Consultation Response

Draft Guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of Regulation (EU) 2024/1689

3 June 2026

The Association for Financial Markets in Europe (AFME) welcomes the opportunity to comment on the **Draft Guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of Regulation (EU) 2024/1689 (the 'AI Act')** (the "Draft GLs"). The Association for Financial Markets in Europe (AFME) is the voice of the leading banks in Europe's financial markets, providing expertise across a broad range of regulatory and capital markets issues. We represent over 150 leading global and European banks and other significant market players. Our members play a vital role in Europe's financial ecosystem, underwriting around 90% of European corporate and sovereign debt, and 85% of European listed equity capital issuances. Importantly, AFME members are market makers, providing liquidity, which is essential for ensuring financial markets can function efficiently. We also represent law firms and other associate members which advise market participants and support AFME's legal and regulatory initiatives.

AFME is registered on the EU Transparency Register, registration number 65110063986-76.

Section I – Background and Objectives

Question A. Are there any aspects / paragraphs of this Section of the draft Guidelines that you believe require further clarification?
Yes

If so, how would you suggest they be improved to ensure effective implementation and compliance?
Para 2: We suggest cross-referencing Section 8.4, para 141–142, as para 2 refers to the 2 August 2026 application date, while para 141 further explains that in-scope AI systems must be compliant regardless of their date of placement on the market or of being put into service, subject to the possible targeted grandfathering rule for Art 50(2).

Association for Financial Markets in Europe
London Office: Level 10, 20 Churchill Place, London E14 5HJ, United Kingdom T: +44 (0)20 3828 2700
Brussels Office: Rue de la Loi 82, 1040 Brussels, Belgium T: +32 (0)2 883 5540
Frankfurt Office: c/o SPACES – Regus, First Floor Reception, Große Gallusstraße 16-18, 60312, Frankfurt am Main, Germany T:+ 49 (0)69 710 456 660
www.afme.eu

Company Registration No: 6996678 Registered Office: Level 10, 20 Churchill Place, London E14 5HJ
AFME is registered on the EU Transparency Register, registration number 65110063986-76

Section II – Overview of Transparency Obligations and Horizontal Topics

Question A. Are there any aspects / sections / paragraphs that require further clarification?
Yes

Please indicate specific parts
Para 10-14: Clarification on responsibilities across complex provider/deployer value chains (particularly beyond the GLs on High-Risk AI and where AI agents may be operators in their own right).

Model providers and upstream AI system developers are best placed to implement and maintain watermarking/provenance/disclosure, while downstream deployers should be able to rely on those to an appropriate extent. The GLs should explicitly recognise, (i) that downstream providers should not be expected to validate provider transparency controls where they do not control underlying model architecture; and (ii) the role of vendor disclosures and technical documentation in downstream compliance.

Further to Art 25, we request clarity that (as with models), once a non-high risk AI system has been developed and put into service/placed on the market for the first time, there is no other provider for that AI system unless one of the events in Art 25 applies. Without this, Art 25 could be interpreted to mean that an entity which uses a third-party GPAI system without development to define a use case/intended purpose can become the provider.

Para 22-24: Interplay with other EU legal acts is only covered in general terms. For banks, this is complex and multi-layered. The GLs should include practical examples showing interaction with sectoral rules under: (a) MiFID II (investor information, suitability); (b) PSD2 (payment transparency); (c) GDPR (automated decision-making disclosures); and (d) AML and Consumer Protection Obligations (particularly re AML/KYC checks).

Para 24: While there is no direct application to GPAI models, Art 50(1-2) could be applied at model level. We request guidance on the criteria for how the model provider should implement transparency, given there is no L1 obligation to do so. Downstream entities may not have sufficient technical visibility into embedded systems and may therefore need to rely on vendor-provided transparency controls and supporting documentation.

Section III – Article 50(1): Transparency for Interactive AI Systems

Question A. Are there aspects requiring further clarification for effective implementation and compliance?
Yes

Please indicate specific parts
Para 28: Disclosure of all likely interaction scenarios raises practical questions for banking AI agents operating across multiple channels (email, chat, voice, in-app messaging),: should disclosure be embedded once in the agent's identity or repeated at each interaction? Clarity is needed on whether persistent, session-level disclosure satisfies Art 50(1) or if per-interaction disclosure is required.

Para 35: We suggest deleting the final bullet, which appears to require also that additional information be provided explaining the system's function or its implications for the user. While this may be justified for general transparency, it seems to go beyond Art 50(1). Moreover, it is not fully consistent with earlier parts of the Draft GLs — e.g. para 34—where there is no additional requirement and where a disclosure like "You are interacting with an AI system", possibly complemented by similar disclosures in additional channels, appears sufficient to meet the transparency obligation.

Para 36: We suggest deleting as it creates uncertainty as to whether a banner may be used and leaves room for differing interpretations at a national level. Alternatively, guidance would be helpful on the risk of habituation effects ("banner blindness"), to clarify how disclosures should be calibrated to remain clear and distinguishable, as required by Art 50(5).

Para 39-42: Further clarity is needed on *"obviousness"*. It is not clear whether a deployer could also benefit from this exception. Finally, it may be necessary to address changes over time – e.g. a voice assistant which started out with a robot-like voice but becomes more human-like with further development.

Question B. Are there additional practical examples or use-cases that should be clarified?
No

Section IV – Article 50(2): Marking and Detection of AI-Generated or Manipulated Content

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation of marking and detection obligations?
Yes

Please specify
Para 28 and 59: In a multi-agent system, where one agent generates text that is ultimately surfaced to a natural person (the instructing user/a third party) does Art 50(2) apply to the intermediate agent's output, the final output, or both? Is one disclosure at system level sufficient?

Para 65-73: We repeat our para 10-14 comments.

That no one technique meets all Art 50(2) criteria matters for banks, where AI content has multiple transformations (e.g. PDF generation, email insertion). Clarity is needed on how providers ensure that at least one method is effective across common transformation pipelines. Deployers shouldn't bear a remarking burden due to poor upstream design.

Para 77-78: Please clarify that, if a provider is compliant in their approach, they need not continuously adapt to *"state-of-the-art"*.

Para 81–82 should be extended to internal enterprise environments. Purely internal outputs (e.g. logs, source code) needn't be marked unless later exposed externally. We assume this covers internal banking uses e.g. risk model outputs, and MI generated by AI tools?

We amend para 81(ii) : *"…output …is **not intended to be shared outside the company or to be usable or verifiable by external persons**, with appropriate safeguards …"* Examples of *"appropriate safeguards…"* are needed, and for para 82, beyond video games.

Para 84–86: *"AI-generated translations …"* below para 86 seems to contradict para 85 as well as para 84 . Generally, an alteration of "structure" or "style" should not be considered "substantial" – e.g. if the AI summary is based only on input data. This has cross-border banking implications. If retained, please clarify if it applies to translation or the final document, and how it fits with legally prescribed formats (e.g. PRIIPs KID).

Para 87 refers to "points 47–48" for law-enforcement, which is in para 43–46. Generally, exceptions introduced for Art 50(4) in case of human review or editorial control and responsibility should be extended to Art 50(2).

Question B. Are there additional examples, use-cases or technical approaches that should be included?
No

Question C. What technical approaches or solutions should be considered for the implementation of the marking and detection obligations under Article 50(2) AI Act, in particular in relation to AI agents and virtual or immersive environments (e.g. virtual reality)? Please provide concrete examples, methodologies or tools where possible.
Para 57, 59 and 69–78: a layered approach should be encouraged. Providers should combine, where technically feasible, metadata, cryptographic provenance, watermarking, content credentials, fingerprinting, logging and detection tools. No single method is likely to satisfy effectiveness, reliability, robustness and interoperability in all contexts.

We note that the Draft GLs define *"technical feasibility"* as an objective notion not dependent on individual providers' resources. While we support a high standard, the GLs should acknowledge that for financial institutions operating legacy document management systems, the technical integration of marking preservation (especially across PDF generation, email archiving and regulatory submission pipelines) may present significant implementation challenges. We recommend that the GLs encourage providers to publish marking-preservation specifications and test suites. This would enable deployers to validate marking survival across their specific processing chains.

Solutions should be tested against copy-paste, compression, cropping, format conversion, PDF generation, email insertion, screenshotting and re-publication. We recommend adding regulatory archiving, MiFID II record-keeping, and long-term document retention (up to 10 years under DORA/MiFID) to the list of transformations against which marking robustness should be validated.

Section V – Article 50(3): Emotion Recognition and Biometric Categorisation

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation?

Yes

Please indicate specific parts
Para 95–99: Further clarification is needed on the distinction between emotion recognition, biometric categorisation, biometric identification, biometric verification/authentication and fraud-prevention controls. Per the Draft GLs, further details on the notion of *"emotion recognition system"* as well as *"biometric categorisation system"* are to be provided as part of the High-Risk AI GLs, currently under development. This distinction is particularly important in financial services, where biometric technologies may be used for onboarding, authentication, liveness detection, anti-fraud checks, call-centre analytics or customer support purposes.

Para 99 states that the AI Act does not require information about the reasons for the system's operation. The Guidelines should clarify how this interacts with GDPR transparency obligations (Art. 13–14 GDPR, Art. 22 on automated decision-making) and that Art 50(3) notices may be combined with privacy notices, provided the AI-specific information remains clear and distinguishable.

Para 99-101: we suggest that the obligations in article 50.3 (notice to natural persons) should not apply to the use of emotion recognition systems analyzing publicly available recordings.

Para 107 reads "*Existing … it is sufficient for simulated persons, objects, places, entities or events to resemble someone or something that can exist or could have existed in reality to be considered a deep fake."* This seems to broaden the scope of the definition which mentions *"existing"*, i.e. *"AI-generated or manipulated image, audio or video content that resembles existing persons, objects, places, entities or events and would falsely appear to a person to be authentic or truthful."*

Question B. Are additional examples or use-cases or clarifications that should be clarified?
Yes

Please specify:
Additional examples could include the use of biometric technology for AML/KYC requirements. For each example, the GLs should clarify the boundary between systems falling within Art 50(3) and those serving a purely security or authentication purpose, which may fall outside its scope.

Section VI - Article 50(4): Labelling of Deep Fakes and Certain Text

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation?
No

Question B. Are there additional examples or use-cases needed that should be clarified?
No

Section VII - Horizontal requirements (Article 50(5))

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation? Please indicate specific parts.
Para 130-133: Further practical guidance is needed on what constitutes *"clear and distinguishable"* information in regulated omnichannel environments. Banks communicate through mobile apps, websites, email, PDFs, chat interfaces, call centres, social media and third-party platforms.

There is also room for interpretation between *"hidden under layers of menu options"* (layers = plural) and *"noticeable and easy to understand."* We assume that a single interaction (e.g. closing a banner or opening an information pane) is still acceptable, and where the threshold to an impermissible "layering" of options is considered to be.

The GLs should also clarify that generic references in terms and conditions, privacy notices or AI policies should not replace contextual disclosure at the point of interaction or exposure.

Para 132 should be further clarified for the regulatory archiving context. Under MiFID II (Art. 16(7)), DORA (Art. 11) and national AML regulations, financial institutions must retain communications records for 5–10 years. AI-generated or AI-manipulated content that is archived must retain its transparency markers throughout the retention period. The GLs should clarify: (a) whether archived content must remain labelled or otherwise marked during the full retention period; (b) whether audit evidence of original marking is sufficient if the marking itself degrades over time; (c) how marking interacts with requirements on document immutability in regulatory archives.

AFME Contacts

Coen ter Wal
Managing Director, Technology & Operations
coen.terwal@afme.eu
+44 (0)203 828 2727

Fiona Willis
Director, Technology & Operations
fiona.willis@afme.eu
+44 (0)203 828 2739