Crypto & TradFi

AI Act Special: Omnibus Agreement, Transparency and Bans

Unpacking the regulations for investors

The AI Act Faces the Test of Implementation

The week of 7 May 2026 will be remembered as a turning point for European regulation of artificial intelligence. In the space of forty-eight hours, three events profoundly reshaped the applicable framework: the trilogue political agreement of 7 May 2026 on the Digital Omnibus on AI, the addition of a new prohibited practice targeting so-called 'nudification' applications and the generation of non-consensual intimate content, and the publication on 8 May 2026 by the European Commission of draft guidelines on the transparency obligations set out in Article 50 of the AI Act.

This sequence follows a clear logic: postponing what is not yet operationally ready—the obligations applicable to high-risk AI systems, which should have come into force on 2 August 2026—whilst accelerating and tightening up what is deemed urgent: protections against misuse (non-consensual intimate deepfakes, AI-generated child pornography), and transparency towards users (labelling of synthetic content, labelling of deepfakes).

For investors and regulated entities alike, the message is mixed. On the one hand, an additional window of up to 16 months to structure compliance with high-risk requirements, which should provide relief to companies engaged in AI transformation programmes, particularly in financial services where credit scoring and insurance pricing are explicitly classified as high-risk. On the other hand, this is by no means a message of deregulation: the transparency timeline has been extended only marginally (four months for systems in place as of 2 August 2026), the scope of prohibitions is widening, and the range of penalties remains substantial, with fines of up to €15 million or 3% of annual global turnover for breaches of Article 50.

This special edition of the Regulatory Brief offers a comprehensive analysis of these three publications and their implications, with a particular focus on financial services, a sector which, according to the annexes to the AI Act, has the highest concentration of use cases explicitly designated as 'high risk'.

The weak signal of the month

The AI Act has just undergone its first major regulatory adjustment even before it has fully entered into force, which is rare in the recent history of European law.

Adopted in June 2024 (Regulation (EU) 2024/1689), the AI Act provided for a phased implementation, culminating on 2 August 2026 for high-risk systems. Seven months before this

deadline, the European legislator, having published the Digital Omnibus on AI proposal in November 2025, has revised its own timetable. This reflects an operational reality: the harmonised standards and technical tools required for compliance will not all be in place in time, and the industry, as highlighted in the Draghi report on European competitiveness, is pressing for the necessary time.

For regulated entities, this heralds what might be called a new **era of post-regulatory rigidity** in Europe: the framework is now being developed through an iterative dialogue with its implementation, via successive Omnibus packages. The Commission has already announced a similar initiative for other legislation (GDPR, ePrivacy, Data Act). The risk for businesses is that **they may delay their compliance efforts** in anticipation of further relaxations; however, as the Council of the EU itself points out, the AI Act has been in force since 1 August 2024, and certain obligations already apply (practices prohibited since 2 February 2025, GPAI since 2 August 2025), and the additional window is precisely a call to **use this time strategically** to identify use cases, map data flows and build the required documentation.

Focus 1: Digital Omnibus on AI, the political agreement of 7 May 2026

Following around nine hours of negotiations, and after the failure of an initial trilogue on 28 April 2026, the European Parliament and the Council of the EU reached a provisional political agreement on the Digital Omnibus on AI in the early hours of 7 May 2026. This text amends the AI Act (Regulation (EU) 2024/1689) and constitutes the first deliverable of the 'One Europe, One Market' roadmap agreed between the three institutions.

Revised timetable for high-risk obligations

The agreement replaces the single deadline of 2 August 2026 with **two fixed application dates**:

- **2 December 2027** for high-risk autonomous AI systems (Annex III) — covering in particular biometrics, critical infrastructure, education, employment, law enforcement, border control… **and creditworthiness assessment as well as life and health insurance pricing.**
- **2 August 2028** for high-risk AI systems incorporated into regulated products (Annex I), such as medical devices, toys, lifts, etc.

The conditional mechanism initially proposed by the Commission (triggered by the availability of standards) has been abandoned in favour of these fixed dates, in the interests of legal certainty.

Labelling of AI-generated content: a short reprieve

The obligation to label and detect content generated or manipulated by AI (Article 50(2) of the AI Act), initially applicable from **2 August 2026**, has been postponed to **2 December 2026**, representing a transition period of 4 months – shorter than the 6 months initially proposed by the Commission and the Council, and well below the 12 months requested by the industry. This

transition period only applies to **systems already on the market by 2 August 2026**: systems placed on the market from that date onwards will have to comply from the moment they are launched.

The other transparency obligations set out in Article 50 (Article 50(1) on interaction with an AI system, 50(3) on emotion recognition and biometric categorisation, and 50(4) on the labelling of deepfakes and certain AI-generated texts) **continue to apply from 2 August 2026** without a transition period.

A new prohibited practice: 'nudification'

An addition introduced by the co-legislators during the trilogue (and not present in the Commission's initial proposal), the agreement provides for the addition to Article 5 of the AI Act of a new prohibition targeting AI systems designed to generate or manipulate non-consensual sexual or intimate content, including AI-generated child sexual abuse material (CSAM).

- The ban will apply from **2 December 2026**.
- According to information made public by the European Parliament and several detailed analyses of the compromise text, the ban targets systems whose **intended purpose** is the generation of such content, or those for which such use is reasonably foreseeable without appropriate safeguards being put in place.
- This addition, which was not included in the Commission's initial proposal, is presented by the Council of the EU as a strengthening of "the protection of children against risks associated with AI systems" (press release of 7 May 2026).

Other adjustments

- **Scope of the AI Office**: clarified for AI systems based on general-purpose AI (GPAI) models where the model and the system are developed by the same provider. **Notable exception**: national authorities retain jurisdiction over law enforcement, border control, judicial authorities and **financial institutions**, a point of considerable practical importance for the sector (see Focus 3).
- **National regulatory sandboxes**: the deadline for their establishment by the competent national authorities has been extended to **2 August 2027**.
- **SMEs and small mid-cap companies (SMCs)**: certain exemptions and facilitations reserved for SMEs are extended to the new category of 'small mid-cap companies'.
- **Interaction with product regulations**: the Machinery Regulation is removed from the direct scope of the AI Act; the Commission must publish sector-specific guidelines (medical devices, toys, lifts, watercraft) by 1 **August 2027**.
- **Sensitive personal data**: the legal basis for processing such data for the purposes of detecting and correcting bias is extended to suppliers and deployers of non-high-risk AI systems, subject to strict necessity.
- **Formal adoption**: the political agreement remains provisional and must be formally endorsed by the Council and the Parliament, with publication in the Official Journal expected within a few months, and in any case **before 2 August 2026**.

Impact for investors

- For **technology and industrial** stocks exposed to high-risk AI, the agreement reduces short-term regulatory uncertainty and provides an additional 16 months to structure compliance. Note: this is not deregulation. Requirements regarding security, fundamental rights and governance remain in place.
- For players **in the platform, media, advertising, cybersecurity and content generation tool sectors**, the short deadline (2 December 2026) for labelling synthetic content creates an immediate compliance cost — but also an opportunity for watermarking, cryptographic provenance and generated content detection solutions.
- For **financial institutions**, maintaining the remit of national authorities (in France, the ACPR and AMF, as appropriate) for the supervision of AI systems built on GPAI models is a key point: financial supervision remains the primary focus, which favours integration into existing prudential frameworks rather than a new parallel regime.

Focus 2: The draft Article 50 guidelines, consultation open until 3 June 2026

On 8 May 2026, one day after the political agreement on the Omnibus, the European Commission launched a public consultation on draft guidelines concerning the transparency obligations set out in Article 50 of the AI Act. The text, some 40 pages long, was prepared by the AI Office on the basis of Article 96(1)(d) of the AI Act.

Scope and legal status

- The guidelines are **non-binding**, as only the CJEU can provide an authoritative interpretation, but they constitute the Commission's first instrument covering **the full scope of Article 50**, i.e. the obligations applicable to both providers and operators of certain AI systems.
- They were prepared in parallel with the Code of Practice on Marking and Labelling of AI-Generated Content (the second version of which was published on 5 March 2026), the latter having a narrower scope, limited to Articles 50(2) and 50(4).
- Consultation open until **3 June 2026**. Final adoption expected before 2 August 2026, the date on which the transparency obligations come into force (with the exception of Article 50(2), which has been postponed to 2 December 2026 for existing systems).

Four obligations to be coordinated

📌 **Article 50(1): Information regarding interaction with an AI system**

Providers must design their AI systems intended to interact directly with natural persons in such a way that the latter are informed that they are interacting with an AI system, except where this is 'obvious' from the perspective of a reasonably well-informed, observant and circumspect person.

The guidelines adopt the 'average consumer' benchmark from European consumer law as the standard. They propose a multi-factor test that takes into account the target audience, the possible presence of vulnerable groups (children, older people, people with disabilities) and the level of AI and digital literacy of the intended users. Concrete examples are provided: code-assistance chatbots intended for professional developers and AI-driven non-player characters (NPCs) in video games may meet the threshold of obviousness.

📌 Article 50(2): Machine-readable labelling and detection of AI-generated content

This obligation, which applies to providers of generative AI systems, requires machine-readable labelling of outputs (audio, image, video, text) and the detectability of artificially generated or manipulated content. Technical solutions must be "effective, interoperable, robust and reliable" to the extent that this is technically feasible.

The technical architecture implicitly expected generally combines watermarking, metadata identifiers, cryptographic provenance and fingerprinting — this is the subject of the Code of Practice, which supplements the guidelines. Signatories to the Code benefit from a presumption of compliance; non-signatories are subject to the same statutory quality criteria without this presumption.

📌 Article 50(3): Emotion recognition and biometric categorisation

Operators of emotion recognition or biometric categorisation systems must inform the data subjects concerned of their exposure to such a system and process their personal data in accordance with the GDPR.

📌 Article 50(4): Labelling of deepfakes and certain generated texts

Operators who generate or manipulate images, audio or video constituting deepfakes must indicate that this content has been artificially generated or manipulated. For texts published with the aim of informing the public on matters of general interest, an exception is provided where the content has been subject to human review or editorial control with editorial responsibility resting with a natural or legal person.

The guidelines interpret this exception narrowly: human review must involve a 'deliberate examination of the substance of the content' by persons with 'relevan al competence and professional judgement'. Superficial checks (spelling checks, procedural validation) are not sufficient.

Penalties and personal liability

- Fines for breaches of Article 50 can reach **€15 million or 3% of global annual turnover**; Article 50 thus falls within the second-highest fine bracket under the AI Act.
- The 'personal, non-professional activity' exception in section 2(10) of the AI Act is interpreted narrowly where it relates to public discourse: a deepfake of a politician shared on social media to criticise a decision cannot be covered by this exception.
- AI systems distributed under a free or open-source licence remain **fully subject to Article 50**.

Impact on investors and regulated entities

- For players **in the digital platform, advertising and media sectors**, transparency obligations create significant compliance costs, but also a market for solutions involving labelling, watermarking, detection of synthetic content, reputation monitoring and information risk management.
- For **news publishers and content producers**, the narrow interpretation of the 'human review' exception in Article 50(4) requires a substantive overhaul of editorial workflows: an editorial workflow relying solely on formal correction will not be sufficient to exempt them from the labelling obligation.
- For **financial institutions** deploying chatbots, robo-advisors or AI-based customer support tools, Article 50(1) requires a review of the user experience: the communication of the AI nature of the interaction must be 'clear and distinguishable' within the meaning of the guidelines.
- For **professional investors**, due diligence on companies exposed to risks of disinformation, image abuse or the generation of illegal content is becoming a key focus. A company's ability to demonstrate its technical safeguards (source verification, detection, moderation) is becoming a quality factor.

Focus 3: Why these developments are of particular relevance to financial services

According to Annex III of the AI Act, the financial services sector is, among all the industries covered, the one with the highest concentration of use cases explicitly designated as 'high-risk':

- **Annex III, point 5(b)**: AI systems intended to assess the creditworthiness of natural persons or to establish their credit score (with the exception of systems used to detect financial fraud).
- **Annex III, point 5(c)**: AI systems intended for risk assessment and pricing for natural persons in relation to life and health insurance.
  Recital 58 of the AI Act justifies this classification: these systems determine individuals' access to financial resources or essential services (housing, electricity, telecommunications services) and present risks of discrimination, particularly on grounds of racial or ethnic origin, gender, disability, age or sexual orientation.

Practical implications of the Omnibus for financial sector players

- **Timeline**: the obligations applicable to high-risk autonomous AI systems (Annex III) have been postponed to **2 December 2027**. This gives banks, insurers and asset managers an additional 16-month window to finalise their mapping, governance, technical documentation and compliance assessment work.
- **Supervision**: for AI systems based on GPAI models developed by the same provider, the political agreement explicitly preserves the competence of national authorities for financial institutions — which, in practice, places the ACPR and the AMF in a central position for the AI supervision of regulated entities in France, in conjunction with existing sectoral requirements.
- **Coordination with DORA**: AI systems deployed in the financial sector are also covered by the Digital Operational Resilience Act (DORA, fully applicable since 17 January 2025), which requires ICT risk management, a register of third-party ICT providers (including AI model providers and ML platforms), a resilience testing framework and incident reporting. Regulated entities must manage these two sets of requirements in an integrated manner.
- **Interaction with MiFID II**: robo-advisors and investment advice chatbots remain subject to MiFID II suitability rules as well as consumer protection requirements. Article 50(1) of the AI Act adds a layer of transparency requirements regarding the AI nature of the interaction.
- **Interaction with the GDPR**: for automated decisions affecting natural persons, Article 22 of the GDPR continues to apply, supplemented by Article 13 of the AI Act and recent case law of the CJEU (Dun & Bradstreet, SCHUFA).
- **Recital 158 of the AI Act**: for regulated credit institutions, the assessment of compliance under the AI Act may be combined with the supervisory review and evaluation process provided for under banking law, enabling banks to integrate AI compliance into the existing SREP.

Impact for investors in the financial sector

- The additional time should not be interpreted as an opportunity to procrastinate. Those who have used this window to **integrate AI compliance into their existing prudential frameworks** (model governance, MRM, compliance, internal control) will have an operational advantage.
- **Suppliers of SupTech and RegTech**, AI governance, model auditing, document management, traceability and bias monitoring stand to benefit significantly from the extended timeline: they now have a clear picture of their commercial trajectory.
- For **fintechs and neobanks**, the extension of SME support measures to small and mid-cap companies may reduce the burden of technical documentation. A thorough verification of SMC status against the new criteria is required.

Main sources

- Council of the European Union, press release *"Artificial Intelligence: Council and Parliament agree to simplify and streamline rules"*, 7 May 2026.

- European Commission, press release IP/26/1024 *"EU agrees to simplify AI rules to boost innovation and ban 'nudification' apps to protect citizens"*, 7 May 2026.
- European Commission (AI Office), *"Draft of the guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of the AI Act"*, 8 May 2026.
- European Commission, *"Consultation on the draft guidelines on transparency obligations under the AI Act*".
- European Commission, *"Regulation on the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)"*, proposal published on 19 November 2025.
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act) – Annex III (points 5(b) and 5(c)), Article 50, recitals 58 and 158.

***