Business Software Alliance

BSA

Additional Recommendations on the Draft Guidelines on Classification of High-Risk Systems
To Respect the AI Act's Risk-Based Approach

November 2026

The AI Act is founded on a risk-based approach. The Draft Guidelines on Classification of High-Risk Systems abandon this approach. Instead, the Draft Guidelines treat any AI system as high risk if other entities in the AI supply chain combine that system with high-risk tools. It also treats general-purpose AI systems as high-risk if any potential downstream high-risk usage is not prevented. These fundamental changes are at odds with the Act and should not be made through secondary guidance. Significant revisions are needed to avoid gutting the Act's risk-based approach.

Our **concerns** are structured around **three pillars**:

1. **The concept of intended purpose**

2. **Horizontal issues pertaining to Annex III**

3. **Employment**

***

** This paper follows the paper sent to the EU AI Office in September and the related meeting held by BSA and member companies' representatives on September 8 at the EU AI Office premises.*

***

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

Business Software Alliance

I. THE CONCEPT OF INTENDED PURPOSE

1) General concerns

Paragraphs 10-14 of the Draft Guidelines on 'intended purpose' require further clarification for Business-to-business (B2B) providers, where it is the business customers, not the AI providers, that determine the actual deployment context.

First, the Draft Guidelines should **clarify whether the provider's documented primary intended purpose**, as reflected in product documentation and contractual terms, controls classification, or whether all technically possible uses must be assumed. Treating general-purpose solutions as high-risk by default would create unworkable compliance burdens and deter innovation.

Second, the Draft Guidelines should **allow providers of General-Purpose AI (GPAI) systems to restrict deployers from engaging in high-risk use cases without having to prescribe every possible purpose and application of the GPAI system**. The Draft Guidelines imply that a GPAI system must have a fixed set of intended purposes, which seems self-contradictory and does not align with existing industry practices and deployer expectations.

Finally, the Draft Guidelines should **clarify that a provider's Acceptable Use Policy (AUP), which contractually restricts deployment in high-risk scenarios, narrows the scope of 'reasonably foreseeable uses' for classification purposes**. For example, where an AUP explicitly prohibits use of an AI agent for hiring decisions without mandatory human review, that restriction should inform the 'intended purposes' analysis. Without this clarification, providers operating through intermediaries cannot determine their obligations with sufficient legal certainty, undermining effective implementation and compliance.

2) Concern over the Responsibilities along the AI value chain (Article 25 of the EU AI Act)

Article 25 of the EU AI Act (responsibilities along the AI value chain) provides the mechanism for shifting responsibility when a system's risk profile changes after it leaves the original provider. The Act recognizes that the third party that makes the substantial modification or changes the intended purpose should be solely responsible, as provider, for the resulting high-risk system.

The Draft Guidelines' treatment of the original provider as responsible for that system, or for the modified portion of it, does not follow the allocation of responsibility currently established by Article 25.

The wording provided in the Draft Commission Guidelines on high-risk AI systems may lead to a scenario where a provider that develops a generic, General-Purpose AI (GPAI) system- including AI systems whose purpose is so broad as to encompass a multitude of uses- could become subject to high-risk obligations if a downstream company integrates that system into a high-risk application. This would effectively reclassify the upstream provider as a high-risk provider on the basis of decisions taken by third parties further down the supply chain, over which it has no meaningful control.

In practice, this risks implying that providers of GPAI systems within the meaning of the AI Act will routinely be treated as providers of high-risk AI systems. This is because GPAI systems - including AI systems whose purpose is so broad as to encompass a multitude of uses- functionality is broad and can be used in a variety of use cases, including high-risk – although this is not its primary intended purpose. The practical effect would be to collapse the distinction the AI Act draws between GPAI models and systems on the one hand and high-risk AI systems on the other. We assume this is not the intended outcome and ask the Commission to clarify the Guidelines accordingly.

Examples (illustrative only; classification in each case depends on the intended purpose declared by the provider)

System	High-risk by virtue of its own intended purpose?	Could become high-risk when used for...
Planning Analytics (forecasting, planning)	No	Creditworthiness assessments, eligibility for public benefits, employment decisions, insurance risk assessments, or other Annex III use cases, where a downstream party adopts such an intended purpose.
Enterprise data analytics	No	Where a downstream party intends the system to support a high-risk use case such as worker evaluation, recruitment, access to essential services, or critical infrastructure operations.
AI assistants	No	Where a downstream party deploys the assistant to perform Annex III functions, for example screening job candidates, evaluating employees, determining eligibility for benefits, or supporting judicial decisions.

These systems are not inherently high-risk under the EU AI Act. Classification depends on the intended purpose and deployment context, particularly whether the system is placed on the market or put into service for an Annex III high-risk use case.

Recommendations:

- **Paragraph 12**: We recommend using the term "explicitly" over "effectively", to provide legal clarity, since multi-purpose/GPAI systems can unintentionally "provide for" a high-risk use. It is also disproportionate to require limitations across all material; one clear reference should suffice.
- **Paragraph 75–76**: A non-high-risk or GPAI component supplied by a third party should not become high-risk merely by being part of a combined high-risk system; only the system as a whole is high-risk. "Contribute to a high-risk purpose" is not an appropriate threshold.

Ultimately, we ask the Commission to align paragraphs 12 and 75–76 with Article 25, so that the party that makes the substantial modification or determines the high-risk intended purpose is

identified as the provider of the resulting high-risk system, with the original provider's obligations limited to those set out in Article 25(2) and (4).

Business Software Alliance

II. HORIZONTAL ISSUES PERTAINING TO ANNEX III

AI Systems as part of complex systems/products and services (Section 2.3, Paragraph 75)

- Concern: The Draft Guidelines do not adequately account for AI systems that execute multi-step workflows, where each step may involve decision-making.
- Recommendation: The evaluation should focus on the final consequential decision instead of evaluating independently each step against Article 6 requirements. “Combined intended purpose” test should only apply to systems offered by the same Provider, while a deployer should solely be responsible for the combined intended purpose of all integrated AI systems.

“On behalf of” (Section 2.5, paragraph 79)

- Concern: The Draft Guidelines lack clear criteria for distinguishing when a private entity acts "on behalf of" a public authority versus independently. Without clarification, private AI providers risk inadvertent high-risk reclassification.
- Recommendation: The Draft Guidelines should define objective criteria, such as formal delegation, statutory mandate, or direct public authority oversight, for determining "on behalf of" status and confirm that a provider merely holding a government contract does not trigger this classification.

Conditions for the ‘filter’ to apply (Section 2.7.1., paragraph 103)

- Concern: The lack of clear guidance on when an AI system is considered preparatory.
- Recommendation: the need to establish objective thresholds for “material influence” to avoid a subjective or overbroad standard. For example, an objective threshold should exclude standard tools like summarizing materials and workflow steps that prepare data for a human-reviewed decision. It should also confirm that AI tools for which outputs require substantive human interpretation do not meet the threshold for “material influence.”

Protection against circumvention (Section 2.7.4., paragraph 117):

- Concern: Because the filter mechanism is both novel and complex, it is inevitable that there will be a learning period for companies in understanding its application.
- Recommendation: The Draft Guidelines establish a transition period during which good-faith misapplication of the filter period is not treated as circumvention. In addition, the Draft Guidelines should identify clear criteria that distinguishes deliberate circumvention from reasonable, good faith interpretative disagreement.

Education and Vocational Training (Section 3.3.2., paragraphs 217-219)

- Concern: the current scope of high-risk AI systems in the access/admissions subsection to AI systems that are being used for substantive decision-making in the admissions or assignment process is too vague. Indeed, the current draft employs vague examples, e.g. systems that “assist” in admissions decisions.
- Recommendation: The scope should be limited to reflect that AI systems that help admissions staff with procedural tasks like understanding applicant data without contributing to the admission outcome itself do not pose a significant risk and should therefore not fall under the high-risk classification.

III. EMPLOYMENT

AI Systems Intended to be Used for the Recruitment or Selection of Natural Persons – Notion of Recruitment (Section 3.4.2.i, paragraph 246))

- Concern: The Draft Guidelines' approach to employment and recruitment is too broad
- Recommendation: (1) limiting the definition of "recruitment" to AI systems that directly support selection, screening, or assessment decisions affecting individual candidates, and (2) confirm that tools used for preliminary or administrative activities, like drafting and distributing job postings or interview scheduling are not high risk and do not require registration.

AI Systems Intended to be Used for the Recruitment or Selection of Natural Persons – Evaluate Candidates (Paragraph 3.4.2.v, paragraph 254)

- Concern: Most AI-assisted recruitment tools necessarily involve processing data about candidates to support human decision-making, this provision effectively removes the filter mechanism from the entire recruitment category. The compliance burden of assuming the high-risk obligations is disproportionate for AI tools that play only an auxiliary role in recruitment workflows. The blanket exclusion of profiling-involved systems from the Article 6(3) filter does not differentiate between (a) AI systems that themselves generate predictive profiles or rankings of candidates (for which the exclusion is appropriate) and (b) AI tools that process candidate data incidentally in performing a preparatory or administrative function (for which the exclusion is disproportionate).
- Recommendation: Consistent with our recommendation in Section 2.7.1. (paragraph 103) regarding the filter exemption and objective "material influence" thresholds, the Draft Guidelines should clarify that the profiling bar applies only where the AI system's primary function is to generate a profile, prediction, or evaluation of a natural person's characteristics or behavior (such as their suitability or performance at work), not where such evaluation is incidental to an otherwise preparatory function (e.g. merely summarizing or organizing candidate input data).

AI Systems Intended to be Used to Manage Work-Related Relationships (Paragraph 3.4.3)

- Concern: The scope of "workforce" for employment use cases is too broad
- Recommendation: The scope should be limited to relationships involving meaningful power asymmetry. The Guidelines should confirm that AI tools used in B2B commercial contracting for professional services — where both parties are acting in a commercial capacity, at arm's length, with negotiated terms — fall outside the scope of Point 4(a) of Annex III. The expansion to sole traders and freelancers is appropriate only where the economic power dynamics resemble employment.

***

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