EU AI Act Article 50

Transparency Obligations for Providers and Deployers

A Premium AI Governance Desk Briefing for Compliance Officers, Legal Counsel, and AI Product Leaders

Published by AI Governance Desk
Editorial Team: AI Governance Desk Regulatory Analysts
July 2026 Edition

Table of Contents

Executive Summary

Article 50 of Regulation (EU) 2024/1689 (the EU AI Act) introduces horizontal transparency duties that apply regardless of an AI system's risk classification. Providers of systems that interact with natural persons, generate synthetic content, or produce deepfakes, and deployers of emotion-recognition, biometric-categorisation or deepfake systems, must disclose specific facts to affected persons in a clear, distinguishable, and timely manner [Source: European Commission, Regulation (EU) 2024/1689, Article 50, eur-lex.europa.eu/eli/reg/2024/1689/oj, Accessed July 2026].

The obligations become applicable on **2 August 2026**, twenty-four months after entry into force on 1 August 2024 [Source: European Commission, Regulation (EU) 2024/1689, Article 113(b), eur-lex.europa.eu/eli/reg/2024/1689/oj, Accessed July 2026]. Non-compliance is sanctioned under Article 99(4) at up to €15,000,000 or 3% of worldwide annual turnover, whichever is higher [Source: European Commission, Regulation (EU) 2024/1689, Article 99(4), Accessed July 2026].

Governance note

Article 50 is a horizontal duty: it applies alongside — not instead of — high-risk obligations in Chapter III and general-purpose AI (GPAI) duties in Chapter V. A single deployment can trigger Article 50, Article 26 (deployer obligations for high-risk systems), and Article 53 (GPAI) simultaneously.

This briefing operationalises Article 50 for regulated entities. It maps each sub-paragraph to concrete provider and deployer actions, crosswalks the obligations to the NIST AI Risk Management Framework 1.0 (Govern/Map/Measure/Manage) and ISO/IEC 42001:2023, and closes with an implementation checklist and evidence-log template that compliance teams can adopt directly.

1. Regulatory Context and Scope of Article 50

The EU AI Act was adopted on 13 June 2024 and published in the Official Journal on 12 July 2024, entering into force on 1 August 2024 [Source: Publications Office of the EU, OJ L, 2024/1689, eur-lex.europa.eu/eli/reg/2024/1689/oj, Accessed July 2026]. Article 50, titled "Transparency obligations for providers and deployers of certain AI systems", sits in Chapter IV. Unlike Chapter III (high-risk systems), Chapter IV applies horizontally, based on the function of the AI system rather than its risk tier.

1.1 Personal, Material and Temporal Scope

Article 50 binds two categories of operator: (a) **providers** that place on the EU market or put into service systems interacting with natural persons or generating synthetic content, and (b) **deployers** that use emotion-recognition, biometric-categorisation, or deepfake-generating systems within the Union [Source: European Commission, Regulation (EU) 2024/1689, Article 50(1)–(4), Accessed July 2026]. Territorial reach follows Article 2, capturing non-EU providers whose output is used in the Union.

Application date: **2 August 2026**. Providers must ship compliant products by that date; deployers must have disclosure procedures operational for any in-scope use from the same date [Source: European Commission, Regulation (EU) 2024/1689, Article 113(b), Accessed July 2026].

1.2 Exemptions

Article 50 does not apply where use is authorised by law to detect, prevent, investigate or prosecute criminal offences, subject to safeguards for third-party rights, or where the content forms part of an evidently artistic, creative, satirical or fictional work — in which case disclosure is limited to a manner that does not hamper enjoyment of the work [Source: European Commission, Regulation (EU) 2024/1689, Article 50(2) and 50(4) subparagraph 3, Accessed July 2026].

2. Obligations for Providers of Generative and Interactive AI

2.1 Interactive Systems — Article 50(1)

Providers must design and develop systems intended to interact directly with natural persons such that affected persons are **informed that they are interacting with an AI system**, unless this is obvious from the perspective of a reasonably well-informed natural person, taking into account the circumstances and context of use [Source: European Commission, Regulation (EU) 2024/1689, Article 50(1), Accessed July 2026].

2.2 Synthetic Content — Article 50(2)

Providers of AI systems, including general-purpose AI systems, that generate synthetic audio, image, video or text content, must ensure that outputs are **marked in a machine-readable format and detectable as artificially generated or manipulated**. Technical solutions must be effective, interoperable, robust and reliable, as far as this is technically feasible, having regard to the specificities and limitations of various types of content, the costs of implementation, and the

generally acknowledged state of the art, as may be reflected in relevant technical standards [Source: European Commission, Regulation (EU) 2024/1689, Article 50(2), Accessed July 2026].

Standards signal

The Commission has requested CEN-CENELEC JTC 21 to develop harmonised standards addressing content provenance and watermarking. Providers should track work items linked to the C2PA specification and ISO/IEC AWI 42005 (AI system impact assessment) as candidate benchmarks of the 'state of the art' referenced in Article 50(2) [Source: European Commission, Standardisation Request C(2023)3215, single-market-economy.ec.europa.eu, Accessed July 2026].

Provider trigger	Required disclosure	Modality
Chatbot / voice assistant (Art. 50(1))	'You are interacting with an AI system'	In-product, first interaction, persistent
Text/image/audio/video generation (Art. 50(2))	Machine-readable marker of synthetic origin	Embedded metadata, watermark, or equivalent
GPAI model with system-level output	Downstream deployer must be equipped to comply	Model card, use-restriction documentation

Provider trigger	Required disclosure	Modality
Accessibility route	Alternative accessible format required	Screen-reader labels, captions, ARIA

Disclosures must be provided to the affected natural persons in a **clear and distinguishable manner at the latest at the time of the first interaction or exposure**, and must conform to the applicable accessibility requirements set out in Directive (EU) 2019/882 and Directive (EU) 2016/2102 [Source: European Commission, Regulation (EU) 2024/1689, Article 50(5), Accessed July 2026].

3. Deployer Obligations: Emotion Recognition, Biometric Categorisation, Deepfakes

3.1 Emotion Recognition and Biometric Categorisation — Article 50(3)

Deployers of an emotion-recognition system or a biometric-categorisation system must **inform the natural persons exposed to it of its operation** and process personal data in accordance with Regulations (EU) 2016/679 (GDPR) and (EU) 2018/1725, and Directive (EU) 2016/680, where applicable [Source: European Commission, Regulation (EU) 2024/1689, Article 50(3), Accessed July 2026]. The obligation does not apply to systems permitted by law for law-enforcement purposes to detect, prevent or investigate criminal offences, subject to appropriate safeguards for third-party rights and freedoms.

3.2 Deepfakes — Article 50(4)

Deployers of an AI system that generates or manipulates image, audio or video content constituting a **deepfake**, as defined in Article 3(60), must disclose that the content has been artificially generated or manipulated. This obligation does not apply where use is authorised by

law to detect, prevent, investigate or prosecute criminal offences. Where content is part of an evidently artistic, creative, satirical, fictional or analogous work, the disclosure is limited to a manner that **does not hamper the display or enjoyment of the work** [Source: European Commission, Regulation (EU) 2024/1689, Article 50(4) subparagraphs 1–3, Accessed July 2026].

3.3 AI-Generated Text on Matters of Public Interest

Deployers of AI systems that generate or manipulate **text published to inform the public on matters of public interest** must disclose that the text has been artificially generated or manipulated. This does not apply where the AI-generated content has undergone a process of **human review or editorial control** and where a natural or legal person holds editorial responsibility for the publication [Source: European Commission, Regulation (EU) 2024/1689, Article 50(4) subparagraph 2, Accessed July 2026].

> **Compliance risk — 'artistic work' carve-out**
>
> The carve-out in Article 50(4) does not eliminate the disclosure duty; it only narrows its manner. Deployers must still label deepfakes forming part of a creative work — for example in credits, in metadata, or in a linked disclosure — sufficient to signal artificial origin without impairing the work. Silent publication is not permitted even for satire.

4. Interaction with Articles 52, 53 and GPAI Duties

Article 50 does not operate in isolation. Providers of general-purpose AI models are separately bound by Article 53, which requires technical documentation, information for downstream providers, a copyright policy, and a summary of training data using the Commission-provided template [Source: European Commission, Regulation (EU) 2024/1689, Article 53(1)(a)–(d), Accessed July 2026]. GPAI models with systemic risk are additionally subject to Article 55 (evaluations, adversarial testing, cybersecurity).

Article 52 governs the classification procedure for GPAI models with systemic risk and the notification duty when the cumulative amount of compute used exceeds **10^{25} floating-point operations** [Source: European Commission, Regulation (EU) 2024/1689, Article 51(2) and Article 52(1), Accessed July 2026]. A single generative platform can therefore trigger Article 50 (transparency to end users), Article 53 (model documentation for deployers), and Article 55 (systemic-risk mitigations) at once.

Provision	Actor	Trigger	Application date
Art. 50(1)	Provider	Interactive AI system	2 Aug 2026
Art. 50(2)	Provider	Synthetic content generation	2 Aug 2026
Art. 50(3)	Deployer	Emotion recognition / biometric categorisation	2 Aug 2026
Art. 50(4)	Deployer	Deepfake or AI text on public-interest matters	2 Aug 2026
Art. 53	GPAI provider	Placing a GPAI model on the EU market	2 Aug 2025
Art. 55	Systemic-risk GPAI provider	Compute ≥ 10^{25} FLOP or Commission designation	2 Aug 2025

5. Enforcement, Penalties, and Supervisory Architecture

Enforcement of Article 50 falls to **national market-surveillance authorities** designated under Article 70, coordinated at Union level by the **AI Office** established within the European Commission and by the **AI Board** composed of Member State representatives [Source: European Commission, Regulation (EU) 2024/1689, Articles 64–70, Accessed July 2026; European Commission, Commission Decision C(2024)1459 establishing the AI Office, digital-strategy.ec.europa.eu, Accessed July 2026].

Penalties are set by Article 99. Non-compliance with the transparency obligations in Article 50 is sanctioned under Article 99(4) at administrative fines of up to **€15,000,000 or 3% of the total worldwide annual turnover for the preceding financial year, whichever is higher** [Source: European Commission, Regulation (EU) 2024/1689, Article 99(4), Accessed July 2026]. Prohibited-practice breaches under Article 5 attract the top tier of €35M / 7% under Article 99(3); supply of incorrect information to authorities is sanctioned under Article 99(5) at up to €7.5M / 1%.

Penalty tier	Trigger	Cap (higher of)
Art. 99(3)	Prohibited practices (Art. 5)	€35M or 7% worldwide turnover
Art. 99(4)	Breach of Art. 50 transparency and most obligations	€15M or 3% worldwide turnover
Art. 99(5)	Incorrect/misleading information to authorities	€7.5M or 1% worldwide turnover
Art. 101	GPAI provider fines (Commission-imposed)	3% worldwide turnover or €15M

For SMEs including start-ups, Article 99(6) requires authorities to consider the lower of the percentage or fixed amount [Source: European Commission, Regulation (EU) 2024/1689, Article 99(6), Accessed July 2026].

6. Crosswalk: Article 50 ↔ NIST AI RMF ↔ ISO/IEC 42001

The following crosswalk maps each Article 50 duty to the corresponding function of the NIST AI Risk Management Framework 1.0 (January 2023) and to clauses of ISO/IEC 42001:2023 (AI Management System). It is intended as an audit aid: satisfying the NIST/ISO controls listed does not by itself discharge the EU obligation, but it evidences the governance maturity that supervisory authorities are expected to weigh under Article 99(7) [Source: NIST, AI Risk Management Framework 1.0, NIST AI 100-1, nist.gov/itl/ai-risk-management-framework, January 2023; ISO/IEC 42001:2023, iso.org/standard/81230.html, Accessed July 2026].

Article 50 duty	NIST AI RMF function / category	ISO/IEC 42001:2023 clause
50(1) Inform user of AI interaction	GOVERN 1.5; MAP 3.4 (context of use)	Clause 6.1.4; Annex A control A.6.2.6
50(2) Machine-readable marking of synthetic output	MEASURE 2.8 (content provenance); MANAGE 4.1	Clause 8.3; Annex A control A.7.4
50(3) Notice for emotion / biometric categorisation	GOVERN 5.1; MAP 5.1 (impacts to individuals)	Clause 6.1.2; Annex A control A.5.4
50(4) Deepfake and AI-text disclosure	MEASURE 2.11 (content authenticity); MANAGE 2.3	Clause 8.4; Annex A control A.7.3
50(5) Accessible, timely disclosure	GOVERN 3.2 (workforce roles); MANAGE 1.3	Clause 7.4; Annex A control A.6.2.8
Evidence and record-keeping	GOVERN 1.6 (documentation); MEASURE 3.3	Clause 7.5 (documented information)

Gaps to note: NIST AI RMF is voluntary US guidance without legal force; ISO/IEC 42001 is a certifiable management-system standard but does not prescribe watermarking techniques. Neither substitutes for the specific labelling and disclosure duties in Article 50, which remain **legally binding under EU law** regardless of certification status.

7. Implementation Checklist for Providers and Deployers

7.1 Provider Checklist (Article 50(1)–(2))

- Inventory every deployed model that interacts with natural persons or produces synthetic audio, image, video, or text output.
- Add a persistent 'AI system' disclosure at first interaction in every UI surface (web, mobile, voice, API playgrounds).
- Implement C2PA-compatible content credentials or an equivalent watermarking scheme for generative outputs; document why the chosen technique is state of the art.
- Publish model documentation covering intended purpose, capabilities and limitations to enable downstream deployer compliance (aligns with Article 53(1)(b)).
- Verify accessibility of the disclosure under Directive (EU) 2019/882 (European Accessibility Act).
- Log a dated compliance decision record for each release; retain for ten years under Article 18 by analogy for high-risk systems, and for the operational lifetime for others.

7.2 Deployer Checklist (Article 50(3)–(4))

- Screen every AI use case for emotion recognition, biometric categorisation, and deepfake generation before go-live.
- Publish a plain-language notice at the point of exposure — physical signage, in-app banner, or pre-roll — describing the AI operation.
- For AI-generated text on matters of public interest, either apply a visible 'AI-generated' label or document the human editorial review process assigning editorial responsibility.
- Coordinate with the Data Protection Officer to ensure the GDPR / LED lawful-basis analysis is on file before processing biometric or emotion data.
- Maintain a deepfake disclosure register: content ID, publication channel, disclosure method, editorial owner, date.
- Train front-line staff on how to answer end-user questions about AI use; log training completion.

7.3 Evidence Log Template

Field	Description	Retention
System ID	Internal identifier and version hash	Lifecycle + 10 years

Field	Description	Retention
Article 50 sub-paragraph	50(1), 50(2), 50(3), or 50(4)	Lifecycle
Disclosure text and format	Screenshot, transcript, or metadata sample	Lifecycle
Accessibility conformance	EAA / WCAG 2.1 AA test result and date	3 years rolling
State-of-the-art justification (50(2))	Reference to standard, paper, or CEN WI	Lifecycle
Approver	Named person and role	Lifecycle

Sources and Disclaimer

A. Official Regulatory and Government

European Commission — Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act), eur-lex.europa.eu/eli/reg/2024/1689/oj, Accessed July 2026.

European Commission — Commission Decision C(2024)1459 establishing the European AI Office, digital-strategy.ec.europa.eu/en/policies/ai-office, Accessed July 2026.

Publications Office of the European Union — OJ L series, 2024/1689, eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689, Accessed July 2026.

European Parliament Think Tank — 'Artificial Intelligence Act' briefing PE 698.792, europarl.europa.eu/thinktank/en/document/EPRS_BRI(2021)698792, Accessed July 2026.

B. Standards and Frameworks

ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system, iso.org/standard/81230.html, Published December 2023.

CEN-CENELEC JTC 21 — Standardisation Request C(2023)3215 on Artificial Intelligence, single-market-economy.ec.europa.eu/single-market/european-standards/standardisation-requests_en, Accessed July 2026.

C. International Organizations

OECD — Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449), oecd.ai/en/ai-principles, Updated May 2024.

D. Established Policy Media and Research

NIST — Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, nist.gov/itl/ai-risk-management-framework, Published January 2023.

Future of Life Institute — EU AI Act Explorer, artificialintelligenceact.eu, Accessed July 2026.

E. Tools and Platforms Referenced

C2PA — Coalition for Content Provenance and Authenticity, c2pa.org — content credentials specification for machine-readable provenance markers.

Disclaimer

This guide is for informational purposes only and reflects publicly available information as of July 2026. AI laws, regulatory interpretations, technical standards, enforcement priorities, and institutional governance practices may change. Readers should verify current requirements through official regulatory sources, standards bodies, legal advisors, compliance professionals, or qualified governance specialists before making decisions. AIGovernanceDesk.com and its editorial team are not liable for decisions made based on this material. This document does not constitute legal, regulatory, technical, or compliance advice.

Website: aigovernancedesk.com | Email: editorial@aigovernancedesk.com | Published: July 2026

Thank you for reading.

For more governance briefings, visit **aigovernancedesk.com**.