BSA's Recommendations on the Second Draft of the EU Code of Practice on Transparency Requirements (Art. 50 of the EU AI Act)

April 7, 2026

The Business Software Alliance (BSA) wanted to recognize the various improvements this raise concerns about the second draft Code of Practice (CoP) on the transparency requirements deriving from Article 50 of the EU AI Act, compared to the previous version. However, we are of the opinion that several concerns remains.

BSA is the global trade association representing the enterprise (B2B) software industry. Our members[1] operate at the forefront of AI, cybersecurity, cloud computing, and data-driven innovation. In particular, BSA members are on the leading edge of providing AI-enabled products and services. As such, they have unique insights into the technology's potential to spur digital transformation and practices that can best support the responsible development and use of AI.

BSA and its members are determined for the Code to be successful with many companies wanting to sign on to it. This is directly linked to the EU's ambition to be a leading AI continent with innovative AI companies. However, the current draft does not meet this competitiveness requirement and, as it stands, would further create complexity, burden and cost for businesses developing and deploying generative AI systems in Europe.

While BSA recognizes and strongly welcomes the significant improvements made from the first draft, we wanted to share some remaining concerns, and related recommendations, which are built around three key pillars:

- Clarify the carve-out for industrial Business-to-Business use-cases;
- Respect strictly the scope of the EU AI Act;
- Focus on the proportionality and technical feasibility of the measures proposed.

Finally, as before, we want to reiterate our insistence on the overall consistency of the current initiatives pertaining to AI.

- The present Code needs to be consistent with the Commission's upcoming Guidelines on Transparency which are being developed in parallel and are expected by June. This would require a dedicated stakeholder consultation.
- There also must be consistency between the Code, the aforementioned Guidelines and the ongoing AI Digital Omnibus proposal, as some of the proposed measures relate to Article 50.
- Finally, the European Commission should also provide consistent guidance on transparency requirements for agentic AI and voice assistants, in particular to stress that that AI agents should only be in scope of Art 50(1) obligations, i.e. no marking requirement but only the disclosure obligation, and how this would work in practice.

[1] *BSA's members include: Adobe, Akamai, Alteryx, Amadeus, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Cohesity, Dassault Systemes, Databricks, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.*

44 Avenue des Arts Brussels 1040 Belgium
+32 (0)2 274 13 10 **www.bsa.org**

Business Software Alliance

***

General. Clarify carve-out for industrial and B2B use cases

Given the ultimate purpose of Article 50 transparency obligations, it is neither necessary nor proportionate for B2B and industrial use cases and should therefore be clearly excluded from the scope of AI Transparency Code of Practice.
Article 50 is grounded in the need to address information asymmetries and risks of deception vis-à-vis natural persons, in particular consumers and citizens whose fundamental rights may be directly affected. In professional, B2B and industrial contexts, AI systems are deployed between informed economic operators, under contractual arrangements, sector-specific regulation and existing product safety and liability regimes. These use-cases do not present the same risks of manipulation or opacity towards end-users that Article 50 is designed to mitigate.
Extending horizontal transparency obligations to such contexts would disregard the risk-based and proportionate logic of the AI Act, create legal uncertainty, and undermine legitimate interests such as the protection of trade secrets and confidential business information, without delivering corresponding fundamental rights benefits.

We therefore recommend clarifying that the Code does not apply to B2B, corporate internal and industrial use cases.

1. **Machine-readable marking techniques** (Section 1, Measure 1.1)

The measures outlined across Section 1, Commitment 1 must ensure they strictly adhere to the scope (language and intent) of Article 50(2) of the AI Act. In particular, obligations should only be imposed where "technically feasible," while giving due regard to the different roles of different companies along the AI value chain, the specific limitations of different content types, implementation costs, and the current state of the art.

Moreover, the measures and requirements suggested in Section 1 still seem to create indirect obligations for model providers and/or third-parties in the business of developing marking techniques. The draft should not make compliance of downstream providers directly conditional on implementation of rules by third-party model providers or the help of third-party marking providers.

We therefore ask for explicit exclusion of model-level obligations from the Code.

We have **identified the following concerns and related recommendations:**

- **Section 1, introduction**: the section starts with the mention “so long as no single marking approach is sufficient”. This approach is too rigid and fails to recognise current technological limitations in implementing marking techniques, specifically for text modality. We would welcome a more flexible approach and avoid wording implying that no single marking approach is sufficient to meet the requirements of Art 50(2).

- **Section 1, Sub-measure 1.1.1. (Digitally signed metadata):** we would welcome clarity as to what this "digitally signed manifest" covers in practice, namely whether it encompasses, for example, audit trails and logging capabilities.
    - The "provenance certificate" concept for free text is technically ambiguous, and we would welcome clarity especially in circumstances where the text is modified even trivially, after download. We would recommend to clarify this measure to cover circumstances in which AI-generated or AI-manipulated text is subsequently reviewed or modified by a human, particularly in low-risk B2B contexts. In such cases, preserving a provenance certificate may be misleading, as it would no longer accurately reflect the final output, since there was a human intervention. If a human substantively modifies the AI content but the AI-generated tag remains, the human editor is effectively stripped of their professional agency, as the tag becomes a false positive for AI generation, undermining the transparency goal by devaluing the human-in-the-loop contribution.
    - The Code of Practice treats "text" as a single modality. Software code is technically text, but it has unique properties that make many of the proposed marking techniques infeasible. Generated code gets auto-formatted by linters and prettifiers, refactored by developers, compiled or transpiled into different languages (TypeScript to JavaScript, Sass to CSS), minified for production, and merged with human-written code in version control. These steps strip metadata and defeat watermarks. We request to explicitly exclude software code from the scope of "AI-generated content" covered by the Code.

- **Section 1, Sub-measure 1.1.3 (Fingerprinting or logging facilities):** The proposed measure regarding fingerprinting and logging lacks clarity concerning its scope and application, risking misinterpretation as a default obligation rather than a supplementary one. Given the significant scalability challenges within complex value chains and the proportionality concerns raised by universal databases, these tools should remain optional mechanisms utilized primarily for internal accountability. As with the sub-measures discussed above, this should also be revised to apply specifically to signatories in the role of providers of AI systems, rather than providers of AI models.
- Indicate "optional supplementary measures" to include logging mechanisms. For internally deployed AI systems (e.g. systems not interacting with external end users), we think logging should meet the 50(2) obligations.

- **Section 1, Sub-measure 1.2 (Non-removal of machine readable techniques):**
    - We would recommend to clarify this measure to cover circumstances in which AI-generated or AI-manipulated content is subsequently reviewed or modified by a human, particularly in low-risk B2B contexts. In such cases, preserving a mark as "AI-generated" or "AI-manipulated" may be misleading, as it would no longer accurately reflect the final output, since there was a human intervention. If a human substantively modifies the AI content but the AI-generated tag remains, the human editor is effectively stripped of their professional agency, as the tag becomes a false positive for AI generation, undermining the transparency goal by devaluing the human-in-the-loop contribution.

      The CoP should therefore clarify that marks do not need to be preserved where subsequent human editorial or expert review has substantively shaped the final content. This is relevant, for instance, for AI-generated text subject to editorial review, or AI-generated image masks used in radiotherapy treatment planning and subsequently reviewed or modified by a human expert.

2. <u>Section 1, Commitment 2: Detection of the marking of AI-generated content</u>

- **Section 1, Measure 2.1 (Measure 2.1: Detection mechanisms for active marking made available to deployers, end-users and other third parties)**
  - This measure should be reconsidered, as it goes beyond the scope of AI Act Art. 50(2). Art. 50(2) focuses on detectability through marking and does not mandate that each provider develops and maintains dedicated detection infrastructure. For this reason, it would impose significant additional burdens, specifically for B2B and internal use cases where such requirements may not be relevant in practice. Requiring signatories to make detection mechanisms available free of charge to a broad range of third parties risks extending the obligations under the AI Act beyond its transparency objective and risks creating a roadmap for adversarial actors to guess and check their way around watermarks until they find a method that bypasses detection, thereby weakening the ecosystem's security.
    The CoP should therefore clarify that detection mechanisms are voluntary, rather than a required element of compliance under this measure. This is particularly important where content is generated or used in closed professional contexts and is not disseminated to the wider public.
  - "Interface": We would welcome clarity as to what the notion entails and, specifically, whether allowing a place for deployers to include AI disclosures to their end users would qualify.
  - Local hosting of the detection tools: We would like clarity as to the rationale of this requirement and, in particular, whether it can be subject to lawful cross-border data transfer mechanisms. We believe this requirement introduces an unjustified data localisation requirement that is not grounded in the AI Act's legal text.
  - "If the detection mechanism requires uploading content . . . Signatories will not retain a verbatim copy of the content or personally identifiable information about the person requesting verification": This requirement conflicts with the fingerprinting option in Measure 1.1.3. If detection uses fingerprinting, it must retain a derivative of the content to function. As this section explicitly indicates that fingerprinting or logging is not excluded from this requirement, we would welcome clarity on how this requirement may be achieved.

- **Section 1, Measure 2.3 (Clear and accessible disclosure of verification and detection results)**
  - Mandating forensic detection mechanisms that function independently of marking is technically unfeasible and exceeds the scope of the AI Act by implying provider liability for downstream misuse beyond their control.
  - Furthermore, the inherent inaccuracy of such tools creates a high risk of false positives and negatives, potentially leading to harmful consequences like the wrongful rejection of legitimate documents (e.g., flagging CVs or cover letters as AI-generated).
  - This measure is not necessarily relevant in the context of markings (AI origin is labelled and does not need to be explained), and not at all for low-risk B2B systems.
  - Moreover, it should be clarified that the measure does not apply where the outputs concerned, or the AI system or product itself, are not intended for laypersons. Similarly, references to the European Accessibility Act (EAA) and the Web Accessibility Directive should apply only where the relevant system, product or interface falls within the scope of those instruments.
  - The CoP should provide concrete examples for disclosing evidence in a manner that does not compromise trade secrets, undermine system security, or facilitate adversarial attacks on detection mechanisms.

- **Section 1, Measure 2.4 (Human-understandable and accessible disclosure of verification and detection results):** The requirement to provide "human-understandable explanations" for detection outcomes is largely redundant for marked content where origin is explicit, and is particularly irrelevant for low-risk B2B systems.
  - While potentially applicable to unmarked content detection, we maintain strong reservations regarding the reliability of forensic methods in this context (see Measure 2.3).
  - Furthermore, any guidance on explanations must carefully balance transparency with security, ensuring that disclosures do not compromise trade secrets, system integrity, or facilitate adversarial circumvention.

3. Section 1, Commitment 3: Measures to meet the requirement for marking and detection techniques

- **Section 1, General** Detection obligations must strictly adhere to a risk-based approach, ensuring that low-risk B2B applications, which face significantly fewer adversarial threats, are not subject to the same standards as AI systems for the general public. Moreover, such detection mechanisms should be voluntary.
- **Section 1, Measure 3.3 (Robustness**): Measure 3.3 must be revised to explicitly acknowledge that absolute error-free performance is unattainable, and to incorporate qualifying language such as "where technically feasible" and "state-of-the-art" to reflect realistic technical limitations.
  Furthermore, robustness obligations must strictly adhere to a risk-based approach, ensuring that low-risk B2B applications, which face significantly fewer adversarial threats, are not subject to the same standards as AI systems for the general public.
- **Section 1, Measure 3.4 (Interoperability):** This measure creates an obligation for signatories to implement technical solutions for marking and detecting AI-generated content that work "across distribution channels and technological environments, regardless of the application domain or context." This must be revised to recognize that interoperability may not always be possible as it would require all parties along the value chain to come together and agree on common solutions that would work across their tech stacks.
  While the mention of "to the extent technically feasible is welcomed, ensuring the cooperation and consensus of every entity along the distribution channel is a challenge separate from the mere technical feasibility.
  - The AI Office could leverage existing international and European standardisation work to support the development of interoperable solutions and standards, without requiring a level of interoperability that is not practical.

4. Section 1, Commitment 4: Testing, verification and compliance

This commitment must also strictly adhere to a risk-based approach, ensuring that low-risk B2B applications, which face significantly fewer adversarial threats, are not subject to the same standards as AI systems for the general public. Moreover, we would welcome guidance from the AI Office in the mater.

5. Avoid mandating prescriptive marking icons

We understand that the goal of the Code is to develop common approach to marking AI-generated content, which will be easily perceptible by users. However, we advise against mandating a standardized icon, at least before conducting extensive research. Existing studies have shown that the same AI label can result in different user perceptions depending on the country and context. For example, one study found that labels such as “AI generated” and “Artificial” elicited negative feelings towards the content and creator in US audiences but positive in Mexico and Brazil. Another study found that AI labels lead to different levels of engagement for political vs. entertainment content. Therefore, we recommend conducting further study as well as considering domain-specific examples or variations of how the icon could be applied.

44 Avenue des Arts Brussels 1040 Belgium
+32 (0)2 274 13 10 **www.bsa.org**