WHITE PAPER · NOVEMBER 2026

EU AI Act deadlines after the Omnibus: what is due in 2026 and 2027

Which dates moved, which did not, and how to find the one that applies to your AI system

Published by TAK Devs
Published November 18, 2026 · Last checked November 17, 2026

Contents

/01 What does a missed EU AI Act deadline cost a software company?

/02 What did the Digital Omnibus change in the EU AI Act, and when?

/03 Which EU AI Act deadlines did the Omnibus not move?

/04 Which EU AI Act deadline applies to your AI system?

/05 What should product teams build before the next AI Act deadline?

/06 What belongs on an EU AI Act deadline checklist?

/07 How did TAK Devs build governance into the RegioKI AI platform?

/08 Why put a verify step between building and publishing AI?

/09 Frequently asked questions

/10 How this paper was checked

/11 References

/00

Executive summary

THE SHORT ANSWER

The Digital Omnibus, Regulation (EU) 2026/1744 , entered into force on 27 July 2026. It moves Annex III high-risk duties from 2 August 2026 to 2 December 2027, and Annex I systems to 2 August 2028. Article 50 transparency rules did not move. Generative systems on the market before 2 August 2026 must mark their output by 2 December 2026.

“The AI Act got delayed” is half true, and the wrong half can cost you. The high-risk rules moved. The transparency rules for chatbots and generated content did not, and one of them lands eight days after this paper goes live.

This paper is for CEOs, CTOs and product leads who ship AI features to EU users and are not sure which date now applies to them. It gives the old and new dates side by side, a decision path for each system, what to build first, an eight-point checklist and RegioKI as proof of how governance can sit inside the product.

Three findings

01 High-risk duties moved, not vanished. Annex III systems now face 2 December 2027 and Annex I systems 2 August 2028, under the amended Article 113[1].

02 Transparency did not move. Article 50 has applied to new systems since 2 August 2026, and generative systems already on the market must mark their output by 2 December 2026, according to the Commission's AI Act timeline[2].

03 Fines are real money. Breaching Article 50 can cost up to EUR 15 million or 3% of worldwide annual turnover, whichever is higher, under Article 99 of the AI Act[3].

/01

What does a missed EU AI Act deadline cost a software company?

Picture the support ticket. An enterprise buyer in Munich asks whether the images your product generates carry a machine-readable marker. Your team says the AI Act was delayed. It was not, not for that feature, and the buyer's procurement team knows it.

That is the first cost of a wrong date: a stalled deal, long before any regulator calls. The second cost is the fine. Article 99 of the AI Act[3] sets fines of up to EUR 15 million or 3% of total worldwide annual turnover, whichever is higher, for breaching the Article 50 transparency duties. Banned practices carry up to EUR 35 million or 7%.

2 Dec 2026
Article 50(2) marking deadline for generative systems already on the market (Regulation (EU) 2026/1744)

EUR 15M or 3%
maximum fine for breaching Article 50, whichever is higher (AI Act, Article 99)

20.0%
of EU enterprises with 10 or more staff used AI in 2025, up from 13.5% in 2024 (Eurostat)

More companies are in scope every year. Eurostat's 2025 enterprise survey[4] found that 20.0% of EU enterprises with 10 or more employees used AI technologies in 2025, up from 13.5% a year earlier. Many of them bought those tools from software companies like yours, which makes you the provider the rules talk about.

The confusion is understandable. The Digital Omnibus on AI changed several dates at once, and many pages still quote 2 August 2026 for everything. A team that reads one headline and parks the topic until 2027 will miss the deadline that matters this year.

The expensive mistake is not missing the high-risk date. It is assuming every date moved with it.

The fix is not a legal project. It is an inventory: which AI features you ship, what each one does, and which date applies to each. The rest of this paper walks through that, one system at a time.

Start now because the fix takes longer than it looks. Marking generated output touches every place your product creates content: the editor, exports, the API, email digests. Each needs its own test. Expect procurement teams to ask for that evidence in writing, and answer with a dated register rather than a promise.

/02

What did the Digital Omnibus change in the EU AI Act, and when?

The Digital Omnibus on AI is an amending regulation. It does not replace the AI Act. It changes dates and a handful of rules inside it. Council and Parliament negotiators reached a provisional deal on 7 May 2026, and the Council gave its final approval on 29 June 2026, according to the Council of the EU's AI timeline[5].

Regulation (EU) 2026/1744[1] was published in the Official Journal on 24 July 2026 and entered into force on 27 July 2026, the third day after publication. Its new Article 113 sets the high-risk dates: 2 December 2027 for systems listed in Annex III, and 2 August 2028 for AI that is part of a product covered by Annex I, such as medical devices or toys.

TABLE 1. EU AI ACT DATES BEFORE AND AFTER THE DIGITAL OMNIBUS. SOURCES: REGULATION (EU) 2024/1689 AND REGULATION (EU) 2026/1744, OFFICIAL JOURNAL OF THE EU.

Obligation	Original date	After the Omnibus
Prohibited practices (Article 5)	2 February 2025	Unchanged
New bans on generating non-consensual intimate imagery and child sexual abuse material	Not in the original Act	2 December 2026
General-purpose AI model duties (Chapter V)	2 August 2025	Unchanged
Article 50 transparency, systems placed on the market from 2 August 2026	2 August 2026	Unchanged
Article 50(2) marking, generative systems placed on the market before 2 August 2026	2 August 2026	2 December 2026
High-risk systems in Annex III areas	2 August 2026	2 December 2027
High-risk AI in products covered by Annex I	2 August 2027	2 August 2028

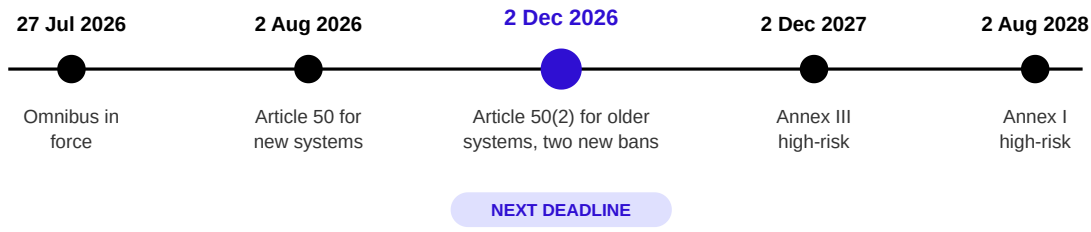


Figure 1. The EU AI Act dates that matter for software companies after the Omnibus. Source: Regulation (EU) 2026/1744.

Two other changes matter for mid-sized software firms. Orrick's July 2026 client alert[6] explains that small mid-cap companies, with fewer than 750 staff and turnover up to EUR 150 million or a balance sheet up to EUR 129 million, now get SME-style relief such as simpler documentation. It also notes a narrower definition of a safety component, which keeps some assistive features out of the high-risk class.

/03

Which EU AI Act deadlines did the Omnibus not move?

Most of the AI Act's general rules started to apply on 2 August 2026, as the Commission's AI Act timeline[2] shows. That includes Article 50, the transparency chapter. Cooley's August 2026 analysis[7] reads the Omnibus the same way: transparency duties for new systems and the duties of general-purpose AI model providers stayed on their original schedule.

Article 50 of the AI Act[3] sets four duties that touch most product teams:

- Chatbots and other systems that talk to people must tell users they are dealing with AI, unless that is obvious.
- Systems that generate synthetic audio, images, video or text must mark the output in a machine-readable format so it can be detected as AI-generated.
- Deployers of emotion recognition or biometric categorisation systems must inform the people exposed to them.
- Deployers who publish deep fakes must disclose that the content was generated or manipulated.

The only relief is narrow. The new Article 111(4), added by Regulation (EU) 2026/1744[1], gives providers of generative systems placed on the market before 2 August 2026 until 2 December 2026 to meet the marking duty in Article 50(2). Everything else in Article 50 already applies.

The banned practices in Article 5 have applied since 2 February 2025, and the two new bans on generating non-consensual intimate imagery and child sexual abuse material apply from 2 December 2026. Rules for general-purpose AI models have applied since 2 August 2025. To help with marking, the Commission has published a Code of Practice on marking and labelling AI-generated content[8].

HELP WITH MARKING IS VOLUNTARY, THE DUTY IS NOT

The European Commission's Code of Practice on marking and labelling AI-generated content, published on 10 June 2026, is voluntary. Signing it is a choice. Meeting Article 50 is not.

/04

Which EU AI Act deadline applies to your AI system?

Run each AI feature through the same path. One product can hit three dates, because the AI Act regulates systems and uses, not companies. Work feature by feature and write the answer down.

1. **Rule out banned uses first**
   If a feature could fall under Article 5, stop and get legal advice. Those bans have applied since 2 February 2025, and two more start on 2 December 2026.
2. **Check for generated content**
   Does the feature generate text, images, audio or video? If it was on the market before 2 August 2026, its Article 50(2) marking deadline is 2 December 2026. If it launched later, the duty already applies.
3. **Check for people-facing AI**
   Chat interfaces, emotion recognition and deep fakes fall under the other Article 50 duties, which have applied since 2 August 2026.
4. **Check the Annex III areas**
   Hiring, credit, education, biometrics, critical infrastructure and the other Annex III areas point to 2 December 2027.
5. **Check for regulated products**
   AI that is part of a product covered by Annex I, such as a medical device, points to 2 August 2028.

One trap sits in the last two steps. Under the amended Article 111(2) in Regulation (EU) 2026/1744[1], a high-risk system already on the market before its date is caught only if its design changes significantly after that date. A major rebuild in 2028 can pull an older system into scope. Log design changes now, so you can show later what changed and when.

More dates follow. The Future of Life Institute's AI Act Explorer[9] lists Commission deliverables still to come, including guidance and a template for post-market monitoring plans due by 2 September 2027. Watch those, because guidance often decides what evidence a reviewer expects.

/05

What should product teams build before the next AI Act deadline?

Most of the work is product work, not paperwork. These five items cover the Article 50 duties and set up the evidence a high-risk review would ask for in 2027.

1. An AI feature register: every feature that uses a model, who supplies the model, what it outputs and which date applies.
2. Disclosure in the interface: a clear line wherever users talk to an AI, built into the design system so new screens inherit it.
3. Machine-readable marking: watermarks or metadata on generated images, audio, video and text, tested on every output type you ship.
4. Change logging: who approved each model or prompt change, when, and what tests ran before release.
5. Human oversight hooks: a way for a named person to review, override or switch off an AI output in the workflows where errors hurt people.

Build these into the release pipeline, not a spreadsheet. A register that updates when code ships stays true. A spreadsheet goes stale the week after the audit.

If your company has fewer than 750 staff, check the small mid-cap thresholds that Orrick summarises[6] before scoping the documentation work. The relief changes how much paperwork a high-risk system needs, not whether Article 50 applies.

Give each item an owner and a date. The register and the marking work come first, because they decide the 2 December 2026 deadline. Logging and oversight can follow early in 2027.

/06

What belongs on an EU AI Act deadline checklist?

Answer these eight questions for each AI feature. Two or more unclear answers mean the feature has no owner for compliance yet.

EU AI ACT DEADLINE CHECKLIST

- [ ] **Was the system on the EU market before or after 2 August 2026?**
  The answer decides whether the Article 50(2) marking deadline is 2 December 2026 or already passed.
- [ ] **Does it generate synthetic text, images, audio or video?**
  If yes, outputs need machine-readable marking. Test it on every output format, including exports and API responses.
- [ ] **Do users know they are talking to an AI?**
  Chat and voice features need a clear disclosure unless it is obvious. This has applied since 2 August 2026.
- [ ] **Is it used for hiring, credit, education, biometrics or critical infrastructure?**
  These are Annex III areas. They point to the 2 December 2027 high-risk date and its documentation duties.
- [ ] **Is it part of a regulated product under Annex I?**
  Medical devices, toys and other covered products point to 2 August 2028. Check the narrower safety component definition first.
- [ ] **Who signs off on model changes, and is that logged?**
  A named approver and a timestamped log are the cheapest evidence you will ever produce for a review.
- [ ] **Can you produce technical documentation on request?**
  Architecture, data sources, model versions, tests and known limits, kept current with each release.
- [ ] **Could a planned rebuild bring an older system into scope?**
  A significant design change after the high-risk date can trigger the full obligations. Plan major rebuilds with that in mind.

Questions one to three decide your exposure this quarter. Questions four to eight decide how calm 2027 will be.

Run the checklist with the product owner and the engineer who knows the model integration, not with legal alone. Legal can read the rule. Only the team can say whether the export endpoint strips the watermark. Record each answer with a date and an owner, then repeat the exercise whenever a feature ships or a model changes.

/07

How did TAK Devs build governance into the RegioKI AI platform?

RegioKI is a Germany-based AI software-as-a-service platform that lets small and medium-sized businesses build, verify and publish AI automations without writing code. According to the RegioKI case study[10], TAK Devs worked as its research and development partner, from specification and architecture through design and production deployment.

6+ months
research, design and full-cycle development

14+ people
engineers, AI scientists and product strategists

5 component types
NLP, OCR, vision, speech and LLM chains, verified before publishing

The platform is multi-tenant and serves both individuals and enterprises. Users assemble automations in a visual drag-and-drop builder from five component types: natural language processing, optical character recognition, vision, speech and large language model chains.

Governance sits in the product, not beside it. The case study describes a Build, Verify, Publish, Subscribe lifecycle, so components are verified before anyone can publish them to the platform's AI marketplace. Enterprise role-based access control decides who can do what, and GDPR-aligned data controls govern how data is handled.

The work started before the build. TAK Devs co-authored an early specification framework for agentic workflows with RegioKI, designed how components and agents communicate, and produced the designs for the builder, marketplace and admin portals. Delivery ran on AWS with CI/CD pipelines and containerised microservices, with Angular, Node.js, PostgreSQL, vector search and LangChain for multi-agent coordination.

RegioKI's founder, Thomas Wehner, called TAK Devs "not less than the top EU companies I have worked with" in the published case study.

This paper does not describe RegioKI as compliant with the AI Act or as a high-risk system, and neither does TAK Devs. The point is narrower and more useful: verification, access control and data controls were built into the platform itself, which is the kind of evidence AI Act reviews ask for.

For a CTO with an AI Act review ahead, the lesson is order. Access, verification and data handling sat in the project scope next to the builder and the marketplace, as product requirements rather than add-ons for later. Read the full RegioKI case study for the stack and the screens.

/08

Why put a verify step between building and publishing AI?

RegioKI's lifecycle puts a gate between what someone builds and what customers can use: per the case study[10], components are verified before they are published. TAK Devs treats that gate as a design default for AI products, not a regulatory extra.

A gate gives you three things the AI Act keeps asking for: a record of what was checked, a named person who approved it, and a point where a change can be stopped. Add role-based access, so only the right roles can publish, and the approval trail writes itself.

Retrofitting that gate after launch is slow, because every existing automation has to pass through it. Building it in from the start costs a few screens and a status field.

The minimum version is small: a status on every AI component (draft, verified, published), a rule that only certain roles can move it forward, and a log entry each time it moves. That is enough to answer the first questions any reviewer, buyer or regulator will ask.

/09

Frequently asked questions

Is the EU AI Act delayed?

Partly. The Digital Omnibus moved the high-risk dates to 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems. Article 50 transparency rules, banned practices and general-purpose AI model duties kept their dates. Generative systems already on the market get until 2 December 2026 for output marking.

When do Article 50 transparency obligations apply?

Article 50 has applied since 2 August 2026. One exception: providers of generative AI systems placed on the market before 2 August 2026 have until 2 December 2026 to mark their output in a machine-readable format, under Article 111(4) as added by Regulation (EU) 2026/1744.

What is the EU AI Act high-risk deadline in 2027?

High-risk obligations for systems in Annex III areas, such as hiring, credit, education and critical infrastructure, apply from 2 December 2027. AI that is part of a product covered by Annex I follows on 2 August 2028. Older systems are caught if their design changes significantly after those dates.

What are the fines for missing an AI Act deadline?

Under Article 99, banned practices can draw fines of up to EUR 35 million or 7% of worldwide annual turnover. Breaching Article 50 or most other operator duties can draw up to EUR 15 million or 3%. For each, the higher figure applies, and lower caps apply to SMEs.

Find out which date applies before 2 December

A generative feature shipped before August 2026 has until 2 December 2026 to mark its output. Book an AI compliance readiness review and we will map each of your AI features to the date that applies, using the EU AI Act deadline checklist.

Book an AI compliance readiness review

Read the full RegioKI case study

TAK Devs is a custom software and AI development company with headquarters in Islamabad and a presence in the US, serving clients in the US and Europe. We own the full software lifecycle: strategy, UI/UX, development, QA, cloud and maintenance. ISO 9001 and ISO 27001 certified. Developing Futures.

/10

How this paper was checked

We opened every source below on November 17, 2026. Every date comes from the Official Journal text of Regulation (EU) 2024/1689 and Regulation (EU) 2026/1744, cross-checked against the Commission's AI Act timeline. Legal readings are attributed to the named law firm. TAK Devs is not a law firm: ask counsel before relying on any reading here.

> General information, not legal advice.

/11

References

[1] **Official Journal of the European Union**. Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 (Digital Omnibus on AI). Jul 24, 2026. https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng

[2] **European Commission**. Timeline for the Implementation of the EU AI Act. 2026. https://ai-act-service-desk.ec.europa.eu/en/ai-act/timeline/timeline-implementation-eu-ai-act

[3] **Official Journal of the European Union**. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Jul 12, 2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

[4] **Eurostat**. 20% of EU enterprises use AI technologies. Dec 11, 2025. https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2

[5] **Council of the European Union**. Timeline: artificial intelligence. 2026. https://www.consilium.europa.eu/en/policies/artificial-intelligence-act/timeline-artificial-intelligence/

[6] **Orrick, Herrington and Sutcliffe LLP**. EU AI Act Update: Digital Omnibus Finalizes 8 Compliance Changes. Jul 29, 2026. https://www.orrick.com/en/Insights/2026/07/EU-AI-Act-Update-Digital-Omnibus-Finalizes-8-Compliance-Changes

[7] **Cooley LLP**. Digital AI Omnibus delays key deadlines, introduces new rules. Aug 3, 2026. https://cdp.cooley.com/digital-ai-omnibus-delays-key-deadlines-introduces-new-rules/

[8] **European Commission, Shaping Europe's digital future**. Commission publishes Code of Practice on marking and labelling AI-generated content. Jun 10, 2026. https://digital-strategy.ec.europa.eu/en/news/commission-publishes-code-practice-marking-and-labelling-ai-generated-content

[9] **Future of Life Institute**. Digital Omnibus on AI (AI Act Explorer). 2026. https://artificialintelligenceact.eu/ai-act-explorer/digital-omnibus/

[10] **TAK Devs**. Pioneering Agentic AI for SME Workflow Automation (RegioKI case study). 2026. https://takdevs.com/case-studies/pioneering-agentic-ai-for-sme-workflow-automation/