

The EU AI Act and the AI Omnibus: A Complete Guide



EXECUTIVE SUMMARY

The AI Omnibus, in force since July 27, 2026, delayed the EU AI Act's high-risk deadlines to December 2027 and August 2028, but left transparency, general purpose AI (GPAI) enforcement, and market surveillance obligations on their original August 2026 date. This guide covers what changed, who's affected, and the practical sequence for staying compliant through 2028.

Table of Contents

Introduction	2
PART ONE: What the EU AI Act Actually Requires	2
PART TWO: What the AI Omnibus Changed	3
PART THREE: What Stayed on the Original Date	3
PART FOUR: Who Needs to Act, and On What	4
PART FIVE: The Runway Comes With Conditions	4
PART SIX: The Practical Sequence	5
PART SEVEN: How BigID Helps	5
PART EIGHT: What This Means for AI Regulation Beyond the EU	6
Conclusion	6
Glossary	7

Introduction

The EU AI Act's compliance timeline changed on July 27, 2026, when the AI Omnibus entered into force. This is the first substantive amendment to the Act since it was adopted in 2024, and it resets the deadline every compliance team had circled in red: August 2, 2026, the date high-risk AI obligations under Annex III were set to become enforceable.

That deadline moved. But not all of it did, and the parts that stayed in place are the parts most organizations still need to act on immediately. This guide walks through what changed, who's affected, what's still due on the original date, and a practical sequence for the next two years.

PART ONE

What the EU AI Act Actually Requires

The Act uses a tiered, risk-based structure. Which tier a given AI system falls into determines which obligations apply and when.

- **Prohibited practices:** certain AI uses are banned outright, including social scoring, exploitation of vulnerable groups, real-time biometric identification in public spaces (with narrow exceptions), and manipulative techniques that cause harm. Enforceable since February 2, 2025, with no extension.
- **General purpose AI (GPAI):** models designed for a wide range of tasks, the kind that power chatbots, coding assistants, and content generation tools. Providers have had compliance obligations, technical documentation, training data summaries, copyright compliance, since August 2, 2025. What changes August 2, 2026 is enforcement.
- **High-risk AI systems (Annex III):** specific use cases carrying meaningful risk to health, safety, or fundamental rights: employment and worker management, credit scoring and insurance access, education, critical infrastructure, law enforcement, migration and border control, and administration of justice. This is the tier most affected by the Omnibus.
- **High-risk AI systems (Annex I):** AI embedded in products already regulated under other EU frameworks, medical devices, machinery, vehicles. This tier follows a separate timeline from standalone Annex III systems.
- **Transparency obligations (Article 50):** requirements that apply regardless of risk tier, disclosing when someone is interacting with an AI system, labeling AI-generated content, and marking deepfakes. The broadest-reaching obligation in the Act.

PART TWO

What the AI Omnibus Changed

The Omnibus is the outcome of a process that started with the European Commission's proposal in November 2025, moved through a provisional agreement in May 2026, and was formally adopted by Parliament on June 16 and Council on June 29, before entering into force on July 27.

- **High-risk deadline, standalone Annex III systems:** pushed from August 2, 2026 to December 2, 2027.
- **High-risk deadline, product-embedded Annex I systems:** pushed from August 2, 2026 to August 2, 2028.
- **AI literacy requirement:** loosened. Instead of requiring providers and deployers to ensure a sufficient level of AI literacy among staff, the amended language requires them to take measures supporting the development of that literacy. A lower bar, but not a removed one.
- **Watermarking obligations:** given a four-month transitional period for systems already on the market before August 2, 2026, recognizing that retrofitting labeling into existing products takes time.
- **New prohibition:** AI systems that generate or manipulate non-consensual intimate images, video, or audio, or generate child sexual abuse material, are now explicitly banned, effective December 2, 2026, on its own separate clock from the rest of the Omnibus.

PART THREE

What Stayed on the Original Date

Three things did not move, and they matter more than the delay does for most organizations in the near term.

- **Article 50 transparency obligations** take effect August 2, 2026, unchanged. Any organization with a customer-facing chatbot, a content generation feature, or synthetic media output needs disclosure and labeling mechanisms working by that date.
- **National market surveillance authorities** gain enforcement power, the ability to investigate and sanction AI Act breaches, on August 2, 2026.
- **The European Commission's enforcement powers over GPAI providers** activate August 2, 2026, turning obligations that have applied since 2025 into obligations with real penalties behind them.

Worth noting: the EU's AI Office has already opened an enforcement action on synthetic media practices before the high-risk phase even arrived. The enforcement infrastructure is operating now, ahead of the biggest compliance wave, not waiting for it.

PART FOUR

Who Needs to Act, and On What

Any organization whose AI output reaches EU users, customers, or citizens is in scope, regardless of where the company is headquartered. Reach, not location, is the test.

- **If you provide or embed a general purpose model:** GPAI obligations already apply, and enforcement teeth arrive August 2, 2026.

- **If you operate in employment, credit and insurance, education, critical infrastructure, law enforcement, or migration:** Annex III obligations apply to you, on the new December 2027 timeline, but the underlying work should start now given how long conformity documentation takes to build.
- **If you make or sell products with embedded AI in a regulated category** (medical devices, machinery, vehicles): Annex I obligations apply, on the August 2028 timeline.
- **If you run any customer-facing AI feature at all:** Article 50 applies to you regardless of any of the above, on the original August 2026 date.



PART FIVE

The Runway Comes With Conditions

- **Standards are still catching up.** This is the first amendment to the Act since it was adopted, and harmonized technical standards for high-risk compliance are still behind schedule, part of why the delay happened in the first place. Standards bodies missing their own timelines means more slippage is plausible, not settled.
- **Conformity work has a long lead time regardless of the deadline.** Data governance documentation, risk management processes, and human oversight design take years to build properly, not months. Organizations that build the foundation now will be in a materially stronger position in 2027 than organizations that wait until the deadline is close again.
- **Article 50 and GPAI enforcement are not part of the delay.** Whatever is true about an organization's high-risk exposure, the transparency and disclosure work is due in days, not years.

PART SIX

The Practical Sequence

- **Now through August 2, 2026:** complete an AI system inventory across the organization, including tools procurement never formally approved. Classify each system against prohibited, high-risk, GPAI, and transparency-obligation categories. Confirm Article 50 disclosure and labeling mechanisms are live for anything customer-facing. Document GPAI usage across the vendor stack.
- **Through end of 2026:** begin building the data governance foundation for Annex III systems, training data documentation, data lineage, access controls. Establish the risk management and human oversight processes Annex III conformity will eventually require. Assign clear ownership of ongoing regulatory monitoring.
- **Through 2027:** complete conformity assessments and technical documentation for standalone high-risk systems ahead of December 2, 2027. Register applicable high-risk systems in the EU database as required.

- **Through 2028:** extend the same conformity work to product-embedded high-risk systems ahead of August 2, 2028.



PART SEVEN

How BigID Helps

Every obligation in this framework, current or deferred, rests on the same underlying question: does the organization actually know what data is feeding its AI systems, where that data came from, and who can access it.

- **Data governance documentation for Annex III:** BigID's discovery and classification engine scans training, validation, and testing datasets across cloud, on-prem, and SaaS sources, then tags what's sensitive or regulated within them, so "the data is representative and error-free" is backed by a scan result instead of a compliance team's assertion.
- **Shadow AI and system inventory:** BigID's AI model discovery capability identifies every AI tool and model connected to organizational data, including ones spun up outside procurement, by scanning for AI-related infrastructure and data flows rather than relying on a self-reported list.
- **Data lineage into and out of AI pipelines:** BigID traces how sensitive data moves from source systems into training sets, fine-tuning jobs, and prompts, and follows where model outputs get stored afterward, building the flow map that both conformity assessments and breach investigations depend on.
- **Access governance for training and output data:** BigID's access intelligence maps who and what, including service accounts and other AI systems, can reach a given dataset, which is the access control evidence Annex III's human oversight requirements ask for directly.
- **Article 50 readiness:** the same AI model discovery that surfaces shadow AI also flags which of those systems are generating customer-facing content, giving compliance teams a starting inventory for transparency disclosures instead of a manual survey of business units.
- **Audit-ready reporting:** BigID consolidates discovery, classification, and lineage results into dashboards and exportable reports mapped to specific regulatory requirements, so a market surveillance authority request gets answered with generated evidence, not a scramble to assemble one.

PART EIGHT

What This Means for AI Regulation Beyond the EU

The EU AI Act was positioned as the definitive template for AI regulation globally, the way GDPR became the reference point for data privacy law. The Omnibus complicates that story without undermining it. It shows a regulator adjusting to the gap between what it legislated and what standards bodies and companies could actually deliver on schedule,

while holding firm on the parts it considers non-negotiable: prohibitions, transparency, enforcement authority.

The pattern worth watching, in the EU and in any other jurisdiction drafting AI rules, is this: obligations tied to fundamental rights and public trust tend to move first and stay fixed, while obligations tied to technical conformity move on whatever timeline reality allows. Organizations operating across multiple jurisdictions should expect more of this pattern, not less, as other regulators watch how Brussels handles the gap between ambition and implementation capacity.

Conclusion

Two dates now answer the question of when the EU AI Act's next phase takes effect. The high-risk deadline moved to December 2027 and August 2028. The transparency and enforcement deadline stayed at August 2, 2026.

Organizations that treat both dates seriously, rather than defaulting to whichever one is more convenient, will be the ones prepared for whatever comes after either of them.

Glossary

- **AI Omnibus:** the package of amendments to the EU AI Act, adopted in mid-2026, that postponed the high-risk system compliance deadlines and made several other targeted changes.
- **Annex III:** the list of high-risk AI use cases under the EU AI Act, covering areas like employment, credit and insurance access, education, critical infrastructure, law enforcement, and migration.
- **Annex I:** the list of product categories, such as medical devices, machinery, and vehicles, where embedded AI is regulated as high-risk under existing EU product safety law.
- **Article 50:** the EU AI Act's transparency provision, requiring disclosure of AI interactions, labeling of AI-generated content, and marking of deepfakes.
- **Conformity assessment:** the formal process a provider must complete to demonstrate a high-risk AI system meets the Act's requirements before it can be placed on the market.
- **GPAI (general purpose AI):** AI models designed for a wide range of tasks rather than a single narrow use, such as the models underlying chatbots and coding assistants; providers of these models carry their own compliance obligations under the Act.
- **AI Office:** the European Commission body responsible for overseeing and enforcing rules for general purpose AI models and coordinating AI Act implementation across member states.
- **Market surveillance authority:** the national body in each EU member state responsible for investigating and sanctioning AI Act violations within its jurisdiction.
- **High-risk AI system:** an AI system that falls into one of the Act's specifically enumerated high-risk categories (Annex I or Annex III), triggering the most demanding set of obligations under the law.
- **Prohibited AI practice:** an AI use banned outright under the Act, including social scoring, exploitation of vulnerable groups, and, as of the AI Omnibus, generating non-consensual intimate imagery or CSAM.

CUSTOMER VALIDATION

Proof in Action

"BigID gives me better visibility into sensitive data, helps prioritize security protections, reduces attack surfaces, strengthens compliance, and increases operational efficiency overall — a strategic pillar of our AI-First Cybersecurity Transformation"

- CISO, Global Healthcare Company

"We needed to automate processes and data management across systems for the data of a few million customers across a lot of systems and data. BigID was the one solution that did this in the most efficient and sophisticated way - and had more use cases we could add on moving forward."

- Chief Privacy Officer, Global Telecoms Company

ANALYST RECOGNITION

Validated by Industry Leaders

- 2025 Company of the Year for AI Governance - Frost & Sullivan
- Named a Leader across all DSPM research including Omdia DSPM Universe, CB Insights for DSPM, TAG, GigaOm, Frost & Sullivan
- Named a Leader in all Privacy Management evals including Forrester Privacy Management Wave & IDC Privacy Compliance MarketScape
- BigID named a Leader across 4 GigaOm evals: DSPM, Data Security Platforms, Unstructured Data Management, Data Access Governance
- Named in Gartner's 2025 DLP Market Guide
- Named in Gartner's 2025 AI TRiSM Market Guide
- Represented 30+ times in Gartner's 2025 Hype Cycles for Security, Privacy, AISPM

FREE RISK ASSESSMENT

Discover Your Data & AI Risk

An agentless, cloud-native assessment to enable security, privacy, and compliance: pinpoint where sensitive data and AI risks live - and how to fix them.

What You Get:

- Complementary 2 week engagement on real data to gain visibility on risk across data and AI
- Executive-ready report identifying high-risk data and AI across cloud environments with next-step recommendations

NEXT STEPS

Ready to Connect the Dots in Data & AI?

Schedule your free risk assessment or ask for a deeper demo across DSPM, DLP, AI SPM, Privacy Automation, AI Risk, and Data Lifecycle Management

► www.bigid.com/demo

► www.bigid.com/risk-assessment

► info@bigid.com



BigID Next

"The features they offer are the moat that they have built around the product. This is what makes them the market leader in Data Privacy and Protection."



IT Security & Risk Management Associate - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.



BigID Next

"Exceptionally powerful and easy to use tool that integrated smoothly into our environment and delivered value from day one. The BigID team is fabulous to work with."



Senior Vice President & Chief Data Officer - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.