannot rely on the Article 6(3) filter on its own. This is the most important position for anyone building agentic systems or assembling systems from services; what gets classified is the functional whole, not the individual module you happen to ship. It matches the Article 50 guidelines' holistic treatment of multi-agent setups, and the fact that agentic AI now has its own designation code (AIH 0401) in the new Annex XIV. The Act's various documents are converging on one view: splitting a system into pieces doesn't split the responsibility.

3. Multi-purpose systems aren't automatically high-risk, but configuring them can be.

A highly capable general-purpose assistant isn't high-risk just because it could be used for an Annex III task. But if you configure or fine-tune it to actually do an Annex III use case, you (the configuring party) usually become the provider of a high-risk system, and the full Chapter III rulebook is attached. For deployers, this is the most likely way to end up as a provider, and it is triggered by everyday business activity: building a screening workflow on a general assistant, tuning a model for credit decisioning, or wiring an assistant into an eligibility determination.

4. Putting a human in the loop doesn't change the classification.

Human oversight is a duty under Article 14 – it's not a way out of being high-risk. The employment materials make the point sharply: a tool can be high-risk where it materially influences employment decisions, even if a person makes the final call. Any governance design that relies on human sign-off as a classification defence doesn't work.

5. The intended purpose is the control point.

Classification turns on the intended purpose the provider gives the system, and the draft requires that purpose to be consistent across the instructions for use, the marketing, and the technical documentation. Disclaimers won't override how a system is realistically used. Practically, this makes your marketing copy and solution-engineering collateral are classification evidence, so they have to be governed as such. It also means the defensive exclusion of downstream high-risk repurposes (discussed in Section 5 of this playbook) must be honoured in practice, not just recited in your terms.

6. The Article 6(3) filter is narrower than people expected.

The filter lets an Annex III system escape high-risk where it doesn't pose a significant risk of harm to health, safety or fundamental rights, including by not materially influencing decision-making. Helpfully, the four conditions are alternative rather than cumulative, so meeting any one is enough. But the draft reads them restrictively, and three limits bite:

- the filter is never available if the system profiles people;
- relying on it carries its own documentation duty under Article 6(4), and you must product that assessment on request; and
- you still have to register in the EU database (though the Omnibus trimmed this by deleting points 7 and 9 of Annex VIII Section B).

So a filtered system isn't an unregulated one; it is a documented, registered exemption.

7. The Annex III use cases come with worked examples, and the boundaries are specific. Take employment: point 4(a) covers recruitment and selection; point 4(b) covers managing the working relationship once a person is in a job or seeking to stay in one – with examples like allocating work, determining pay, assessing performance, supporting promotion decisions, monitoring behaviour, and affecting whether someone is let go or kept on. The draft also gives examples that fall outside: a general employer branding tool may be outside point 4(a) because it isn't tied to a specific recruitment process. The same worked example treatment runs across biometrics, education, credit and essential services, critical infrastructure, law enforcement, migration, and justice. And (as noted in Section 12.4 of this playbook), the definitions of emotion recognition and biometric categorisation systems developed here are cross-referenced by the Article 50 guidelines and apply for transparency purposes too.

8. A dedicated Article 25 value-chain guidelines instrument is promised for the future. Since the Omnibus made Article 25(2) and (4) separately fineable, that document will matter. Because it isn't here yet, draft your supplier agreements so they can be amended later, rather than locking in fixed schedules now.

What to do now that the consultation has closed.

The 23 July window is gone, but three things are still open to you:

- Re-check all of your systems against the draft now, in writing. The draft is the best available prediction of the final text, so any change between draft and final will just refine a position you've already taken, not force a fresh one.
- Document every time you rely on the Article 6(3) filter, to the standard the draft sets, and register accordingly.
- Use the channels that are still open. These include the national competent authorities' guidance functions (which the Omnibus strengthened in Article 70(8)), the AI Act Service Desk, industry associations and the standards process, where the technical detail behind classification arguments is still being written.

Part VI | •

The rest of the living instrument

The amended AI Act, through the AI Omnibus, is only the frame. The detail that actually governs you arrives later – through codes of practice, guidelines, delegated acts, templates, and standards. And each of these needs a named person inside your organisation who is responsible for tracking it and putting it into practice.

Part VI | • Section 16

The General-Purpose AI (GPAI) Code of Practice

Published 10 July 2025; approved as adequate by the Commission and the AI Board on 1 August 2025.

It has three chapters:

- **transparency** and **copyright** apply to all model providers; and
- **safety and security** to systemic-risk models

Open-source models are largely spared the heavier documentation unless they pose systemic risk, but they're still bound by the copyright chapter.

Following the code shows compliance, but there is no **presumption of conformity** - only cited harmonised standards can do that, and none exist yet for GPAI. So throughout the delay, the Code is the practical way to show you comply. The Omnibus removed the one-time formal signoff from Article 56(6), which turns “adequate” into an ongoing, published assessment - in effect, a rolling bulletin of enforcement expectations.

Two dates matter:

2 August 2026:

the AI Office's good-faith posture toward signatories still working towards compliance ended, and this is also when Article 101 switches on the Commission's power to fine.

2 August 2027

models placed on the market before August 2025 must be compliant by this date (Article 111(3)).

Part VI | • Section 17

The European Commission guidelines inventory

Already issued guidance (six, all still evolving)

In terms of already existing guidance issued from the European Commission, there are 6 pieces evolving:

- the AI system definition guidelines ([C\(2025\) 5053 final](#));
- the prohibited practices guidelines ([C\(2025\) 5052 final](#)) - which are silent on the two new prohibitions;
- the [scope of obligations for GPAI models guidelines](#);
- the draft Article 73 serious incident guidance of 26 September 2025 - setting reporting windows of two to fifteen days, treating an indirect causal link between system and harm as enough, and mapping the overlap with DORA and the medical devices regulation;
- the final Article 50 transparency guidelines (see Part III of this playbook); and
- the draft Article 6 high risk classification guidelines, with the final version due at the end of 2026 (see Part V of this playbook).

Five more are anticipated, two of them significant:

- Guidelines on Articles 8(2), 9(10) and 17(3), due by 1 August 2027 - these will determine whether product manufacturers face one compliance story or two.
- Joint Commission and EDPB guidelines on the AI Act and data protection - being prepared for the end of 2026. They'll cover the legal basis for using training data, Article 4a (use of special category data for bias work) in practice, and how to integrate the data protection impact assessment (DPIA) and the fundamental rights impact assessment (FRIA).

Watch for the upcoming consultations on these.

Also pending: research-exemption guidance and Article 25 value-chain guidelines, among others.

Part VI | • Section 18

Delegated acts, implementing acts and templates

These are the detailed follow-on rules and forms the Commission writes to fill in the Act -

delegated and implementing acts carry binding legal force, while templates are standard forms you use to comply. Most of the concrete detail that will actually affect you lands here, not in the amended articles, so each item below is worth tracking.

Delegated acts:

- The **Article 2(13) sectoral-equivalence** delegated acts (due by 2 August 2027) may limit specific Articles 9 to 15 and 17 to 25 requirements for Annex I Section A systems. They're not guaranteed until adopted, so treat them as possible upside, not something to count on.
- The **Machinery Regulation delegated** acts apply by 2 August 2028 and will set out the entire compliance content for AI in machinery.

Templates: Three templates matter:

- the **simplified technical documentation form** that notified bodies must accept from SMEs, start-ups and SMCs;
- the **fundamental rights impact assessment (FRIA) questionnaire and automated tool**, which must support DPIA cross-referencing, and which you should use rather than build your own; and
- the **post-market monitoring guidance with a voluntary template** by 2 September 2027.

Implementing acts are also due on sandbox governance and on AI Office procedure under (Article 75d) - the latter matters directly for how ready you need to be for an inspection.

Part VI | • Section 19

European harmonised standards

This is the only layer that reverses the burden of proof.

Under Article 40, a provider that follows a standard **cited in the Official Journal** has a presumption of conformity. A regulator that wants to challenge that has to show the standard was inadequate or wasn't applied properly - the provider doesn't have to prove conformity from scratch.

The standardisation request has slipped repeatedly. The original request (M/593, Implementing Decision C(2023) 3215) was cancelled and replaced by C(2025) 3871 final of 23 June 2025 (M/613), which reset the deadline to 31 August 2025 - and that date was also missed.

The first standard is here on quality management - but not yet “live”. EN 18286:2026, the quality management system (QMS) standard for EU AI Act regulatory purposes, was approved on 12 July 2026. It's the first standard home-grown by the European committee (JTC 21) specifically for the Act. This one matters more than it sounds: because the QMS is the spine you use to evidence every Chapter III requirement, meeting this standard effectively means complying with the whole Act. Its annexes map to ISO 9001 and ISO/IEC 42001 (Annex A), but one way only, because 42001 certification does not demonstrate Article 17 conformity.

The gate has not opened. Approval makes EN 18286 a European harmonised standard, but it gives you no presumption of compliance until the Commission cites it in the Official Journal of the EU.

Other standards in process: Behind EN 18286 on QMS, still in draft sit:

- prEN 18229-1 - logging, transparency, human oversight;
- prEN 18228 - risk management;
- prEN 18282 - cybersecurity and
- a conformity assessment framework that isn't expected to carry the presumption.

The new Article 40(2) requires the Commission to ask for joint deliverables that let you comply with the AI Act and Annex I product legislation at the same time - this is the standards-side equivalent of the Article 2(13) equivalence mechanism, and the most likely source of real burden reduction. If the standards pipeline fails, Article 41 lets the Commission impose its own “common specifications”, replacing an industry-consensus text with a Commission-drafted one.

Bottom line:

Adopt EN 18286 now, so the presumption attaches the day the citation lands, and track the Official Journal citations, not press releases.

Part VI | • Section 20

Institutional build-out to monitor

Keep an eye on:

1

The **AI Office**. It holds the Articles 75a to 75d powers from 27 July 2026 and has sole authority under the amended Article 75(1). But its funding rule (Article 64(3)) only promises “adequate resources” subject to the normal budget process – there is no guaranteed money. This points to concentrated enforcement (a few high-profile cases) rather than broad enforcement.

2

The **EU sandbox**. The EU-level sandbox gives priority access to SMEs, start-ups and SMCs. **National sandboxes** are due by 2 August 2027 and can now take in real-world testing plans.

3

The **Article 43(3) notified body bridge**. Under Article 43(3), sectoral notified bodies must apply for AI Act approval by 28 January 2028. That flow of applications (tracked by Annex XIV code) is the leading indicator of whether there will be enough assessment capacity.

4

The **published code adequacy assessments**. These are the cheapest way to learn what the AI Office actually considers sufficient.

Part VII •

Using the playbook

Part VII • Section 21

Action checklist by horizon

Specific sections of this playbook are listed for more details related to each action.

Immediate action (from 2 August):

- ☐ Inventory every interactive and generative system against Article 50(1) to (4) and record your reasoning, including why anything is excluded. (Section 12 of this playbook)
- ☐ Add disclosure at the point of interaction; avoid the five approaches that aren't enough on their own; and double check anywhere you're relying on the obviousness exception. (Section 12 of this playbook)
- ☐ For AI agents, disclose both that the agent is artificial and who it acts for. (Section 15 of this playbook)
- ☐ Put marking in place for AI generated content, together with a working way to detect the mark. (Section 12 of this playbook)
- ☐ Notify people wherever emotion recognition or biometric categorisation is running. (Section 12 of this playbook)
- ☐ Label deep fakes and public-interest text, use the EU icon set, and confirm the labels survive when the content is syndicated. (Sections 12 and 13 of this playbook)
- ☐ Put the editorial review process in place, and check your backlog of already-published material. (Sections 12 and 13 of this playbook)
- ☐ Decide whether to sign the Transparency Code - and handle consumer-law pre-contract disclosure as a separate task. (Section 13 of this playbook)
- ☐ **GPAI providers:** 2 August is when enforcement switches on. (Section 16 of this playbook)

2 August to 2 December 2026:

- ☐ Bring legacy generative systems into line with Article 50(2). (Section 12 and 13 of this playbook)
- ☐ Build your Article 5 safeguards stack – across the data, model, system and operational layers, tailored to how you release (including open releases), and update your acceptable-use and monitoring frameworks to match. (Section 14 of this playbook)
- ☐ Re-classify your whole portfolio against the draft Article 6 high risk classification guidelines and the amended Article 3(14) and Article 6(1a) to (1c) – documenting and registering every time you rely on the Article 6(3) filter. (Section 15 of this playbook)
- ☐ Lock in a design-change process keyed to the Article 11(2) grandfathering trigger and start your Article 25 contracting programme. (Section 15 of this playbook)

2027 and beyond:

- ☐ Get your value-chain agreements under Article 25(4) – with hand-over clauses, or the defensive exclusion written into your own terms.
- ☐ Build the Chapter III backbone: risk management, data governance under the reworked Articles 10 and 4a, technical documentation, and quality management system (QMS) scaled to your size (anchored on EN 18286, with owners and evidence assigned to each part), and set up your DPIA process so it feeds right into the FRIA.
- ☐ **Key deadlines:** legacy GPAI models comply by 2 August 2027; Annex III systems by 2 December 2027; Annex I systems by 2 August 2028; and public-sector legacy systems by 2 August 2030.

Part VII • Section 22

The judgments that matter

1

The delay is preparation time, not a pause. The obligations haven't changed, and the ecosystem of standards, bodies and guidance is being built to be ready by the new dates. Businesses who use 2026 and 2027 to lock down classification, contracts, and architecture will take in the follow-on rules as refinements. Those who wait will have to rebuild under them.

2

A large part of the Act is already in force and untouched. The original prohibitions have applied since February 2025 and the GPAI rules since August 2025. Neither was changed in substance. Any story that the rules were "rolled back" is simply wrong where it reaches these parts, and your compliance programmes here have to keep running.

3

2026 is a transparency and prohibitions year, not a quiet one. Two firm deadlines, 2 August and 2 December, fall inside the supposed relief period. This first carries the transparency-level fines ; the second carries the highest fines of all.

4

Classification is the gateway decision, and the Commission reads it broadly. The safety-component clarifications narrow the scope in some places, while the enhanced-scrutiny and holistic-agentic readings widen it in others. Re-classify now, in writing.

5

Soft law is the real rulebook, and standards are its hardest edge. Two codes of practice (each with a published adequacy scorecard), two sets of guidelines, a pipeline of delegated-acts, and the harmonised-standards catalogue – with EN 18286 on QMS first among them – will carry more day-to-day compliance detail than the amended articles do. Standards are the only layer that reverses the burden of proof, which is why the Official Journal citation calendar deserves the same attention as the legislative one. Give each instrument a named owner.

6

Enforcement has a new centre of gravity. For centrally supervised businesses, the AI Office is now a regulator with powers to enter premises, seal them, levy daily-penalties, and publish its decisions. For everyone else, the referral channels and the stronger access given to fundamental-rights authorities mean your exposure at the national level reaches further than before.

Taken together rather than as a list, the six judgments are really about one thing: sequencing. The delay is time allocated for a purpose. The original prohibitions and the GPAI rules are obligations that arrived before the ones the market is watching. The transparency duties and the two new prohibitions fall inside a stretch widely described as quiet. Classification has to be settled before the guidelines governing it are final, because the alternative is settling it afterwards, under pressure, against a fixed deadline. The standards layer will arrive late and matter most. And the enforcement machinery switched on sixteen months before the bulk of what it exists to police. This is the point the playbook began with – that most commentary fixated on the postponement, when the sequence mattered more – now made six times, with the evidence attached.

The practical upshot: what the Act asks of an organisation between now and 2028 is not a project with a finish line. Every one of the six judgments, once you act on it, produces something you have to maintain, not something you deliver.

A classification position holds only as long as it matches the current reading of Article 6. A safeguards stack is adequate only against the state of the art the last time you tested it. Signing a code of practice is an ongoing commitment to something the Commission now re-assesses and re-publishes continuously. The deliverables in Section 21 aren't an end state. They are the first version of a file you'll rewrite several times before any regulator ever reads it.

That is also the limit of what a playbook can do. What's here is a snapshot reading of a law that's still moving – taken in August 2026, when four of the five layers are still incomplete, the Article 6 guidelines are in draft, no AI Act standard has been cited yet, and the two new prohibitions have no guidance at all. The overall structure will hold. But several of the specific positions won't survive once the final texts land, and the passages most at risk are flagged as such throughout. No document can solve that, because the difficulty isn't that the analysis might be wrong. It's that the analysis dates, and the organisation relying on it needs to know when.

Where does Credo AI fit?

The sections above describe two things that are true at the same time, and it's the combination, not either one alone, that creates the governance problem.

First: in the EU AI Act, you comply by producing a record, not by understanding the rules. A classification written down. A registered reliance on the Article 6(3) filter. A documented and tested safeguards stack. An open-weights decision recorded at the time you made it – each one tied to who owns it, the procedure behind it, and where the evidence lives. A regulator arriving in 2028 reads that file. It doesn't assess what you intended, and it gives no credit for an analysis you did but never wrote down.

Second: the four layers beneath the Act keep moving the whole time you're building that file. So the file that satisfies a regulator in December 2027 is not the file that would have satisfied one in August 2026, and the gap between them isn't something you close just by adding pages.

So the real question isn't whether your organisation understands the Act. It's whether your file moves when the layer beneath the Act moves, and whether you can show when it moved, on what basis, and on whose authority.

Credo AI fits at four points in this document, and it's worth being precise about which.

- 1. Classification:** the gateway decision (judgment 4) that everything else hangs on. The call has to be made for each system, recorded with its reasoning, and re-run when the final Article 6 guidelines land at year end and – and again against the amended Articles 3(14) and 6(1a) to (1c). Doing that once across a portfolio is a project. Doing it every time the reading shifts, with each previous position and its rationale still on hand, is a different kind of capability – and it's the one December 2027 actually calls for.
- 2. The moving-target challenge:** the “living instrument” (judgment 5), now stretched across the whole structure. Final guidelines, revised codes, and newly-cited standards change what you have to prove without changing a single article of the Act. An organisation that takes in that movement in one central place is in a very different position from one that re-reads the whole landscape team by team every time a new instrument lands. This is where the distinction between a programme built to the structure and one built to a single reading at a single moment stops being talked and becomes real.

3. Inventory: where the architecture is blunt – an inventory you haven’t risk-classified in writing isn’t a compliance programme. The draft Article 6 guidelines sharpen this, because they assess modular and agentic systems as one whole. That turns knowing exactly which individual agents, models, and vendors you have from a nice-to-have into something the classification itself depends on.

4. Evidence held against fixed dates. The Article 50 checks now underway, the safeguards stack due on 2 December, and the Chapter III backbone due in December 2027 all come down to the same thing: evidence tied to named owners, in known locations, in a form that holds up when a regulator reads it cold and without context.

Two limits are worth stating plainly, because a playbook that oversells its own tools won’t survive the first board meeting where it’s put to the test.

First, none of this makes the judgment calls for you. Whether to rely on the Article 6(3) filter, whether to sign the codes of practice, and how to set safeguards for an open-weights release – these stay legal calls made by named people who may be asked to justify them. A platform keeps those positions consistent across your whole portfolio, records them at the time they were made, and lets you re-run them when the basis changes. It doesn’t make them for you, and any claim that it won’t survive contact the first lawyer who looks at it.

Second, no tool shortens the queue for a notified body. Assessment capacity is the real bottleneck for December 2027. The leading indicator of this capacity is how many bodies apply for approval before the bridge expires on 28 January 2028 – and that’s a conversation you have to start early, through a named responsible person.

The bottom line

The Omnibus bought time, not relief. And that time only helps the organisations that spend it building. Classify in writing, paper the value chain, evidence your safeguards, and give every moving instrument a named owner, so the final texts land as refinements rather than a rebuild. None of that is a one-off project; it’s a file that has to move as the Act moves beneath it. That is the problem Credo AI is built to hold – and the moment to start is now, while the layers are still forming and the runway is still yours. **Credo AI is where that record is built, kept current, and reviewed again each time the rules underneath it change.**

[Talk to us](#)

-

About the Authors

About the Authors



Vassilis Rovilos

EU Policy Director

Credo AI

Vasileios Rovilos is EU Policy Director at Credo AI, leading its EU regulatory strategy and AI Act implementation. He previously worked at the Future of Privacy Forum, Microsoft and the European Commission (DG JUST and DG GROW). He is a designated expert member of the Commission's AI Advisory Forum and of UNESCO's AI Ethics Experts Without Borders Group. He taught the EU AI Act segment of UC Law San Francisco's LexLab Law & AI certificate programme. He is a Globethics Emerging Leaders in Ethical AI Governance Fellow and a GLOBSEC Young Leader. He also conducts academic research on AI and is a member of the Advanced Research Group on Law and AI at Universidade Católica Portuguesa. He holds degrees from Maastricht, Leiden and the College of Europe, with executive education at Stanford and Oxford, and is a PhD candidate at the NKUA Law School (Athens, Greece).



[View on LinkedIn](#)

Reviewers



Akanksha Ray

Global Policy Director

Credo AI

Akanksha Ray is Director of Global Policy at Credo AI. She was previously a senior manager for Digital Trust and AI in Intel’s Government, Policy and Trade group, where she managed global programs that upskill governments, the workforce, and citizens on technology issues such as AI governance, digital trust, supply chain resiliency, and more. She also sits on Intel’s Responsible AI Council, conducting ethical impact assessments and driving corporate initiatives for AI accessibility, sustainability, and digital inclusion. Outside of Intel, Akanksha is a team lead for a research group at the Center for AI and Digital Policy advancing global AI governance that promote human rights. She also sits on the Board of Governors of Lakewood University—a non-profit, distance education school in Cleveland Heights, Ohio. She holds a bachelor’s degree in public policy and economics from Duke University with a focus on foreign policy and international development.



[View on LinkedIn](#)

Reviewers



Linsey Krolik

Senior Director, AI Governance and Product Legal

Credo AI

Linsey Krolik is Product, Privacy and Regulatory Counsel at Credo AI, an AI governance technology company, and an Assistant Clinical Professor at Santa Clara University School of Law. Linsey teaches and practices in the areas of privacy, AI governance, startup, and technology law.



[View on LinkedIn](#)

Contributors



Design and Layout

Lucy van der Meer

Brand and Design Manager

Credo AI

Lucy van der Meer leads brand and design at Credo AI, where she helps translate complex ideas around trusted AI Governance into visual experiences that leave their impact in the hearts of people.

Her background spans traditional arts, digital design, and everything in between. But what's always driven her isn't the medium — it's the mark. The moment a brand stops being something people see and starts being something they remember. Lucy works at the intersection of human psychology and design strategy, building brands that don't just look good, but mean something to communities all around the world.

Before Credo AI, Lucy launched and relaunched brands internationally, working across user research, web design, and brand identity to create experiences that stuck around long after the first impression. Like graffiti on a wall, weathered, real, and impossible to ignore.



[View on LinkedIn](#)

Contributors



Content Development

Jerome J. Sanders

Senior Content Marketing Manager

Credo AI

Jerome leads content marketing at Credo AI, where he helps Fortune 500 organizations build trusted AI Governance through data-driven strategic storytelling and thought leadership. With over a decade of experience in content strategy, product marketing, and product development at Cisco, he specializes in translating complex AI concepts into narratives that drive business outcomes, clarity, and industry influence for senior executives.

A 2024 Aspen Institute First Movers Fellow, United Nations Global Compact SDG Innovation Accelerator and Board Chair at Honor Your Dreams, Jerome focuses on expanding equitable access to foundational AI technologies. He holds a business degree from Santa Clara University's Leavey School of Business.



[View on LinkedIn](#)

Regulation (EU) 2024/1689 (aka the EU AI Act)
as amended by Regulation (EU) 2026/1744
(aka the AI Omnibus)

Sources: Regulation (EU) 2026/1744 of 8 July 2026 (OJ L, 24.7.2026); Regulation (EU) 2024/1689 as amended. ELL: data.europa.eu/eli/reg/2026/1744/oj.

Date	Milestone	What it covers	Who it affects	Status	Reference
2 February 2025	Chapters I and II apply	General provisions and the original Article 5 prohibitions. Eight prohibited practices including social scoring, untargeted facial scraping, emotion inference at work and in education, and real-time remote biometric identification. Also the Article 4 AI literacy duty, since recast by the Omnibus as an obligation of means.	Everyone placing on the market, putting into service or using AI systems in the Union	In application	Art. 113(3)(a)
2 August 2025	GPAI model regime and governance apply	Articles 51 to 55 for general-purpose AI model providers. Also Chapter III Section 4 on notifying authorities and notified bodies, Chapter VII governance, Chapter XII penalties other than Article 101, and Article 78 confidentiality. The Omnibus left Articles 51 to 55 untouched, but did amend other provisions in this group (Articles 28, 29 and 30, Article 56(6), the new Article 64(3) and Article 99) which take effect in amended form on 27 July 2026.	GPAI model providers, Member States, notified bodies, and enforcement structures	In application	Chapter V; Art. 113(3)(b)
27 July 2026	AI Omnibus enters into force	Third day after publication. Articles 102 to 110 of the AI Act, as amended, apply from this date. The five-year clock on the new delegated powers in Articles 2(13) and 30(2) starts running, as does the 18-month notified-body bridge under Article 43(3).	AI Office, national authorities, product sectors affected by the Annex I reshuffle	In application	Art. 4 Omnibus; Art. 113(3)(d)
27 July 2026	Article 4a bias legal basis available	Narrow legal basis for processing special categories of personal data for bias detection and correction, subject to six cumulative safeguards. Extended by the Omnibus to deployers of high-risk systems and to providers and deployers of other AI systems and models. Article 4a sits in Chapter I, which has applied since 2 February 2025, so the provision takes effect on entry into force rather than at the general application date. Recital 9 confirms it is deliberately available ahead of the Chapter III deadlines, so that providers can carry out bias work in preparation for Article 10(2)(f) and (g). It is a permission, not a duty.	Providers and deployers conducting bias work	New legal basis	New Art. 4a
2 August 2026	General application of the AI Act	The date the Omnibus did not move. Everything not expressly deferred or already applying starts here.	All operators	Unchanged	Art. 113, first para
2 August 2026	Article 50 transparency duties apply in full	Interaction disclosure under 50(1), marking and detection under 50(2) for systems newly placed on the market, emotion recognition and biometric categorisation notice under 50(3), deep fake and public-interest text labelling under 50(4), and the horizontal requirements in 50(5). Not deferred. The Article 111(2) grandfathering for legacy high-risk systems does not reach Article 50.	Providers and deployers of interactive and generative systems; publishers and broadcasters	Unchanged	Art. 50
2 August 2026	Chapter III Section 5 applies	Harmonised standards and common specifications, conformity assessment under Article 43, the EU declaration of conformity, CE marking and registration in the EU database under Articles 49 and 71. Only Sections 1 to 3 of Chapter III were deferred; this Section was not.	Providers of high-risk systems, notified bodies	Unchanged	Arts. 40 to 49
2 August 2026	Chapter IX market surveillance switches on	Post-market monitoring under Article 72, the Article 73 serious incident regime, the Article 85 complaint right, and the AI Office's new supervisory and enforcement powers under Articles 75a to 75d: entry and inspection, sealing of premises, binding commitments, fines, and periodic penalty payments of up to 5 per cent of average daily income or of worldwide annual turnover in the preceding financial year, per day.	All operators; centrally supervised providers in particular	Unchanged	Ch. IX; new Arts. 75a-75d
2 August 2026	Article 101 fining power becomes available	The Commission can fine general-purpose AI model providers up to 3 per cent of worldwide annual turnover or 15 million euros, whichever is higher. Always scheduled for this date and untouched by the Omnibus. The AI Office's first-year good-faith posture toward code signatories was tied to the period ending here.	GPAI model providers	Unchanged	Art. 101; Art. 113
2 August 2026	Member States notify penalty rules	Rules on penalties and other enforcement measures communicated to the Commission, without delay and at the latest by the date of entry into application. The amended Article 99(1) now requires Member States to weigh the economic viability of SMEs and small mid-caps, and permits warnings and non-monetary measures.	Member States	In application	Art. 99(2)
2 December 2026	Prohibitions on NCII and CSAM generators apply	Providers are caught where prohibited generation is the intended purpose, or where it is a reasonably foreseeable and reproducible outcome and the system lacks reasonable and adequate technical safeguards. Deployers are caught only where they use the system for that purpose. Top fines tier: 35 million euros or 7 per cent. No Commission guidance exists yet.	All providers and deployers of generative systems, including open-weight releases	New obligation	Art. 113, third para., point (a); Art. 5(1)(ba), (bb); 5(1a), (1b)
2 December 2026	Legacy generative systems meet marking rules	Four-month transitional window for providers of systems generating synthetic audio, image, video or text placed on the market before 2 August 2026 to comply with Article 50(2). Note the text says placed on the market, not put into service.	Providers of pre-existing generative systems	New obligation	New Art. 111(4)
1 August 2027	Guidelines on minimising duplicate compliance	Commission guidance on the practical implementation of Articles 8(2), 9(10) and 17(3): folding AI-specific risk and quality management into existing sectoral systems. The date is 1 August, a day before the cluster that	Annex I providers facing two compliance narratives	Deadline	Art. 96(1)(g)
2 August 2027 was 2 August 2026	National regulatory sandboxes operational	At least one per Member State, alone or jointly with others. Sandbox plans may now absorb real-world testing plans under Articles 60 and 60a. Separately, the AI Office may run a Union-level sandbox with priority access for SMEs, start-ups and small mid-caps.	Member States; smaller providers seeking supervised routes	Postponed	Art. 57(1)
2 August 2027 was 2 August 2026	National regulatory sandboxes operational	At least one per Member State, alone or jointly with others. Sandbox plans may now absorb real-world testing plans under Articles 60 and 60a. Separately, the AI Office may run a Union-level sandbox with priority access for SMEs, start-ups and small mid-caps.	Member States; smaller providers seeking supervised routes	Postponed	Art. 57(1)
2 August 2027	Delegated acts on equivalent sectoral protection	Commission specifies which high-risk systems may have Articles 9 to 15 and 17 to 25 requirements limited where Annex I Section A legislation already provides equivalent or higher protection. Discretionary until adopted: treat as potential upside, not as a planning assumption.	Annex I Section A providers	Deadline	New Art. 2(13); Art. 97
2 August 2027	Legacy GPAI models comply	Providers of general-purpose AI models placed on the market before 2 August 2025 meet the Articles 53 and 55 obligations: documentation, copyright policy, training-content summary, and for systemic-risk models the safety and security framework. Unchanged by the Omnibus.	GPAI model providers	Unchanged	Art. 111(3)
2 September 2027	Post-market monitoring guidance published	Commission guidance including a voluntary template, replacing the implementing-act power the Omnibus removed. Providers gain flexibility to tailor monitoring to their organisation.	Providers of high-risk systems	Deadline	Art. 72(3)
2 December 2027 was 2 August 2026	Annex III high-risk requirements apply	Chapter III Sections 1 to 3 for systems classified as high-risk under Article 6(2): classification, the Articles 9 to 15 requirements, and the Articles 16 to 27 provider and deployer obligations. Covers recruitment, credit scoring, education, essential services, biometrics, critical infrastructure, law enforcement, migration and justice. Article 6(5) is carved out and keeps its original timing.	Providers, deployers, importers and distributors of Annex III systems	Postponed	Art. 113(3)(c)(i)
28 January 2028	Notified body designation deadline	Bodies notified under Annex I Section A legislation, relying on the Article 43(3) bridge to assess high-risk AI systems, must apply for designation under the AI Act. Applications remain possible after this date under Article 28. Watch this flow as the leading indicator of assessment capacity.	Notified bodies; operators depending on their capacity	Deadline	Art. 43(3), 2nd subpara
2 August 2028 was 2 August 2027	Annex I high-risk requirements apply	Chapter III Sections 1 to 3 for systems classified as high-risk under Article 6(1): AI as a safety component in products under Union harmonisation legislation, or AI that is itself such a product.	Product manufacturers acting as providers; sectoral value chains	Postponed	Art. 113(3)(c)(ii)
2 August 2028	Machinery delegated acts apply	Annex III to the Machinery Regulation gains health and safety requirements mirroring AI Act Chapter III Section 2 and Articles 17, 19, 72 and 73. Machinery has moved from Annex I Section A to Section B, so AI in machinery migrates to a sectoral conformity track.	Machinery manufacturers	New obligation	Art. 8 Reg. (EU) 2023/1230
2 August 2028	European Commission evaluation and review	Article 112 evaluation and periodic review, including the annual assessment of whether Annex III and the Article 5 prohibitions need amending.	European Commission	Unchanged	Art. 112
2 August 2030	Public authority legacy systems comply	Providers and deployers of high-risk systems intended for use by public authorities take the necessary steps, whether or not the system has been significantly changed. Public buyers will demand forward compliance in tenders well before this backstop.	Public-sector suppliers and deployers	Unchanged	Art. 111(2)
31 December 2030	Annex X large-scale IT systems comply	AI systems that are components of the large-scale EU information systems established by the legal acts listed in Annex X. The true end of the calendar.	Operators of EU large-scale IT systems	Unchanged	Art. 111(1)