ACM ETPC response to targeted EC consultation: Draft Commission Guidelines on the Classification of High-Risk AI Systems

June 2026
Authors: Gaston Besanson, Paolo Giudici, Francisco Medeiros, Ahmed Nagy, Alejandro Saucedo, and Neil Yorke-Smith. Reviewer: Gerhard Schimpf. Editor: Tom Romanoff

Executive Summary

The Association for Computing Machinery (ACM) is the world's longest-established professional society of individuals involved in all aspects of computing. It annually bestows the ACM A.M. Turing Award, often popularly referred to as the "Nobel Prize of Computing." ACM's Europe Technology Policy Committee ("Europe TPC") is charged with and committed to providing sound **technical information** to policymakers and the general public in the service of sound public policymaking. Europe TPC has previously responded to European Union stakeholder consultations in the context of the AI Act[1], the Data Act[2], the Digital Services Act[3][4], the Digital Citizen Principles[5], the Cyber Resilience Act[6], amongst others[7]. ACM and Europe TPC are non-profit, non-political, and non-lobbying organisations.

This submission provides architectural and operational feedback regarding the draft guidelines published on May 19, 2026. EuropeTPC fully supports the foundation of the AI Act to ensure safe and trustworthy AI. However, the interpretive boundaries proposed within these draft Guidelines risk misclassifying foundational data infrastructure and stifling the adoption of secure, composable architectures.

1. General Principles: End-to-End Assessment and *Modular AI Architectures*

The Commission's Draft Position:

The guidelines establish that multi-component *or composite* systems must be assessed "end-to-end." The regulatory intent is to prevent the artificial splitting of high-risk workflows into unclassified micro-components to evade compliance.

The Technical Reality & Enterprise Friction:

Modern enterprise systems are aggressively moving away from monolithic models toward

[1] https://www.acm.org/binaries/content/assets/public-policy/europe-tpc-comments-ai-consultation.pdf
[2] https://www.acm.org/binaries/content/assets/public-policy/acm-eur-tpc-data-act-comments-13may22a.pdf
[3] https://www.acm.org/binaries/content/assets/public-policy/europetpc-digital-services-act-comments.pdf
[4] https://www.acm.org/binaries/content/assets/public-policy/acm-europe-tpc-dsa-comments.pdf
[5] https://www.acm.org/binaries/content/assets/public-policy/europetpc-comments-digital-principles.pdf
[6] https://www.acm.org/binaries/content/assets/public-policy/acm-europe-tpc-cyber-reslience-comments-pdf
[7] https://www.acm.org/public-policy/public-policy-statements

composable, *modular architectures*. In these setups, a *primary orchestrating model* routes tasks to highly specialised, narrow-scope *micro-models or services*

Under the current draft's broad "end-to-end" mandate, a benign, narrowly scoped *sub-component* inherits the high-risk classification of the broader workflow simply by being part of the topology. This creates an untenable regulatory burden and directly penalises robust architectural design, forcing enterprises to build risk-heavy monoliths rather than modular, auditable microservices.

Proposed Amendment to the Guidelines:

We urge the Commission to explicitly introduce a **"Bounded Sub-Component Exemption"** within the General Principles.

- **Suggested Textual Addition:** *"Where an AI system utilises a multi-component or modular architecture, the high-risk classification and its associated obligations shall apply to the primary orchestrating system or the final-mile decision node. Sub-components, micro-models, or discrete services that (a) lack independent decision-making authority, (b) are deterministic in their output, and (c) operate strictly within technical guardrails defined by the orchestrator, shall not independently inherit the high-risk classification, provided overall system accountability is maintained by the provider of the orchestrating system."*

2. Annex I: Harmonisation Frameworks and the MLOps Lifecycle

The Commission's Draft Position:

The guidelines clarify the application of the AI Act to systems that act as safety components under existing EU harmonisation legislation (e.g., the Machinery Directive and the Medical Devices Directive), confirming the 2 August 2028 compliance deadline.

The Technical Reality & Enterprise Friction:

The draft fails to reconcile the static nature of traditional product safety certifications with the dynamic reality of enterprise Machine Learning Operations (MLOps). Enterprise AI relies on Continuous Integration and Continuous Deployment (CI/CD) pipelines. To prevent model drift and maintain accuracy, models require automated data refresh cycles (data hydration) and periodic weight updates.

The guidelines leave critical ambiguity regarding the definition of a "substantial modification." If routine MLOps pipeline updates are treated as substantial modifications, enterprises will be forced into continuous, paralysing reassessments of third-party conformity, effectively halting the deployment of critical safety and operational updates.

Proposed Amendment to the Guidelines:

The guidelines must define the threshold for substantial modifications in the context of automated machine learning lifecycles.

- **Suggested Textual Addition:** *"For systems operating under Annex I, routine operational updates executed via governed MLOps pipelines, including periodic model retraining, minor weight adjustments, and data hydration cycles, shall not constitute a 'substantial modification' triggering a new conformity assessment, provided that such updates (a) do not alter the system's originally intended purpose, and (b) remain strictly within the mathematical safety and performance envelope established during the initial conformity assessment."*

3. Annex III: Standalone Systems, Data Foundations, and the Art. 6(3) Filter

The Commission's Draft Position:

Annex III details the eight high-risk areas in exhaustive detail. Crucially, it applies a rigid interpretation of the AI Act Article 6(3) filter conditions, specifically the exemption for systems that perform a "narrow procedural task." The draft implies that if a system materially influences a human's final decision, the exemption is voided.

The Technical Reality & Enterprise Friction:

This interpretation threatens to misclassify the entirety of enterprise data foundation programmes. A robust data foundation utilises AI upstream to ingest raw data, perform complex entity resolution, and hydrate data pipelines. This structured data is then surfaced to a human operator (e.g., in HR or credit contexts) who makes the actual qualitative judgment.
If the Commission equates "providing structured, high-quality data" with "materially influencing a decision," then foundational data engineering itself becomes a high-risk activity. This conflates data structuring with decision recommendation, placing an impossible compliance burden on backend enterprise data architecture.

Proposed Amendment to the Guidelines:

1. Draw a definitive technical demarcation between "Decision Recommendation" and "Upstream Data Structuring".
The final guidelines must draw a definitive technical demarcation between "Decision Recommendation" and "Upstream Data Structuring". To this aim, real use cases should be provided.

- **Suggested Textual Addition:** *"The exemption for 'narrow procedural tasks' under AI Act Article 6(3) shall explicitly include AI systems utilised for foundational data operations, including but not limited to data quality enforcement, semantic mapping, entity resolution, and pipeline hydration. If an AI system operates strictly to structure and presents factual data without providing a qualitative assessment, scoring, or automated recommendation regarding the specific high-risk use case, and provided that the rules governing data selection, sorting, and presentation are transparent, objective, and auditable to prevent hidden anchoring or framing biases, it shall be deemed to not materially replace or influence the human assessment and is therefore exempt from the high-risk classification."*

2. Clarify Intended Purpose for General Purpose and Configurable AI Systems
High risk classification should not apply solely because a system is technically capable of use in

an Annex III context. Classification should depend on the provider's documented intended purpose, deployment constraints, contractual restrictions, technical safeguards, and reasonably foreseeable use.

- **Suggested text Addition:**
  "Where an AI system is technically capable of use in a high-risk context but is not marketed, configured, documented, or supplied for such use, and where the provider implements meaningful technical and organisational restrictions, capability alone should not trigger high-risk classification. Where product documentation, examples, sales materials, APIs, templates, or deployment support facilitate Annex III use cases, the intended purpose should be considered to include those use cases."

3. Establish an Evidence-Based Article 6(3) Assessment Framework
The Article 6(3) filter would benefit from a common evidence framework to support consistent interpretation across sectors and Member States. This section could benefit from incorporating real-world use cases as references for effective implementation.

- **Suggested text Addition:**
  "Providers relying on Article 6(3) should document the intended purpose, relevant Annex III use case, applicable exemption condition, reasons why the system does not materially influence decision making, reasons why it does not perform profiling, the level of human review, the nature of the system output, and safeguards preventing progression from procedural support to decision recommendation."

4. Clarify Meaningful Human Review
The Guidelines should further clarify the distinction between meaningful human review and formal human involvement. The determining factor should be whether the human retains genuine authority and independent judgment.

- **Suggested text Addition:**
  "Human review should be considered meaningful only where the reviewer has sufficient authority, competence, time, and access to underlying evidence to challenge, reject, or modify the system output. Human review should not be considered meaningful where AI outputs are routinely accepted without substantive independent assessment."

Strengthen Procurement Transparency
Public authorities and regulated organisations need practical mechanisms to assess provider classification decisions.

- **Suggested text Addition:**
  "Providers should supply a classification statement indicating whether the system is high risk under Article 6(1), high risk under Article 6(2), exempt under Article 6(3), or outside the scope of Article 6. Where Article 6(3) is applied, supporting documentation should be made available to facilitate due diligence and oversight."

5. Create a Review Mechanism for Modular Systems

Modular and multi-component AI architectures are evolving rapidly. Periodic updates would improve legal certainty and consistent implementation.

- **Suggested text Addition:**
  "The Commission should periodically review and update guidance relating to multi-component and modular AI systems, drawing on input from technical experts, standardisation bodies, civil society, deployers, and market surveillance authorities."

References

European Commission. (2026). Draft Commission guidelines on the classification of high-risk AI systems under Article 6 of Regulation (EU) 2024/1689.

European Commission. (2026). Targeted consultation on the draft guidelines for the classification of high-risk artificial intelligence systems.

Artificial Intelligence Act. European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence.

European Commission. (2022). The Blue Guide on the implementation of EU product rules 2022.