BSA's Input
to
the European Commission's Public Consultation
on
the Draft Implementing Regulation on detailed arrangements on evaluations and proceedings
for
the Artificial Intelligence Act

April 9, 2026

BSA welcomes the opportunity to comment on the European Commission's public consultation on the Draft Implementing Regulation on detailed arrangements on evaluations and proceedings with regards the Artificial Intelligence Act.

BSA is the leading advocate for the global software industry and the voice of Business-to-Business (B2B) global software companies[1] that are leaders in artificial intelligence, cybersecurity, cloud computing, and other cutting-edge technologies. We work in over 20 markets in the US, Europe, and Asia, advocating for policies that build trust in technology so that every industry sector and the public can benefit from innovation.

We are raising the **following concerns** to the draft:

- **Scope of Access Powers (Article 2(1) of the Draft)**
  Article 92(3) of the AI Act allows the Commission to "request access to the general-purpose AI model concerned through APIs or further appropriate technical means and tools, including source code." This is deliberately open-ended, but it is anchored to a specific purpose: evaluating the model to assess compliance or investigate systemic risks under Article 92(1).
  The Draft implementing act goes considerably further. Article 2(1) specifies that access "may include" a non-exhaustive list that extends well beyond the AI Act's language: API access, internal access, source code, model weights, hosting infrastructure, the ability to "inspect and modify system state interaction with the model," and "all levels of access granted to employees of the provider." Several of these measures merit closer examination, both in terms of their legal basis under Article 92 and their practical implications for model providers.
  - **Mandatory access to the source code** (Article 2 (1) of the Draft). This measure does not present added value for such access, making it an overreaching measure compared to its actual benefits.
    BSA recommends removing that element as it would trigger concerns with regards competition as it could lead to the divulgation of proprietary information to competitors.
  - **Model weights.** Granting third party access to model weights raises acute concerns across security, safety, and intellectual-property dimensions.
    - Emerging policy and regulatory frameworks increasingly treat frontier model weights as sensitive assets that developers are responsible for securing, controlling, and

[1] BSA's members include: Adobe, Akamai, Alteryx, Amadeus, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Cohesity, Dassault Systèmes, 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.

monitoring across the entire model lifecycle. Mandating weight access for regulatory investigations directly contradicts the strict custody and access controls that regulators otherwise require developers to maintain. Compelling such access also sets a precedent for similar demands from jurisdictions with weaker security, rule-of-law, or adversarial interests.

- Model weights are core proprietary assets protected under the EU Trade Secrets Directive and represent high-value targets for state-level actors and organized crime.[2]
- This approach also goes beyond existing frameworks: the GPAI Code of Practice does not require signatories to provide evaluators with weight access, instead permitting API access, on-premise testing, or secure hardware environments (Measure 3.5). The implementing act should remain consistent with this standard.
- Weight access grants capabilities far exceeding what regulatory assessment requires, including the ability to fine-tune and repurpose models. The AI Office's objectives can be fully achieved through controlled API access with audit trails, joint testing protocols, or secure execution environments that preserve developer custody.

BSA recommends that Article 2(1) be revised to remove all references to third-party access to model weights, recognizing that weight access exceeds what is necessary for regulatory assessment, contradicts developers' existing custody obligations, and creates security and precedent risks that outweigh any investigatory benefit.

- **"Modify system state."** Article 92 is an evaluation power, not an intervention power. The inclusion of "modify" goes beyond read-only assessment and could be interpreted as permitting the Commission or its appointed experts to alter the operational behavior of a live model.
  BSA recommends revising Article 2(1), to replace "*modify* system state*"* with "*inspect* system state."

- **Providing third-party evaluators and regulators with all levels of access granted to employees.** This wording is too open-ended and could potentially encompass any type of internal access. There is no necessity test linking the scope of access to the evaluation's stated purpose. This sits in direct tension with Article 92(4) of the AI Act, which requires each access request to state its purpose and reasons.
  BSA recommends Article 2(1) be revised to remove all references to "all levels of access granted to employees of the provider."

- **Removal of any logging measures.** The final sentence of Article 2(1) empowers the Commission to require providers to "disable and remove any logging measures" that could track the Commission's access. This has no textual basis in Article 92. Its practical effect is to eliminate the provider's ability to verify that access (e.g., to source code, model weights, production infrastructure) was used only for the stated evaluation purpose.
  BSA recommends removing this provision entirely. If the concern is that provider-side logging could compromise evaluation integrity, the appropriate solution is a jointly agreed logging protocol that protects both the confidentiality of the evaluation and the provider's legitimate security interests.

---

[2] *RAND's 2024 report on securing AI model weights[2] catalogued 38 distinct attack vectors, noting that frontier model weights are high-value targets for state-level actors and organized crime.* Nevo, S., Lahav, D., Karpur, A., Bar-On, Y., Bradley, H. A., & Alstott, J. (2024). *Securing AI Model Weights: Preventing Theft and Misuse of Frontier Models*. RAND Corporation. https://www.rand.org/pubs/research_reports/RRA2849-1.html

- **Infrastructure Access, Third-Party Hosting, and Open-Source Models (Article 2 (1) of the Draft)**
  - Article 2(1) includes "access to the infrastructure used for hosting the general-purpose AI model." In practice, a single GPAI model may be distributed across multiple platforms at once, leaving it unclear what "access to the infrastructure" means. Many GPAI models are not hosted on infrastructure the provider owns or exclusively controls. GPAI models are routinely placed on the market through third-party cloud platforms, hosted on model repositories, or deployed via API gateway services operated by downstream providers.In such instances, the model provider may lack contractual or legal authority to grant the Commission access to third-party infrastructure it does not own.
    The problem becomes particularly acute for open-source and open-weight models. These are hosted and deployed by a diffuse ecosystem of independent parties. A model released under an open licence may be running on thousands of deployments across the EU, with the original provider no longer hosting it at all. When the implementing act refers to "the infrastructure used for hosting the general-purpose AI model," it is unclear whether it refers to the provider's or the downstream deployers' infrastructure. The implementing act draws no distinction between closed and open-source distribution models. Article 92(3) contemplates access to "the general-purpose AI model concerned," not to third-party infrastructure.
  - BSA recommends to clarify that infrastructure access under Article 2(1) is limited to infrastructure owned or controlled by the provider. For open-source models, the access framework should reflect the reality that the provider may not control where the model is hosted. Where access to third-party infrastructure is genuinely necessary for an evaluation, the implementing act should establish a separate mechanism for the Commission to engage directly with the hosting provider, rather than routing the obligation through the model provider.

- **Introduce a new paragraph in Article 2 to require a structured dialogue prior to access requests, and ensuring any access request remains limited to what is *necessary and proportionate for the stated evaluation purpose***
  Article 92(7) provides for a preliminary exchange whereby the AI Office may engage with the provider of the AI model to understand the provider's internal testing practices and internal safeguards for managing systemic risks, before proceeding to formal access requests to obtain additional information. This step serves an important function: it allows the AI Office to determine whether existing provider processes and documentation already address its concerns, potentially rendering more intrusive access unnecessary. However, this step is not reflected in the current wording of Article 2 of the present draft.
  BSA recommends:
  - larifying that any requested access should be *limited to what is necessary and proportionate for the stated evaluation purpose*, consistent with the purpose-specificity requirement in Article 92(4).
  - Adding a new paragraph Article 92(3) reflecting the concerns outlined above.
    Article 92(3): "Prior to adopting a decision requesting access to a general-purpose AI model pursuant to Article 92(3) of Regulation (EU) 2024/1689, the Commission shall engage in structured dialogue with the provider pursuant to Article 92(7) to determine the level of access necessary and proportionate to achieve the evaluation objectives".

- **No oral hearing or hearing officer (Article 7 of the Draft).**
  - The right to be heard is limited to written submissions, with no oral hearing and no independent hearing officer. Both are standard in competition law enforcement. For

disputes involving model architecture, training methodology, and systemic risk assessments, written submissions alone may not be sufficient for a meaningful defence.

- BSA recommends the introduction of a right to request an oral hearing and establishing an independent procedural arbiter, consistent with the framework under Regulation 773/2004 relating to the conduct of proceedings by the Commission pursuant to Articles 81 and 82 of the EC Treaty.

***

For further information, please contact:
Hadrien Valembois, Director, Policy – EMEA hadrienv@bsa.org or
+32.474.72.34.59